



CHAPTER 26

IKE

IKE, also called ISAKMP, is the negotiation protocol that lets two hosts agree on how to build an IPsec security association. To configure the security appliance for virtual private networks, you set global IKE parameters that apply system wide, and you also create IKE policies that the peers negotiate to establish a VPN tunnel.

Certificate Group Matching

Certificate Group Matching lets you define rules to match a user's certificate to a permission group based on fields in the distinguished name (DN). You can use any field of the certificate or you can have all certificate users share a permission group.

To match user permission groups based on fields of the certificate, you must define rules that specify the fields to match for a group and then enable each rule for that selected group. A group must already exist in the configuration before you can create a rule for it. If the group does not exist in the configuration, you must define it by using **Configuration > VPN > General > Tunnel Group**.

Once you have defined rules, you must configure a certificate group matching policy to define the method to use for identifying the permission groups of certificate users: match the group from the rules, match the group from the OU field, or use a default group for all certificate users. You can use any or all of these methods.

Policy

A certificate group matching policy defines the method to use for identifying the permission groups of certificate users. You can use any or all of these methods.

Fields

- **Use the configured rules to match a certificate to a group**—Lets you use the rules you have defined under **Rules**.
- **Use the certificate OU field to determine the group**—Lets you use the organizational unit field to determine the group to which to match the certificate. This is selected by default.
- **Use the IKE identity to determine the group**—Lets you use the identity you previously defined under **Configuration > VPN > IKE > Global Parameters**. The IKE identity can be hostname, IP address, key ID, or automatic.
- **Use the peer IP address to determine the group**—Lets you use the peer's IP address. This is selected by default.

- **Default to group**—Lets you select a default group for certificate users that is used when none of the preceding methods resulted in a match. This is selected by default. Click the default group in the **Default to group** list. The group must already exist in the configuration. If the group does not appear in the list, you must define it by using **Configuration > VPN > General > Tunnel Group**.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Rules

The **Certificate Group Matching Rules** panel lets you add, edit, and delete rules to/from your configuration.

To match user permission groups based on fields of the certificate, you must define rules that specify which fields to match for a group and then enable each rule for that selected group. Rules cannot be longer than 255 characters. A group must already exist in the configuration before you can create a rule for it.

You can assign multiple rules to the same group. To do that, you add the rule priority and group first. Then you define as many criteria statements as you need for each group. When multiple rules are assigned to the same group, a match results for the first rule that tests true.

To match user permission groups based on multiple fields in the certificate so that all the criteria must match for the user to be assigned to a permission group, create a single rule with multiple matching criteria. To match user permission groups based on one criterion or another so that successfully matching any of the criteria identifies the member of the group, create multiple rules.

Fields

- **Map Name**—Displays the names of configured certificate-to-group maps.
- **Rule Priority**—Displays the importance of the rule as a matching criterion. The lower the value, the higher the priority. The default value is 10.
- **Mapped to Group**—Displays the group to which the rule is mapped.
- **Field**—Displays the type of distinguished name (Subject or Issuer) used in the rule.
- **Component**—Displays the distinguished name component used in the rule. For a list of possibilities and their definitions, see **Add Certificate Matching Rule Criterion Help**.
- **Operator**—Displays the operator used in the rule: Equals, Not Equals, Contains, and Does Not Contain.
- **Value**—Displays the value to be matched against.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	—	—

Add/Edit Certificate Matching Rule

Use the **Add/Edit Certificate Matching Rule** dialog box to define a certificate matching rule.

Fields

- **Map Existing**—Select the name of the map to include the rule.
- **Map New**—Enter a new map name for a rule.
- **Rule Priority**—Specify a number that indicates the level of priority for this rule. For the first rule defined, the default priority is 10. Each rule must have a unique priority. The lower the value, the higher the priority.
- **Mapped to Group**—Select the tunnel group to map to this rule. Groups must already exist in the configuration. If the group does not appear in the list, you must define it by using **Configuration > VPN > General > Tunnel Group**.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	—	—

Add/Edit Certificate Matching Rule Criterion

Use the **Add/Edit Certificate Matching Rule Criterion** dialog box to configure a certificate matching rule criterion for the selected group.

Fields

- **Field**—Specify the type of distinguished name to be used in the rule: Subject and Issuer, which consist of a specific-to-general identification hierarchy: CN, OU, O, L, SP, and C. These labels and acronyms conform to X.500 terminology.
 - **Subject**—The person or system that uses the certificate. For a CA root certificate, the Subject and Issuer are the same.
 - **Issuer**—The CA or other entity (jurisdiction) that issued the certificate.

- **Component**—Select the distinguished name component used in the rule:

DN Field	Definition
Whole Field	The entire DN.
Country (C)	The two-letter country abbreviation. These codes conform to ISO 3166 country abbreviations.
Common Name (CN)	The name of a person, system, or other entity. This is the lowest (most specific) level in the identification hierarchy.
DN Qualifier (DNQ)	A specific DN attribute.
E-mail Address (EA)	The e-mail address of the person, system or entity that owns the certificate.
Generational Qualifier (GENQ)	A generational qualifier such as Jr., Sr., or III.
Given Name (GN)	The first name of the certificate owner.
Initials (I)	The first letters of each part of the certificate owner's name.
Locality (L)	The city or town where the organization is located.
Name (N)	The name of the certificate owner.
Organization (O)	The name of the company, institution, agency, association, or other entity.
Organizational Unit (OU)	The subgroup within the organization.
Serial Number (SER)	The serial number of the certificate.
Surname (SN)	The family name or last name of the certificate owner.
State/Province (S/P)	The state or province where the organization is located.
Title (T)	The title of the certificate owner, such as Dr.
User ID (UID)	The identification number of the certificate owner.

- **Operator**—Select the operator used in the rule:
 - **Equals**—The distinguished name field must exactly match the value.
 - **Contains**—The distinguished name field must include the value within it.
 - **Does Not Equal**—The distinguished name field must not match the value
 - **Does Not Contain**—The distinguished name field must not include the value within it.
- **Value**—Specify the value to match against.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Here is some text marked print. Print is hidden.

Global Parameters

This panel lets you set system wide values for VPN tunnels. The following sections describe each of the options.

Enabling IKE on Interfaces

You must enable IKE for each interface that you want to use for VPN tunnels.

Enabling IPsec over NAT-T

NAT-T lets IPsec peers establish both remote access and LAN-to-LAN connections through a NAT device. It does this by encapsulating IPsec traffic in UDP datagrams, using port 4500, thereby providing NAT devices with port information. NAT-T auto-detects any NAT devices, and only encapsulates IPsec traffic when necessary. This feature is disabled by default.

- The security appliance can simultaneously support standard IPsec, IPsec over TCP, NAT-T, and IPsec over UDP, depending on the client with which it is exchanging data.
- When both NAT-T and IPsec over UDP are enabled, NAT-T takes precedence.
- When enabled, IPsec over TCP takes precedence over all other connection methods.

The security appliance implementation of NAT-T supports IPsec peers behind a single NAT/PAT device as follows:

- One LAN-to-LAN connection.
- Either a LAN-to-LAN connection or multiple remote access clients, but not a mixture of both.

To use NAT-T you must:

- Open port 4500 on the security appliance.
- Enable IPsec over NAT-T globally in this panel.
- Select the second or third option for the Fragmentation Policy parameter in the **Configuration > VPN > IPsec > Pre-Fragmentation** panel. These options let traffic travel across NAT devices that do not support IP fragmentation; they do not impede the operation of NAT devices that do support IP fragmentation.

Enabling IPsec over TCP

IPsec over TCP enables a VPN client to operate in an environment in which standard ESP or IKE cannot function, or can function only with modification to existing firewall rules. IPsec over TCP encapsulates both the IKE and IPsec protocols within a TCP packet, and enables secure tunneling through both NAT and PAT devices and firewalls. This feature is disabled by default.



Note

This feature does not work with proxy-based firewalls.

IPsec over TCP works with remote access clients. It works on all physical and VLAN interfaces. It is a client to security appliance feature only. It does not work for LAN-to-LAN connections.

- The security appliance can simultaneously support standard IPsec, IPsec over TCP, NAT-Traversal, and IPsec over UDP, depending on the client with which it is exchanging data.
- The VPN 3002 hardware client, which supports one tunnel at a time, can connect using standard IPsec, IPsec over TCP, NAT-Traversal, or IPsec over UDP.
- When enabled, IPsec over TCP takes precedence over all other connection methods.

You enable IPsec over TCP on both the security appliance and the client to which it connects.

You can enable IPsec over TCP for up to 10 ports that you specify. If you enter a well-known port, for example port 80 (HTTP) or port 443 (HTTPS), the system displays a warning that the protocol associated with that port will no longer work. The consequence is that you can no longer use a browser to manage the security appliance through the IKE-enabled interface. To solve this problem, reconfigure the HTTP/HTTPS management to different ports.

You must configure TCP port(s) on the client as well as on the security appliance. The client configuration must include at least one of the ports you set for the security appliance.

Determining ID Method

During IKE negotiations the peers must identify themselves to each other. You can choose the identification methods from the following options:

Address	Uses the IP addresses of the hosts exchanging ISAKMP identity information.
Hostname	Uses the fully-qualified domain name of the hosts exchanging ISAKMP identity information (default). This name comprises the hostname and the domain name.
Key ID	Uses the string the remote peer uses to look up the preshared key.
Automatic	Determines IKE negotiation by connection type: <ul style="list-style-type: none"> • IP address for preshared key • Cert DN for certificate authentication.

Disabling Inbound Aggressive Mode Connections

Phase 1 IKE negotiations can use either Main mode or Aggressive mode. Both provide the same services, but Aggressive mode requires only two exchanges between the peers, rather than three. Aggressive mode is faster, but does not provide identity protection for the communicating parties. It is therefore necessary that they exchange identification information prior to establishing a secure SA in which to encrypt information. This feature is disabled by default.

Alerting Peers Before Disconnecting

Client or LAN-to-LAN sessions may be dropped for several reasons, such as: a security appliance shutdown or reboot, session idle timeout, maximum connection time exceeded, or administrator cut-off.

The security appliance can notify qualified peers (in LAN-to-LAN configurations), VPN Clients and VPN 3002 Hardware Clients of sessions that are about to be disconnected, and it conveys to them the reason. The peer or client receiving the alert decodes the reason and displays it in the event log or in a pop-up panel. This feature is disabled by default.

This panel lets you enable the feature so that the security appliance sends these alerts, and conveys the reason for the disconnect.

Qualified clients and peers include the following:

- Security appliance devices with Alerts enabled.
- VPN clients running 4.0 or later software (no configuration required).
- VPN 3002 hardware clients running 4.0 or later software, and with Alerts enabled.
- VPN 3000 Series Concentrators running 4.0 or later software, with Alerts enabled.

Waiting for Active Sessions to Terminate Prior to Reboot

You can schedule a security appliance reboot to occur only when all active sessions have terminated voluntarily. This feature is disabled by default.

Fields

- **Enable IKE**—Shows IKE status for all configured interfaces.
 - **Interface**—Displays names of all configured security appliance interfaces.
 - **IKE Enabled**—Shows whether IKE is enabled for each configured interface.
 - **Enable/Disable**—Click to enable or disable IKE for the highlighted interface.
- **NAT Transparency**—Lets you enable or disable IPsec over NAT-T and IPsec over TCP.
 - **Enable IPsec over NAT-T**—Select to enable IPsec over NAT-T.
 - **NAT Keepalive**—Type the number of seconds that can elapse with no traffic before the security appliance terminates the NAT-T session. The default is 20 seconds. The range is 10 to 3600 seconds (one hour).
 - **Enable IPsec over TCP**—Select to enable IPsec over TCP.
 - **Enter up to 10 comma-separated TCP port values**—Type up to 10 ports on which to enable IPsec over TCP. Use a comma to separate the ports. You do not need to use spaces. The default port is 10,000. The range is 1 to 65,635.
- **Identity to Be Sent to Peer**—Lets you set the way that IPsec peers identify themselves to each other.
 - **Identity**—Select one of the following methods by which IPsec peers identify themselves:

Address	Uses the IP addresses of the hosts.
Hostname	Uses the fully-qualified domain names of the hosts. This name comprises the hostname and the domain name.
Key ID	Uses the string the remote peer uses to look up the preshared key.
Automatic	Determines IKE negotiation by connection type: IP address for preshared key or cert DN for certificate authentication.

- **Key Id String**—Type the alpha-numeric string the peers use to look up the preshared key.
- **Disable inbound aggressive mode connections**—Select to disable aggressive mode connections.
- **Alert peers before disconnecting**—Select to have the security appliance notify qualified LAN-to-LAN peers and remote access clients before disconnecting sessions.
- **Wait for all active sessions to voluntarily terminate before rebooting**—Select to have the security appliance postpone a scheduled reboot until all active sessions terminate.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Policies

Each IKE negotiation is divided into two sections called Phase1 and Phase 2.

Phase 1 creates the first tunnel, which protects later IKE negotiation messages. Phase 2 creates the tunnel that protects data.

To set the terms of the IKE negotiations, you create one or more IKE policies, which include the following:

- A unique priority (1 through 65,543, with 1 the highest priority).
- An authentication method, to ensure the identity of the peers.
- An encryption method, to protect the data and ensure privacy.
- An HMAC method to ensure the identity of the sender, and to ensure that the message has not been modified in transit.
- A Diffie-Hellman group to establish the strength of the of the encryption-key-determination algorithm. The security appliance uses this algorithm to derive the encryption and hash keys.
- A limit for how long the security appliance uses an encryption key before replacing it.

If you do not configure any IKE policies, the security appliance uses the default policy, which is always set to the lowest priority, and which contains the e default value for each parameter. If you do not specify a value for a specific parameter, the default value takes effect.

When IKE negotiation begins, the peer that initiates the negotiation sends all of its policies to the remote peer, and the remote peer searches for a match with its own policies, in priority order.

A match between IKE policies exists if they have the same encryption, hash, authentication, and Diffie-Hellman values, and an SA lifetime less than or equal to the lifetime in the policy sent. If the lifetimes are not identical, the shorter lifetime—from the remote peer policy—applies. If no match exists, IKE refuses negotiation and the IKE SA is not established.

Fields

- **Policies**—Displays parameter settings for each configured IKE policy.
 - **Priority #**—Shows the priority of the policy.
 - **Encryption**—Shows the encryption method.
 - **Hash**—Shows the has algorithm.
 - **D-H Group**—Shows the Diffie-Hellman group.
 - **Authentication**—Shows the authentication method.
 - **Lifetime (secs)**—Shows the SA lifetime in seconds.
- **Add/Edit/Delete**—Click to add, edit, or delete an IKE policy.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Add/Edit IKE Policy

Fields

Priority #—Type a number to set a priority for the IKE policy. The range is 1 to 65,543, with 1 the highest priority.

Encryption—Select an encryption method. This is a symmetric encryption method that protects data transmitted between two IPsec peers. The choices follow:

des	56-bit DES-CBC. Less secure but faster than the alternatives. The default.
3des	168-bit Triple DES.
aes	128-bit AES.
aes-192	192-bit AES.
aes-256	256-bit AES.

Hash—Select the hash algorithm that ensures data integrity. It ensures that a packet comes from whom you think it comes from, and that it has not been modified in transit.

sha	SHA-1	The default is SHA-1. MD5 has a smaller digest and is considered to be slightly faster than SHA-1. A successful (but extremely difficult) attack against MD5 has occurred; however, the HMAC variant IKE uses prevents this attack.
md5	MD5	

Authentication—Select the authentication method the security appliance uses to establish the identity of each IPsec peer. Pre-shared keys do not scale well with a growing network but are easier to set up in a small network. The choices follow:

pre-share	Pre-shared keys.
rsa-sig	A digital certificate with keys generated by the RSA signatures algorithm.
crack	IKE Challenge/Response for Authenticated Cryptographic Keys protocol for mobile IPsec-enabled clients which use authentication techniques other than certificates.

D-H Group—Select the Diffie-Hellman group identifier, which the two IPsec peers use to derive a shared secret without transmitting it to each other.

1	Group 1 (768-bit)	The default, Group 2 (1024-bit Diffie-Hellman) requires less CPU time to execute but is less secure than Group 2 or 5.
2	Group 2 (1024-bit)	
5	Group 5 (1536-bit)	
7	Group 7 (Elliptical curve field size is 163 bits.)	Group 7 is for use with the Movian VPN client, but with any peer that supports Group 7 (ECC).

Lifetime (secs)—Either select Unlimited or type an integer for the SA lifetime. The default is 86,400 seconds or 24 hours. With longer lifetimes, the security appliance sets up future IPsec security associations more quickly. Encryption strength is great enough to ensure security without using very fast rekey times, on the order of every few minutes. We recommend that you accept the default.

Time Measure—Select a time measure. The security appliance accepts the following values:

- 120 - 86,400 seconds
- 2 - 1440 minutes
- 1 - 24 hours
- 1 day

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

IP Address Management

IP addresses make internetwork connections possible. They are like telephone numbers: both the sender and receiver must have an assigned number to connect. But with VPNs, there are actually two sets of addresses: the first set connects client and server on the public network; and once that connection is made, the second set connects client and server through the VPN tunnel.

In security appliance address management, we are dealing with the second set of IP addresses: those private IP addresses that connect a client with a resource on the private network, through the tunnel, and let the client function as if it were directly connected to the private network. Furthermore, we are dealing only with the private IP addresses that get assigned to clients. The IP addresses assigned to other resources on your private network are part of your network administration responsibilities, not part of security appliance management.

Therefore, when we discuss IP addresses here, we mean those IP addresses available in your private network addressing scheme, that let the client function as a tunnel endpoint.

Assignment

The Assignment panel lets you choose a way to assign IP addresses to remote access clients.

Fields

- **Use authentication server**—Select to assign IP addresses retrieved from an authentication server on a per-user basis. If you are using an authentication server (external or internal) that has IP addresses configured, we recommend using this method. Configure AAA servers on the **Configuration > Properties > AAA Setup > AAA Servers** and **AAA Server Group** panels.
- **Use DHCP**— Select to obtain IP addresses from a DHCP server. If you use DHCP, configure the server on the **Configuration > Properties > DHCP Services > DHCP Server** panel.

- **Use internal address pools**—Select to have the security appliance assign IP addresses from an internally configured pool. Internally configured address pools are the easiest method of address pool assignment to configure. If you use this method, configure the IP address pools on **Configuration > VPN > IP Address Management > IP Pools** panel.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

IP Pools

The IP Pool box shows each configured address pool by name, and with their IP address range, for example: 10.10.147.100 to 10.10.147.177. If no pools exist, the box is empty. The security appliance uses these pools in the order listed: if all addresses in the first pool have been assigned, it uses the next pool, and so on.

If you assign addresses from a non-local subnet, we suggest that you add pools that fall on subnet boundaries to make adding routes for these networks easier.

Fields

- **Pool Name**—Displays the name of each configured address pool.
- **Starting Address**—Shows first IP address available in each configured pool.
- **Ending Address**—Shows the last IP address available in each configured pool.
- **Subnet Mask**—Shows the subnet mask for addresses in each configured pool.
- **Add**—Click to add a new address pool.
- **Edit/Delete**—Click to edit or delete an already configured address pool.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Add/Edit IP Pool

These panels let you:

- Add a new pool of IP addresses from which the security appliance assigns addresses to clients.
- Modify an IP address pool that you have previously configured.

The IP addresses in the pool range must not be assigned to other network resources.

Fields

- **Name**—Assign an alpha-numeric name to the address pool. Limit 64 characters
- **Starting IP Address**—Enter the first IP address available in this pool. Use dotted decimal notation, for example: 10.10.147.100.
- **Ending IP Address**—Enter the last IP address available in this pool. Use dotted decimal notation, for example: 10.10.147.100.
- **Subnet Mask**—Select the subnet mask for the IP address pool.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

IPsec

The security appliance uses IPsec for LAN-to-LAN VPN connections, and provides the option of using IPsec for client-to-LAN VPN connections. In IPsec terminology, a “peer” is a remote-access client or another secure gateway.



Note

The ASA supports LAN-to-LAN IPsec connections with Cisco peers, and with third-party peers that comply with all relevant standards.

During tunnel establishment, the two peers negotiate security associations that govern authentication, encryption, encapsulation, and key management. These negotiations involve two phases: first, to establish the tunnel (the IKE SA); and second, to govern traffic within the tunnel (the IPsec SA).

A LAN-to-LAN VPN connects networks in different geographic locations. In IPsec LAN-to-LAN connections, the security appliance can function as initiator or responder. In IPsec client-to-LAN connections, the security appliance functions only as responder. Initiators propose SAs; responders accept, reject, or make counter-proposals—all in accordance with configured SA parameters. To establish a connection, both entities must agree on the SAs.

This security appliance supports these IPsec attributes:

- Main mode for negotiating phase one ISAKMP security associations when using digital certificates for authentication
- Aggressive mode for negotiating phase one ISAKMP Security Associations (SAs) when using preshared keys for authentication
- Authentication Algorithms:
 - ESP-MD5-HMAC-128
 - ESP-SHA1-HMAC-160

- Authentication Modes:
 - Preshared Keys
 - X.509 Digital Certificates
- Diffie-Hellman Groups 1, 2, 5, and 7
- Encryption Algorithms:
 - AES-128, -192, and -256
 - 3DES-168
 - DES-56
 - ESP-NULL
- Extended Authentication (XAuth)
- Mode Configuration (also known as ISAKMP Configuration Method)
- Tunnel Encapsulation Mode
- IP compression (IPCOMP) using LZS

IPsec Rules

This pane shows the currently configured IPsec rules. Use it to add, edit, delete and move up, move down, cut, copy, and paste an IPsec rule.

Fields



Note

You cannot edit, delete, or copy an implicit rule. The security appliance implicitly accepts the traffic selection proposal from remote clients when configured with a dynamic tunnel policy. You can override it by giving a specific traffic selection.

- **Type: Priority**—Displays the type of rule (static or dynamic) and its priority.
- **Traffic Selection**
 - **#**—Indicates the rule number.
 - **Source**—Indicates the IP addresses that are subject to this rule when traffic is sent to the IP addresses listed in the **Remote Side Host/Network** column. In detail mode (see the **Show Detail** button), an address column might contain an interface name with the word **any**, such as **inside:any**. **any** means that any host on the inside interface is affected by the rule.
 - **Destination**—Lists the IP addresses that are subject to this rule when traffic is sent from the IP addresses listed in the **Security Appliance Side Host/Network** column. In detail mode (see the **Show Detail** button), an address column might contain an interface name with the word **any**, such as **outside:any**. **any** means that any host on the outside interface is affected by the rule. Also in detail mode, an address column might contain IP addresses in square brackets, for example, [209.165.201.1-209.165.201.30]. These addresses are translated addresses. When an inside host makes a connection to an outside host, the security appliance maps the inside host's address to an address from the pool. After a host creates an outbound connection, the security appliance maintains this address mapping. This address mapping structure is called an xlate, and remains in memory for a period of time.
 - **Service**—Specifies the service and protocol specified by the rule (TCP, UDP, ICMP, or IP).

- **Action**—Specifies the type of IPsec rule (protect or do not protect).
- **Transform Set**—Displays the transform set for the rule.
- **Peer**—Identifies the IPsec peer.
- **PFS**—Displays Perfect Forward Secrecy settings for the rule.
- **NAT-T Enabled**—Indicates whether NAT Traversal is enabled for the policy.
- **Reverse Route Enabled**—Indicates whether Reverse Route Injection is enabled for the policy.
- **Connection Type**—(Meaningful only for static tunnel policies.) Identifies the connection type for this policy as bidirectional, originate-only, or answer-only).
- **SA Lifetime**—Displays the SA lifetime for the rule.
- **Trustpoint**—Displays the trust point for the policy. This applies to static connections only.
 - **Chain Enabled**—Displays whether the policy transmits the entire trust point chain.
- **IKE Negotiation Mode**—Displays whether IKE negotiations use main or aggressive mode.
- **Description**—(Optional) Specifies a brief description for this rule. For an existing rule, this is the description you typed when you added the rule. An implicit rule includes the following description: “Implicit rule.” To edit the description of any but an implicit rule, right-click this column, and choose Edit Description or double-click the column.
- **Enable Anti-replay window size**—Sets the anti-replay window size, between 64 and 1028 in multiples of 64. One side-effect of priority queueing in a hierarchical QoS policy with traffic shaping (see the [“Rule Actions > QoS Tab”](#) section on page 21-26) is packet re-ordering. For IPsec packets, out-of-order packets that are not within the anti-replay window generate warning syslog messages. These warnings becomes false alarms in the case of priority queueing. Configuring the anti-replay window size helps you avoid possible false alarms.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Tunnel Policy (Crypto Map) - Basic

Use this pane to define a new Tunnel Policy for an IPsec rule. The values you define here appear in the IPsec Rules table after you click OK. All rules are enabled by default as soon as they appear in the IPsec Rules table.

The Tunnel Policy panel lets you define a tunnel policy that is used to negotiate an IPsec (Phase 2) security association (SA). ASDM captures your configuration edits, but does not save them to the running configuration until you click Apply.

Every tunnel policy must specify a transform set and identify the security appliance interface to which it applies. The transform set identifies the encryption and hash algorithms that perform IPsec encryption and decryption operations. Because not every IPsec peer supports the same algorithms, you might want to specify a number of policies and assign a priority to each. The security appliance then negotiates with the remote IPsec peer to agree on a transform set that both peers support.

Tunnel policies can be *static* or *dynamic*. A static tunnel policy identifies one or more remote IPsec peers or subnetworks to which your security appliance permits IPsec connections. A static policy can be used whether your security appliance initiates the connection or receives a connection request from a remote host. A static policy requires you to enter the information necessary to identify permitted hosts or networks.

A dynamic tunnel policy is used when you cannot or do not want to provide information about remote hosts that are permitted to initiate a connection with the security appliance. If you are only using your security appliance as a VPN client in relation to a remote VPN central-site device, you do not need to configure any dynamic tunnel policies. Dynamic tunnel policies are most useful for allowing remote access clients to initiate a connection to your network through a security appliance acting as the VPN central-site device. A dynamic tunnel policy is useful when the remote access clients have dynamically assigned IP addresses or when you do not want to configure separate policies for a large number of remote access clients.

Fields

- **Interface**—Select the interface name to which this policy applies.
- **Policy Type**—Select the type, static or dynamic, of this tunnel policy.
- **Priority**—Enter the priority of the policy.
- **Transform Set to Be Added**—Select the transform set for the policy and click **Add** to move it to the list of active transform sets. Click **Move Up** or **Move Down** to rearrange the order of the transform sets in the list box. You can add a maximum of 11 transform sets to a crypto map entry or a dynamic crypto map entry.
- **Peer Settings - Optional for Dynamic Crypto Map Entries**—Configure the peer settings for the policy.
 - **Connection Type**—(Meaningful only for static tunnel policies.) Select bidirectional, originate-only, or answer-only to specify the connection type of this policy. For LAN-to-LAN connections, select bidirectional or answer-only (not originate-only). Select answer-only for LAN-to-LAN redundancy.
 - **IP Address of Peer to Be Added**—Enter the IP address of the IPsec peer you are adding.
- **Enable Perfect Forward Secrecy**—Check to enable Perfect Forward Secrecy for the policy. PFS is a cryptographic concept where each new key is unrelated to any previous key. In IPsec negotiations, Phase 2 keys are based on Phase 1 keys unless you specify Perfect Forward Secrecy.
- **Diffie-Hellman Group**—When you enable PFS you must also select a Diffie-Hellman group which the security appliance uses to generate session keys. The choices are as follows:
 - Group 1 (768-bits) = Use Perfect Forward Secrecy, and use Diffie-Hellman Group 1 to generate IPsec session keys, where the prime and generator numbers are 768 bits. This option is more secure but requires more processing overhead.
 - Group 2 (1024-bits) = Use Perfect Forward Secrecy, and use Diffie-Hellman Group 2 to generate IPsec session keys, where the prime and generator numbers are 1024 bits. This option is more secure than Group 1 but requires more processing overhead.
 - Group 5 (1536-bits)
 - Group 7 (ECC) = Use Perfect Forward Secrecy, and use Diffie-Hellman Group 7 (ECC) to generate IPsec session keys, where the elliptic curve field size is 163 bits. This option is the fastest and requires the least overhead. It is intended for use with the Movian VPN client, but you can use it with any peers that support Group 7 (ECC).

Tunnel Policy (Crypto Map) - Advanced

Fields

- **Security Association Lifetime** parameters—Configures the duration of a Security Association (SA). This parameter specifies how to measure the lifetime of the IPsec SA keys, which is how long the IPsec SA lasts until it expires and must be renegotiated with new keys.
 - **Time**—Specifies the SA lifetime in terms of hours (hh), minutes (mm) and seconds (ss).
 - **Traffic Volume**—Defines the SA lifetime in terms of kilobytes of traffic. Enter the number of kilobytes of payload data after which the IPsec SA expires. Minimum is 100 KB, default is 10000 KB, maximum is 2147483647 KB.
- **Enable NAT-T**— Enables NAT Traversal (NAT-T) for this policy.
- **Enable Reverse Route Injection**—Enables Reverse Route Injection for this policy.
- **Static Type Only Settings**—Specifies parameters for static tunnel policies.
 - **Trust Point Name**—Selects the trust point to use. If you select something other than None (Use Preshared Keys), which is the default, the Enable entire chain transmission check box becomes active.
 - **Enable entire chain transmission**—Enables transmission of the entire trust point chain.
 - **IKE Negotiation Mode**—Selects the IKE negotiation mode, Main or Aggressive. This parameter sets the mode for exchanging key information and setting up the SAs. It sets the mode that the initiator of the negotiation uses; the responder auto-negotiates. Aggressive Mode is faster, using fewer packets and fewer exchanges, but it does not protect the identity of the communicating parties. Main Mode is slower, using more packets and more exchanges, but it protects the identities of the communicating parties. This mode is more secure and it is the default selection. If you select Aggressive, the Diffie-Hellman Group list becomes active.
 - **Diffie-Hellman Group**—Select the Diffie-Hellman group to apply. The choices are as follows: Group 1 (768-bits), Group 2 (1024-bits), Group 5 (1536-bits), Group 7 (ECC).

Tunnel Policy (Crypto Map) -Traffic Selection

This pane lets you define what traffic to protect.

Fields

- **Interface and Action**
 - **Interface**—Identify the interface for the rule.
 - **Action**—Specify the action for this rule to take. The selections are protect and do not protect.
- **Source and Destination**—Specify the IP address, network object group or interface IP address for the source host or network. A rule cannot use the same address as both the source and destination.
 - **IP Address**—Indicates that the parameters that follow specify the interface, IP address, and subnet mask of the source host or network.
 - **Name**—Indicates that the parameters that follow specify the name of the source host or network.
 - **Group**—Indicates that the parameters that follow specify the interface and group name of the source host or network.
 - **Interface**—Selects the interface name for the IP address. This parameter appears when you select the IP Address option button.

- **IP address**—Specifies the IP address of the interface to which this policy applies. This parameter appears when you select the IP Address option button.
- **...** —Displays the Select host/network panel, on which you can select an interface and display and select the associated hosts. Your selection appears in the Source Host/Network IP address and Mask boxes on the Add Rule panel. This selection also fills in the Mask field. This parameter appears when you select the IP Address option button.
- **Mask**—Selects a standard subnet mask to apply to the IP address. This parameter appears when you select the IP Address option button.
- **Name**—Selects the interface name to use as the source or destination host or network. This parameter appears when you select the Name option button. This is the only parameter associated with this option.
- **Interface**—Selects the interface name for the IP address. This parameter appears when you select the Group option button.
- **Group**—Selects the name of the group on the specified interface for the source or destination host or network. If the list contains no entries, you can enter the name of an existing group. This parameter appears when you select the Group option button.
- **Rule Flow Diagram**—Shows the results of your selections for the source and destination host/network group boxes.
- **Protocol and Service**—Specifies protocol and service parameters relevant to this rule.

**Note**

“Any - any” IPsec rules are not allowed. This type of rule would prevent the device and its peer from supporting multiple LAN -to-LAN tunnels.

- **TCP**—Specifies that this rule applies to TCP connections. This selection also displays the **Source Port** and **Destination Port** group boxes.
- **UDP**—Specifies that this rule applies to UDP connections. This selection also displays the **Source Port** and **Destination Port** group boxes.
- **ICMP**—Specifies that this rule applies to ICMP connections. This selection also displays the **ICMP Type** group box.
- **IP**—Specifies that this rule applies to IP connections. This selection also displays the **IP Protocol** group box.
- **Manage Service Groups**—Displays the Manage Service Groups panel, on which you can add, edit, or delete a group of TCP/UDP services/ports.
- **Source Port** and **Destination Port** —Contains TCP or UDP port parameters, depending on which option button you selected in the Protocol and Service group box.
- **Service**—Indicates that you are specifying parameters for an individual service. Specifies the name of the service and a boolean operator to use when applying the filter.
- **Boolean operator** (unlabeled)—Lists the boolean conditions (equal, not equal, greater than, less than, or range) to use in matching the service specified in the service box.
- **Service** (unlabeled)—Identifies the service (such as https, kerberos, or any) to be matched. If you specified the range service operator this parameter becomes two boxes, into which you enter the start and the end of the range.
- **...** —Displays a list of services from which you can select the service to display in the Service box.

- **Service Group**—Indicates that you are specifying the name of a service group for the source port.
- **Service** (unlabeled)—Selects the service group to use.
- **ICMP Type**—Specifies the ICMP type to use. The default is any. Click the ... button to display a list of available types.
- **Options**
 - **Time Range**—Specify the name of an existing time range or create a new range.
 - ... —Displays the Add Time Range pane, on which you can define a new time range.
 - **Please enter the description below (optional)**—Provides space for you to enter a brief description of the rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Pre-Fragmentation

Use this panel to set the IPsec pre-fragmentation policy and do-not-fragment (DF) bit policy for any interface.

The IPsec pre-fragmentation policy specifies how to treat packets that exceed the maximum transmission unit (MTU) setting when tunneling traffic through the public interface. This feature provides a way to handle cases where a router or NAT device between the security appliance and the client rejects or drops IP fragments. For example, suppose a client wants to FTP get from an FTP server behind a security appliance. The FTP server transmits packets that when encapsulated would exceed the security appliance's MTU size on the public interface. The selected options determine how the security appliance processes these packets. The pre-fragmentation policy applies to all traffic travelling out the security appliance public interface.

The security appliance encapsulates all tunneled packets. After encapsulation, the security appliance fragments packets that exceed the MTU setting before transmitting them through the public interface. This is the default policy. This option works for situations where fragmented packets are allowed through the tunnel without hindrance. For the FTP example, large packets are encapsulated and then fragmented at the IP layer. Intermediate devices may drop fragments or just out-of-order fragments. Load-balancing devices can introduce out-of-order fragments.

When you enable pre-fragmentation, the security appliance fragments tunneled packets that exceed the MTU setting before encapsulating them. If the DF bit on these packets is set, the security appliance clears the DF bit, fragments the packets, and then encapsulates them. This action creates two independent non-fragmented IP packets leaving the public interface and successfully transmits these packets to the peer site by turning the fragments into complete packets to be reassembled at the peer site. In our example, the security appliance overrides the MTU and allows fragmentation by clearing the DF bit.

**Note**

Changing the MTU or the pre-fragmentation option on *any* interface tears down *all* existing connections. For example, if 100 active tunnels terminate on the public interface, and you change the MTU or the pre-fragmentation option on the external interface, all of the active tunnels on the public interface are dropped.

Fields

- **Pre-Fragmentation**—Shows the current pre-fragmentation configuration for every configured interface.
 - **Interface**—Shows the name of each configured interface.
 - **Pre-Fragmentation Enabled**—Shows, for each interface, whether pre-fragmentation is enabled.
 - **DF Bit Policy**—Shows the DF Bit Policy for each interface.
- **Edit**—Displays the Edit IPsec Pre-Fragmentation Policy dialog box.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Edit IPsec Pre-Fragmentation Policy

Use this panel to modify an existing IPsec pre-fragmentation policy and do-not-fragment (DF) bit policy for an interface selected on the parent panel, **Configuration > VPN > IPsec > Pre-Fragmentation**

Fields

- **Interface**—Identifies the selected interface. You cannot change this parameter using this dialog box.
- **Enable IPsec pre-fragmentation**—Enables or disables IPsec pre-fragmentation. The security appliance fragments tunneled packets that exceed the MTU setting before encapsulating them. If the DF bit on these packets is set, the security appliance clears the DF bit, fragments the packets, and then encapsulates them. This action creates two independent, non-fragmented IP packets leaving the public interface and successfully transmits these packets to the peer site by turning the fragments into complete packets to be reassembled at the peer site.
- **DF Bit Setting Policy**—Selects the do-not-fragment bit policy: Copy, Clear, or Set.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Transform Sets

Use this panel to view and add or edit transform sets. A transform is a set of operations done on a data flow to provide data authentication, data confidentiality, and data compression. For example, one transform is the ESP protocol with 3DES encryption and the HMAC-MD5 authentication algorithm (ESP-3DES-MD5).

Fields

- **Transform Sets**—Shows the configured transform sets.
 - **Name**—Shows the name of the transform sets.
 - **Mode**—Shows the mode, Tunnel, of the transform set. This parameter specifies the mode for applying ESP encryption and authentication; in other words, what part of the original IP packet has ESP applied. Tunnel mode applies ESP encryption and authentication to the entire original IP packet (IP header and data), thus hiding the ultimate source and destination addresses.
 - **ESP Encryption**—Shows the Encapsulating Security Protocol (ESP) encryption algorithms for the transform sets. ESP provides data privacy services, optional data authentication, and anti-replay services. ESP *encapsulates* the data being protected.
 - **ESP Authentication**—Shows the ESP authentication algorithms for the transform sets.
- **Add**—Opens the Add Transform Set dialog box, in which you can add a new transform set.
- **Edit**—Opens the Edit Transform Set dialog box, in which you can modify an existing transform set.
- **Delete**—Removes the selected transform set. There is no confirmation or undo.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Add/Edit Transform Set

Use this panel to add or modify a transform set. A transform is a set of operations done on a data flow to provide data authentication, data confidentiality, and data compression. For example, one transform is the ESP protocol with 3DES encryption and the HMAC-MD5 authentication algorithm (ESP-3DES-MD5).

Fields

- **Set Name**—Specifies a name for this transform set.
- **Properties**—Configures properties for this transform set. These properties appear in the Transform Sets table.
 - **Mode**—Shows the mode, Tunnel, of the transform set. This field shows the mode for applying ESP encryption and authentication; in other words, what part of the original IP packet has ESP applied. Tunnel mode applies ESP encryption and authentication to the entire original IP packet (IP header and data), thus hiding the ultimate source and destination addresses.
 - **ESP Encryption**—Selects the Encapsulating Security Protocol (ESP) encryption algorithms for the transform sets. ESP provides data privacy services, optional data authentication, and anti-replay services. ESP *encapsulates* the data being protected.
 - **ESP Authentication**—Selects the ESP authentication algorithms for the transform sets.



Note The IPsec ESP (Encapsulating Security Payload) protocol provides both encryption and authentication. Packet authentication proves that data comes from whom you think it comes from; it is often referred to as “data integrity.”

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Load Balancing



Note VPN load balancing runs only on security appliance models ASA 5520 and higher.

VPN load balancing requires an active 3DES/AES license. The security appliance checks for the existence of this crypto license before enabling load balancing. If it does not detect an active 3DES or AES license, the security appliance prevents the enabling of load balancing and also prevents internal configuration of 3DES by the load balancing system unless the license permits this usage.

This panel lets you enable load balancing on the security appliance. Enabling load balancing involves:

- Configuring the load-balancing cluster by establishing a common virtual cluster IP address, UDP port (if necessary), and IPsec shared secret for the cluster. These values are identical for every device in the cluster.
- Configuring the participating device by enabling load balancing on the device and defining device-specific properties. These values vary from device to device.

If you have a remote-client configuration in which you are using two or more security appliances connected to the same network to handle remote sessions, you can configure these devices to share their session load. This feature is called *load balancing*. Load balancing directs session traffic to the least loaded device, thus distributing the load among all devices. It makes efficient use of system resources and provides increased performance and availability.

**Note**

Load balancing is effective only on remote sessions initiated with the Cisco VPN Client (Release 3.0 and later), the Cisco VPN 3002 Hardware Client (Release 3.5 and later), or the ASA 5505 operating as an Easy VPN Client. All other clients, including LAN-to-LAN connections, can connect to a security appliance on which load balancing is enabled, but they cannot participate in load balancing.

To implement load balancing, you group together logically two or more devices on the same private LAN-to-LAN network into a *virtual cluster*.

All devices in the virtual cluster carry session loads. One device in the virtual cluster, the *virtual cluster master*, directs incoming calls to the other devices, called *secondary devices*. The virtual cluster master monitors all devices in the cluster, keeps track of how busy each is, and distributes the session load accordingly. The role of virtual cluster master is not tied to a physical device; it can shift among devices. For example, if the current virtual cluster master fails, one of the secondary devices in the cluster takes over that role and immediately becomes the new virtual cluster master.

The virtual cluster appears to outside clients as a single *virtual cluster IP address*. This IP address is not tied to a specific physical device. It belongs to the current virtual cluster master; hence, it is virtual. A VPN client attempting to establish a connection connects first to this virtual cluster IP address. The virtual cluster master then sends back to the client the public IP address of the least-loaded available host in the cluster. In a second transaction (transparent to the user) the client connects directly to that host. In this way, the virtual cluster master directs traffic evenly and efficiently across resources.

**Note**

All clients other than the Cisco VPN client, the Cisco VPN 3002 Hardware Client, or the ASA 5505 operating as an Easy VPN Client connect directly to the security appliance as usual; they do not use the virtual cluster IP address.

If a machine in the cluster fails, the terminated sessions can immediately reconnect to the virtual cluster IP address. The virtual cluster master then directs these connections to another active device in the cluster. Should the virtual cluster master itself fail, a secondary device in the cluster immediately and automatically takes over as the new virtual session master. Even if several devices in the cluster fail, users can continue to connect to the cluster as long as any one device in the cluster is up and available.

Prerequisites

Load balancing is disabled by default. You must explicitly enable load balancing.

You must have first configured the public and private interfaces and also have previously configured the interface to which the virtual cluster IP address refers.

All devices that participate in a cluster must share the same cluster-specific values: IP address, encryption settings, encryption key, and port.

Fields

- **VPN Load Balancing**—Configures virtual cluster device parameters.
 - **Participate in Load Balancing Cluster**—Specifies that this device is a participant in the load-balancing cluster.

- **VPN Cluster Configuration**—Configures device parameters that must be the same for the entire virtual cluster. All servers in the cluster must have an identical cluster configuration.
- **Cluster IP Address**—Specifies the single IP address that represents the entire virtual cluster. Choose an IP address that is within the public subnet address range shared by all the security appliances in the virtual cluster.
- **UDP Port**—Specifies the UDP port for the virtual cluster in which this device is participating. The default value is 9023. If another application is using this port, enter the UDP destination port number you want to use for load balancing.
- **Enable IPsec Encryption**—Enables or disables IPsec encryption. If you select this check box, you must also specify and verify a shared secret. The security appliances in the virtual cluster communicate via LAN-to-LAN tunnels using IPsec. To ensure that all load-balancing information communicated between the devices is encrypted, select this check box.

**Note**

When using encryption, you must have previously configured the load-balancing inside interface. If that interface is not enabled on the load-balancing inside interface, you get an error message when you try to configure cluster encryption.

If the load-balancing inside interface was enabled when you configured cluster encryption, but was disabled before you configured the participation of the device in the virtual cluster, you get an error message when you select the Participate in Load Balancing Cluster check box, and encryption is not enabled for the cluster.

- **IPsec Shared Secret**—Specifies the shared secret to between IPsec peers when you have enabled IPsec encryption. The value you enter in the box appears as consecutive asterisk characters.
- **Verify Secret**—Confirms the shared secret value entered in the IPsec Shared Secret box.
- **VPN Server Configuration**—Configures parameters for this specific device.
 - **Interfaces**—Configures the public and private interfaces and their relevant parameters.
 - **Public**—Specifies the name or IP address of the public interface for this device.
 - **Private**—Specifies the name or IP address of the private interface for this device.
 - **Priority**—Specifies the priority assigned to this device within the cluster. The range is from 1 to 10. The priority indicates the likelihood of this device becoming the virtual cluster master, either at start-up or when an existing master fails. The higher you set the priority (for example, 10), the more likely this device becomes the virtual cluster master.

**Note**

If the devices in the virtual cluster are powered up at different times, the first device to be powered up assumes the role of virtual cluster master. Because every virtual cluster requires a master, each device in the virtual cluster checks when it is powered-up to ensure that the cluster has a virtual master. If none exists, that device takes on the role. Devices powered up and added to the cluster later become secondary devices. If all the devices in the virtual cluster are powered up simultaneously, the device with the highest priority setting becomes the virtual cluster master. If two or more devices in the virtual cluster are powered up simultaneously, and both have the highest priority setting, the one with the lowest IP address becomes the virtual cluster master.

- **NAT Assigned IP Address**—Specifies the IP address that this device's IP address is translated to by NAT. Enter 0.0.0.0 if NAT is not being used or if the device is not behind a firewall using NAT.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

NAC

Configuration > VPN > NAC

The security appliance uses Extensible Authentication Protocol (EAP) over UDP (EAPoUDP) messaging to validate the posture of remote hosts. Posture validation involves the checking of a remote host for compliancy with safety requirements before the assignment of a network access policy. An Access Control Server must be configured for Network Admission Control before you configure NAC on the security appliance.

The NAC window lets you set attributes that apply to all NAC communications. The following global attributes at the top of the window apply to EAPoUDP messaging between the security appliance and remote hosts:

- **Retransmission Timer**—The security appliance starts this timer when it sends an EAPoUDP message to the host. A response from the host clears the timer. If the timer expires before the security appliance receives a response, it resends the message. The setting is in seconds. Enter a value in the range 1 to 60. The default setting is 3.
- **Hold Timer**—The security appliance starts this timer when it places the NAC session for a remote host into a hold state. It places a session in a hold state if it does not receive a response after sending EAPoUDP messages equal to the number of EAPoUDP Retries. The security appliance also starts this timer after it receives an Access Reject message from the ACS server. When the timer expires, the security appliance tries to initiate a new EAP over UDP association with the remote host. The setting is in seconds. Enter a value in the range 60 to 86400. The default setting is 180.
- **EAPoUDP Retries**—Number of times the security appliance resends an EAP over UDP message. This attribute limits the number of consecutive retries sent in response to Retransmission Timer expirations. The setting is in seconds. Enter a value in the range 1 to 3. The default setting is 3.
- **EAPoUDP Port**—Port number for EAP over UDP communication with the Cisco Trust Agent (CTA) on the host. This number must match the port number configured on the CTA. Enter a value in the range 1024 to 65535. The default setting is 21862.

The Clientless Authentication area of the NAC window lets you configure settings for hosts that are not responsive to the EAPoUDP requests. Hosts for which there is no CTA running do not respond to these requests.

- **Enable Clientless Authentication**—Check to enable clientless authentication. The security appliance sends the configured clientless username and password to the Access Control Server in the form of a user authentication request. The ACS in turn requests the access policy for clientless hosts. If you uncheck this attribute, the security appliance applies the default ACL for clientless hosts.

- Username—Username configured for clientless hosts on the ACS. The default setting is clientless. Enter 1 to 64 ASCII characters, excluding leading and trailing spaces, pound signs (#), question marks (?), single and double quotation marks (“ ” and "), asterisks (*), and angle brackets (< and >).
- Password—Password configured for clientless hosts on the ACS. The default setting is clientless. Enter 4 – 32 ASCII characters.
- Confirm Password—Password configured for clientless hosts on the ACS repeated for validation.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

