Section 3

ASA Access Control

Now that you have connectivity to the ASA and have configured basic networking settings on the ASA, you can start to look at making the ASA operate as a true firewall.

This section covers the services offered by the ASA such as Network Address Translation (NAT) and Access Lists that make up the main foundation of protection offered by the Cisco ASA.

Network Address Translation

Network Address Translation (NAT) is a key concept and technology used by the ASA. The main purpose of NAT is to translate one IP address into another. It is commonly used to translate private IP addresses into publicly routable IP addresses for use over the Internet.

One of the main uses of NAT is at the perimeter of the corporate network. Good network design dictates that corporate networks use a private IP addressing scheme as defined by RFC 1918. This includes the following addresses for use within corporate networks:

- 10.0.0.0/8
- 172.16.0.0/12
- 192.168.0.0/16

The preceding networks are not publically routable on the Internet. Therefore to communicate on the public Internet you need to translate these private addresses into a public address. A public address is usually assigned by an Internet service provider and is an address outside of the RFC 1918 private range. For example, 194.73.134.1 is considered a public IP address.

Figure 21 shows how NAT would be used in the simple network you configured on the ASA.

FIGURE 21
Simple NAT with the ASA

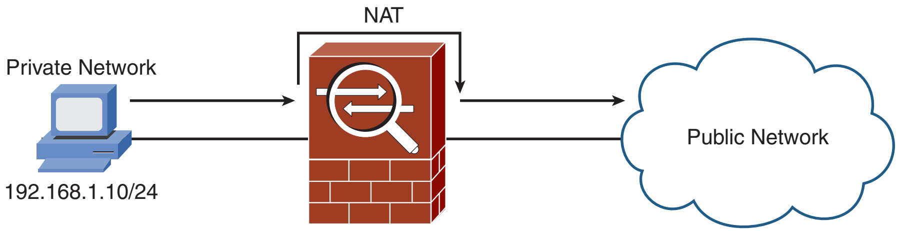

NAT, in the true sense, translates one address to another address. There is also a subfunction of NAT called Port Address Translation (PAT), which is where multiple internal addresses are translated into a single external address. Different source ports are used on the external address to differentiate between the internal addresses, and this information is held by the device performing the translation so that it can work out where to send the return packets. PAT is also commonly referred to as NAT-Overload.

On the ASA, NAT is required when traffic is flowing from a lower-security interface to a higher-security interface. For example, the outside interface has a security level of 0, and the inside interface has a security level of 100. Therefore, NAT is required for hosts on the outside to communicate with hosts on the inside.

NAT is not required by default for traffic flowing from a higher-security interface to a lower-security interface. This has been the case since Cisco released PIX and ASA 7.0. You can enable this setting by issuing the **nat-control** command, which then forces the use of NAT on all interfaces in all directions.

Simple NAT Configuration

To configure NAT on the ASDM, choose Configuration from the toolbar and then Firewall. NAT rules are created on the ASA and perform the translations depending on the configuration. By default, no NAT rules are configured on the ASA.

You can add three main types of NAT rules: static, dynamic, and exempt.

Start these examples by adding a static NAT rule.

Adding a Static NAT Rule with ASDM

Static NAT is where you can perform a one-to-one NAT translation; a single internal IP address is translated to a single external IP address. This is normally used for inbound access where external users are accessing resources such as a corporate web or email server.

Figure 22 shows the inclusion of a web server in your simple network topology.

FIGURE 22
Web Server

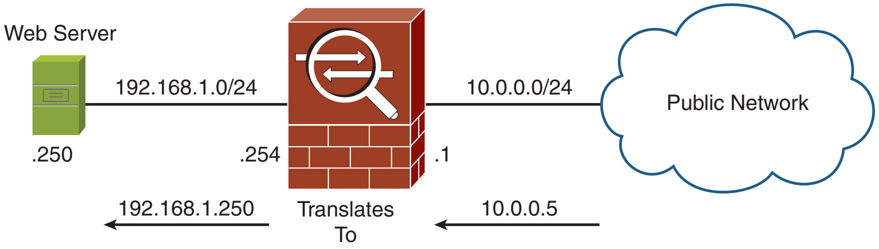

The web server has an internal IP address of 192.168.1.250/24. Hosts from the outside cannot access this server on this address because it is not routable via the outside interface (because a NAT translation is required).

What we need to configure is a static NAT entry from 192.168.1.250 that translates to 10.0.0.5. Doing so then enables external users to access the web server on 10.0.0.5.

Figure 23 shows a static NAT translation and the settings.

FIGURE 23
Static NAT Configuration

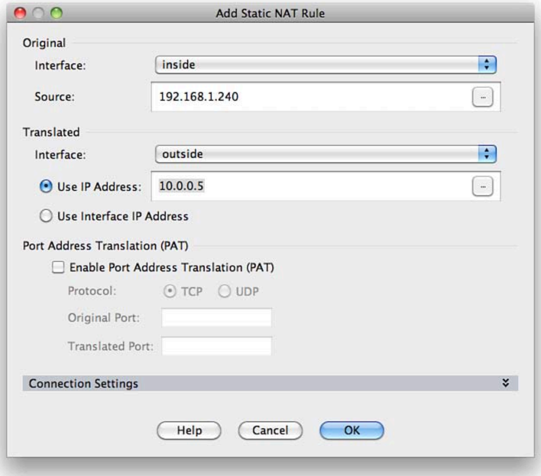

Add Static NAT Rule

Original
Interface: inside
Source: 192.168.1.240

Translated
Interface: outside
● Use IP Address: 10.0.0.5
○ Use Interface IP Address

Port Address Translation (PAT)
□ Enable Port Address Translation (PAT)
Protocol: ● TCP ○ UDP
Original Port:
Translated Port:

You can see in Figure 23 the entered original and translated addresses into the ASDM. This setting can translate 192.168.1.250 to 10.0.0.5 on the outside interface. When applying this setting, you then go back to the NAT Rules configuration screen shown in Figure 24. This screen lists all the configured NAT rules on the ASA, which in this case is only the single static NAT rule.

Now that you have configured a static NAT rule through the ASDM, the next rule to look at is a dynamic NAT rule.

FIGURE 24
NAT Rules Screen

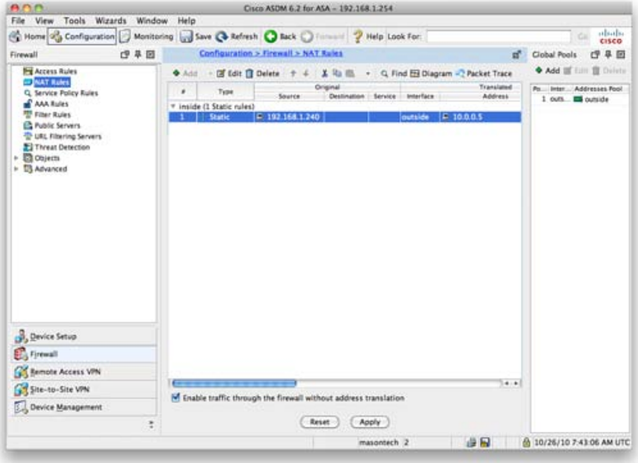

Cisco ASDM 6.2 for ASA - 192.168.1.254

Configuration > Firewall > NAT Rules

#	Type	Original			Translated	
		Source	Destination	Service	Interface	Address
inside (1 Static rules)						
1	Static	192.168.1.240			outside	10.0.0.5

Global Pools

Po...	Inter...	Addresses Pool
1	outs...	outside

☑ Enable traffic through the firewall without address translation

Adding a Dynamic NAT Rule with ASDM

Static NAT translations provide a one-to-one translation of IP addresses. A dynamic NAT rule creates a one-to-many translation of IP addresses.

The most common use of dynamic NAT is when the ASA is placed at the network perimeter between the corporate network and the Internet. Users on the corporate network want Internet access, so they require a NAT translation to translate their private IP address to a publicly routable IP address. If there are 50 internal users, you would require 50 public addresses for the translation if you use a static NAT translation.

Dynamic NAT uses a single public IP address and enables all the internal users access to the Internet. They all use the same public IP address and are tracked by the ASA by using different source ports for each internal client.

When you add a dynamic NAT rule, you need to link it to an address pool. You can create these address pools via **Objects > Global Pool**, or you can create them while adding the dynamic NAT rule.

Figure 25 shows a dynamic NAT rule that applies to 192.168.1.0/24 and uses the address of the outside interface for the translation.

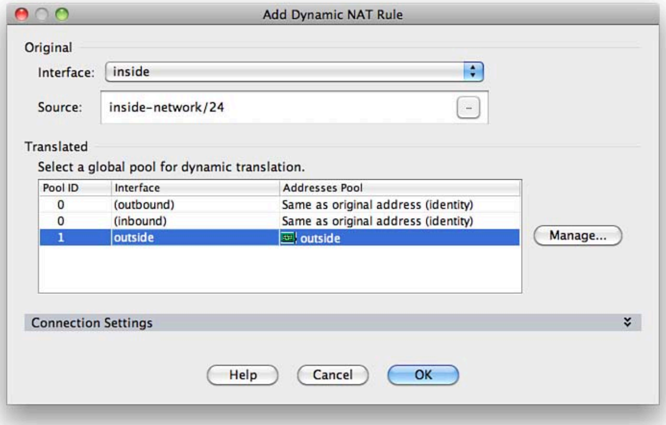

Add Dynamic NAT Rule

Original

Interface: inside

Source: inside-network/24

Translated

Select a global pool for dynamic translation.

Pool ID	Interface	Addresses Pool
0	(outbound)	Same as original address (identity)
0	(inbound)	Same as original address (identity)
1	outside	outside

FIGURE 25
Adding a Dynamic NAT Rule

When you apply this rule, you go back to the NAT Rules screen. On this screen, as shown in Figure 26, you see a configured static NAT rule and a dynamic NAT rule.

The last type of NAT rule to look at is a NAT exempt rule.

FIGURE 26
NAT Rules Screen

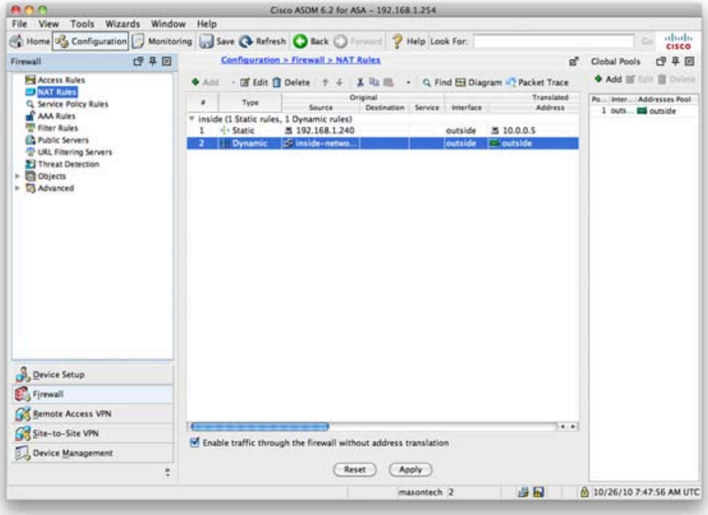

#	Type	Original			Translated	
		Source	Destination	Service	Interface	Address
inside (1 Static rules, 1 Dynamic rules)						
1	Static	192.168.1.240			outside	10.0.0.5
2	Dynamic	inside-netwo...			outside	outside

Po...	Inter...	Addresses Pool
1	outs...	outside

☑ Enable traffic through the firewall without address translation

Adding a NAT Exempt Rule with ASDM

You have just configured static and dynamic NAT on the ASA via ASDM. Now look at the third option available when adding a NAT rule: a NAT exempt rule.

NAT exemption exempts addresses from NAT translation. When NAT is configured on an interface, you sometimes might need a specific host to bypass NAT and be exempt from the NAT rules. A common use of this is when configuring VPNs and you want the local private network to communicate with the remote private network without being translated.

Access Lists

You just looked at configuring NAT on the ASA. Now that you have configured NAT, the next element to look at is ACLs, which are the restrictive lists that define the firewall. These are also called a firewall rule set or rule base. The ACL is one of the most important aspects of the firewall because it permits and denies traffic through the firewall. The incorrect configuration of an ACL can result in a security hole that a potential attacker may use to exploit an internal system. Without an ACL, a firewall is not much more than a standard router.

Configuring ACLs with ASDM

In the example, you have shown a web server on the inside of the network. This web server has a static NAT translation. Now provide an ACL that enables inbound traffic matching the web server address to access the web server.

To configure ACLs, from the toolbar select **Configuration** and then **Firewall > Access Rules**.

You see that by default some rules are already applied to the ASA, which are the implicit rules configured by default on the ASA device. These rules cannot be removed; they are the catchall rules matched if no other rule is matched first.

Figure 27 shows these implicit rules. Looking at Figure 27, you can see three implicit access rules. Two are applied to the inside interface and one to the outside interface.

The implicit access rules applied to the inside interface are as follows:

- Permit traffic from anywhere destined to a lower-security interface.
- Deny any traffic from anywhere to anywhere.

FIGURE 27
Implicit Access Rules

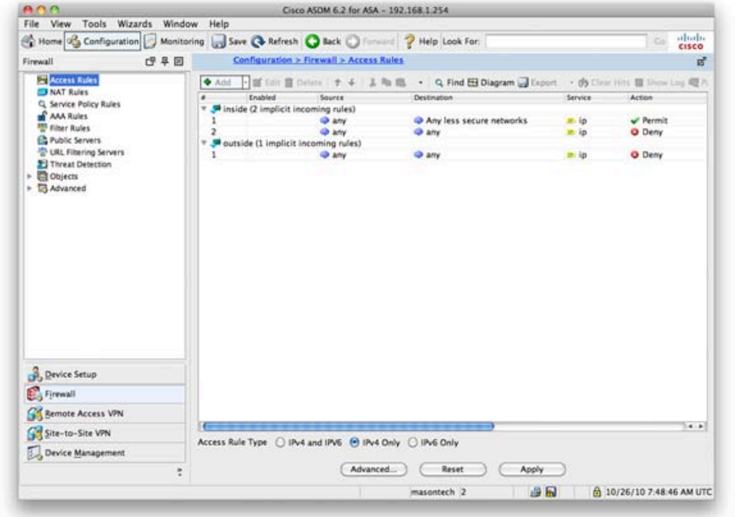

#	Enabled	Source	Destination	Service	Action
inside (2 implicit incoming rules)					
1		any	Any less secure networks	ip	✓ Permit
2		any	any	ip	Deny
outside (1 implicit incoming rules)					
1		any	any	ip	Deny

Access Rule Type ○ IPv4 and IPV6 ● IPv4 Only ○ IPv6 Only

This rule implements the ASA mentioned earlier. Any traffic from the inside interface is permitted only to lower-security interfaces. All other traffic is denied.

The implicit access rule applied to the outside interface is as follows:

- Deny any traffic from anywhere to anywhere.

Because the outside interface has the lowest available security level (0), all traffic is by default denied unless a more specific access rule permits it. This default ensures that nothing enters the firewall from the outside without previous configuration.

Clicking Add brings up the Add an Access Rule screen. You need to add a rule on the outside interface to enable access to the internal web server on TCP port 80.

Figure 28 shows the completed screen to create this access rule.

FIGURE 28
Add an Access Rule

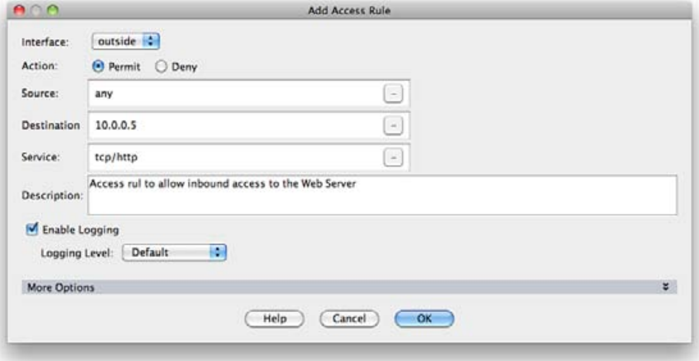

Add Access Rule

Interface: outside
Action: ● Permit ○ Deny
Source: any
Destination 10.0.0.5
Service: tcp/http
Description: Access rul to allow inbound access to the Web Server
☑ Enable Logging
Logging Level: Default

You can see from Figure 28 that you created a rule that permits TCP port 80/HTTP traffic from anywhere to access the outside interface and the address 10.0.0.5. This is the address you used for the static NAT translation for the web server.

Applying this access rule takes you back to the Access Rules screen of the ASDM, as shown in Figure 29. You can now see the new rule that has been configured on the outside interface and has been placed above the implicit rule that denies all other traffic. Traffic can now access the website from anywhere on the Internet.

FIGURE 29
Access Rules Screen

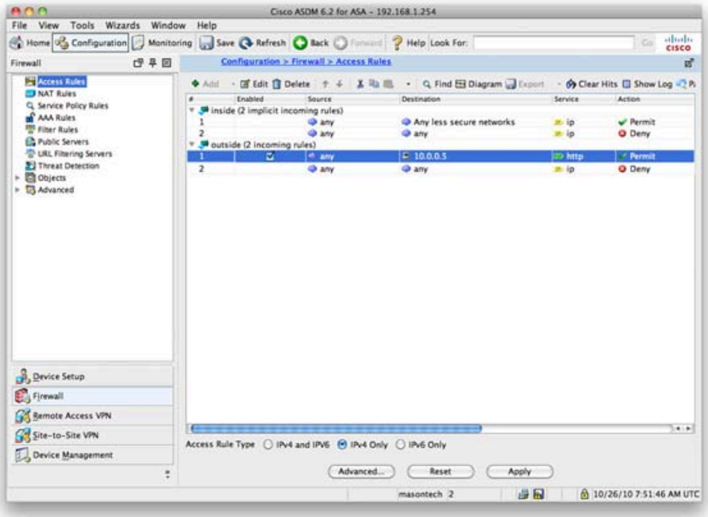

#	Enabled	Source	Destination	Service	Action
inside (2 implicit incoming rules)					
1		any	Any less secure networks	ip	Permit
2		any	any	ip	Deny
outside (2 incoming rules)					
1	☑	any	10.0.0.5	http	Permit
2		any	any	ip	Deny

Access Rule Type ○ IPv4 and IPv6 ● IPv4 Only ○ IPv6 Only

Using Object Groups Within ACLs

You now want to extend this further and permit HTTPS into the web server. You also want to use the name **webserver-public** rather than the public IP address of the web server. You can achieve both of these goals by configuring objects on the ASA.

Start by defining the web server as an object within the ASA. Navigate to **Firewall > Objects > Network Objects/Groups**. Add an entry using the public IP address of the web server and the name of **webserver-public**. You can see this completed in Figure 30.

FIGURE 30
Configure an IP Name

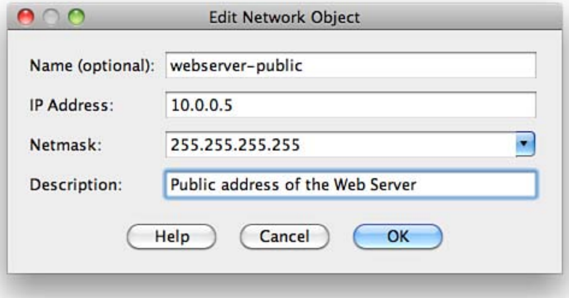

Name (optional): webserver-public

IP Address: 10.0.0.5

Netmask: 255.255.255.255

Description: Public address of the Web Server

You have now configured a more meaningful name for the public IP address of the web server. Next configure an access rule to also enable HTTPS access to the web server. You could just add another access rule that permits HTTPS from anywhere to the web server, as you did with the initial access rule that permitted HTTP. However, you can achieve this is by creating a service group. Navigate to **Firewall > Objects > Service Groups**, and then click **Add** to add a new service group.

Because both HTTP and HTTPS are TCP protocols, you need to select **TCP Service Group**. Call the group **webaccess** and add both HTTP and HTTPS to the group.

Figure 31 shows that you created a single TCP service group called **webaccess** that now contains these two protocols.

The next step is to go back to the Access Rules configuration screen and change the existing access rule to use the new group.

When you navigate to **Firewall > Access Rules**, the destination field for the access rule that you created has now changed from 10.0.0.5 to **webserver-public**. This is because you added the IP name object for the web server. From now on, the ASA will know 10.0.0.5 as **webserver-public**, making it easier to read the rule base and understand it.

FIGURE 31
Edit TCP Service Group

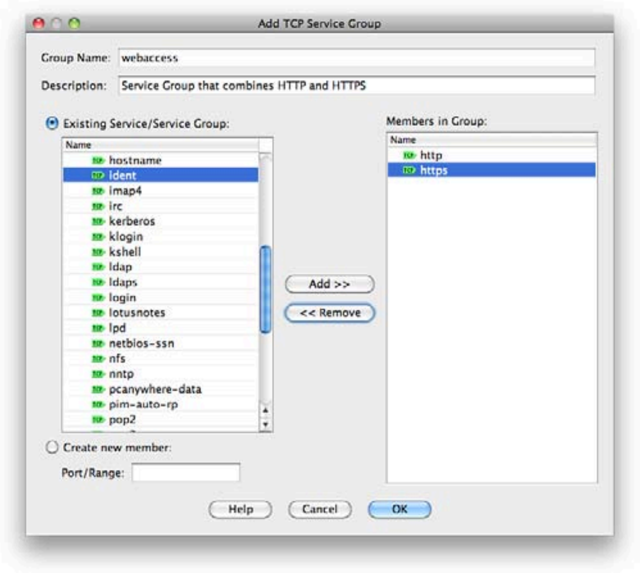

Add TCP Service Group

Group Name: webaccess

Description: Service Group that combines HTTP and HTTPS

● Existing Service/Service Group:

Members in Group: http, https

Highlight the rule you created earlier, and click **Edit**. Select the service by clicking the ellipsis (**...**) button next to the service. Now add the TCP service group that you have just created. The service group should be at the top of the list. Select it and click **OK** to return to the main Access Rules screen, as shown in Figure 32.

Figure 32 shows that the destination name is now **webserver-public** and the service is now **webaccess** and not HTTP.

You just created a single rule that enables both HTTP and HTTPS access inbound on the outside interface to the public address of the web server.

FIGURE 32
Access Rules Screen

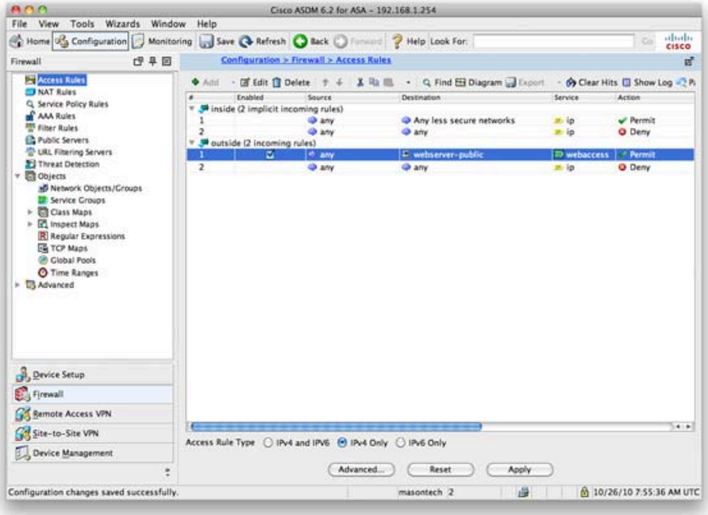

Configuration > Firewall > Access Rules

#	Enabled	Source	Destination	Service	Action
inside (2 implicit incoming rules)					
1		any	Any less secure networks	ip	Permit
2		any	any	ip	Deny
outside (2 incoming rules)					
1	☑	any	webserver-public	webaccess	Permit
2		any	any	ip	Deny

Access Rule Type ○ IPv4 and IPv6 ● IPv4 Only ○ IPv6 Only

Configuration changes saved successfully. masontech 2 10/26/10 7:55:36 AM UTC

Modular Policy Framework

The Modular Policy Framework (MPF) is an advanced feature of the ASA that provides the security administrator with greater granularity and more flexibility when configuring network policies. The security administrator can do the following:

- Define flows of traffic.
- Associate security policies to traffic flows.
- Enable a set of security policies on an interface or globally.

Modular policies consist of the following components:

- Class maps
- Policy maps
- Service policies

Class Maps

A class map is a configuration element used to match something. A class map is similar in operation to an access control list (ACL), but with class maps you can match other items that ACLs cannot match.

Class maps can define a class of traffic by matching via the follow command keywords:

- **access list**: An entry in an ACL.
- **any:** Any packet.
- **default inspection traffic**: The default TCP and UDP ports used by all applications that the security appliance can inspect. You can specify an ACL-based class along with the default inspection traffic class to narrow the matched traffic.
- **dscp**: A differentiated services code point (DSCP) value in the IP header defined by the Internet Engineering Task Force (IETF).
- **flow**: All traffic going to a unique IP destination address.

- **port**: Traffic using the TCP or UDP destination port or a contiguous range of ports.
- **precedence**: The precedence value represented by the Type of Service (ToS) byte in the IP header.
- **rtp:** Real-Time Transport Protocol (RTP) destination port.
- **tunnel-group**: VPN tunnel traffic. If you use this criterion, you can also configure the class to match a specific destination IP address within the tunnel group.

Class maps are assigned to policy maps.

Policy Maps

The class map determines what is matched, and the policy map associates one or more actions with a class of traffic.

The policy actions that can be configured are as follows:

- Forward the traffic flow to the Security Services Module (when present) for intrusion protection or content security and control services by creating an intrusion prevention system (IPS) or a content security and control (CSC) policy.
- Perform a specified protocol inspection or inspections by creating an inspection policy.
- Police the bandwidth used by the specified flow by creating a quality of service (QoS) police policy.
- Direct the flow to the low-latency queue by creating a QoS priority policy.
- Set connection parameters on the flows by creating a set connection policy.

Service Policies

The service policy activates a policy map on a targeted interface or globally on all interfaces. Service policies are represented as service policy rules in the ASDM.

To configure a service policy rule, you first need to navigate to **Firewall > Service Policy Rules**. Figure 33 displays the Service Policy rules screen.

Clicking **Add** launches the Add Service Policy Rule Wizard. Three steps to this wizard configure a service policy rule:

Step 1: Configure a service policy.

Step 2: Configure the traffic classification criteria for the service policy rule.

Step 3: Configure actions on the traffic classified by the service policy rule.

FIGURE 33
Service Policy Rules

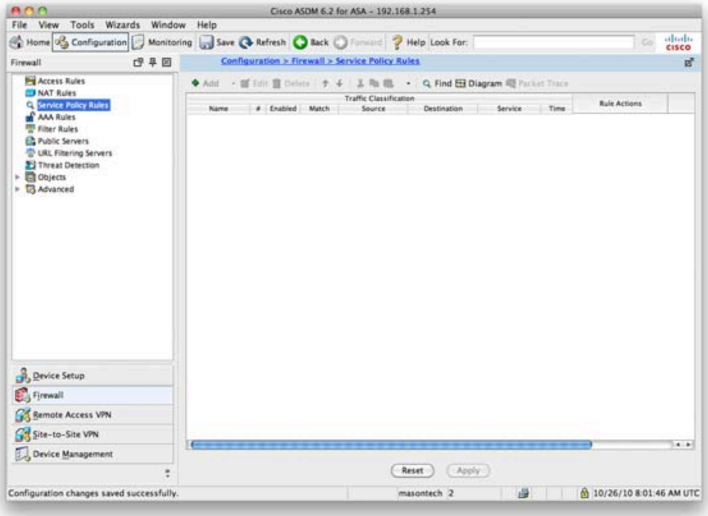
Cisco ASDM 6.2 for ASA - 192.168.1.254
File View Tools Wizards Window Help
Home Configuration Monitoring Save Refresh Back Forward Help Look For: Go
Firewall
Configuration > Firewall > Service Policy Rules
Access Rules
NAT Rules
Service Policy Rules
AAA Rules
Filter Rules
Public Servers
URL Filtering Servers
Threat Detection
Objects
Advanced
Add Edit Delete Find Diagram Packet Trace
Traffic Classification
Name # Enabled Match Source Destination Service Time
Rule Actions
Device Setup
Firewall
Remote Access VPN
Site-to-Site VPN
Device Management
Reset Apply
masontech 2
10/26/10 8:01:46 AM UTC

Configuration changes saved successfully.

Step 1: Configure a Service Policy

In Step 1, you need to give the service policy a name and either apply it to a specific interface or apply it globally, which applies the policy on all interfaces. You can also provide a description of the service policy. You can see the screen in Figure 34.

FIGURE 34
Service Policy Step 1

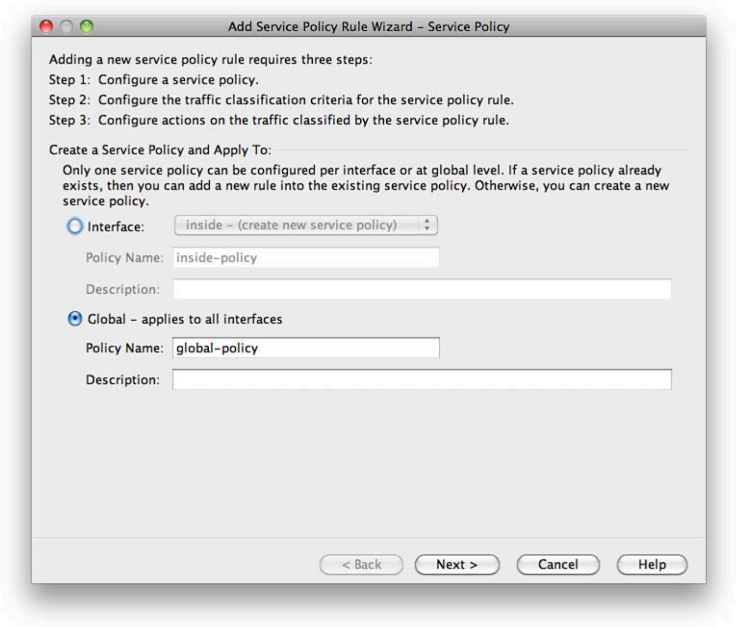

Add Service Policy Rule Wizard – Service Policy

Adding a new service policy rule requires three steps:
Step 1: Configure a service policy.
Step 2: Configure the traffic classification criteria for the service policy rule.
Step 3: Configure actions on the traffic classified by the service policy rule.

Create a Service Policy and Apply To:

Only one service policy can be configured per interface or at global level. If a service policy already exists, then you can add a new rule into the existing service policy. Otherwise, you can create a new service policy.

○ Interface: inside – (create new service policy)
Policy Name: inside–policy
Description:

● Global – applies to all interfaces
Policy Name: global–policy
Description:

Step 2: Configure the Traffic Classification Criteria for the Service Policy Rule

You are now asked to either create a new traffic class or use an existing traffic class. When creating a new traffic class, you must enter the name for the new traffic class and supply a description. You have the option to match traffic against the criteria covered earlier in this section about class maps. This is shown in Figure 35.

When you select one of the traffic-match criteria, the next screen is the configuration screen for that criterion. You chose Tunnel Group as the traffic-match criteria, and Figure 36 shows that you have the option now to select a tunnel group to match.

FIGURE 35
Traffic Classification

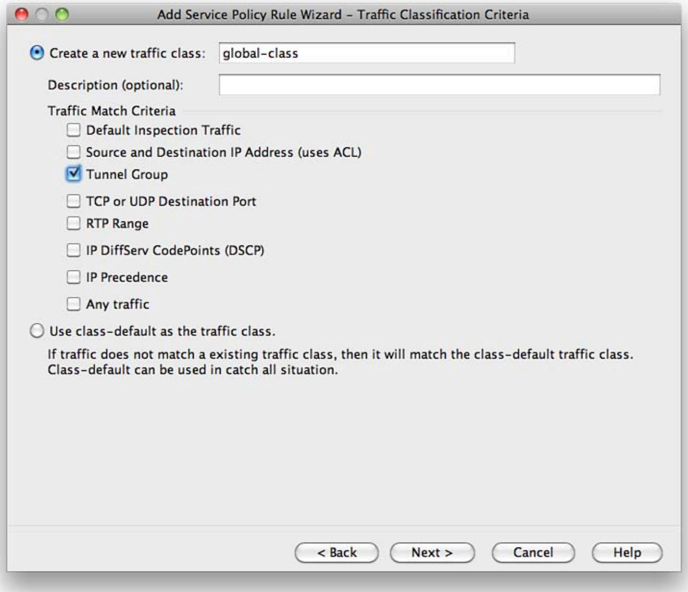

Add Service Policy Rule Wizard – Traffic Classification Criteria

● Create a new traffic class: global-class

Description (optional):

Traffic Match Criteria

- □ Default Inspection Traffic
- □ Source and Destination IP Address (uses ACL)
- ☑ Tunnel Group
- □ TCP or UDP Destination Port
- □ RTP Range
- □ IP DiffServ CodePoints (DSCP)
- □ IP Precedence
- □ Any traffic

○ Use class-default as the traffic class.

If traffic does not match a existing traffic class, then it will match the class-default traffic class. Class-default can be used in catch all situation.

FIGURE 36
Tunnel Group

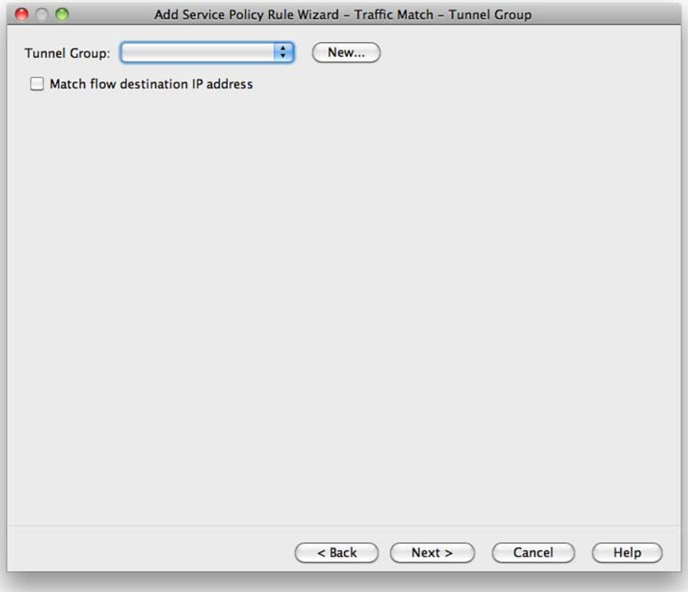

Add Service Policy Rule Wizard – Traffic Match – Tunnel Group

Tunnel Group:

□ Match flow destination IP address

Step 3: Configure Actions on the Traffic Classified by the Service Policy Rule

The next screen is the Rule Actions screen. Three tabs display at the top of the screen:

- Protocol Inspection
- Connection Settings
- QoS

The Protocols Inspection tab enables you to configure protocol-specific inspections if the traffic-match criteria allow it, as shown in Figure 37.

FIGURE 37
Protocol Inspection

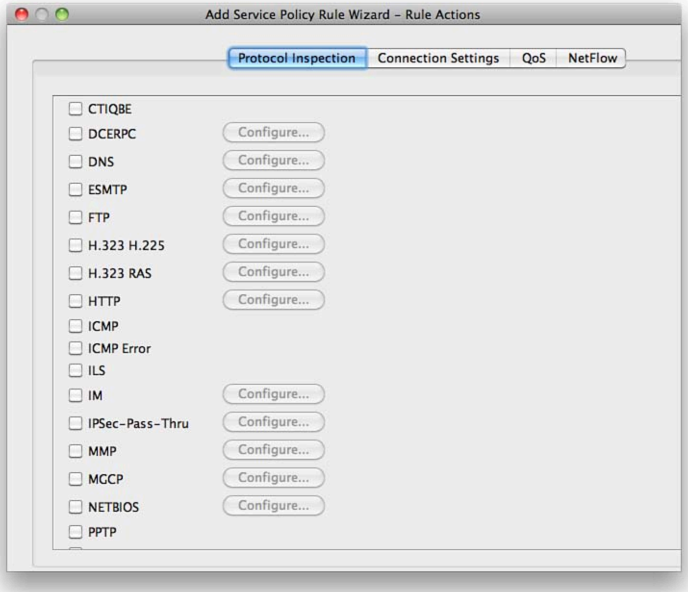

☐ CTIQBE
☐ DCERPC
☐ DNS
☐ ESMTP
☐ FTP
☐ H.323 H.225
☐ H.323 RAS
☐ HTTP
☐ ICMP
☐ ICMP Error
☐ ILS
☐ IM
☐ IPSec-Pass-Thru
☐ MMP
☐ MGCP
☐ NETBIOS
☐ PPTP

The Connection Settings tab enables you to set the maximum connections for TCP and UDP connections and the TCP timeout. You can also choose to randomize the TCP sequence number and enable TCP normalization, as shown in Figure 38.

FIGURE 38
Connection Settings

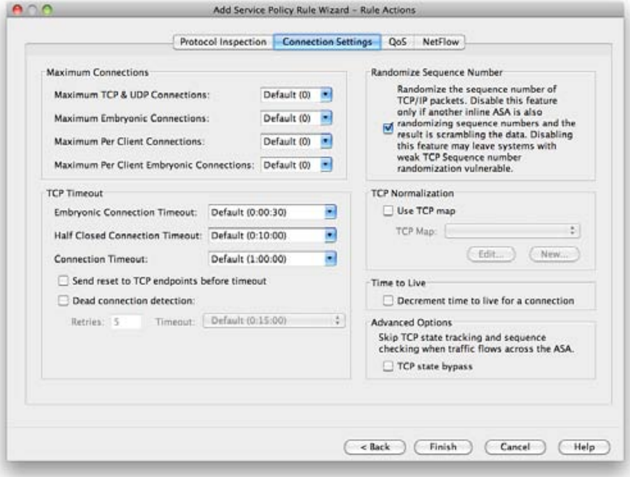

Maximum Connections

Maximum TCP & UDP Connections: Default (0)

Maximum Embryonic Connections: Default (0)

Maximum Per Client Connections: Default (0)

Maximum Per Client Embryonic Connections: Default (0)

Randomize Sequence Number

☑ Randomize the sequence number of TCP/IP packets. Disable this feature only if another inline ASA is also randomizing sequence numbers and the result is scrambling the data. Disabling this feature may leave systems with weak TCP Sequence number randomization vulnerable.

TCP Timeout

Embryonic Connection Timeout: Default (0:00:30)

Half Closed Connection Timeout: Default (0:10:00)

Connection Timeout: Default (1:00:00)

□ Send reset to TCP endpoints before timeout

□ Dead connection detection:

Retries: 5 Timeout: Default (0:15:00)

TCP Normalization

□ Use TCP map

TCP Map:

Time to Live

□ Decrement time to live for a connection

Advanced Options

Skip TCP state tracking and sequence checking when traffic flows across the ASA.

□ TCP state bypass

You can use the QoS tab to enable priority and policing for the traffic flow. When policing is selected, you can apply QoS settings to the flow to restrict the amount of bandwidth the flow is provided when traversing the interfaces of the firewall.

You can use this setting to reduce potential denial-of-service (DoS) attempts, because you can limit the amount of bandwidth allocated to a protocol, as shown in Figure 39.

FIGURE 39
QoS

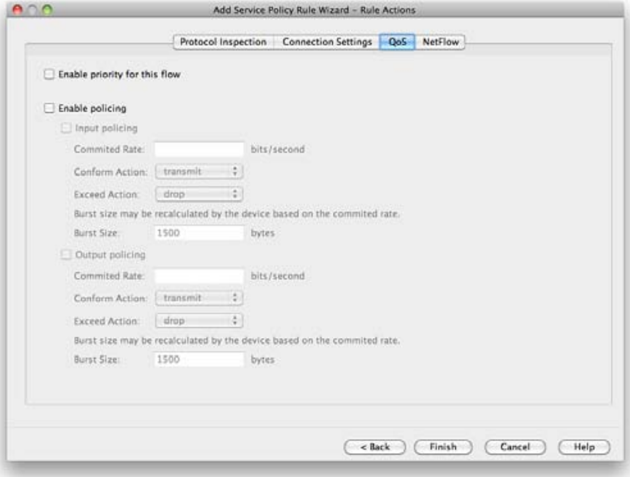

□ Enable priority for this flow

□ Enable policing

□ Input policing

Commited Rate: bits/second

Conform Action: transmit

Exceed Action: drop

Burst size may be recalculated by the device based on the commited rate.

Burst Size: 1500 bytes

□ Output policing

Commited Rate: bits/second

Conform Action: transmit

Exceed Action: drop

Burst size may be recalculated by the device based on the commited rate.

Burst Size: 1500 bytes

You click **Finish** to apply the service policy rule. It is added to the ASA when you click **Apply** from the main Service Policy Rules window on the ASDM.

Application Layer Policies

The Cisco ASA Application Inspection and Control (AIC) features provide advanced application layer (OSI Layers 5 to 7) filtering that provides a more granular level of control for modern day applications.

This advanced inspection helps to prevent malicious content from being delivered to endpoints protected by the ASA that would normally bypass traditional Layers 3 and 4 Access Control Lists. AIC inspection can be performed against protocols such as HTTP, FTP, DNS, ESMTP, and other common application protocols.

The following features are provided by Application Inspection and Control (AIC) on the Cisco ASA:

- **Protocol Minimization**: Enables a minimal required set of protocol features through the ASA
- **Payload Minimization**: Enables transport of minimally required payloads over the application session
- **Application Layer Signatures**: Enables and drops known malicious payloads in application layer sessions
- **Protocol Verification**: Detects and drops anomalous application layer protocol units

Configuring HTTP Inspection

Now take a look at configuring AIC for HTTP inspection. The ASA HTTP AIC inspection can granularly parse HTTP request and responses and enable specific value and regular expression matching against this traffic. The HTTP inspector also verifies adherence to the HTTP protocol and performs URL filtering and checking against several built in HTTP signatures.

For this example configure a HTTP protection policy that filters application layer traffic from the outside to the web server previously configured that needs protecting. Create a protection policy that verifies adherence to the HTTP protocol and enables only the HTTP GET method.

Following are two steps involved in creating the HTTP protection policy:

Step 1: Create an HTTP inspection policy map.

Step 2: Apply the HTTP inspection policy map.

Step 1: Create an HTTP Inspection Policy Map

You need to create the special HTTP inspection policy map to include all the inspection rules and their associated actions. The HTTP inspection policy map is the configuration container that contains a set of inspections that will be applied to a specific set of flows and is reusable in multiple traffic classes.

To create the HTTP inspection policy map using ASDM navigate to Configuration > Firewall > Objects > Inspect Maps > HTTP. This screen can be seen in figure 40.

FIGURE 40
HTTP Inspection
Policy Map

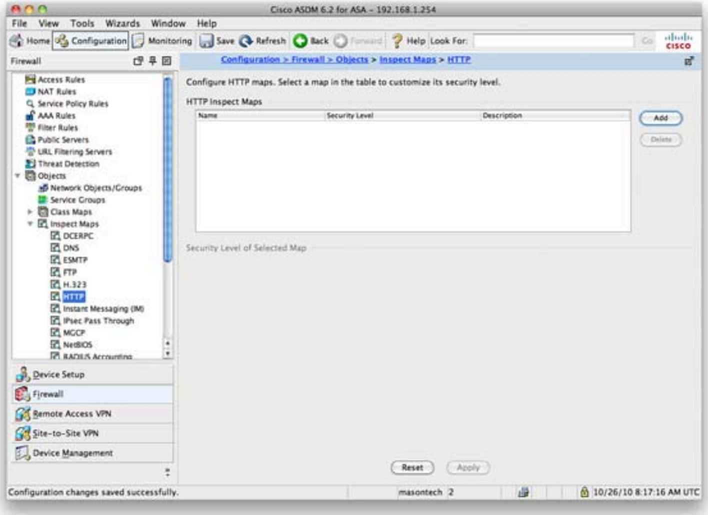

Configure HTTP maps. Select a map in the table to customize its security level.

HTTP Inspect Maps

Name	Security Level	Description

Configuration changes saved successfully.

Clicking Add opens a new HTTP inspection policy map. You can either use one of the built in Security Levels of Low, Medium, or High, or you can click the Details button to configure individual inspection.

For this example select the Details button to configure individual inspection for the HTTP inspection policy map.

Before you do, give the inspection policy map a name of **HTTP-Policy** and a description of **Example HTTP Inspection Policy Map**.

The first inspection to add is to implement adherence to the HTTP protocol. This is done from the Parameters tab selected by default. Ensure the Check For Protocol Violations radio button is selected, and also ensure that the traffic is getting dropped and logging is enabled.

The completed screen can be seen in Figure 41.

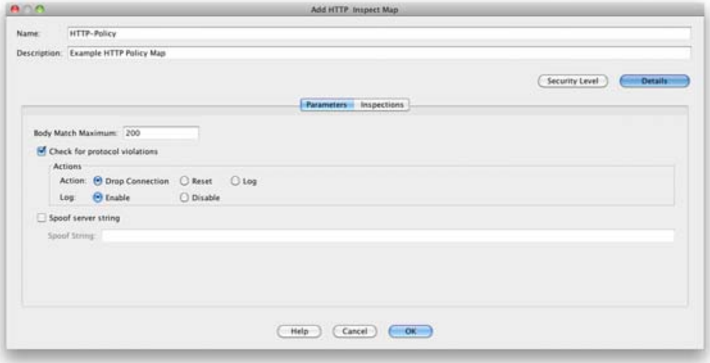

FIGURE 41
HTTP Protocol Adherence

Name: HTTP-Policy
Description: Example HTTP Policy Map

Body Match Maximum: 200
☑ Check for protocol violations
Actions
Action: ● Drop Connection ○ Reset ○ Log
Log: ● Enable ○ Disable
□ Spoof server string
Spoof String:

To configure the next inspection, select Inspections tab to start adding the manual inspections to the HTTP policy map.

Select Add and start to configure the inspection. The first inspection to add is to allow only the GET request method HTTP protocol to the server.

In the Match Criteria section, select the Single Match option to specify a single condition. In the Match Criteria section, specify No Match as the Match Type. This drops all traffic except traffic matching the specific criteria.

In the Criterion drop-down list, specify Request Method to filter traffic based on the HTTP request method. Select GET from the Method drop-down list.

In the Actions section, ensure that Drop Connection is chosen and ensure Logging is enabled.

The completed screen can be seen in Figure 42.

FIGURE 42
HTTP GET method

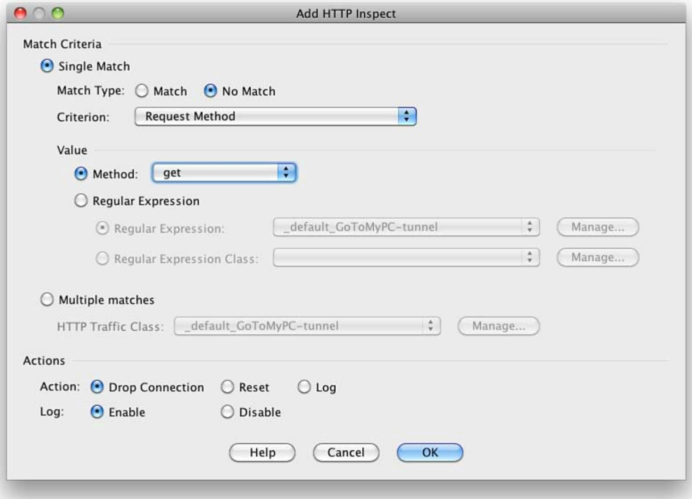

Add HTTP Inspect

Match Criteria

● Single Match

Match Type: ○ Match ● No Match

Criterion: Request Method

Value

● Method: get

○ Regular Expression

● Regular Expression: _default_GoToMyPC-tunnel

○ Regular Expression Class:

○ Multiple matches

HTTP Traffic Class: _default_GoToMyPC-tunnel

Actions

Action: ● Drop Connection ○ Reset ○ Log

Log: ● Enable ○ Disable

Click OK when the preceding is completed to return to the Inspection Policy Map screen.

You have now configured the HTTP inspection policy map to adhere to the HTTP protocol and enable only HTTP GET requests. The inspection policy is called HTTP-Policy and can be seen at the HTTP Inspect Maps screen, as shown in Figure 43.

FIGURE 43
HTTP Inspect Maps

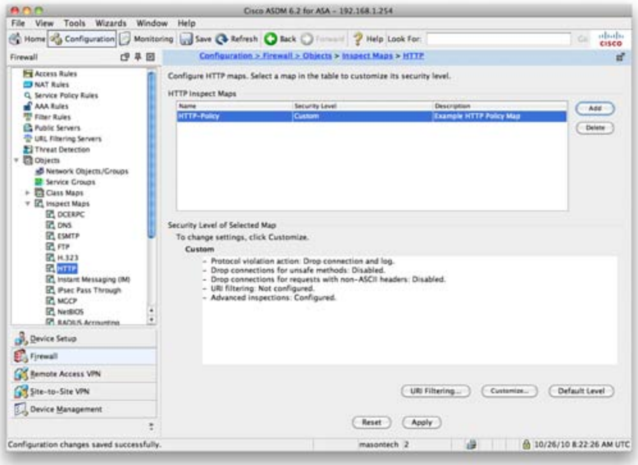

Configuration > Firewall > Objects > Inspect Maps > HTTP

Configure HTTP maps. Select a map in the table to customize its security level.

HTTP Inspect Maps

Name	Security Level	Description
HTTP-Policy	Custom	Example HTTP Policy Map

Security Level of Selected Map

To change settings, click Customize.

Custom

- Protocol violation action: Drop connection and log.
- Drop connections for unsafe methods: Disabled.
- Drop connections for requests with non-ASCII headers: Disabled.
- URI filtering: Not configured.
- Advanced inspections: Configured.

Configuration changes saved successfully. masontech 2 10/26/10 8:22:26 AM UTC

Step 2: Apply the HTTP Inspection Policy Map

In Step 1 you configured an HTTP Inspection Policy Map. The next step is to apply this policy map to the ASA so that traffic is inspected against the policy map and the correct action is taken.

The implementation of the HTTP inspection policy map is identical to the way that service policies are applied under the Modular Policy Framework covered earlier in this section.

When configuring the Service Policy rule, be sure to select the correct HTTP inspection policy map, as shown in Figure 44.

FIGURE 44
HTTP Inspection
Policy Map

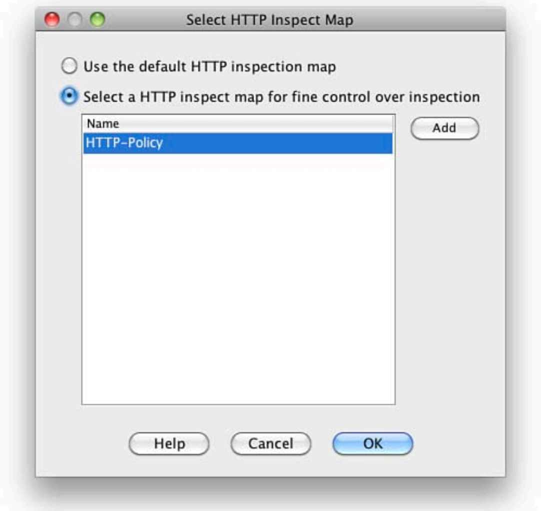

○ Use the default HTTP inspection map

● Select a HTTP inspect map for fine control over inspection

Name
HTTP-Policy

NOTE
Basic threat detection is enabled by default on the ASA. There is a minimal impact on performance when there are drops or potential threats on the ASA.

Threat Detection

Threat detection on the ASA is similar in operation to an IPS.

Two types of threat detection are available on the ASA:

- Basic threat detection
- Scanning threat detection

Basic threat detection is enabled by default. You can enable both basic and scanning threat detection independently of each other. One is not dependent on the other, and therefore you can have one, both, or neither configured on your ASA.

Basic Threat Detection

The security appliance basic threat detection feature provides threat-related drop statistics by monitoring the rate of dropped packets and security events per second (eps).

When the rate of dropped packets or security events exceeds established thresholds, basic threat detection generates a syslog message.

This enables you to detect activity that might be related to an attack, such as a DoS attack.

The ASA basic threat detection provides threat-related drop statistics by monitoring the following events:

- Access list denials
- Bad packet format
- Exceeded connection limits

NOTE
The scanning threat detection feature can significantly affect the performance and memory use of the ASA while it creates and gathers the host- and subnet-based data structure and information. Performance impact varies depending on the ASA platform.

- Detection of DoS attacks
- Failed basic firewall checks
- Detection of suspicious Internet Control Message Protocol (ICMP) packets
- Packets failing application inspection
- Interface overload
- Detection of scanning attacks
- Detection of incomplete sessions, such as TCP SYN attacks or no data UDP session attacks

The ASA tracks two types of rates for each monitored events: the average rate and burst rate. The average rate is the average rate over a time interval, and the burst rate is the one-tenth of the average rate or 10 seconds, whichever is the highest.

Syslog messages are generated when either of the rates for the monitored events is exceeded.

The following table shows the default threshold rates for basic threat detection.

Default Threshold Rates for Basic Threat Detection

Packet Drop Reason	Average Rate	Burst Rate
DoS attack detected Bad packet format Connection limits exceeded Suspicious ICMP packets	100 drops per second over the last 600 seconds	400 drops per second over the last 10-second period
Scanning attack detected	5 drops per second over the last 600 seconds	10 drops per second over the last 10-second period
Incomplete session	100 drops per second over the last 600 seconds	200 drops per second over the last 10-second period
Denial by access list	400 drops per second over the last 600 seconds	800 drops per second over the last 10-second period
Basic firewall checks failed Packet failed application inspection	400 drops per second over the last 600 seconds	1600 drops per second over the last 10-second period
Interface overload	2000 drops per second over the last 600 seconds	8000 drops per second over the last 10-second period

Basic threat detection is configured from the **Firewall > Threat Detection** screen. This is shown in Figure 45.

FIGURE 45
Threat Detection

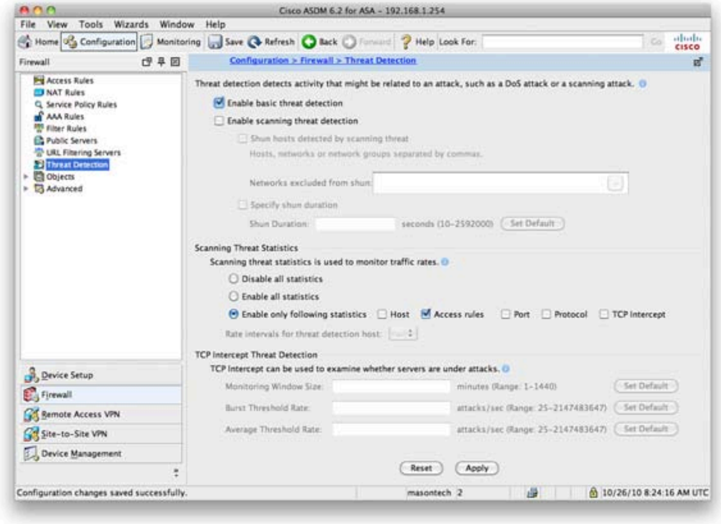

Threat detection detects activity that might be related to an attack, such as a DoS attack or a scanning attack.

☑ Enable basic threat detection

□ Enable scanning threat detection

□ Shun hosts detected by scanning threat

Hosts, networks or network groups separated by commas.

Networks excluded from shun:

□ Specify shun duration

Shun Duration: seconds (10-2592000)

Scanning Threat Statistics

Scanning threat statistics is used to monitor traffic rates.

○ Disable all statistics

○ Enable all statistics

● Enable only following statistics □ Host ☑ Access rules □ Port □ Protocol □ TCP Intercept

Rate intervals for threat detection host:

TCP Intercept Threat Detection

TCP Intercept can be used to examine whether servers are under attacks.

Monitoring Window Size: minutes (Range: 1-1440)

Burst Threshold Rate: attacks/sec (Range: 25-2147483647)

Average Threshold Rate: attacks/sec (Range: 25-2147483647)

Configuration changes saved successfully. masontech 2 10/26/10 8:24:16 AM UTC

You can see from Figure 45 that basic threat detection is enabled on this ASA. To disable it, uncheck the check box.

Tuning of the basic threat detection is performed in the CLI configuration with the **threat-detection** command. This is beyond the scope of the ASAF exam.

Scanning Threat Detection

The scanning threat detection feature of the ASA is concerned with hosts performing network scans against networks protected by the ASA.

Network reconnaissance scans, or port scans as they are commonly known, are normally a precursor to an attacker launching a full-blown attack on a system. The first step is normally to identify which ports and services are available on a system before enumerating and fingerprinting these ports to check for known vulnerabilities. A known vulnerability is always the preferred route in for attackers because they can use simple attack scripts to gain access and then escalate privileges.

Please see page 110 for more details.

When performing scanning threat detection, the ASA uses an extensive database of host statistics to generate syslog messages when a host is identified as either an attacker, or a target.

As with basic threat detection, scanning threat detection is configured from the **Firewall > Threat Detection** screen.

You have now enabled scanning threat detection and selected to shun hosts detected by scanning threat.

Figure 46 shows the Threat Detection configuration window with both basic threat detection and scanning threat detection enabled.

FIGURE 46
Threat Detection: Basic and Scanning

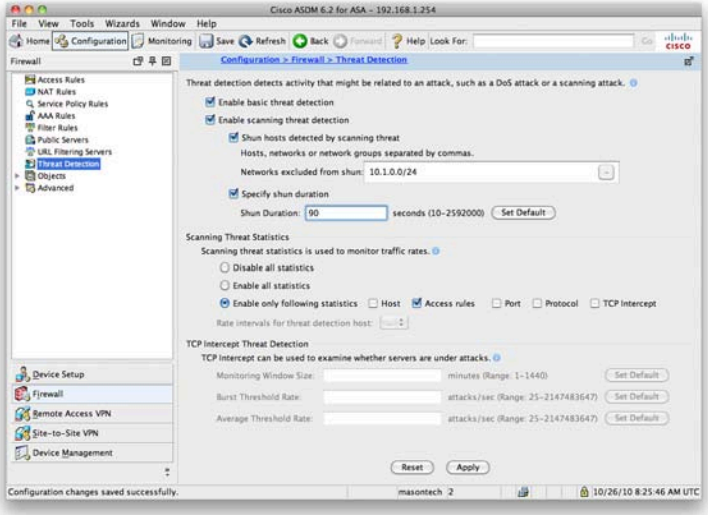

Threat detection detects activity that might be related to an attack, such as a DoS attack or a scanning attack.

☑ Enable basic threat detection
☑ Enable scanning threat detection
☑ Shun hosts detected by scanning threat
Hosts, networks or network groups separated by commas.
Networks excluded from shun: 10.1.0.0/24
☑ Specify shun duration
Shun Duration: 90 seconds (10-2592000)

Scanning Threat Statistics
Scanning threat statistics is used to monitor traffic rates.
○ Disable all statistics
○ Enable all statistics
● Enable only following statistics □ Host ☑ Access rules □ Port □ Protocol □ TCP Intercept
Rate intervals for threat detection host:

TCP Intercept Threat Detection
TCP Intercept can be used to examine whether servers are under attacks.
Monitoring Window Size: minutes (Range: 1-1440)
Burst Threshold Rate: attacks/sec (Range: 25-2147483647)
Average Threshold Rate: attacks/sec (Range: 25-2147483647)

Configuration changes saved successfully. masontech 2 10/26/10 8:25:46 AM UTC

When a shun is activated, all current connections from the malicious host are dropped, and all future connections are blocked at the outside interface of the ASA. Shuns are dynamic in nature and are not stored as a part of the configuration. If the security appliance loses power or reloads, any active shuns are lost.

You can specify a network, or network object, that will *not* be shunned. In the example, you set that 10.1.0.0/24 will not be shunned. This setting is useful for entering networks that should never be blocked, such as testing partners or third-party support organizations. You also set the ASA so that devices are shunned for 90 seconds.

Summary

This section expanded on the previous section and provided a simple network that addressed the interfaces, configured NAT for outbound access, and enabled a static translation to an internal web server. You also created an access rule that enabled inbound HTTP and HTTPS traffic to the web server using an object group rather than individual access control list entries.

You then looked at more advanced topics such as the Modular Policy Framework (MPF), Application Inspection Policies (AIP), and Threat Detection.

The next section looks at integrating the ASA into your network and covers routing, switching, and operating the firewall in Transparent mode.