Introduction to the ASAv

The Adaptive Security Virtual Appliance (ASAv) brings full firewall functionality to virtualized environments to secure data center traffic and multitenant environments.

You can manage and monitor the ASAv using ASDM or CLI. Other management options may be available.

Hypervisor Support

For hypervisor support, see Cisco ASA Compatibility.

Licensing for the ASAv

The ASAv uses Cisco Smart Software Licensing. For complete information, see Smart Software Licensing.

Note You must install a smart license on the ASAv. Until you install a license, throughput is limited to 100 Kbps so you can perform preliminary connectivity tests. A smart license is required for regular operation.

Beginning with 9.13(1), any ASAv license can be used on any supported ASAv vCPU/memory configuration. This allows you to deploy an ASAv on a wide variety of VM resource footprints. Session limits for AnyConnect Client and TLS Proxy are determined by the ASAv platform entitlement installed rather than a platform limit tied to a model type.

See the following sections for information about ASAv licensing entitlements and resource specifications for the supported private and public deployment targets.

About Smart License Entitlements

Any ASAv license can be used on any supported ASAv vCPU/memory configuration. This allows you to run the ASAv on a wide variety of VM resource footprints. This also increases the number of supported AWS and Azure instances types. When configuring the ASAv machine, the maximum supported number of vCPUs is 8; and the maximum supported memory is 64GB for ASA virtual deployed on all platforms other than AWS and OCI. For ASA virtual deployed on AWS and OCI, the maximum supported memory is 128GB.

Important It is not possible to change the resource allocation (memory, CPUs, disk space) of an ASAv instance once it is deployed. If you need to increase your resource allocations for any reason, for example to change your licensed entitlement from the ASAv30/2Gbps to the ASAv50/10Gbps, you need to create a new instance with the necessary resources.

- vCPUs—The ASAv supports 1 to 8 vCPUs.
- Memory—The ASAv supports 2GB to 64GB of RAM for ASA virtual deployed on all platforms other than AWS and OCI. For ASA virtual deployed on AWS and OCI, the maximum supported memory is 128GB.
- Disk storage—The ASAv supports a minimum virtual disk of 8GB by default. Depending on the type of platform, the virtual disk support varies between 8GB to 10GB. Keep this in mind when you provision your VM resources.

Important The minimum memory requirement for the ASAv is 2 GB. If your current ASAv runs with less than 2 GB of memory, you cannot upgrade to version 9.13(1) or greater from an earlier version without increasing the memory of your ASAv machine. You can also redeploy a new ASAv machine with the latest version.

The minimum memory requirement for deploying ASAv with more than 1 vCPU is 4 GB.

For upgrading from ASAv version 9.14 and later to a latest version, the ASA virtual machine requires a minimum memory of 4 GB and 2 vCPU.

Session Limits for Licensed Features

Session limits for AnyConnect Client and TLS Proxy are determined by the installed ASAv platform entitlement tier, and enforced via a rate limiter. The following table summarizes the session limits based on the entitlement tier and rate limiter.

Table 1: ASAv Session Limits by Entitlement

Entitlement	AnyConnect Client Premium Peers	Total TLS Proxy Sessions	Rate Limiter
Standard Tier, 100M	50	500	150 Mbps
Standard Tier, 1G	250	500	1 Gbps
Standard Tier, 2G	750	1000	2 Gbps
Standard Tier, 10G	10,000	10,000	10 Gbps

The session limits granted by an entitlement, as shown in the previous table, cannot exceed the session limits for the platform. The platform session limits are based on the amount of memory provisioned for the ASAv.

Table 2: ASAv Session Limits by Memory Requirement

Provisioned Memory	AnyConnect Client Premium Peers	Total TLS Proxy Sessions
2 GB to 7.9 GB	250	500
8 GB to 15.9 GB	750	1000
16 GB - 64 GB	10,000	10,000
64 GB to 128 GB	20,000	20,000

Platform Limits

Firewall connections, concurrent and VLANs are platform limits based on the ASAv memory.

Note We limit the firewall connections to 100 when the ASAv is in an unlicensed state. Once licensed with any entitlement, the connections go to the platform limit. The minimum memory requirement for the ASAv is 2GB.

Table 3: Platform Limits

ASAv Memory	Firewall Conns, Concurrent	VLANs
2 GB to 7.9 GB	100,000	50
8 GB to 15.9 GB	500,000	200
16 GB to 64	2,000,000	1024

ASAv Private Cloud Entitlements (VMware, KVM, Hyper-V)

Because any ASAv license can be used on any supported ASAv vCPU/memory configuration, you have greater flexibility when you deploy the ASAv in a private cloud environment (VMware, KVM, Hyper-V).

Session limits for AnyConnect Client and TLS Proxy are determined by the installed ASAv platform entitlement tier, and enforced via a rate limiter. The following table summarizes the session limits based on the entitlement tier for the ASAv deployed to a private cloud environment, with the enforced rate limiter.

Note ASAv session limits are based on the amount of memory provisioned for the ASAv; see Table 2: ASAv Session Limits by Memory Requirement, on page 3.

Table 4: ASAv on VMware/KVM/HyperV Private Cloud - Licensed Feature Limits Based on Entitlement

RAM (GB)		Entitlement Support*			
Min	Max	Standard Tier, 100M	Standard Tier, 1G	Standard Tier, 2G	Standard Tier, 10G
2	7.9	50/500/100M	250/500/1G	250/500/2G	250/500/10G
8	15.9	50/500/100M	250/500/1G	750/1000/2G	750/1000/10G
16	64	50/500/100M	250/500/1G	750/1000/2G	10K/10K/10G
*AnyConnect Client Sessions / TLS Proxy Sessions / Rate Limiter per entitlement/instance.					

ASAv Public Cloud Entitlements (AWS)

Because any ASAv license can be used on any supported ASAv vCPU/memory configuration, you can deploy the ASAv on a wide variety AWS instances types. Session limits for AnyConnect Client and TLS Proxy are determined by the installed ASAv platform entitlement tier, and enforced via a rate limiter.

The following table summarizes the session limits and rate limiter based on the entitlement tier for AWS instance types. See "About ASAv Deployment On the AWS Cloud" for a breakdown of the AWS VM dimensions (vCPUs and memory) for the supported instances.

Table 5: ASAv on AWS - Licensed Feature Limits Based on Entitlement

Instance	BYOL Entitlement Support*				PAYG**
	Standard Tier, 100M	Standard Tier, 1G	Standard Tier, 2G	Standard Tier, 10G	
c5.xlarge	50/500/100M	250/500/1G	750/1000/2G	750/1000/10G	750/1000
c5.2xlarge	50/500/100M	250/500/1G	750/1000/2G	10K/10K/10G	10K/10K
c4.large	50/500/100M	250/500/1G	250/500/2G	250/500/10G	250/500
c4.xlarge	50/500/100M	250/500/1G	250/500/2G	250/500/10G	250/500
c4.2xlarge	50/500/100M	250/500/1G	750/1000/2G	10K/10K/10G	750/1000
c3.large	50/500/100M	250/500/1G	250/500/2G	250/500/10G	250/500
c3.xlarge	50/500/100M	250/500/1G	250/500/2G	250/500/10G	250/500
c3.2xlarge	50/500/100M	250/500/1G	750/1000/2G	10K/10K/10G	750/1000
m4.large	50/500/100M	250/500/1G	250/500/2G	250/500/10G	250/500
m4.xlarge	50/500/100M	250/500/1G	250/500/2G	250/500/10G	10K/10K
m4.2xlarge	50/500/100M	250/500/1G	750/1000/2G	10K/10K/10G	10K/10K
*AnyConnect Client Sessions / TLS Proxy Sessions / Rate Limiter per entitlement/instance. **AnyConnect Client Sessions / TLS Proxy Sessions. The Rate Limiter is not employed in PAYG mode.					

Pay-As-You-Go (PAYG) Mode

The following table summarizes the Smart Licensing entitlements for each tier for the hourly billing (PAYG) mode, which is based on the allocated memory.

Table 6: ASAv on AWS - Smart License Entitlements for PAYG

RAM (GB)	Hourly Billing Mode Entitlement
< 2 GB	Standard Tier, 100M (ASAv5)
2 GB to < 8 GB	Standard Tier, 1G (ASAv10)
8 GB to < 16 GB	Standard Tier, 2G (ASAv30)
16 GB < 32 GB	Standard Tier, 10G (ASAv50)
30 GB and higher	Standard Tier, 20G (ASAv100)

ASAv Public Cloud Entitlements (Azure)

Because any ASAv license can be used on any supported ASAv vCPU/memory configuration, you can deploy the ASAv on a wide variety Azure instances types. Session limits for AnyConnect Client and TLS Proxy are determined by the installed ASAv platform entitlement tier, and enforced via a rate limiter.

The following table summarizes the session limits and rate limiter based on the entitlement tier for the Azure instance types. See "About ASAv Deployment On the Microsoft Azure Cloud" for a breakdown of the Azure VM dimensions (vCPUs and memory) for the supported instances.

Note Pay-As-You-Go (PAYG) Mode is currently not supported for the ASAv on Azure.

Table 7: ASAv on Azure - Licensed Feature Limits Based on Entitlement

Instance	BYOL Entitlement Support*				
	Standard Tier, 100M	Standard Tier, 1G	Standard Tier, 2G	Standard Tier, 10G	Standard Tier, 20G
D1, D1_v2DS1, DS1_v2	50/500/100M	250/500/1G	250/500/2G	250/500/10G	250/500/20G
D2, D2_v2, DS2, DS2_v2	50/500/100M	250/500/1G	250/500/2G	250/500/10G	250/500/20G
D3, D3_v2, DS3, DS3_v2	50/500/100M	250/500/1G	750/1000/2G	750/1000/10G	750/1000/20G
D4, D4_v2, DS4, DS4_v2	50/500/100M	250/500/1G	750/1000/2G	10K/10K/10G	10K/10K/20G
D2_v3	50/500/100M	250/500/1G	750/1000/2G	750/1000/10G	750/1000/20G
D4_v3	50/500/100M	250/500/1G	750/1000/2G	10K/10K/10G	10K/10K/20G

Instance	BYOL Entitlement Support*				
	Standard Tier, 100M	Standard Tier, 1G	Standard Tier, 2G	Standard Tier, 10G	Standard Tier, 20G
D8_v3	50/500/100M	250/500/1G	750/1000/2G	10K/10K/10G	10K/10K/20G
F4, F4s	50/500/100M	250/500/1G	750/1000/2G	750/1000/10G	750/1000/20G
F8, F8s	50/500/100M	250/500/1G	750/1000/2G	10K/10K/10G	10K/20K/20G
F16, F16s	50/500/100M	250/500/1G	750/1000/2G	10K/10K/10G	10K/20K/20G
*AnyConnect Client Sessions / TLS Proxy Sessions / Rate Limiter per entitlement/instance.					

Guidelines and Limitations

The ASAv firewall functionality is very similar to the ASA hardware firewalls, but with the following guidelines and limitations.

Guidelines and Limitations for the ASAv (all entitlements)

Smart Licensing Guidelines

- The maximum supported number of vCPUs is 16. The maximum supported memory is 64GB for ASA virtual deployed on all platforms other than AWS and OCI. For ASA virtual deployed on AWS and OCI, the maximum supported memory is 128GB. Any ASAv license can be used on any supported ASAv vCPU/memory configuration.
- Session limits for licensed features and unlicensed platform capabilities are set based on the amount of VM memory.
- Session limits for AnyConnect Client and TLS Proxy are determined by the ASAv platform entitlement; session limits are no longer associated with an ASAv model type (ASAv5/10/30/50).
- Session limits have a minimum memory requirement; in cases where the VM memory is below the minimum requirement, the session limits will be set for the maximum number supported by the amount of memory.
- There are no changes to existing entitlements; the entitlement SKU and display name will continue to include the model number (ASAv5/10/30/50).
- The entitlement sets the maximum throughput via a rate limiter.
- There is no change to customer ordering process.

Disk Storage

The ASAv supports a maximum virtual disk of 8 GB by default. You cannot increase the disk size beyond 8 GB. Keep this in mind when you provision your VM resources.

Context Mode Guidelines

Supported in single context mode only. Does not support multiple context mode.

Failover for High Availability Guidelines

For failover deployments, make sure that the standby unit has the same license entitlement; for example, both units should have the 2Gbps entitlement.

Important When creating a high availability pair using ASAv, it is necessary to add the data interfaces to each ASAv in the same order. If the exact same interfaces are added to each ASAv, but in different order, errors may be presented at the ASAv console. Failover functionality may also be affected.

Unsupported ASA Features

The ASAv does not support the following ASA features:

- Clustering (for all entitlements, except KVM and VMware)
- Multiple context mode
- Active/Active failover
- EtherChannels
- Shared AnyConnect Premium Licenses

Limitations

- The ASAv is not compatible with the 1.9.5 i40en host driver for the x710 NIC. Older or newer driver versions will work. (VMware only)

Guidelines and Limitations for the 1 GB Entitlement

Performance Guidelines

- Jumbo frame reservation on the 1 GB platform with 9 or more configured e1000 interfaces may cause the device to reload. If **jumbo-frame reservation** is enabled, reduce the number of interfaces to 8 or less. The exact number of interfaces will depend on how much memory is needed for the operation of other features configured, and could be less than 8.

Guidelines and Limitations for the 10 GB Entitlement

Performance Guidelines

- Supports 10Gbps of aggregated traffic.
- Supports the following practices to improve ASAv performance:
  - Numa nodes

  - Multiple RX queues
  - SR-IOV provisioning
  - See Performance Tuning and Performance Tuning for more information.
- CPU pinning is recommended to achieve full throughput rates; see Increasing Performance on ESXi Configurations and Increasing Performance on KVM Configurations.
- Jumbo frame reservation with a mix of e1000 and i40e-vf interfaces may cause the i40e-vf interfaces to remain down. If **jumbo-frame reservation** is enabled, do not mix interface types that use e1000 and i40e-vf drivers.

Limitations

- Transparent mode is not supported.
- The ASAv is not compatible with the 1.9.5 i40en host driver for the x710 NIC. Older or newer driver versions will work. (VMware only)
- Not supported on Hyper-V.

ASAv Interfaces and Virtual NICs

As a guest on a virtualized platform, the ASAv uses the network interfaces of the underlying physical platform. Each ASAv interface maps to a virtual NIC (vNIC).

- ASAv Interfaces
- Supported vNICs

ASAv Interfaces

The ASAv includes the following Gigabit Ethernet interfaces:

- Management 0/0

  For AWS and Azure, Management 0/0 can be a traffic-carrying “outside” interface.
- GigabitEthernet 0/0 through 0/8. Note that the GigabitEthernet 0/8 is used for the failover link when you deploy the ASAv as part of a failover pair.

Note To simply configuration migration, Ten GigabitEthernet interfaces, like those available on the VMXNET3 driver, are labeled GigabitEthernet. This has no impact on the actual interface speed and is cosmetic only.

The ASAv defines GigabitEthernet interfaces using the E1000 driver as 1Gbps links. Note that VMware no longer recommends using the E1000 driver.

- Hyper-V supports up to eight interfaces. Management 0/0 and GigabitEthernet 0/0 through 0/6. You can use GigabitEthernet 0/6 as a failover link.

Supported vNICs

The ASAv supports the following vNICs. Mixing vNICs, such as e1000 and vmxnet3, on the same ASAv is not supported.

Table 8: Supported vNics

vNIC Type	Hypervisor Support		ASAv Version	Notes
	VMware	KVM		
vmxnet3	Yes	No	9.9(2) and later	VMware default When using vmxnet3, you need to disable Large Receive Offload (LRO) to avoid poor TCP performance. See Disable LRO for VMware and VMXNET3, on page 9.
e1000	Yes	Yes	9.2(1) and later	Not recommended by VMware.
virtio	No	Yes	9.3(2.200) and later	KVM default
ixgbe-vf	Yes	Yes	9.8(1) and later	AWS default; ESXi and KVM for SR-IOV support.
i40e-vf	No	Yes	9.10(1) and later	KVM for SR-IOV support.

Disable LRO for VMware and VMXNET3

Large Receive Offload (LRO) is a technique for increasing inbound throughput of high-bandwidth network connections by reducing CPU overhead. It works by aggregating multiple incoming packets from a single stream into a larger buffer before they are passed higher up the networking stack, thus reducing the number of packets that have to be processed. However, LRO can lead to TCP perfomance problems where network packet delivery may not flow consistently and could be "bursty" in congested networks.

Important VMware enables LRO by default to increase overall throughput. It is therefore a requirement to disable LRO for ASAv deployments on this platform.

You can disable LRO directly on the ASAv machine. Power off the virtual machine before you make any configuration changes.

1. Find the ASAv machine in the vSphere Web Client inventory.
   a. To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
   b. Click the **Related Objects** tab and click **Virtual Machines**.

2. Right-click the virtual machine and select **Edit Settings**.
3. Click **VM Options**.
4. Expand **Advanced**.
5. Under Configuration Parameters, click the **Edit Configuration** button.
6. Click **Add Parameter** and enter a name and value for the LRO parameters:
   - Net.VmxnetSwLROSL | 0
   - Net.Vmxnet3SwLRO | 0
   - Net.Vmxnet3HwLRO | 0
   - Net.Vmxnet2SwLRO | 0
   - Net.Vmxnet2HwLRO | 0

Note Optionally, if the LRO parameters exist, you can examine the values and change them if needed. If a parameter is equal to 1, LRO is enabled. If equal to 0, LRO is disabled.

7. Click **OK** to save your changes and exit the **Configuration Parameters** dialog box.
8. Click **Save**.

See the following VMware support articles for more information:

- VMware KB 1027511
- VMware KB 2055140

ASAv and SR-IOV Interface Provisioning

Single Root I/O Virtualization (SR-IOV) allows multiple VMs running a variety of guest operating systems to share a single PCIe network adapter within a host server. SR-IOV allows a VM to move data directly to and from the network adapter, bypassing the hypervisor for increased network throughput and lower server CPU burden. Recent x86 server processors include chipset enhancements, such as Intel VT-d technology, that facilitate direct memory transfers and other operations required by SR-IOV.

The SR-IOV specification defines two device types:

- Physical Function (PF)—Essentially a static NIC, a PF is a full PCIe device that includes SR-IOV capabilities. PFs are discovered, managed, and configured as normal PCIe devices. A single PF can provide management and configuration for a set of virtual functions (VFs).
- Virtual Function (VF)—Similar to a dynamic vNIC, a VF is a full or lightweight virtual PCIe device that provides at least the necessary resources for data movements. A VF is not managed directly but is derived from and managed through a PF. One or more VFs can be assigned to a VM.

SR-IOV is defined and maintained by the Peripheral Component Interconnect Special Interest Group (PCI SIG), an industry organization that is chartered to develop and manage the PCI standard. For more information about SR-IOV, see PCI-SIG SR-IOV Primer: An Introduction to SR-IOV Technology.

Provisioning SR-IOV interfaces on the ASAv requires some planning, which starts with the appropriate operating system level, hardware and CPU, adapter types, and adapter settings.

Guidelines and Limitations for SR-IOV Interfaces

The specific hardware used for ASAv deployment can vary, depending on size and usage requirements. Licensing for the ASAv, on page 1 explains the compliant resource scenarios that match license entitlement for the different ASAv platforms. In addition, SR-IOV Virtual Functions require specific system resources.

Host Operating System and Hypervisor Support

SR-IOV support and VF drivers are available for:

- Linux 2.6.30 kernel or later

The ASAv with SR-IOV interfaces is currently supported on the following hypervisors:

- VMware vSphere/ESXi
- QEMU/KVM
- AWS

Hardware Platform Support

Note You should deploy the ASAv on any *server class* x86 CPU device capable of running the supported virtualization platforms.

This section describes hardware guidelines for SR-IOV interfaces. Although these are guidelines and not requirements, using hardware that does not meet these guidelines may result in functionality problems or poor performance.

A server that supports SR-IOV and that is equipped with an SR-IOV-capable PCIe adapter is required. You must be aware of the following hardware considerations:

- The capabilities of SR-IOV NICs, including the number of VFs available, differ across vendors and devices.
- Not all PCIe slots support SR-IOV.
- SR-IOV-capable PCIe slots may have different capabilities.

Note You should consult your manufacturer's documentation for SR-IOV support on your system.

- For VT-d enabled chipsets, motherboards, and CPUs, you can find information from this page of virtualization-capable IOMMU supporting hardware. VT-d is a required BIOS setting for SR-IOV systems.
- For VMware, you can search their online Compatibility Guide for SR-IOV support.
- For KVM, you can verify CPU compatibility. Note that for the ASAv on KVM we only support x86 hardware.

Note We tested the ASAv with the Cisco UCS C-Series Rack Server. Note that the Cisco UCS-B server does not support the ixgbe-vf vNIC.

Supported NICs for SR-IOV

- Intel Ethernet Network Adapter X710

Attention The ASAv is not compatible with the 1.9.5 i40en host driver for the x710 NIC. Older or newer driver versions will work. (VMware only)

- Intel Ethernet Server Adapter X520 - DA2

CPUs

- x86_64 multicore CPU

  Intel Sandy Bridge or later (Recommended)

Note We tested the ASAv on Intel's Broadwell CPU (E5-2699-v4) at 2.3GHz.

- Cores
  - Minimum of 8 physical cores per CPU socket
  - The 8 cores must be on a single socket.

Note CPU pinning is recommended to achieve full throughput rates on the ASAv50 and ASAv100; see Increasing Performance on ESXi Configurations and Increasing Performance on KVM Configurations.

BIOS Settings

SR-IOV requires support in the BIOS as well as in the operating system instance or hypervisor that is running on the hardware. Check your system BIOS for the following settings:

- SR-IOV is enabled
- VT-x (Virtualization Technology) is enabled
- VT-d is enabled
- (Optional) Hyperthreading is disabled

We recommend that you verify the process with the vendor documentation because different systems have different methods to access and change BIOS settings.

Limitations

Be aware of the following limitations when using ixgbe-vf interfaces:

- The guest VM is not allowed to set the VF to promiscuous mode. Because of this, transparent mode is not supported when using ixgbe-vf.
- The guest VM is not allowed to set the MAC address on the VF. Because of this, the MAC address is not transferred during HA like it is done on other ASA platforms and with other interface types. HA failover works by transferring the IP address from active to standby.

Note This limitation is applicable to the i40e-vf interfaces too.

- The Cisco UCS-B server does not support the ixgbe-vf vNIC.
- In a failover setup, when a paired ASAv (primary unit) fails, the standby ASAv unit takes over as the primary unit role and its interface IP address is updated with a new MAC address of the standby ASAv unit. Thereafter, the ASAv sends a gratuitous Address Resolution Protocol (ARP) update to announce the change in MAC address of the interface IP address to other devices on the same network. However, due to incompatibility with these types of interfaces, the gratuitous ARP update is not sent to the global IP address that is defined in the NAT or PAT statements for translating the interface IP address to global IP addresses.