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, on page 1
- Licensing for the ASAv, on page 1
- Licensing for the ASAv, on page 1
- Guidelines and Limitations, on page 4
- ASAv Interfaces and Virtual NICs, on page 7
- ASAv and SR-IOV Interface Provisioning, on page 9

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.

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

Licensing for the ASAv

See the following tables for information about ASAv licensing entitlements, licensing states, requried resources, and model specifications:

- Table 1: ASAv Smart License Entitlements—Shows the compliant resources scenarios that match license entitlement for the ASAv platform.

Note The ASAv uses Cisco Smart Software Licensing. A smart license is required for regular operation. Until you install a license, throughput is limited to 100 Kbps so you can perform preliminary connectivity tests.

- Table 2: ASAv Licensing States—Shows the ASAv states and messages connected to resources and entitlement for the ASAvs.
- Table 3: ASAv Model Descriptions and Specifications—Shows the ASAv models and associated specifications, resource requirements, and limitations.

Smart License Entitlements

The ASAv uses Cisco Smart Software Licensing. For detailed information, see Smart Software Licensing for the ASAv and ASA.

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.

Table 1: ASAv Smart License Entitlements

License Entitlement	vCPU/RAM	Throughput	Rate Limiter Enforced
Lab Edition Mode (no license)	All Platforms	100Kbps	Yes
ASAv5 (100M)	1vCPU/1 GB to 1.5 GB (2 GB recommended)	100Mbps	Yes
ASAv10 (1 GB)	1vCPU/2 GB	1Gbps	Yes
ASAv30 (2 GB)	4vCPU/8 GB	2Gbps	Yes
ASAv50 (10 GB)	8vCPU/16 GB	10Gbps	Yes

Licensing States

Table 2: ASAv Licensing States

State	Resources vs. Entitlement	Actions and Messages
Compliant	Resource = Entitlement limits (vCPU, GB of RAM)	Appliances optimally resourced ASAv5 (1vCPU,1G), ASAv10 (1vCPU,2G), ASAv30 (4vCPU,8G), ASAv50 (8vCPU, 16G) No actions, no messages
	Resources < Entitlement limits Under-provisioned	No actions while Warning messages are logged that ASAv cannot run at licensed throughput.
Non-compliant	Resources > Entitlement limits Over-provisioned	ASAv rate limiter engages to limit performance and log Warnings on the console.
		ASAv10, ASAv30, ASAv50 reboot after logging Error messages on the console.

Model Descriptions and Specifications

Table 3: ASAv Model Descriptions and Specifications

Model	License Requirement
ASAv5	Smart License See the following specifications: • 100 Mbps throughput • 1 vCPU • 1 GB RAM (adjustable to 1.5 GB) **Note** For optimum performance we recommend 2 GB of memory for the ASAv5. . When deploying ASAv5 with 2 GB RAM for optimal performance, you may encounter a warning message that appears due to the default entitlement of 1.5 GB RAM associated with the ASAv5 licensing. It is advisable to ignore the warning message. • 50,000 concurrent firewall connections • Does not support AWS • Supports Azure on a Standard D3 and Standard D3_v2 instances

Model	License Requirement
ASAv10	Smart License See the following specifications: • 1 Gbps throughput • 1 vCPU • 2 GB RAM • 100,000 concurrent firewall connections • Supports AWS on c3.large, c4.large, and m4.large instances • Supports Azure on a Standard D3 and Standard D3_v2 instances
ASAv30	Smart License See the following specifications: • 2 Gbps throughput • 4 vCPUs • 8 GB RAM • 500,000 concurrent firewall connections • Supports AWS on c3.xlarge, c4.xlarge, and m4.xlarge instances • Supports Azure on a Standard D3 and Standard D3_v2 instances
ASAv50	Smart License See the following specifications: • 10 Gbps throughput • 8 vCPUs (minimum of 8 physical cores per CPU socket required; cannot be provisioned across multiple CPU sockets) • 16 GB RAM • 2,000,000 concurrent firewall connections • Does not support AWS, Microsoft Azure, or Hyper-V

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 models)

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 model license; for example, both units should be ASAv30s.

Important

- When creating a High Availability (HA) 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.
- HA can be configured between two ASA virtual instances even if there is a resource mismatch (Example: One instance with 8GB RAM and another with 16GB RAM). This configuration is supported to facilitate hitless upgrades. However, it is not recommended running HA with resource disparities beyond the necessary duration until the resource allocation changes is complete.

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 ASAv5

Performance Guidelines

- Supports 8000 connections per second, 25 maximum VLANs, 50,000 concurrent session, and 50 VPN sessions.

- The ASAv5 is intended for users who require a small memory footprint and small throughput, so that you can deploy larger numbers of ASAv5s without using unnecessary memory.
- Beginning with 9.5(1.200), the memory requirement for the AVAv5 was reduced to 1GB. Downgrading the available memory on an ASAv5 from 2 GB to 1 GB is not supported. To run with 1 GB of memory, the ASAv5 VM must be redeployed with version 9.5(1.200) or later. Similarly, if you try to downgrade to a version earlier than 9.5(1.200), you must increase the memory to 2 GB.

Note For optimum performance we recommend 2 GB of memory for the ASAv5.

- In some situations, the ASAv5 may experience memory exhaustion. This can occur during certain resource heavy applications, such as enabling AnyConnect or downloading files.
  - Console messages related to spontaneous reboots or critical syslogs related to memory usage are symptoms of memory exhaustion.
  - In these cases, you can enable the ASAv5 to be deployed in a VM with 1.5 GB of memory. To change from 1GB to 1.5 GB, power down your VM, modify the memory, and power the VM back on.
  - You can display a summary of the maximum memory and current free memory available to the system using the `show memory` command from the CLI.
- The ASAv5 will begin to drop packets soon after the threshold of 100 Mbps is reached (there is some headroom so that you get the full 100 Mbps).

Limitations

- ASAv5 is not compatible with AnyConnect HostScan 4.8, which requires 2 GB of RAM.
- ASAv5 is not supported on Amazon Web Services (AWS).
- Jumbo frames are not supported.

Guidelines and Limitations for the ASAv50

Performance Guidelines

- Supports 10Gbps of aggregated traffic.
- Supported only on ESXi and KVM.
- 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.
-

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 Amazon Web Services (AWS), Microsoft Azure, and Hyper-V.
- The ixgbe vNIC is not supported in this release.

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.

Table 4: Supported vNics

Table 5: Supported vNics

vNIC Type	Hypervisor Support		ASAv Version	Notes
	VMware	KVM		
e1000	Yes	Yes	9.2(1) and later	VMware default
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

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 Server Adapter X520 - DA2
- Intel Ethernet Server Adapter X540

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.