



Testing and Troubleshooting

This chapter describes how to troubleshoot the ASA and test basic connectivity.

- [Recover Enable and Telnet Passwords, on page 1](#)
- [Configure and Run Captures with the Packet Capture Wizard, on page 5](#)
- [CPU Usage and Reporting, on page 10](#)
- [Test Your Configuration, on page 15](#)
- [Monitoring Performance and System Resources, on page 23](#)
- [Monitoring Connections, on page 25](#)
- [History for Testing and Troubleshooting , on page 25](#)

Recover Enable and Telnet Passwords

If you forget the enable or Telnet passwords, you can recover them for ASAv and ISA 3000 models. The procedure differs by device type. You must perform the task using the CLI.



Note For other platforms, you cannot recover lost passwords. You can only restore the factory default configuration, and reset the passwords to the default. For Firepower 4100/9300, see the [FXOS configuration guide](#). For Firepower 1000 and 2100, see the [FXOS troubleshooting guide](#).

Recover Passwords on the ASA 5506-X, ASA 5508-X, ASA 5516-X, and ISA 3000

To recover passwords for the ASA 5506-X, ASA 5508-X, ASA 5516-X, and ISA 3000 perform the following steps:

Procedure

- Step 1** Connect to the ASA console port.
- Step 2** Power off the ASA, then power it on.
- Step 3** After startup, press the **Escape** key when you are prompted to enter ROMMON mode.
- Step 4** To update the configuration register value, enter the following command:

```
rommon #1> confreg 0x41
```

You must reset or power cycle for new config to take effect

The ASA displays the current configuration register value and a list of configuration options. Record the current configuration register value, so you can restore it later.

```
Configuration Register: 0x00000041
```

```
Configuration Summary
[ 0 ] password recovery
[ 1 ] display break prompt
[ 2 ] ignore system configuration
[ 3 ] auto-boot image in disks
[ 4 ] console baud: 9600
boot: ..... auto-boot index 1 image in disks
```

Step 5 Reload the ASA by entering the following command:

```
rommon #2> boot
Launching BootLoader...
Boot configuration file contains 1 entry.

Loading disk0:/asa932-226-k8.bin... Booting...Loading...
```

The ASA loads the default configuration instead of the startup configuration.

Step 6 Access the privileged EXEC mode by entering the following command:

```
ciscoasa# enable
```

Step 7 When prompted for the password, press **Enter**.

The password is blank.

Step 8 Load the startup configuration by entering the following command:

```
ciscoasa# copy startup-config running-config
```

Step 9 Access the global configuration mode by entering the following command:

```
ciscoasa# configure terminal
```

Step 10 Change the passwords, as required, in the default configuration by entering the following commands:

```
ciscoasa(config)# password password
ciscoasa(config)# enable password password
ciscoasa(config)# username name password password
```

Step 11 Load the default configuration by entering the following command:

```
ciscoasa(config)# no config-register
```

The default configuration register value is 0x1. See the [command reference](#) for more information about the configuration register.

Step 12 Save the new passwords to the startup configuration by entering the following command:

```
ciscoasa(config)# copy running-config startup-config
```

Recover Passwords or Images on the ASAv

To recover passwords or images on the ASAv, perform the following steps:

Procedure

Step 1 Copy the running configuration to a backup file on the ASAv:

```
copy running-config filename
```

Example:

```
ciscoasa# copy running-config backup.cfg
```

Step 2 Restart the ASAv:

```
reload
```

Step 3 From the GNU GRUB menu, press the down arrow, choose the <filename> with no configuration load option, then press **Enter**. The filename is the default boot image filename on the ASAv. The default boot image is never automatically booted through the **fallback** command. Then load the selected boot image.

```
GNU GRUB version 2.0(12)4
bootflash:/asa100123-20-smp-k8.bin
bootflash: /asa100123-20-smp-k8.bin with no configuration load
```

Example:

```
GNU GRUB version 2.0(12)4
bootflash: /asa100123-20-smp-k8.bin with no configuration load
```

Step 4 Copy the backup configuration file to the running configuration.

```
copy filename running-config
```

Example:

```
ciscoasa (config)# copy backup.cfg running-config
```

Step 5 Reset the password.

enable password *password*

Example:

```
ciscoasa(config)# enable password cisco123
```

Step 6 Save the new configuration.

write memory

Example:

```
ciscoasa(config)# write memory
```

Disable Password Recovery for ASA or ISA 3000 Hardware



Note You cannot disable password recovery on the ASAv, Firepower models.

To disable password recovery to ensure that unauthorized users cannot use the password recovery mechanism to compromise the ASA, perform the following steps.

Before you begin

On the ASA, the **no service password-recovery** command prevents you from entering ROMMON mode with the configuration intact. When you enter ROMMON mode, the ASA prompts you to erase all Flash file systems. You cannot enter ROMMON mode without first performing this erasure. If you choose not to erase the Flash file system, the ASA reloads. Because password recovery depends on using ROMMON mode and maintaining the existing configuration, this erasure prevents you from recovering a password. However, disabling password recovery prevents unauthorized users from viewing the configuration or inserting different passwords. In this case, to restore the system to an operating state, load a new image and a backup configuration file, if available.

The **service password-recovery** command appears in the configuration file for information only. When you enter the command at the CLI prompt, the setting is saved in NVRAM. The only way to change the setting is to enter the command at the CLI prompt. Loading a new configuration with a different version of the command does not change the setting. If you disable password recovery when the ASA is configured to ignore the startup configuration at startup (in preparation for password recovery), then the ASA changes the setting to load the startup configuration as usual. If you use failover, and the standby unit is configured to ignore the startup configuration, then the same change is made to the configuration register when the **no service password-recovery** command replicates to the standby unit.

Procedure

Disable password recovery.

no service password-recovery

Example:

```
ciscoasa (config)# no service password-recovery
```

Configure and Run Captures with the Packet Capture Wizard

You can use the Packet Capture Wizard to configure and run captures for troubleshooting errors. The captures can use ACLs to limit the type of traffic captured, the source and destination addresses and ports, and one or more interfaces. The wizard runs one capture on each of the ingress and egress interfaces. You can save the captures on your PC to examine them in a packet analyzer.



Note This tool does not support clientless SSL VPN capture.

To configure and run captures, perform the following steps:

Procedure

- Step 1** Choose **Wizards > Packet Capture Wizard**.
- The **Overview of Packet Capture** screen appears, with a list of the tasks through which the wizard will guide you to complete. Those tasks include the following:
- Selecting an ingress interface.
 - Selecting an egress interface.
 - Setting the buffer parameters.
 - Running the captures.
 - Saving the captures to your PC (optional).
- Step 2** Click **Next**.
- In a clustering environment, the **Cluster Option** screen appears. Go to Step 3.
- In a non-clustering environment, the **Ingress Traffic Selector** screen appears. Go to Step 4.
- Step 3** Choose one of the following options in the **Cluster Option** screen for running a capture: **This device only** or **The whole cluster**, then click **Next** to display the **Ingress Selector** screen.
- Step 4** Click the **Select Interface** radio button to capture packets on an interface.
- In a clustering environment, to capture only the cluster control plane packets, select the **CP-Cluster** check box.
- Step 5** Click the **Use backplane channel** radio button to capture packets on the ASA CX dataplane.

- Step 6** Do one of the following in the **Packet Match Criteria** area:
- Click the **Specify access-list** radio button to specify the ACL to use for matching packets, then choose the ACL from the **Select ACL** drop-down list. Click **Manage** to display the **ACL Manager** pane to add a previously configured ACL to the current drop-down list. Choose an ACL, then click **OK**.
 - Click the **Specify Packet Parameters** radio button to specify packets parameters.
- a) Do one of the following in the **ICMP Capture** drop-down list:
- Note** The **ICMP Capture** field is populated only when you select **The whole cluster** as the cluster option in the previous window.
- Select **include-decrypt** to capture decrypted IPsec packets which contain both normal and decrypted traffic once they enter the firewall device.
 - Select **persist** to capture persistent packets on cluster units.
- Step 7** To continue, see [Ingress Traffic Selector, on page 8](#).
- Step 8** Click **Next** to display the **Egress Traffic Selector** screen.
- Step 9** Click the **Select Interface** radio button to capture packets on an interface.
- In a clustering environment, to capture the cluster control plane packets, select the **CP-Cluster** check box.
- Note** To know more details on the Egress Traffic Selector fields, see [Egress Traffic Selector, on page 9](#).
- To know more details on the Egress Traffic Selector fields, see [Egress Traffic Selector, on page 9](#).
- Step 10** Click **Next** to display the **Buffers & Captures** screen. To continue, see [Buffers](#).
- Step 11** Check the **Get capture every 10 seconds** check box in the **Capture Parameters** area to obtain the latest capture every 10 seconds automatically. By default, this capture uses the circular buffer.
- Step 12** You specify the buffer size and packet size in the **Buffer Parameters** area. The buffer size is the maximum amount of memory that the capture can use to store packets. The packet size is the longest packet that the capture can hold. We recommend that you use the longest packet size to capture as much information as possible.
- a) Enter the packet size. The valid size ranges from 14 - 1522 bytes.
 - b) Enter the buffer size. The valid size ranges from 1534 - 33554432 bytes.
 - c) Check the **Use circular buffer** check box to store captured packets.
- Note** When you choose this setting, if all the buffer storage is used, the capture starts overwriting the oldest packets.
- Step 13** Click **Next** to display the **Summary** screen, which shows the cluster options for all units in the cluster (if you are using clustering), traffic selectors, and buffer parameters that you have entered. To continue, see [Summary](#).
- Step 14** Click **Next** to display the **Run Captures** screen, then click **Start** to begin capturing packets. Click **Stop** to end the capture. To continue, see [Run Captures, on page 9](#). If you are using clustering, go to Step 16.
- Step 15** Click **Get Capture Buffer** to determine how much buffer space you have remaining. Click **Clear Buffer on Device** to remove the current content and allow room in the buffer to capture more packets.
- Step 16** In a clustering environment, on the **Run Captures** screen, perform one or more of the following steps:
- Click **Get Cluster Capture Summary** to view a summary of packet capture information for all units in the cluster, followed by packet capture information for each unit.

- Click **Get Capture Buffer** to determine how much buffer space you have remaining in each unit of the cluster. The **Capture Buffer from Device** dialog box appears.
- Click **Clear Capture Buffer** to remove the current content for one or all of the units in a cluster and allow room in the buffer to capture more packets.

- Step 17** Click **Save captures** to display the **Save Capture** dialog box. You have the option of saving either the ingress capture, the egress capture, or both. To continue, see [Save Captures](#).
- Step 18** Click **Save Ingress Capture** to display the **Save capture file** dialog box. Specify the storage location on your PC, then click **Save**.
- Step 19** Click **Launch Network Sniffer Application** to start the packet analysis application specified in **Tools > Preferences** for analyzing the ingress capture.
- Step 20** Click **Save Egress Capture** to display the **Save capture file** dialog box. Specify the storage location on your PC, then click **Save**.
- Step 21** Click **Launch Network Sniffer Application** to start the packet analysis application specified in **Tools > Preferences** for analyzing the egress capture.
- Step 22** Click **Close**, then click **Finish** to exit the wizard.
-

Guidelines for Packet Capture

Context Mode

- You can configure captures on the cluster control link within a context; only the packet that is associated with the context sent in the cluster control link is captured.
- You can only configure one capture for a shared VLAN; if you configure a capture in multiple contexts on the shared VLAN, then only the last capture that was configured is used.
- If you remove the last-configured (active) capture, no captures become active, even if you have previously configured a capture in another context; you must remove the capture and add it again to make it active.
- All traffic that enters the interface to which the capture is attached is captured, including traffic to other contexts on the shared VLAN. Therefore, if you enable a capture in Context A for a VLAN that is also used by Context B, both Context A and Context B ingress traffic are captured.
- For egress traffic, only the traffic of the context with the active capture is captured. The only exception is when you do not enable the ICMP inspection (therefore the ICMP traffic does not have a session in the accelerated path). In this case, both ingress and egress ICMP traffic for all contexts on the shared VLAN is captured.

Additional Guidelines

- If the ASA receives packets with an incorrectly formatted TCP header and drops them because of the *invalid-tcp-hdr-length* ASP drop reason, the **show capture** command output on the interface where those packets are received does not show those packets.
- You can only capture IP traffic; you cannot capture non-IP packets such as ARPs.
- For inline SGT tagged packets, captured packets contain an additional CMD header that your PCAP viewer might not understand.

- Packet captures include packets that the system modifies or injects into the connection due to inspection, NAT, TCP normalization, or other features that adjust the content of a packet.
- The trace of the lifespan of an injected virtual packet in a datapath does not exactly reflect how the datapath handles the physical packets. This difference depends on the software version, configuration, and type of the injected virtual packets. Following are configuration settings that might lead to the disparity:
 - at least 2 NAT statements for the same host exist.
 - forward and reverse flows of a connection having different protocols. For example, forward flow is UDP or TCP, reverse flow is ICMP.
 - ICMP error inspection being enabled.

Ingress Traffic Selector

To configure the ingress interface, source and destination hosts or networks, and the protocol for packet capture, perform the following steps:

Procedure

-
- Step 1** Choose the ingress interface name from the drop-down list.
- Step 2** Enter the ingress source host and network. Click the **Use backplane channel** radio button to capture packets on the ASA CX dataplane.
- Step 3** Enter the ingress destination host and network.
- Step 4** Enter the protocol type to capture. Available protocols are ah, eigrp, esp, gre, icmp, icmp6, igmp, igmp, ip, ipinip, nos, ospf, pcp, pim, snp, tcp, or udp.
- Enter the ICMP type for ICMP only. Available types include all, alternate address, conversion-error, echo, echo-reply, information-reply, information-request, mask-reply, mask-request, mobile-redirect, parameter-problem, redirect, router-advertisement, router-solicitation, source-quench, time-exceeded, timestamp-reply, timestamp-request, traceroute, or unreachable.
 - Specify the source and destination port services for the TCP and UDP protocols only. Available options include the following:
 - Choose **All Services** to include all services.
 - Choose **Service Groups** to include a service group.

To include a specific service, choose one of the following: aol, bgp, chargen, cifx, citrix-ica, ctique, daytime, discard, domain, echo, exec, finger, ftp, ftp-data, gopher, h323, hostname, http, https, ident, imap4, irc, kerberos, klogin, kshell, ldap, ldaps, login, lotusnotes, lpd, netbios-ssn, nntp, pcanypwhere-data, pim-auto-rp, pop2, pop3, pptp, rsh, rtsp, sip, smtp, sqlnet, ssh, sunrpc, tacacs, talk, telnet, uucp, or whois.
- Step 5** To enable packet capture for the Cisco TrustSec service, in the **Security Group Tagging** area, check the **SGT number** check box and enter the security group tag number. Valid security group tag numbers range from 2 to 65519.
-

Egress Traffic Selector

To configure the egress interface, source and destination hosts/networks, and source and destination port services for packet capture, perform the following steps:

Procedure

- Step 1** Click the **Select Interface** radio button to capture packets on an interface. Click the **Use backplane channel** radio button to capture packets on the ASA CX dataplane.
- Step 2** Choose the egress interface name from the drop-down list.
- Step 3** Enter the egress source host and network.
- Step 4** Enter the egress destination host and network.

The protocol type selected during the ingress configuration is already listed.

Buffers

To configure the packet size, buffer size, and use of the circular buffer for packet capture, perform the following steps:

Procedure

- Step 1** Enter the longest packet that the capture can hold. Use the longest size available to capture as much information as possible.
 - Step 2** Enter the maximum amount of memory that the capture can use to store packets.
 - Step 3** Use the circular buffer to store packets. When the circular buffer has used all of the buffer storage, the capture will overwrite the oldest packets first.
-

Summary

The **Summary** screen shows the cluster options (if you are using clustering), traffic selectors, and the buffer parameters for the packet capture selected in the previous wizard screens.

Run Captures

To start and stop the capture session, view the capture buffer, launch a network analyzer application, save packet captures, and clear the buffer, perform the following steps:

Procedure

- Step 1** Click **Start** to begin the packet capture session on a selected interface.

- Step 2** Click **Stop** to stop the packet capture session on a selected interface.
 - Step 3** Click **Get Capture Buffer** to obtain a snapshot of the captured packets on the interface.
 - Step 4** Click **Ingress** to show the capture buffer on the ingress interface.
 - Step 5** Click **Egress** to show the capture buffer on the egress interface.
 - Step 6** Click **Clear Buffer on Device** to clear the buffer on the device.
 - Step 7** Click **Launch Network Sniffer Application** to start the packet analysis application for analyzing the ingress capture or the egress capture specified in **Tools > Preferences**.
 - Step 8** Click **Save Captures** to save the ingress and egress captures in either ASCII or PCAP format.
-

Save Captures

To save the ingress and egress packet captures to ASCII or PCAP file format for further packet analysis, perform the following steps:

Procedure

- Step 1** Click **ASCII** to save the capture buffer in ASCII format.
 - Step 2** Click **PCAP** to save the capture buffer in PCAP format.
 - Step 3** Click **Save ingress capture** to specify a file in which to save the ingress packet capture.
 - Step 4** Click **Save egress capture** to specify a file in which to save the egress packet capture.
-

CPU Usage and Reporting

The CPU Utilization report summarizes the percentage of the CPU used within the time specified. Typically, the core operates at approximately 30 to 40 percent of total CPU capacity during non-peak hours and approximately 60 to 70 percent capacity during peak hours.

vCPU Usage in the ASAv

Use the **show cpu usage** command on the ASAv to display CPU utilization statistics. The ASAv vCPU usage shows the amount of vCPUs used for the data path, control point, and external processes.

The vCPU usage reported by the cloud service provider (such as VMware, Azure, OCI, and so on) includes the ASAv usage as described plus:

- ASAv idle time
- %SYS overhead used for the ASA virtual VM
- Overhead of moving packets between vSwitches, vNICs, and pNICs. This overhead can be quite significant.

CPU Usage Example

The following is an example in which the reported vCPU usage is substantially different:

- ASAv reports: 40%
- DP: 35%
- External Processes: 5%
- vSphere reports: 95%
- ASA (as ASAv 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.

Usage can exceed 100% because the ESXi server can use additional compute resources for overhead on behalf of the ASAv.

VMware CPU Usage Reporting

In vSphere, click the **VM Performance** tab, then click **Advanced** to display the **Chart Options** drop-down list, which shows vCPU usage for each state (%USER, %IDLE, %SYS, and so on) of the VM. This information is useful for understanding VMware's perspective on where CPU resources are being used.

On the ESXi server shell (you access the shell by using SSH to connect to the host), `esxtop` is available. `Esxtop` has a similar look and feel to the Linux `top` command and provides VM state information for vSphere performance, including the following:

- Details on vCPU, memory, and network usage
- vCPU usage for each state of each VM.
- Memory (type M while running) and network (type N while running), as well as statistics and the number of RX drops

ASAv and vCenter Graphs

There are differences in the CPU % numbers between the ASAv and vCenter:

- The vCenter graph numbers are always higher than the ASAv numbers.
- vCenter calls it %CPU usage; the ASAv 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.

vCenter 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 Usage in MHz / number of virtual CPUs x core frequency

When you compare the usage in MHz, both the vCenter and ASAv numbers match. According to the vCenter graph, MHz % CPU usage is calculated as $60 / (2499 \times 1 \text{ vCPU}) = 2.4$

Amazon CloudWatch CPU Usage Reporting

You can view the metrics explorer to monitor resources by their tags and properties. Perform the following steps to view CPU utilization statistics for a specific instance:

Procedure

-
- Step 1** Open the **CloudWatch** console and choose **Metrics** in the navigation pane.
 - Step 2** Select the **EC2** metric namespace and select **Per-instance Metrics** dimension.
 - Step 3** Enter **CPUtilization** in the search field and press Enter. Select the row for the required instance to display a graph for the **CPUtilization** metric for that instance.
- See [Amazon CloudWatch documentation](#) for more information.
-

ASAv and Amazon CloudWatch Graphs

The Amazon CloudWatch graph numbers are higher than the numbers because of the different ways in which CPU usage is calculated on the ASAv and the CloudWatch.

When ASAv is running in poll mode, each CPU runs a loop of lightweight commands instead of entering power saving mode or any other idle state. This improves performance by keeping each core active at all times instead of having to be turned on/off or have their clocks adjusted by the Intel power states.

Inside ASAv, this activity is understood to be idling behaviour, and the CPU usage is correctly calculated. However, on the Amazon CloudWatch, the idle behaviour appears like normal CPU activity because all CPU cycles have instructions to run, and this causes the CloudWatch to display a high CPU usage percentage (85-90%).

Azure CPU Usage Reporting

Perform the following steps to view CPU utilization % across all monitored VMs by using VM Insights from the Azure Monitor:

Procedure

- Step 1** Go to the Azure portal, select **Monitor** and choose **Virtual Machines** in the **Solutions** section.
- Step 2** Select the **Performance** tab to display the **CPU Utilization %** chart. This chart displays the top five machines with the highest average processor utilization.
-

Perform the following steps to view the CPU utilization % chart directly from a specific Azure VM:

Procedure

- Step 1** Go to the Azure portal and select **Virtual Machines**.
- Step 2** From the list of VMs, choose a VM.
- Step 3** In the **Monitoring** section, select **Insights**.
- Step 4** Select the **Performance** tab.
- See [How to chart performance with VM Insights](#) for more information.
-

ASAv and Azure Graphs

There are differences in the CPU % numbers between the ASAv and Azure. The Azure graph numbers are always higher than the ASAv numbers because Azure calculates the CPU % usage as the 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 $\text{Usage in MHz} / \text{number of virtual CPUs} \times \text{core frequency}$.

Azure also rate limits the amount of CPU that is requested by the guest OS. Consider a scenario in which ASAv is reporting CPU usage at 40% and the hypervisor is reporting CPU usage at 90%. Now, if the ASAv needs more processing power, the CPU usage may go above 80% and the hypervisor may then report CPU usage at more than 95%. This leads to the hypervisor throttling the ASAv CPU even though the ASAv is just running a loop of lightweight commands in poll mode exhibiting idling behavior.

Hyper-V CPU Usage Reporting

In addition to viewing CPU, RAM, and disk space configuration information for available Cloud Servers, you can also view disk, I/O, and networking information. Use this information to help you decide which Cloud Server is right for your needs. You can view the available servers through either the command-line nova client or the [Cloud Control Panel](#) interface.

On the command line, run the following command:

```
nova flavor-list
```

All available server configurations are displayed. The list contains the following information:

- ID - The server configuration ID
- Name - The configuration name, labeled by RAM size and performance type
- Memory_MB - The amount of RAM for the configuration
- Disk - The size of the disk in GB (for general purpose Cloud Servers, the size of the system disk)
- Ephemeral - The size of the data disk
- Swap - The size of the swap space
- VCPUs - The number of virtual CPUs associated with the configuration
- RXTX_Factor - The amount of bandwidth, in Mbps, allocated to the PublicNet ports, ServiceNet ports, and isolated networks (cloud networks) attached to a server
- Is_Public - Not used

ASA Virtual and Hyper-V Graphs

There are differences in the CPU % numbers between the ASA Virtual and Hyper-V:

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

Hyper-V 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 Usage in MHz / number of virtual CPUs x core frequency



Note It is recommended to look into ASA Virtual reports to get an accurate CPU usage percentage.

OCI CPU Usage Reporting

You can view the CPU utilization % in OCI by using the compute instance metric: **oci_computeagent**. The CpuUtilization metric displays the activity level from the CPU and is expressed as a percentage of total time. Perform the following steps to view metric charts for a single compute instance:

Procedure

- Step 1** Open the navigation menu and click **Instances** under **Compute**.
- Step 2** Click an instance and click **Metrics** under **Resources**.
- Step 3** Select **oci_computeagent** in the Metric namespace list.
See [Compute Instance Metrics](#) for more information.
-

ASAv and OCI Graphs

The OCI graph numbers are always higher than the ASAv numbers because OCI calculates the CPU % usage as the 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 Usage in MHz / number of virtual CPUs x core frequency.

Test Your Configuration

This section describes how to test connectivity for the single mode ASA or for each security context, how to ping the ASA interfaces, and how to allow hosts on one interface to ping through to hosts on another interface.

Test Basic Connectivity: Pinging Addresses

Ping is a simple command that lets you determine if a particular address is alive and responsive. The following topics explain more about the command and what types of testing you can accomplish with it.

What You Can Test Using Ping

When you ping a device, a packet is sent to the device and the device returns a reply. This process enables network devices to discover, identify, and test each other.

You can use ping to do the following tests:

- Loopback testing of two interfaces—You can initiate a ping from one interface to another on the same ASA, as an external loopback test to verify basic “up” status and operation of each interface.
- Pinging to an ASA—You can ping an interface on another ASA to verify that it is up and responding.
- Pinging through an ASA—You can ping through an intermediate ASA by pinging a device on the other side of the ASA. The packets will pass through two of the intermediate ASA’s interfaces as they go in each direction. This action performs a basic test of the interfaces, operation, and response time of the intermediate unit.

- Pinging to test questionable operation of a network device—You can ping from an ASA interface to a network device that you suspect is functioning incorrectly. If the interface is configured correctly and an echo is not received, there might be problems with the device.
- Pinging to test intermediate communications—You can ping from an ASA interface to a network device that is known to be functioning correctly. If the echo is received, the correct operation of any intermediate devices and physical connectivity is confirmed.

Choosing Between ICMP and TCP Ping

The ASA includes the traditional ping, which sends ICMP Echo Request packets and gets Echo Reply packets in return. This is the standard tool and works well if all intervening network devices allow ICMP traffic. With ICMP ping, you can ping IPv4 or IPv6 addresses, or host names.

However, some networks prohibit ICMP. If this is true of your network, you can instead use TCP ping to test network connectivity. With TCP ping, the ping sends TCP SYN packets, and considers the ping a success if it receives a SYN-ACK in response. With TCP ping, you can ping IPv4 addresses or host names, but you cannot ping IPv6 addresses.

Keep in mind that a successful ICMP or TCP ping simply means that the address you are using is alive and responding to that specific type of traffic. This means that basic connectivity is working. Other policies running on a device could prevent specific types of traffic from successfully getting through a device.

Enable ICMP

By default, you can ping from a high security interface to a low security interface. You just need to enable ICMP inspection to allow returning traffic through. If you want to ping from low to high, then you need to apply an ACL to allow traffic.

When pinging an ASA interface, any ICMP rules applied to the interface must allow Echo Request and Echo Response packets. ICMP rules are optional: if you do not configure them, all ICMP traffic to an interface is allowed.

This procedure explains all of ICMP configuration you might need to complete to enable ICMP pinging of ASA interfaces, or for pinging through an ASA.

Procedure

Step 1 Ensure ICMP rules allow Echo Request/Echo Response.

ICMP rules are optional and apply to ICMP packets sent directly to an interface. If you do not apply ICMP rules, all ICMP access is allowed. In this case, no action is required.

However, if you do implement ICMP rules, ensure that you include rules that permit any address for the Echo and Echo-Reply messages on each interface. Configure ICMP rules on the **Configuration > Device Management > Management Access > ICMP** pane.

Step 2 Ensure access rules allow ICMP.

When pinging a host through an ASA, access rules must allow ICMP traffic to leave and return. The access rule must at least allow Echo Request/Echo Reply ICMP packets. You can add these rules as global rules.

If you do not have access rules, you will need to also allow the other type of traffic you want, because applying any access rules to an interface adds an implicit deny, so all other traffic will be dropped.

Configure access rules on the **Configuration > Firewall > Access Rules** pane. If you are simply adding the rules for testing purposes, you can delete them after completing the tests.

Step 3 Enable ICMP inspection.

ICMP inspection is needed when pinging through the ASA, as opposed to pinging an interface. Inspection allows returning traffic (that is, the Echo Reply packet) to return to the host that initiated the ping, and also ensures there is one response per packet, which prevents certain types of attack.

You can simply enable ICMP inspection in the default global inspection policy.

- a) Choose **Configuration > Firewall > Service Policy Rules**.
- b) Edit the **inspection_default** global rule.
- c) On the **Rule Actions > Protocol Inspection** tab, select ICMP.
- d) Click **OK**, then **Apply**.

Ping Hosts

To ping any device, you simply choose **Tools > Ping**, enter the IP address or host name of the destination you are pinging, and click **Ping**. For TCP ping, you select **TCP** and also include the destination port. That is usually the extent of any test you need to run.

Example output for a successful ping:

```

Sending 5, 100-byte ICMP Echos to out-pc, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

```

If the ping fails, the output indicates ? for each failed attempt, and the success rate is less than 100 percent (complete failure is 0 percent):

```

Sending 5, 100-byte ICMP Echos to 10.132.80.101, timeout is 2 seconds:
????
Success rate is 0 percent (0/5)

```

However, you can also add parameters to control some aspects of the ping. Following are your basic options:

- ICMP ping—You can select the interface through which the destination host is connected. If you do not select an interface, the routing table is used to determine the correct interface. You can ping IPv4 or IPv6 addresses or host names.
- TCP ping—You must also select the TCP port for the destination you are pinging. For example, **www.example.com 80** to ping the HTTP port. You can ping IPv4 addresses or host names, but not IPv6 addresses.

You also have the option to specify the source address and port that is sending the ping. In this case, optionally select the interface through which the source sends the ping (the routing table is used when you do not select an interface).

Finally, you can specify how often to repeat the ping (the default is 5 times) or the timeout for each attempt (the default is 2 seconds).

Test ASA Connectivity Systematically

If you want to do a more systematic test of ASA connectivity, you can use the following general procedure.

Before you begin

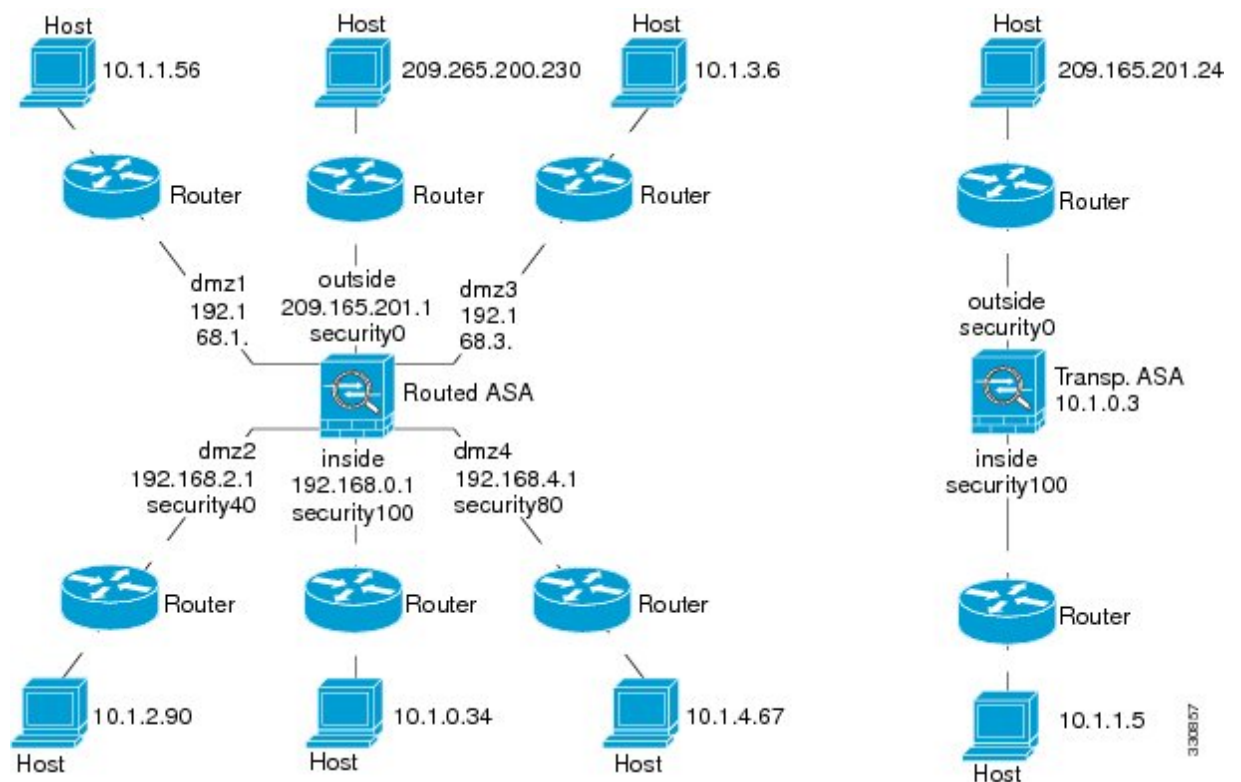
If you want to see the syslog messages mentioned in the procedure, enable logging (the **logging enable** command, or **Configuration > Device Management > Logging > Logging Setup** in ASDM).

Procedure

Step 1

Draw a diagram of your single-mode ASA or security context that shows the interface names, security levels, and IP addresses. The diagram should also include any directly connected routers and a host on the other side of the router from which you will ping the ASA.



Figure 1: Network Diagram with Interfaces, Routers, and Hosts



Step 2

Ping each ASA interface from the directly connected routers. For transparent mode, ping the BVI IP address. This test ensures that the ASA interfaces are active and that the interface configuration is correct.

A ping might fail if the ASA interface is not active, the interface configuration is incorrect, or if a switch between the ASA and a router is down (see the following figure). In this case, no debugging messages or syslog messages appear, because the packet never reaches the ASA.

Figure 2: Ping Failure at the ASA Interface**Figure 3: Ping Failure Because of IP Addressing Problems**

If the ping reply does not return to the router, then a switch loop or redundant IP addresses might exist (see the following figure).

- Step 3** Ping each ASA interface from a remote host. For transparent mode, ping the BVI IP address. This test checks whether the directly connected router can route the packet between the host and the ASA, and whether the ASA can correctly route the packet back to the host.

A ping might fail if the ASA does not have a return route to the host through the intermediate router (see the following figure). In this case, the debugging messages show that the ping was successful, but syslog message 110001 appears, indicating a routing failure has occurred.

Figure 4: Ping Failure Because the ASA Has No Return Route

- Step 4** Ping from an ASA interface to a network device that you know is functioning correctly.
- If the ping is not received, a problem with the transmitting hardware or interface configuration may exist.
 - If the ASA interface is configured correctly and it does not receive an echo reply from the “known good” device, problems with the interface hardware receiving function may exist. If a different interface with “known good” receiving capability can receive an echo after pinging the same “known good” device, the hardware receiving problem of the first interface is confirmed.
- Step 5** Ping from the host or router through the source interface to another host or router on another interface. Repeat this step for as many interface pairs as you want to check. If you use NAT, this test shows that NAT is operating correctly.

If the ping succeeds, a syslog message appears to confirm the address translation for routed mode (305009 or 305011) and that an ICMP connection was established (302020). You can also enter either the **show xlate** or **show conns** command to view this information.

The ping might fail because NAT is not configured correctly. In this case, a syslog message appears, showing that the NAT failed (305005 or 305006). If the ping is from an outside host to an inside host, and you do not have a static translation, you get message 106010.

Figure 5: Ping Failure Because the ASA is Not Translating Addresses



Trace Routes to Hosts

If you are having problems sending traffic to an IP address, you can trace the route to the host to determine if there is a problem on the network path.

Procedure

-
- Step 1** [Make the ASA Visible on Trace Routes, on page 20.](#)
Step 2 [Determine Packet Routes, on page 21.](#)
-

Make the ASA Visible on Trace Routes

By default, the ASA does not appear on traceroutes as a hop. To make it appear, you need to decrement the time-to-live on packets that pass through the ASA, and increase the rate limit on ICMP unreachable messages.

Procedure

-
- Step 1** Decrement the TTL using a service policy.
- Choose **Configuration > Firewall > Service Policy Rules**.
 - Add or edit a rule. For example, if you already have a rule to which you can add the option to decrement TTL, you do not need to create a new one.
 - Progress through the wizard to the Rule Actions page, applying the rule globally or to an interface, and specifying the traffic match. For example, you could create a global match any rule.
 - On the Rule Actions page, click the **Connection Settings** tab, and select **Decrement time to live for a connection**.
 - Click **OK** or **Finish**, then **Apply**.
- Step 2** Increase the ICMP unreachable rate limit.
- Choose **Configuration > Device Management > Management Access > ICMP**.
 - Increase the **IPv4 ICMP Unreachable Message Limits > Rate Limit** value at the bottom of the page. For example, increase it to 50.
 - Click **Apply**.
-

Determine Packet Routes

Use Traceroute to help you to determine the route that packets will take to their destination. A traceroute works by sending UDP packets or ICMPv6 echo to a destination on an invalid port. Because the port is not valid, the routers along the way to the destination respond with an ICMP or ICMPv6 Time Exceeded Message, and report that error to the ASA.

The traceroute shows the result of each probe sent. Every line of output corresponds to a TTL value in increasing order. The following table explains the output symbols.

Output Symbol	Description
*	No response was received for the probe within the timeout period.
U	No route to the destination.
<i>m</i> msec	For each node, the round-trip time (in milliseconds) for the specified number of probes.
!N.	ICMP network unreachable. For ICMPv6, address is out of scope.
!H	ICMP host unreachable.
!P	ICMP unreachable. For ICMPv6, port not reachable.
!A	ICMP administratively prohibited.
?	Unknown ICMP error.

Procedure

-
- Step 1** Choose **Tools > Traceroute**.
- Step 2** Enter the destination hostname or IP address to which you are tracing the route. Configure a DNS server to use a host name.
- Step 3** (Optional) Configure the characteristics of the trace. The defaults are appropriate in most cases.
- **Timeout**—How long to wait for a response before timing out. The default is 3 seconds.
 - **Port**—The UDP port to use. The default is 33434.
 - **Probe**—How many probes to send at each TTL level. The default is 3.
 - **TTL**—The minimum and maximum time-to-live values for the probes. The minimum default is one, but you can set it to a higher value to suppress the display of known hops. The maximum default is 30. The traceroute terminates when the packet reaches the destination or when the maximum value is reached.
 - **Specify source interface or IP address**—The interface to use as the source of the trace. You can specify the interface by name or by IP address. For IPv6, you cannot specify the source interface; you can only specify the source IP address. An IPv6 address is valid only if you enabled IPv6 on an ASA interface. In transparent mode, you must use the management address.
 - **Reverse Resolve**—Whether to have the output display the names of hops encountered if DNS name resolution is configured. Deselect the option to show IP addresses only.
 - **Use ICMP**—Whether to send ICMP probe packets instead of UDP probe packets.

Step 4 Click **Trace Route** to start the traceroute.

The **Traceroute Output** area displays detailed messages about the traceroute results.

Using the Packet Tracer to Test Policy Configuration

You can test your policy configuration by modeling a packet based on source and destination addressing and protocol characteristics. The trace does policy lookup to test access rules, NAT, routing, and so forth, to see if the packet would be permitted or denied.

By testing packets this way, you can see the results of your policies and test whether the types of traffic you want to allow or deny are handled as desired. Besides verifying your configuration, you can use the tracer to debug unexpected behavior, such as packets being denied when they should be allowed.

Procedure

Step 1 Choose **Tools > Packet Tracer**.

Step 2 Choose the source **Interface** for the packet trace.

Step 3 Specify the **Packet Type** for the packet trace. Available protocol types include: ICMP, IP, TCP, UDP, SCTP.

Step 4 (Optional.) If you want to trace a packet where the security group tag value is embedded in the Layer 2 CMD header (Trustsec), check **SGT number** and enter the security group tag number, 0-65533.

Step 5 (Transparent mode) If you want the packet tracer to enter a parent interface, which is later redirected to a subinterface, check **VLAN ID** and enter the ID, 1- 4096. VLAN ID is available only when the input interface is not a subinterface.

Step 6 (Transparent mode) Specify the **Destination MAC Address**.

Step 7 Specify the source and destination for the packets.

You can specify IPv4 or IPv6 addresses, fully-qualified domain names (FQDN), or security group names or tags if you use Cisco Trustsec. For the source address, you can also specify a username in the format Domain\username.

Step 8 Specify the protocol characteristics:

- ICMP—Enter the ICMP type, ICMP code (0-255), and optionally, the ICMP identifier.
- TCP/UDP/SCTP—Enter the source and destination port numbers.
- Raw IP—Enter the protocol number, 0-255.

Step 9 Use packet tracer to debug packets across cluster units. From the **Cluster Capture** drop-down list, select:

- a) **decrypted**—Injects a decrypted packet in a VPN tunnel and also simulates a packet that comes across a VPN tunnel.
- b) **persist**—Injects the packet you want to track across cluster units.
- c) **bypass-checks**—Skips security checks like ACL, VPN filters, IPsec spoof, and uRPF.
- d) **transmit**—Allows simulated packets to egress the ASA.

Step 10 Click **Start** to trace the packet.

The **Information Display Area** shows detailed messages about the results of the packet trace.

Monitoring Performance and System Resources

You can monitor a variety of system resources to identify performance or other potential problems.

Monitoring Performance

You can view ASA performance information in a graphical or tabular format.

Procedure

- Step 1** Choose **Monitoring > Properties > Connection Graphs > Perfmon**.
- Step 2** You can give the graph window a title by entering it in **Graph Window Title**, or you can choose an existing title.
- Step 3** Select up to four entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. The available options are the following:
- AAA Perfmon—Requests per second for authentication, authorization, and accounting requests.
 - Inspection Perfmon—Packets per second for HTTP, FTP, and TCP inspection.
 - Web Perfmon—Requests per second for URL access and URL server requests.
 - Connections Perfmon—Connections per second for all connections, UDP connections, TCP connections, and TCP Intercept.
 - Xlate Perfmon—NAT xlates per second.
- Step 4** Click **Show Graphs**.
- You can toggle each graph between graph and table views. You can also change how often the data refreshes, and export or print the data.
-

Monitoring Memory Blocks

You can view free and used memory blocks information in a graphical or tabular format.

Procedure

- Step 1** Choose **Monitoring > Properties > System Resources Graphs > Blocks**.
- Step 2** You can give the graph window a title by entering it in **Graph Window Title**, or you can choose an existing title.

- Step 3** Select entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. The available options are the following:
- Blocks Used—Displays the ASA used memory blocks.
 - Blocks Free—Displays the ASA free memory blocks.
- Step 4** Click **Show Graphs**.
- You can toggle each graph between graph and table views. You can also change how often the data refreshes, and export or print the data.
-

Monitoring CPU

You can view CPU utilization.

Procedure

- Step 1** Choose **Monitoring > Properties > System Resources Graphs > CPU**.
- Step 2** You can give the graph window a title by entering it in **Graph Window Title**, or you can choose an existing title.
- Step 3** Add CPU Utilization to the Selected Graphs list.
- Step 4** Click **Show Graphs**.
- You can toggle the graph between graph and table views. You can also change how often the data refreshes, and export or print the data.
-

Monitoring Memory

You can view memory utilization information in a graphical or tabular format.

Procedure

- Step 1** Choose **Monitoring > Properties > System Resources Graphs > Memory**.
- Step 2** You can give the graph window a title by entering it in **Graph Window Title**, or you can choose an existing title.
- Step 3** Select entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. The available options are the following:
- Free Memory—Displays the ASA free memory.
 - Used Memory—Displays the ASA used memory.
- Step 4** Click **Show Graphs**.

You can toggle each graph between graph and table views. You can also change how often the data refreshes, and export or print the data.

Monitoring Per-Process CPU Usage

You can monitor the processes that run on the CPU. You can obtain information about the percentage of CPU that is used by a certain process. CPU usage statistics are sorted in descending order to display the highest consumer at the top. Also included is information about the load on the CPU per process, at 5 seconds, 1 minute, and 5 minutes before the log time. This information is updated automatically every 5 seconds to provide real-time statistics. In ASDM, it is updated every 30 seconds.

To view CPU usage on a per-process basis, choose **Monitoring > Properties > Per-Process CPU Usage**.

You can stop the auto refresh, manually refresh the information, or save it to a file. You can also click **Configure CPU Usage Colors** to choose background and foreground colors based on usage percentages, to make it easier to scan for high-usage processes.

Monitoring Connections

To view current connections in a tabular format, in the ASDM main window, choose **Monitoring > Properties > Connections**. Information for each connection includes the protocol, source and destination address characteristics, idle time since the last packet was sent or received, and the amount of traffic in the connection.

History for Testing and Troubleshooting

Feature Name	Platform Releases	Description
IPv6 support for traceroute	9.7(1)	The traceroute command was modified to accept address. We modified the following screen: Tools > Traceroute
Support for the packet tracer for bridge group member interfaces	9.7(1)	You can now use the packet tracer for bridge group member interfaces. We added VLAN ID and Destination MAC Address to the packet-tracer screen: Tools > Packet Tracer
Manually start and stop packet captures	9.7(1)	You can now manually stop and start the captures. Added/Modified screens: Wizards > Packet Capture > Run Captures Added/Modified options: Start button, Stop button

Feature Name	Platform Releases	Description
Enhanced packet tracer and packet capture capabilities	9.9(1)	<p>The packet tracer has been enhanced with the following:</p> <ul style="list-style-type: none"> • Trace a packet when it passes between cluster u • Allow simulated packets to egress the ASA. • Bypass security checks for a simulated packet. • Treat a simulated packet as an IPsec/SSL decrypt <p>The packet capture has been enhanced with the following features:</p> <ul style="list-style-type: none"> • Capture packets after they are decrypted. • Capture traces and retain them in the persistent <p>New or modified screens:</p> <p>Tools > Packet Tracer</p> <p>We added Cluster Capture field to support these options: decrypted, persist, bypass-checks, transmit</p> <p>We added two new options in the Filter By view under Sessions drop-down list: Origin and Origin-ID</p> <p>Monitoring > VPN > VPN Statistics > Packet Tracer > Capture</p> <p>We added ICMP Capture field in the Packet Capture screen: Wizards > Packet Capture Wizard</p> <p>We added two options include-decrypted and persist to the ICMP Capture.</p>
Packet capture support for matching IPv6 traffic without using an ACL	9.10(1)	<p>If you use the match keyword for the capture command, the any keyword only matches IPv4 traffic. You can now use any4 and any6 keywords to capture either IPv4 or IPv6 traffic. The any keyword continues to match only IPv4 traffic.</p> <p>New/Modified commands: capture match</p> <p>No ASDM support.</p>
New debug telemetry command for Forepower 9300/4100.	9.14(1)	<p>If you use the debug telemetry command, debug messages related to telemetry are displayed. The debugs help to identify the cause for errors when generating the telemetry reports.</p> <p>No modified screens.</p>