Basic Settings

This chapter describes how to configure basic settings on the ASA that are typically required for a functioning configuration.

Set the Hostname, Domain Name, and the Enable and Telnet Passwords

To set the hostname, domain name, and the enable and Telnet passwords, perform the following steps.

Before you begin

Before you set the hostname, domain name, and the enable and Telnet passwords, check the following requirements:

- In multiple context mode, you can configure the hostname and domain name in both the system and context execution spaces.

- For the enable and Telnet passwords, set them in each context; they are not available in the system.

- To change from the system to a context configuration, in the **Configuration > Device List** pane, double-click the context name under the active device IP address.

Procedure

Step 1   Choose **Configuration** > **Device Setup** > **Device Name/Password**.

Step 2   Enter the hostname. The default hostname is "ciscoasa."

The hostname appears in the command line prompt, and if you establish sessions to multiple devices, the hostname helps you keep track of where you enter commands. The hostname is also used in syslog messages.

For multiple context mode, the hostname that you set in the system execution space appears in the command line prompt for all contexts. The hostname that you optionally set within a context does not appear in the command line; it can be used for a banner.

Step 3 Enter the domain name. The default domain name is default.domain.invalid.

The ASA appends the domain name as a suffix to unqualified names. For example, if you set the domain name to "example.com" and specify a syslog server by the unqualified name of "jupiter," then the ASA qualifies the name to "jupiter.example.com."

Step 4 Change the privileged mode (enable) password. The default password is blank, but you are prompted to change it the first time you enter the **enable** command at the CLI.

The enable password lets you enter privileged EXEC mode if you do not configure enable authentication. The enable password also lets you log into ASDM with a blank username if you do not configure HTTP authentication. ASDM does not enforce the enable password change like CLI access does.

a) Check the **Change the privileged mode password** check box.
b) Enter the new password, and then confirm the new password. Set a case-sensitive password of 8 to 127 characters long. It can be any combination of ASCII printable characters (character codes 32-126), with the following exceptions:

- No spaces

- No question marks

- You cannot use three or more consecutive sequential or repetitive ASCII characters. For example, the following passwords will be rejected:

  - **abc**user1

  - user**543**

  - user**aaaa**

  - user2**666**

You cannot reset the password to a blank value.

Step 5 Set the login password for Telnet access. There is no default password.

The login password is used for Telnet access when you do not configure Telnet authentication.

a) Check the **Change the password to access the console of the security appliance** check box.
b) Enter the old password (for a new ASA, leave this field blank), new password, then confirm the new password. The password can be up to 16 characters long. It can be any combination of ASCII printable characters (character codes 32-126), with the exception of spaces and the question mark.

Step 6 Click **Apply** to save your changes.

Set the Date and Time

✎

Note   Do not set the date and time for the Firepower 2100 in Platform mode, 4100, or 9300; the ASA receives these settings from the chassis.

Set the Date and Time Using an NTP Server

NTP is used to implement a hierarchical system of servers that provide a precisely synchronized time among network systems. This kind of accuracy is required for time-sensitive operations, such as validating CRLs, which include a precise time stamp. You can configure multiple NTP servers. The ASA chooses the server with the lowest stratum—a measure of how reliable the data is.

Time derived from an NTP server overrides any time set manually.
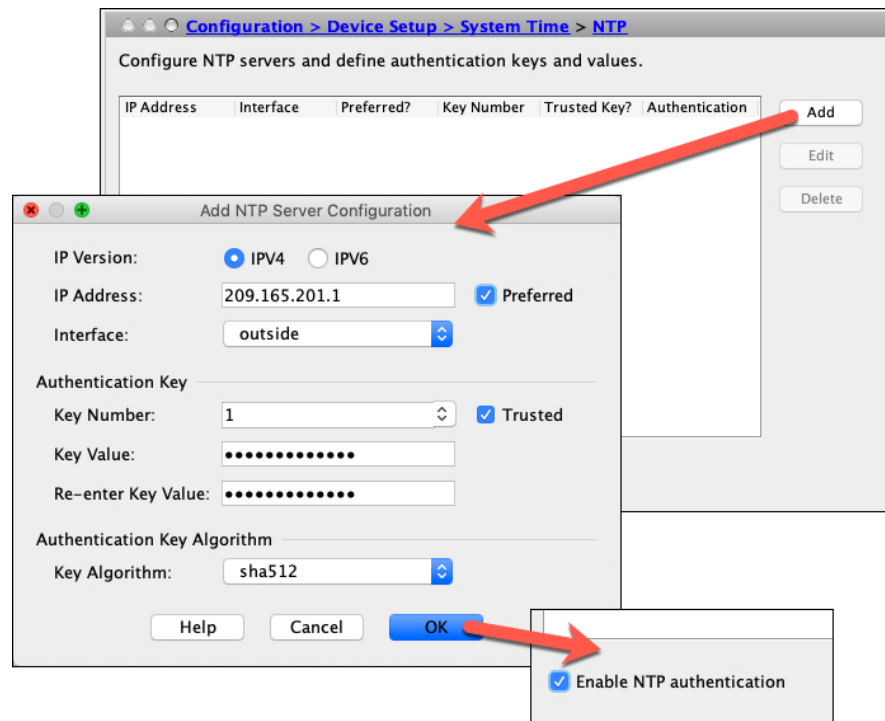
The ASA supports NTPv4.

Before you begin

In multiple context mode, you can set the time in the system configuration only.

Procedure

Step 1   Choose **Configuration** > **Device Setup** > **System Time** > **NTP**.

Step 2   Click **Add** to display the **Add NTP Server Configuration** dialog box.

Step 3   Enter the NTP server **IPv4** or **IPv6 IP Address**.

You cannot enter a hostname for the server; the ASA does not support DNS lookup for the NTP server.

Step 4   (Optional) Check the **Preferred** check box to set this server as a preferred server.

NTP uses an algorithm to determine which server is the most accurate and synchronizes to it. If servers are of similar accuracy, then the preferred server is used. However, if a server is significantly more accurate than the preferred one, the ASA uses the more accurate one.

Step 5   (Optional) Choose the **Interface** from the drop-down list.

This setting specifies the outgoing interface for NTP packets. If the interface is blank, then the ASA uses the default admin context interface according to the management routing table.

Step 6   (Optional) Configure NTP authentication.

a)   Enter a **Key Number** between 1 and 4294967295, or choose an existing key number from the drop-down list if you previously created a key for another NTP server that you want to reuse.

This setting specifies the key ID for this authentication key, which enables you to use authentication to communicate with the NTP server. The NTP server packets must also use this key ID.

b)   Check the **Trusted** check box.

c)   Enter the **Key Value**, which is a string up to 32 characters long, and then re-enter the key value.

d)   Choose a **Key Algorithm** from the drop-down list.

e)   Click **OK**.

Step 7   Check the **Enable NTP authentication** check box to turn on NTP authentication.

Step 8   Click **Apply** to save your changes.

Set the Date and Time Manually

To set the date and time manually, perform the following steps:

Before you begin

In multiple context mode, you can set the time in the system configuration only.

Procedure

Step 1   Choose **Configuration** > **Device Setup** > **System Time** > **Clock**.

Step 2   Choose the time zone from the drop-down list. This setting specifies the time zone as GMT plus or minus the appropriate number of hours. If you select the Eastern Time, Central Time, Mountain Time, or Pacific Time zone, then the time adjusts automatically for daylight savings time, from 2:00 a.m. on the second Sunday in March to 2:00 a.m. on the first Sunday in November.

Note   Changing the time zone on the ASA may drop the connection to intelligent SSMs.

Step 3   Click the **Date** drop-down list to display a calendar. Then find the correct date using the following methods:

- Click the name of the month to display a list of months, then click the desired month. The calendar updates to that month.

- Click the year to change the year. Use the up and down arrows to scroll through the years, or enter a year in the entry field.

- Click the arrows to the right and left of the month and year to scroll the calendar forward and backward one month at a time.

- Click a day on the calendar to set the date.

Step 4 Enter the time manually in hours, minutes, and seconds.

Step 5 Click **Update Display Time** to update the time shown in the bottom right corner of the ASDM pane. The current time updates automatically every ten seconds.

Configure Precision Time Protocol (ISA 3000)

The Precision Time Protocol (PTP) is a time-synchronization protocol developed to synchronize the clocks of various devices in a packet-based network. These device clocks are generally of varying precision and stability. The protocol is designed specifically for industrial, networked measurement and control systems, and is optimal for use in distributed systems because it requires minimal bandwidth and little processing overhead.

A PTP system is a distributed, networked system consisting of a combination of PTP and non-PTP devices. PTP devices include ordinary clocks, boundary clocks and transparent clocks. Non-PTP devices include network switches, routers and other infrastructure devices.

You can configure the ASA device to be a transparent clock. The ASA device does not synchronize its clock with the PTP clocks. The ASA device will use the PTP default profile, as defined on the PTP clocks.

When you configure the PTP devices, you define a domain number for the devices that are meant to function together.Thus, you can configure multiple PTP domains, and then configure each non-PTP device to use the PTP clocks for one specific domain.

Before you begin

- This feature is only available on the ISA 3000.

- Use of PTP is supported in single context mode only.

- Cisco PTP supports multicast PTP messages only.

- PTP is enabled on all ISA 3000 interfaces in transparent mode by default. In routed mode, you must add the necessary configuration to ensure that the PTP packets are allowed to flow through the device.

- PTP is available only for IPv4 networks, not for IPv6 networks.

- PTP configuration is supported on physical Ethernet interfaces, whether stand-alone or bridge group members. It is not supported on:

    - Management interface.

    - Subinterfaces, EtherChannels, BVIs. or any other virtual interfaces.

• PTP flows on VLAN subinterfaces are supported, assuming the appropriate PTP configuration is present on the parent interface.

• You must ensure that PTP packets are allowed to flow through the device. In transparent firewall mode, the access list configuration to allow PTP traffic is configured by default. PTP traffic is identified by UDP ports 319 and 320, and destination IP address 224.0.1.129, so in routed firewall mode any ACL that allows this traffic should be acceptable.

• In routed firewall mode, you must also enable multicast routing for PTP multicast groups:

    • Enter the global configuration mode command **multicast-routing**.

    • And for each interface that is not a bridge group member, and on which PTP is enabled, enter the interface configuration command **igmp join-group 224.0.1.129** to statically enable PTP multicast group membership. This command is not supported or needed for bridge group members.

Procedure

Step 1	Select **Configuration** > **Device Management** > **PTP**.
Step 2	Enter the **Domain value**.

This is the domain number for all ports on the device. Packets received on a different domain are treated like regular multicast packets and will not undergo any PTP processing. This value can be from zero to 255; the default value is zero. Enter the domain number that is configured on the PTP devices in your network.

Step 3 (Optional) Select **Enable End-to-End Transparent Clock Mode** to enable End-to-End Transparent mode on all PTP-enabled interfaces.

A transparent clock is a clock which compensates for its delays by measuring the residence times and updating the correctionField in the PTP packet.

Step 4 Enable PTP on one or more device interfaces by selecting an interface and clicking **Enable** or **Disable**.

Enable PTP on each interface through which the system can contact a PTP clock in the configured domain.

Step 5 Click **Apply**.

What to do next

You can choose **Monitoring** > **Properties** > **PTP** to view PTP clock and interface/port information.

Configure the Master Passphrase

The master passphrase allows you to securely store plain text passwords in encrypted format and provides a key that is used to universally encrypt or mask all passwords, without changing any functionality. Features that use the master passphrase include the following:

• OSPF

• EIGRP

- VPN load balancing

- VPN (remote access and site-to-site)

- Failover

- AAA servers

- Logging

- Shared licenses

Add or Change the Master Passphrase

To add or change the master passphrase, perform the following steps.

Before you begin

- This procedure will only be accepted in a secure session, for example by console, SSH, or ASDM via HTTPS.

- If failover is enabled but no failover shared key is set, an error message appears if you change the master passphrase, informing you that you must enter a failover shared key to protect the master passphrase changes from being sent as plain text.

  Choose **Configuration > Device Management > High Availability > Failover**, enter any character in the **Shared Key** field or 32 hexadecimal numbers (0-9A-Fa-f) if a failover hexadecimal key is selected, except a backspace. Then click **Apply**.

- Enabling or changing password encryption in Active/Standby failover causes a **write standby**, which replicates the active configuration to the standby unit. Without this replication, the encrypted passwords on the standby unit will differ even though they use the same passphrase; configuration replication ensures that the configurations are the same. For Active/Active failover, you must manually enter **write standby**. A **write standby** can cause traffic interruption in Active/Active mode, because the configuration is cleared on the secondary unit before the new configuration is synced. You should make all contexts active on the primary ASA using the **failover active group 1** and **failover active group 2** commands, enter **write standby**, and then restore the group 2 contexts to the secondary unit using the **no failover active group 2** command.

Procedure

Step 1 Choose one of the following options:

- In single context mode, choose **Configuration > Device Management > Advanced > Master Passphrase**.

- In multiple context mode, choose **Configuration > Device Management > Device Administration > Master Passphrase**.

Step 2 Check the **Advanced Encryption Standard (AES) password encryption** check box.

If no master passphrase is in effect, a warning message appears when you click Apply. You can click OK or Cancel to continue.

If you later disable password encryption, all existing encrypted passwords are left unchanged, and as long as the master passphrase exists, the encrypted passwords will be decrypted, as required by the application.

Step 3   Check the **Change the encryption master passphrase** check box to enable you to enter and confirm your new master passphrases. By default, they are disabled.

Your new master passphrase must be between 8 and 128 characters long.

If you are changing an existing passphrase, you must enter the old passphrase before you can enter a new one.

Leave the **New** and **Confirm master passphrase** fields blank to delete the master passphrase.

Step 4   Click **Apply**.

Disable the Master Passphrase

Disabling the master passphrase reverts encrypted passwords into plain text passwords. Removing the passphrase might be useful if you downgrade to a previous software version that does not support encrypted passwords.

Before you begin

- You must know the current master passphrase to disable it.

- This procedure works only in a secure session; that is, by Telnet, SSH, or ASDM via HTTPS.

  To disable the master passphrase, perform the following steps:

Procedure

Step 1   Choose one of the following options:

- In single context mode, choose **Configuration > Device Management > Advanced > Master Passphrase**.

- In multiple context mode, choose **Configuration > Device Management > Device Administration > Master Passphrase**.

Step 2   Check the **Advanced Encryption Standard (AES) password encryption** check box.

If no master passphrase is in effect, a warning statement appears when you click Apply. Click OK or Cancel to continue.

Step 3   Check the Change the encryption master passphrase check box.

Step 4   Enter the old master passphrase in the **Old master passphrase** field. You must provide the old master passphrase to disable it.

Step 5   Leave the **New master passphrase** and the **Confirm master passphrase** fields empty.

Step 6   Click **Apply**.

Configure the DNS Servers

You need to configure a DNS server so that the ASA can resolve host names to IP addresses. You also must configure a DNS server to use fully qualified domain names (FQDN) network objects in access rules.

Some ASA features require use of a DNS server to access external servers by domain name. Other features, such as the **ping** or **traceroute** command, let you enter a name that you want to ping or traceroute, and the ASA can resolve the name by communicating with a DNS server. Many SSL VPN and certificate commands also support names.

By default, there is a default DNS server group called DefaultDNS. You can create multiple DNS server groups: one group is the default, while other groups can be associated with specific domains. A DNS request that matches a domain associated with a DNS server group will use that group. For example, if you want traffic destined to inside eng.cisco.com servers to use an inside DNS server, you can map eng.cisco.com to an inside DNS group. All DNS requests that do not match a domain mapping will use the default DNS server group, which has no associated domains. For example, the DefaultDNS group can include a public DNS server available on the outside interface. Other DNS server groups can be configured for VPN tunnel groups. See the **tunnel-group** command in the command reference for more information.

Note The ASA has limited support for using the DNS server, depending on the feature.

Before you begin

Make sure that you configure the appropriate routing and access rules for any interface on which you enable DNS domain lookup so you can reach the DNS server.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **DNS** > **DNS Client**.

Step 2 Choose one of the following options in the **DNS Setup** area:

- **Configure one DNS server group**—This option defines the servers in the DefaultDNS group.
- **Configure multiple DNS server groups**—With this option, you can configure the DefaultDNS group as well as other groups that you can associate with specific domains, and groups for use with remote access SSL VPN group policies. Even if you configure the DefaultDNS group only, you must select this option if you want to alter the timeout and other characteristics used with the group.

Step 3 If you select **Configure one DNS server group**, configure the servers in the DefaultDNS group.

a) In **Primary DNS Server**, enter the IP address of the DNS server that should be used whenever it is available. For this server and each secondary server, optionally specify the *interface_name* through which the ASA communicates with the server. If you do not specify the interface, the ASA checks the data routing table; if there are no matches, it then checks the management-only routing table.

b) Click **Add** to add secondary DNS servers.

You can add up to six DNS servers. The ASA tries each DNS server in order until it receives a response. Use the **Move Up**/**Move Down** buttons to put the servers in priority order.

c) Enter a DNS domain name appended to the hostname if it is not fully-qualified.

Step 4   If you select **Configure multiple DNS server groups**, define the server group properties.

a)   Click **Add** to create a new group, or select a group and click **Edit**.

The DefaultDNS group is always listed.

b)   Configure the group properties.

• **Server IP Address to Add**, **Source Interface**—Enter the IP address of a DNS server and click **Add>>**. For each server, optionally specify the *interface_name* through which the ASA communicates with the server. If you do not specify the interface, the ASA checks the management-only routing table; if there are no matches, it then checks the data routing table.

You can add up to six DNS servers. The ASA tries each DNS server in order until it receives a response. Use the **Move Up**/**Move Down** buttons to put the servers in priority order.

• **Timeout**—The number of seconds, from 1 to 30, to wait before trying the next DNS server. The default is 2 seconds. Each time the ASA retries the list of servers, this timeout doubles.

• **Retries**—The number of times, from 0 to 10, to retry the list of DNS servers when the ASA does not receive a response.

• **Expire Entry Timer** (DefaultDNS or active group only)—The minimum TTL for the DNS entry, in minutes. If the expiration timer is longer than the entry's TTL, the TTL is increased to the expire entry time value. If the TTL is longer than the expiration timer, the expire entry time value is ignored: no additional time is added to the TTL in this case. Upon expiration, the entry is removed from the DNS lookup table. Removing an entry requires that the table be recompiled, so frequent removals can increase the processing load on the device. Because some DNS entries can have very short TTL (as short as three seconds), you can use this setting to virtually extend the TTL. The default is 1 minute (that is, the minimum TTL for all resolutions is 1 minute). The range is 1 to 65535 minutes. This option is used when resolving FQDN network objects only.

• **Poll Timer** (DefaultDNS or active group only)—The time, in minutes, of the polling cycle used to resolve FQDN network/host objects to IP addresses. FQDN objects are resolved only if they are used in a firewall policy. The timer determines the maximum time between resolutions; the DNS entry's time-to-live (TTL) value is also used to determine when to update to IP address resolution, so individual FQDNs might be resolved more frequently than the polling cycle. The default is 240 (four hours). The range is 1 to 65535 minutes.

• **Domain Name** (DefaultDNS or active group only)—The domain name appended to the hostname if it is not fully-qualified.

c)   Click **OK**.

d)   If you have multiple groups, you can change the default group by selecting it and clicking **Set Active**.

You can only use a group as the default if it does not have any domains mapped to it (see).

Step 5   Ensure that DNS lookup is enabled on at least one interface. In the **DNS Lookup** interface list, below the DNS server group table, click in the **DNS Enabled** column and select **True** to enable lookup on the interface.

Make sure to enable DNS lookup on all interfaces that will be used to access DNS servers.

If you do not enable DNS lookup on an interface, then the DNS server **Source Interface** or the interface found using the routing table cannot be used.

Step 6 (Optional) Under **Trusted DNS Server**, configure the options for determining which servers to trust when resolving domain names in network-service objects.

a) (Optional) Add or remove explicitly-configured trusted DNS servers.

- Click **Add** to add a new server, then select the IP type (IPv4 or IPv6), enter the IP address of the server, and click **OK**.

- Select a server and click **Edit** to change the address.

- Select a server and click **Delete** to remove it from the trusted server list.

b) Select or deselect the following options:

- **Any**—Trust every DNS server, snoop them all. This option is disabled by default.

- **Configured-Servers**—Whether servers configured in DNS server groups should be trusted. This option is enabled by default.

- **DHCP-Client**—Whether the servers that are learned by snooping messages between a DHCP client and DHCP server are considered trusted DNS servers. This option is enabled by default.

- **DHCP-Pools**—Whether the DNS servers that are configured in the DHCP pools for clients that obtain addresses through DHCP servers running on the device interfaces should be trusted. This option is enabled by default.

- **DHCP-Relay**—Whether the servers that are learned by snooping DHCP relay messages between a DHCP client and DHCP server are considered trusted DNS servers. This option is enabled by default.

Step 7 (Optional) Check the **Enable DNS Guard on all interfaces** check box to enforce one DNS response per query.

You can also set DNS Guard when configuring DNS inspection. For a given interface, the DNS Guard setting configured in DNS inspection takes precedence over this global setting. By default, DNS inspection is enabled on all interfaces with DNS Guard enabled.

Step 8 (Optional) Map domains to specific DNS server groups.

You can map up to 30 domains. You cannot map the same domain to multiple DNS server groups, but you can map multiple domains to the same server group. Do not map any domains to the group you want to use for the default (for example, DefaultDNS).

a) In the **DNS Group Map** area, check **Enable DNS Group Map**.
b) Click **Add**.

The **Add Domains to DNS Server Group** dialog box appears.

c) In the **DNS server group to domain name mapping** drop-down list, choose the DNS server group name.
d) In the **Domain Name** field, enter the domain name that you want to map to the DNS group.
e) Click **OK**.
f) Repeat these steps to add more mappings.

Step 9 Click **Apply** to save your changes.

Configure the Hardware Bypass and Dual Power Supply (Cisco ISA 3000)

You can enable the hardware bypass so that traffic continues to flow between an interface pair during a power outage. Supported interface pairs are copper GigabitEthernet 1/1 & 1/2; and GigabitEthernet 1/3 & 1/4. When the hardware bypass is active, no firewall functions are in place, so make sure you understand the risks of allowing traffic through. See the following hardware bypass guidelines:

- This feature is only available on the Cisco ISA 3000 appliance.

- If you have a fiber Ethernet model, only the copper Ethernet pair (GigabitEthernet 1/1 & 1/2) supports hardware bypass.

- When the ISA 3000 loses power and goes into hardware bypass mode, only the supported interface pairs can communicate; when using the default configuration, inside1 <---> inside2, and outside1 <---> outside2 can no longer communicate. Any existing connections between these interfaces will be lost.

- We suggest that you disable TCP sequence randomization (as described in this procedure). If randomization is enabled (the default), then when the hardware bypass is activated, TCP sessions will need to be re-established. By default, the ISA 3000 rewrites the initial sequence number (ISN) of TCP connections passing through it to a random number. When the hardware bypass is activated, the ISA 3000 is no longer in the data path and does not translate the sequence numbers; the receiving client receives an unexpected sequence number and drops the connection. Even with TCP sequence randomization disabled, some TCP connections will have to be re-established because of the link that is temporarily down during the switchover.

- Cisco TrustSec connections on hardware bypass interfaces are dropped when hardware bypass is activated. When the ISA 3000 powers on and hardware bypass is deactivated, the connections are renegotiated.

- When the hardware bypass is deactivated, and traffic resumes going through the ISA 3000 data path, some existing TCP sessions need to be re-established because of the link that is temporarily down during the switchover.

- When hardware bypass is active, the Ethernet PHYs are disconnected, so the ASA is unable to determine the interface status. Interfaces may appear to be in a down state.

For dual power supplies in the ISA 3000, you can establish dual power supplies as the expected configuration in the ASA OS. If one power supply fails, the ASA issues an alarm. By default, the ASA expects a single power supply and won't issue an alarm as long as it includes one working power supply.

Before you begin

- You must attach the hardware bypass interfaces to access ports on the switch. Do not attach them to trunk ports.

Procedure

Step 1   To configure hardware bypass, choose **Configuration** > **Device Management** > **Hardware Bypass**.

Step 2   Configure the hardware bypass to activate for each interface pair by checking the **Enable Bypass during Power Down** check box.

Step 3 (Optional) Configure each interface pair to remain in hardware bypass mode after the power comes back and the appliance boots up by checking the **Stay in Bypass after Power Up** check box.

When the hardware bypass is deactivated, there is a brief connection interruption as the ASA takes over the flows. In this case, you need to manually turn off the hardware bypass when you are ready; this option lets you control when the brief interruption occurs.

Step 4 For an interface pair, manually activate or deactivate the hardware bypass by checking the **Bypass Immediately** check box.

Step 5 (Optional) Configure the hardware bypass to remain active until after the ASA FirePOWER module boots up by checking the **Stay in Bypass Mode until after the ASA Firepower Module Boots Up** check box.

You must enable hardware bypass without the **Stay in Bypass after Power Up** option for the boot delay to operate. Without this option, the hardware bypass is likely to become inactive before the ASA FirePOWER module finishes booting up. This scenario can cause traffic to be dropped if you configured the module to fail-close, for example.

Step 6 Click **Apply**.

Step 7 Disable TCP randomization. This example shows how to disable randomization for all traffic by adding the setting to the default configuration.

   a) Choose **Configuration** > **Firewall** > **Service Policy**.
   b) Select the **sfrclass** rule, and click **Edit**.
   c) Click **Rule Actions**, and then click **Connection Settings**.
   d) Uncheck the **Randomize Sequence Number** check box.
   e) Click **OK**, and then **Apply**.

Step 8 To establish dual power supplies as the expected configuration, choose **Configuration** > **Device Management** > **Power Supply**, check the **Enable Redundant Power Supply** check box, and click **Apply**.

This screen also shows the available power supplies.

Step 9 Click **Save**.

The behavior of hardware bypass after the system comes online is determined by the configuration setting in the startup configuration, so you must save your running configuration.

Adjust ASP (Accelerated Security Path) Performance and Behavior

The ASP is an implementation layer that puts your policies and configurations into action. It is not of direct interest except during troubleshooting with the Cisco Technical Assistance Center. However, there are a few behaviors related to performance and reliability that you can adjust.

Choose a Rule Engine Transactional Commit Model

By default, when you change a rule-based policy (such as access rules), the changes become effective immediately. However, this immediacy comes with a slight cost in performance. The performance cost is

more noticeable for very large rule lists in a high connections-per-second environment, for example, when you change a policy with 25,000 rules while the ASA is handling 18,000 connections per second.

The performance is affected because the rule engine compiles rules to enable faster rule lookup. By default, the system also searches uncompiled rules when evaluating a connection attempt so that new rules can be applied; because the rules are not compiled, the search takes longer.

You can change this behavior so that the rule engine uses a transactional model when implementing rule changes, continuing to use the old rules until the new rules are compiled and ready for use. With the transactional model, performance should not drop during the rule compilation. The following table clarifies the behavioral difference.

Model	Before Compilation	During Compilation	After Compilation
Default	Matches old rules.	Match new rules. (The rate for connections per second decreases.)	Matches new rules.
Transactional	Matches old rules.	Match old rules. (The rate for connections per second is unaffected.)	Matches new rules.

An additional benefit of the transactional model is that, when replacing an ACL on an interface, there is no gap between deleting the old ACL and applying the new one. This feature reduces the chances that acceptable connections may be dropped during the operation.

Tip If you enable the transactional model for a rule type, syslogs to mark the beginning and the end of the compilation are generated. These syslogs are numbered 780001 through 780004.

Use the following procedure to enable the transactional commit model for the rule engine.

Procedure

Choose **Configuration > Device Management > Advanced > Rule Engine** and select the desired options:

- **Access group**—Access rules applied globally or to interfaces.
- **NAT**—Network address translation rules.

Enable ASP Load Balancing

The ASP load balancing mechanism helps avoid the following issues:

- Overruns caused by sporadic traffic spikes on flows
- Overruns caused by bulk flows oversubscribing specific interface receive rings
- Overruns caused by relatively heavily overloaded interface receive rings, in which a single core cannot sustain the load.

ASP load balancing allows multiple cores to work simultaneously on packets that were received from a single interface receive ring. If the system drops packets, and the **show cpu** command output is far less than 100%, then this feature may help your throughput if the packets belong to many unrelated connections.

Note   ASP load balancing is disabled on the ASA virtual. With the integration of DPDK (Dataplane Development Kit) into the ASA virtual's accelerated security path (ASP), the ASA virtual shows better performance with this feature disabled.

Procedure

Step 1   To enable the automatic switching on and off of ASP load balancing, choose **Configuration** > **Device Management** > **Advanced** > **ASP Load Balancing**, and check the **Dynamically enable or disable ASP load balancing based on traffic monitoring** check box .

Step 2   To manually enable or disable ASP load balancing, check or uncheck the **Enable ASP load balancing** check box.

When you manually enable ASP load balancing, it is enabled until you manually disable it, even if you also have the Dynamic option enabled. Manually disabling ASP load balancing applies only if you manually enabled ASP load blancing. If you also enabled the Dynamic option, then the system reverts to automatically enabling or disabling ASP load balancing.

Monitoring the DNS Cache

The ASA provides a local cache of DNS information from external DNS queries that are sent for certain clientless SSL VPN and certificate commands. Each DNS translation request is first looked for in the local cache. If the local cache has the information, the resulting IP address is returned. If the local cache can not resolve the request, a DNS query is sent to the various DNS servers that have been configured. If an external DNS server resolves the request, the resulting IP address is stored in the local cache with its corresponding hostname.

See the following command for monitoring the DNS cache:

- **show dns-hosts**

   This command shows the DNS cache, which includes dynamically learned entries from a DNS server as well as manually entered name and IP addresses using the name command.

History for Basic Settings

Feature Name	Platform Releases	Description
Multiple DNS server groups	9.18(1)	You can now use multiple DNS server groups: one group is the default, while other groups can be associated with specific domains. A DNS request that matches a domain associated with a DNS server group will use that group. For example, if you want traffic destined to inside eng.cisco.com servers to use an inside DNS server, you can map eng.cisco.com to an inside DNS group. All DNS requests that do not match a domain mapping will use the default DNS server group, which has no associated domains. For example, the DefaultDNS group can include a public DNS server available on the outside interface. New/Modified screens: **Configuration** > **Device Management** > **DNS** > **DNS Client**
Trusted DNS servers for network-service object domain resolution.	9.17(1)	You can specify which DNS servers the system should trust when resolving domain names in network-service objects. This feature ensures that any DNS domain name resolutions acquire IP addresses from trusted sources. New/Modified screens: **Configuration** > **Device Management** > **DNS** > **DNS Client**

Stronger local user and enable password requirements	9.17(1)	For local users and the enable password, the following password requirements were added:
		
		• Password length—Minimum 8 characters. Formerly, the minimum was 3 characters.
		
		• Repetitive and sequential characters—Three or more consecutive sequential or repetitive ASCII characters are disallowed. For example, the following passwords will be rejected:
		
		    • **abc**user1
		
		    • user**543**
		
		    • user**aaaa**
		
		    • user2**666**
		
		New/Modified screens:
		
		• **Configuration** > **Device Management** > **Users/AAA** > **User Accounts**
		
		• **Configuration** > **Device Setup** > **Device Name/Password**

NTPv4 support	9.14(1)	The ASA now supports NTPv4. No modified screens.

Feature Name	Platform Releases	Description
Additional NTP authentication algorithms	9.13(1)	Formerly, only MD5 was supported for NTP authentication. The ASA now supports the following algorithms: • MD5 • SHA-1 • SHA-256 • SHA-512 • AES-CMAC New/Modified screens: **Configuration** > **Device Setup** > **System Time** > **NTP** > **Add** button **> Add NTP Server Configuration** dialog box **> Key Algorithm** drop-down list
NTP support on IPv6	9.12(1)	You can now specify an IPv6 address for the NTP server. New/Modified screens: **Configuration** > **Device Setup** > **System Time** > **NTP** > **Add** button **> Add NTP Server Configuration** dialog box
enable password change now required on login	9.12(1)	The default **enable** password is blank. When you try to access privileged EXEC mode on the ASA, you are now required to change the password to a value of 3 to 127 characters. You cannot keep it blank. The **no enable password** command is no longer supported. At the CLI, you can access privileged EXEC mode using the **enable** command, the **login** command (with a user at privilege level 2+), or an SSH or Telnet session when you enable **aaa authorization exec auto-enable**. All of these methods require you to set the enable password. This password change requirement is not enforced for ASDM logins. In ASDM, by default you can log in without a username and with the **enable** password. No modified screens.
ASP load balancing is disabled on the ASA virtual	9.10(1)	With the recent integration of DPDK (Dataplane Development Kit) into the ASA virtual's accelerated security path (ASP), the ASA virtual shows better performance with this feature disabled.
Automatic ASP load balancing now supported for the ASA virtual	9.8(1)	Formerly, you could only manually enable and disable ASP load balancing. We modified the following screen: **Configuration > Device Management > Advanced > ASP Load Balancing**

Feature Name	Platform Releases	Description
PBKDF2 hashing for all local **username** and **enable** passwords	9.7(1)	Local **username** and **enable** passwords of all lengths are stored in the configuration using a PBKDF2 (Password-Based Key Derivation Function 2) hash using SHA-512. Previously, passwords 32 characters and shorter used the MD5-based hashing method. Already existing passwords continue to use the MD5-based hash unless you enter a new password. See the "Software and Configurations" chapter in the General Operations Configuration Guide for downgrading guidelines. We modified the following screens: **Configuration** > **Device Setup** > **Device Name/Password** > **Enable Password** **Configuration** > **Device Management** > **Users/AAA** > **User Accounts** > **Add/Edit User Account** > **Identity**
Dual power supply support for the ISA 3000	9.6(1)	For dual power supplies in the ISA 3000, you can establish dual power supplies as the expected configuration in the ASA OS. If one power supply fails, the ASA issues an alarm. By default, the ASA expects a single power supply and won't issue an alarm as long as it includes one working power supply. We introduced the following screen: **Configuration** > **Device Management** > **Power Supply**
Longer password support for local **username** and **enable** passwords (up to 127 characters)	9.6(1)	You can now create local **username** and **enable** passwords up to 127 characters (the former limit was 32). When you create a password longer than 32 characters, it is stored in the configuration using a PBKDF2 (Password-Based Key Derivation Function 2) hash. Shorter passwords continue to use the MD5-based hashing method. We modified the following screens: **Configuration** > **Device Setup** > **Device Name/Password** > **Enable Password** **Configuration** > **Device Management** > **Users/AAA** > **User Accounts** > **Add/Edit User Account** > **Identity**
ISA 3000 hardware bypass	9.4(1.225)	The ISA 3000 supports a hardware bypass function to allow traffic to continue flowing through the appliance when there is a loss of power. We introduced the following screen: **Configuration** > **Device Management** > **Hardware Bypass** *This feature is not available in Version 9.5(1).*
Automatic ASP Load Balancing	9.3(2)	You can now enable automatic switching on and off of the ASP load balancing feature. **Note** The automatic feature is not supported on the ASA virtual; only manual enabling and disabling is supported. We modified the following screen: **Configuration > Device Management > Advanced > ASP Load Balancing**

Feature Name	Platform Releases	Description
Removal of the default Telnet password	9.0(2)/9.1(2)	To improve security for management access to the ASA, the default login password for Telnet was removed; you must manually set the password before you can log in using Telnet. **Note** The login password is only used for Telnet if you do not configure Telnet user authentication. Previously, when you cleared the password, the ASA restored the default of "cisco." Now when you clear the password, the password is removed. The login password is also used for Telnet sessions from the switch to the ASASM (see the **session** command). For initial ASASM access, you must use the **service-module session** command, until you set a login password. We did not modify any ASDM screens.
Password Encryption Visibility	8.4(1)	We modified the **show password encryption** command.
Master Passphrase	8.3(1)	We introduced this feature. The master passphrase allows you to securely store plain text passwords in encrypted format and provides a key that is used to universally encrypt or mask all passwords, without changing any functionality. We introduced the following screens: **Configuration > Device Management > Advanced > Master Passphrase** **Configuration > Device Management > Device Administration > Master Passphrase**