Prefiltering and Prefilter Policies

About Prefiltering

Prefiltering is the first phase of access control, before the system performs more resource-intensive evaluation. Prefiltering is simple, fast, and early. Prefiltering uses limited outer-header criteria to quickly handle traffic. Compare this to subsequent evaluation, which uses inner headers and has more robust inspection capabilities.

Configure prefiltering to:

- Improve performance— The sooner you exclude traffic that does not require inspection, the better. You can fastpath or block certain types of plaintext, passthrough tunnels based on their outer encapsulation headers, without inspecting their encapsulated connections. You can also fastpath or block any other connections that benefit from early handling.
- Tailor deep inspection to encapsulated traffic—You can rezone certain types of tunnels so that you can later handle their encapsulated connections using the same inspection criteria. Rezoning is necessary because after prefiltering, access control uses inner headers.

About Prefilter Policies

Prefiltering is a policy-based feature. To assign it to a device, you assign it to the access control policy that is assigned to the device.

Policy Components: Rules and Default Action

In a prefilter policy, *tunnel rules*, *prefilter rules*, and a *default action* handle network traffic:

- Tunnel and prefilter rules—First, rules in a prefilter policy handle traffic in the order you specify. Tunnel rules match specific tunnels only and support rezoning. Prefilter rules have a wider range of constraints and do not support rezoning. For more information, see Tunnel vs Prefilter Rules, on page 2.
- Default action (tunnels only)—If a tunnel does not match any rules, the default action handles it. The default action can block these tunnels, or continue access control on their individual encapsulated connections. You cannot rezone tunnels with the default action.

  There is no default action for nonencapsulated traffic. If a nonencapsulated connection does not match any prefilter rules, the system continues with access control.

Connection Logging

You can log connections fastpathed and blocked by the prefilter policy. See *Other Connections You Can Log* in the Cisco Secure Firewall Management Center Administration Guide for more information.

Connection events contain information on whether and how logged connections—including entire tunnels—were prefiltered. You can view this information in event views (workflows), dashboards, and reports, and use it as correlation criteria. Keep in mind that because fastpathed and blocked connections are not subject to deep inspection, associated connection events contain limited information.

Default Prefilter Policy

Every access control policy has an associated prefilter policy.

The system uses a default policy if you do not configure custom prefiltering. Initially, this system-provided policy passes all traffic to the next phase of access control. You can change the policy's default action and configure its logging options, but you cannot add rules to it or delete it.

Prefilter Policy Inheritance and Multitenancy

Access control uses a hierarchical implementation that complements multitenancy. Along with other advanced settings, you can lock a prefilter policy association, enforcing that association in all descendant access control policies. For more information, see Access Control Policy Inheritance.

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain. The default prefilter policy belongs to the Global domain.

Tunnel vs Prefilter Rules

Whether you configure a tunnel or prefilter rule depends on the specific type of traffic you want to match and the actions or further analysis you want to perform.

Characteristic	Tunnel Rules	Prefilter Rules
Primary function	Quickly fastpath, block, or rezone plaintext, passthrough tunnels.	Quickly fastpath or block any other connection that benefits from early handling.
Encapsulation and port/protocol criteria	Encapsulation conditions match only plaintext tunnels over selected protocols, listed in Encapsulation Rule Conditions, on page 13.	Port conditions can use a wider range of port and protocol constraints than tunnel rules; see Port, Protocol, and ICMP Code Rule Conditions.

Characteristic	Tunnel Rules	Prefilter Rules
Network criteria	Tunnel endpoint conditions constrain the endpoints of the tunnels you want to handle; see Network Rule Conditions.	Network conditions constrain the source and destination hosts in each connection; see Network Rule Conditions.
Direction	Bidirectional or unidirectional (configurable). Tunnel rules are bidirectional by default, so they can handle all traffic between tunnel endpoints.	Unidirectional only (nonconfigurable). Prefilter rules match source-to-destination traffic only.
Rezone sessions for further analysis	Supported, using tunnel zones; see Tunnel Zones and Prefiltering, on page 14.	Not supported.

Prefiltering vs Access Control

Prefilter and access control policies both allow you to block and trust traffic, though the prefiltering "trust" functionality is called "fastpathing" because it skips more inspection. The following table explains this and other differences between prefiltering and access control, to help you decide whether to configure custom prefiltering.

If you do not configure custom prefiltering, you can only approximate—not replicate—prefilter functionality with early-placed Block and Trust rules in the access control policy.

Characteristic	Prefiltering	Access Control	For more information, see...
Primary function	Quickly fastpath or block certain types of plaintext, passthrough tunnels (see Encapsulation Rule Conditions, on page 13), or tailor subsequent inspection to their encapsulated traffic. Fastpath or block any other connections that benefit from early handling.	Inspect and control all network traffic, using simple or complex criteria, including contextual information and deep inspection results.	About Prefiltering, on page 1
Implementation	Prefilter policy. The prefilter policy is invoked by the access control policy.	Access control policy. The access control policy is a main configuration. In addition to invoking subpolicies, access control policies have their own rules.	About Prefilter Policies, on page 1 Associating Other Policies with Access Control
Sequence within access control	First. The system matches traffic to prefilter criteria before all other access control configurations.	—	—

Characteristic	Prefiltering	Access Control	For more information, see...
Rule actions	Fewer. You can stop further inspection (Fastpath and Block) or allow further analysis with the rest of access control (Analyze).	More. Access control rules have a larger variety of actions, including monitoring, deep inspection, block with reset, and interactive blocking.	Tunnel and Prefilter Rule Components, on page 9 Access Control Rule Actions
Bypass capability	Fastpath rule action. Fastpathing traffic in the prefilter stage bypasses all further inspection and handling, including: • Security Intelligence • authentication requirements imposed by an identity policy • SSL decryption • access control rules • deep inspection of packet payloads • discovery • rate limiting	Trust rule action. Traffic trusted by access control rules is only exempt from deep inspection and discovery.	Introduction to Access Control Rules
Rule criteria	Limited. Rules in the prefilter policy use simple network criteria: IP address, VLAN tag, port, and protocol. For tunnels, tunnel endpoint conditions specify the IP address of the routed interfaces of the network devices on either side of the tunnel.	Robust. Access control rules use network criteria, but also user, application, requested URL, and other contextual information available in packet payloads. Network conditions specify the IP address of source and destination hosts.	Tunnel vs Prefilter Rules, on page 2 Prefilter Rule Conditions, on page 11 Tunnel Rule Conditions, on page 13
IP headers used (tunnel handling)	Outermost. Using outer headers allows you to handle entire plaintext, passthrough tunnels. For nonencapsulated traffic, prefiltering still uses "outer" headers—which in this case are the only headers.	Innermost possible. For a nonencrypted tunnel, access control acts on its individual encapsulated connections, not the tunnel as a whole.	Passthrough Tunnels and Access Control, on page 5

Characteristic	Prefiltering	Access Control	For more information, see...
Rezone encapsulated connections for further analysis	Rezones tunneled traffic. Tunnel zones allow you to tailor subsequent inspection to prefiltered, encapsulated traffic.	Uses tunnel zones. Access control uses the tunnel zones you assign during prefiltering.	Tunnel Zones and Prefiltering, on page 14
Connection logging	Fastpathed and blocked traffic only. Allowed connections may still be logged by other configurations.	Any connection.	*Other Connections You Can Log* in the Cisco Secure Firewall Management Center Administration Guide
Supported devices	Secure Firewall Threat Defense only.	All.	—

Passthrough Tunnels and Access Control

Plaintext (nonencrypted) tunnels can encapsulate multiple connections, often flowing between discontinuous networks. These tunnels are especially useful for routing custom protocols over IP networks, IPv6 traffic over IPv4 networks, and so on.

An outer *encapsulation header* specifies the source and destination IP addresses of the *tunnel endpoints*—the routed interfaces of the network devices on either side of the tunnel. Inner *payload headers* specify the source and destination IP addresses of the encapsulated connections' actual endpoints.

Often, network security devices handle plaintext tunnels as *passthrough* traffic. That is, the device is not one of the tunnel endpoints. Instead, it is deployed between the tunnel endpoints and monitors the traffic flowing between them.

Some network security devices enforce security policies using outer IP headers. Even for plaintext tunnels, these devices have no control over or insight into individual encapsulated connections and their payloads.

By contrast, the system leverages access control as follows:

- Outer header evaluation—First, prefiltering uses outer headers to handle traffic. You can block or fastpath entire plaintext, passthrough tunnels at this stage.
- Inner header evaluation—Next, the rest of access control (and other features such as QoS) use the innermost detectable level of headers to ensure the most granular level of inspection and handling possible.

  If a passthrough tunnel is not encrypted, the system acts on its individual encapsulated connections at this stage. You must *rezone* a tunnel (see Tunnel Zones and Prefiltering, on page 14) to act on all its encapsulated connections.

Access control has no insight into encrypted passthrough tunnels. For example, access control rules see a passthrough VPN tunnel as one connection. The system handles the entire tunnel using only the information in its outer, encapsulation header.

Best Practices for Fastpath Prefiltering

When you use the fastpath action in a prefilter rule, the matching traffic bypasses inspection and is simply transmitted through the device. Use this action for traffic that you can trust and that would not benefit from any of the security features available.

The following types of traffic are ideal for fastpathing. For example, you could configure the rules to fastpath any traffic from or to the IP addresses of the endpoints or servers. You can further limit the rule based on ports used.

- VPN traffic that is going through the device. That is, the device is not an endpoint in the VPN topology.
- Scanner traffic. Scanner probes can create a lot of false-positive responses from intrusion policies.
- Voice/video.
- Backups.
- Management traffic (sftunnel) that traverses threat defense devices. Performing deep inspection on management traffic (using access control policies) can cause issues. You can prefilter based on port TCP/8305 between the management center and managed devices.

Best Practices for Encapsulated Traffic Handling

This topic discusses guidelines for the following types of encapsulated traffic:

- Generic Routing Encapsulation (GRE)
- Point-to-Point Protocol (PPTP)
- IPinIP
- IPv6inIP
- Teredo

GRE Tunnel Limitations

GRE tunnel processing is limited to IPv4 and IPv6 passenger flows. Other protocols, such as PPTP and WCCP, are not supported within the GRE tunnel.

Understand Snort version support for your managed devices

The inspection engine used by managed devices is known as Snort. Snort 3 supports more features than Snort 2. To understand how these affect managed devices on your network, you must know:

- Which versions of Snort your device supports.

  Snort version support can be found in the section on bundled components in the *Cisco Firepower Compatibility Guide*.

- How the management center and threat defense software support Snort 2 and Snort 3

Limitations of Snort 2 and Snort 3 can be found in the *Feature Limitations of Snort 3 for Management Center-Managed Threat Defense* topic in the Cisco Secure Firewall Management Center Snort 3 Configuration Guide.

GRE v1 and PPTP bypass outer flow processing

GRE v1 (sometimes referred to as *stateful GRE*) and PPTP traffic bypass outer flow processing.

Passenger flow processing is supported for IPv6inIP and Teredo but the following limitations apply:

- Sessions are over a single tunnel that is not load-balanced
- There is no HA or clustering replication
- Primary and secondary flow relationships are not maintained
- Prefilter policy white and black lists are not supported

GRE v0 sequence number field must be optional

All endpoints sending traffic on the network must send GREv0 traffic with the sequence number field as optional; otherwise, the sequence number field is removed. RFC 1701 and RFC 2784 both specify the sequence field as optional.

How tunnels work with interfaces

Prefilter and access control policy rules are applied to all tunnel types on routed, transparent, inline-set, inline-tap, and passive interfaces.

References

For more information about the GRE and PPTP protocols, see the following:

- RFC 1701, RFC 2784, and RFC 2890 (GRE protocol v0)
- RFC 2637 (PPTP and GRE protocol v1)

Requirements and Prerequisites for Prefilter Policies

Model Support

Threat Defense

Supported Domains

Any

User Roles

- Admin
- Access Admin

- Network Admin

Configure Prefiltering

To perform custom prefiltering, configure prefilter policies and assign the policies to access control policies. It is through the access control policy that prefilter policies get assigned to managed devices.

Only one person should edit a policy at a time, using a single browser window. If multiple users save the same policy, the last saved changes are retained. For your convenience, the system displays information on who (if anyone) is currently editing each policy. To protect the privacy of your session, a warning appears after 30 minutes of inactivity on the policy editor. After 60 minutes, the system discards your changes.

Procedure

Step 1 Choose **Policies** > **Access Control** > **Prefilter**.

Step 2 Click **New Policy** to create a custom prefilter policy.

A new prefilter policy has no rules and a default action of Analyze all tunnel traffic. It performs no logging or tunnel rezoning. You can also **Copy** ( ) or **Edit** ( ) an existing policy.

Step 3 Configure the prefilter policy's default action and its logging options.

- Default action—Choose a default action for supported plaintext, passthrough tunnels: **Analyze all tunnel traffic** (with access control) or **Block all tunnel traffic**.
- Default action logging—Click **Logging** ( ) next to the default action; see *Logging Connections with a Policy Default Action* in the Cisco Secure Firewall Management Center Administration Guide. You can configure default action logging for blocked tunnels only.

Step 4 Configure tunnel and prefilter rules.

In a custom prefilter policy, you can use both kinds of rule, in any order. Create rules depending on the specific type of traffic you want to match and the actions or further analysis you want to perform; see Tunnel vs Prefilter Rules, on page 2.

Caution Exercise caution when using tunnel rules to assign tunnel zones. Connections in rezoned tunnels may not match security zone constraints in later evaluation. For more information, see Tunnel Zones and Prefiltering, on page 14.

For detailed information on configuring rule components, see Tunnel and Prefilter Rule Components, on page 9.

Step 5 Evaluate rule order. To move a rule, click and drag or use the right-click menu to cut and paste.

Properly creating and ordering rules is a complex task, but one that is essential to building an effective deployment. If you do not plan carefully, rules can preempt other rules or contain invalid configurations. For more information, see Best Practices for Access Control Rules.

Step 6 Save the prefilter policy.

Step 7 For configurations that support tunnel zone constraints, appropriately handle rezoned tunnels.

Match connections in rezoned tunnels by using tunnel zones as source zone constraints.

Step 8 Associate the prefilter policy with the access control policy deployed to your managed devices.

See Associating Other Policies with Access Control.

Step 9 Deploy configuration changes; see Deploy Configuration Changes.

Note When you deploy a prefilter policy, its rules are not applied on the existing tunnel sessions. Hence, traffic on an existing connection is not bound by the new policy that is deployed. In addition, the policy hit count is incremented only for the first packet of a connection that matches a policy. Thus, the traffic on an existing connection that could match a policy is omitted from the hit count. To have the policy rules effectively applied, clear the existing tunnel sessions, and then deploy the policy.

What to do next

If you will deploy time-based rules, specify the time zone of the device to which the policy is assigned. See Time Zone.

Tunnel and Prefilter Rule Components

State (Enabled/Disabled)

By default, rules are enabled. If you disable a rule, the system does not use it and stops generating warnings and errors for that rule.

Position

Rules are numbered, starting at 1. The system matches traffic to rules in top-down order by ascending rule number. The first rule that traffic matches is the rule that handles that traffic, regardless of rule type (tunnel vs prefilter).

Action

A rule's action determines how the system handles and logs matching traffic:

- Fastpath—Exempts matching traffic from all further inspection and control, including access control, identity requirements, and rate limiting. Fastpathing a tunnel fastpaths all encapsulated connections.
- Block—Blocks matching traffic without further inspection of any kind. Blocking a tunnel blocks all encapsulated connections.
- Analyze—Allows traffic to continue to be analyzed by the rest of access control, using inner headers. If passed by access control and any related deep inspection, this traffic may also be rate limited. For tunnel rules, enables rezoning with the Assign Tunnel Zone option.

Direction (Tunnel Rules Only)

A tunnel rule's direction determines how the system source and destination criteria:

- Match tunnels only from source (unidirectional)—Match source-to-destination traffic only. Matching traffic must originate from one of the specified source interfaces or tunnel endpoints, and leave through one of the destination interfaces or tunnel endpoints.
- Match tunnels from source and destination (bidirectional)—Match both source-to-destination traffic and destination-to-source traffic. The effect is identical to writing two unidirectional rules, one the mirror of the other.

Prefilter rules are always unidirectional.

Assign Tunnel Zone (Tunnel Rules Only)

In a tunnel rule, assigning a tunnel zone (whether existing or created on the fly), *rezones* matching tunnels. Rezoning requires the Analyze action.

Rezoning a tunnel allows other configurations—such as access control rules—to recognize all the tunnel's encapsulated connections as belonging together. By using a tunnel's assigned tunnel zone as an interface constraint, you can tailor inspection to its encapsulated connections. For more information, see Tunnel Zones and Prefiltering, on page 14.

Caution Exercise caution when assigning tunnel zones. Connections in rezoned tunnels may not match security zone constraints in later evaluation. See Using Tunnel Zones, on page 14 for a brief walkthrough of a tunnel zone implementation, and a discussion of the implications of rezoning without explicitly handling rezoned traffic.

Conditions

Conditions specify the specific traffic the rule handles. Traffic must match all a rule's conditions to match the rule. Each condition type has its own tab in the rule editor.

You can prefilter traffic using the following *outer-header* constraints. You must constrain tunnel rules by encapsulation protocol.

- Interface—Interface Rule Conditions
- Network (prefilter rule)/Tunnel Endpoints (tunnel rule)—Network Rule Conditions
- VLAN—VLAN Tags Rule Conditions
- Ports (prefilter rule)/Encapsulation and Ports (tunnel rule)—Port Rule Conditions for Prefilter Rules, on page 12 or Encapsulation Rule Conditions, on page 13
- Time Range—Time and Day Rule Conditions

Logging

A rule's logging settings govern the records the system keeps of the traffic it handles.

In tunnel and prefilter rules, you can log fastpathed and blocked traffic (the Fastpath and Block actions). For traffic subject to further analysis (the Analyze action), logging in the prefilter policy is disabled, although matching connections may still be logged by other configurations. Logging is performed on inner flows, not on the encapsulating flow. For more information, see *Logging Connections with Tunnel and Prefilter Rules* in the Cisco Secure Firewall Management Center Administration Guide.

Comments

Each time you save changes to a rule you can add comments. For example, you might summarize the overall configuration for the benefit of other users, or note when you change a rule and the reason for the change.

You cannot edit or delete these comments after you save the rule.

Related Topics

Best Practices for Access Control Rules

Prefilter Rule Conditions

Rule conditions enable you to fine-tune your prefilter policy to target the networks you want to control. See one of the following sections for more information.

Interface Rule Conditions

Interface rule conditions control traffic by its source and destination interfaces.

Depending on the rule type and the devices in your deployment, you can use predefined *interface objects* called *security zones* or *interface groups* to build interface conditions. Interface objects segment your network to help you manage and classify traffic flow by grouping interfaces across multiple devices; see Interface.

Tip Constraining rules by interface is one of the best ways to improve system performance. If a rule excludes all of a device's interfaces, that rule does not affect that device's performance.

Just as all interfaces in an interface object must be of the same type (all inline, passive, switched, routed, or ASA FirePOWER), all interface objects used in an interface condition must be of the same type. Because devices deployed passively do not transmit traffic, in passive deployments you cannot constrain rules by destination interface.

Network Rule Conditions

Network rule conditions control traffic by its source and destination IP address, using inner headers. Tunnel rules, which use outer headers, have tunnel endpoint conditions instead of network conditions.

You can use predefined objects to build network conditions, or manually specify individual IP addresses or address blocks.

Note You *cannot* use FDQN network objects in identity rules.

Leave matching criteria empty whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against every combination of the contents of the criteria you specify.

VLAN Tags Rule Conditions

Note VLAN tags in access rules only apply to inline sets. Access rules with VLAN tags do not match traffic on firewall interfaces.

VLAN rule conditions control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic. The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note the following Q-in-Q support:

- Threat Defense on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).
- Threat Defense on all other models:
  - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
  - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.

You can specify a maximum of 50 VLAN conditions.

In a cluster, if you encounter problems with VLAN matching, edit the access control policy advanced options, Transport/Network Preprocessor Settings, and select the **Ignore the VLAN header when tracking connections** option.

Note The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal VLAN tags to constrain this configuration can have unexpected results. Using override-enabled objects allows descendant domain administrators to tailor Global configurations to their local environments.

Port Rule Conditions for Prefilter Rules

Port conditions match traffic based on the source and destination ports. Depending on the rule type, “port” can represent any of the following:

- TCP and UDP—You can control TCP and UDP traffic based on the port. The system represents this configuration using the protocol number in parentheses, plus an optional associated port or port range. For example: TCP(6)/22.
- ICMP—You can control ICMP and ICMPv6 (IPv6-ICMP) traffic based on its internet layer protocol plus an optional type and code. For example: ICMP(1):3:3.
- Protocol—You can control traffic using other protocols that do not use ports.

Using Source and Destination Port Constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as destination port conditions in a single access control rule.

Matching Non-TCP Traffic with Port Conditions

You can match non-port-based protocols. By default, if you do not specify a port condition, you are matching IP traffic. Although you can configure port conditions to match other protocols in prefilter rules, you should use tunnel rules instead when matching GRE, IP in IP, IPv6 in IP, and Toredo Port 3544.

Time and Day Rule Conditions

You can specify a continuous time range or a recurring time period.

For example, a rule can apply only during weekday working hours, or every weekend, or during a holiday shutdown period.

Time-based rules are applied based on the local time of the device that processes the traffic.

Time-based rules are supported only on threat defense devices. If you assign a policy with a time-based rule to a different type of device, the time restriction associated with the rule is ignored on that device. You will see warnings in this case.

Tunnel Rule Conditions

Rule conditions enable you to fine-tune your tunnel policy to target the networks you want to control. For tunnel rules, you can use the following conditions:

- **Interface Objects**—The security zones or interface groups that define the device interfaces through which the connections pass. See Interface Rule Conditions.
- **Tunnel Endpoints**—The network objects that define the source and destination IP addresses of the tunnel.
- **VLAN Tags**—The outermost VLAN tag in the tunnel. See VLAN Tags Rule Conditions.
- **Encapsulation and Ports**—The encapsulation protocol of the tunnel. See Encapsulation Rule Conditions, on page 13.
- **Time Range**—The days and times when the rule is active. If you do not specify a time range, the rule is always active. See Time and Day Rule Conditions.

Encapsulation Rule Conditions

Encapsulation conditions are specific to tunnel rules.

These conditions control certain types of plaintext, passthrough tunnels by their encapsulation protocol. You must choose at least one protocol to match before you can save the rule. You can choose:

- GRE (47)
- IP-in-IP (4)
- IPv6-in-IP (41)
- Teredo (UDP (17)/3455)

Tunnel Zones and Prefiltering

Tunnel zones allow you to use prefiltering to tailor subsequent traffic handling to encapsulated connections.

A special mechanism is required because usually, the system handles traffic using the innermost detectable level of headers. This ensures the most granular level of inspection possible. But it also means that if a passthrough tunnel is not encrypted, the system acts on its individual encapsulated connections; see Passthrough Tunnels and Access Control, on page 5.

Tunnel zones solve this problem. During the first phase of access control (prefiltering) you can use outer headers to identify certain types of plaintext, passthrough tunnels. Then, you can *rezone* those tunnels by assigning a custom *tunnel zone*.

Rezoning a tunnel allows other configurations—such as access control rules—to recognize all the tunnel's encapsulated connections as belonging together. By using a tunnel's assigned tunnel zone as an interface constraint, you can tailor inspection to its encapsulated connections.

Despite its name, a tunnel zone is not a security zone. A tunnel zone does not represent a set of interfaces. It is more accurate to think of a tunnel zone as a tag that, in some cases, replaces the security zone associated with an encapsulated connection.

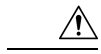

Caution For configurations that support tunnel zone constraints, connections in rezoned tunnels do **not** match security zone constraints. For example, after you rezone a tunnel, access control rules can match its encapsulated connections with their newly assigned *tunnel* zone, but not with any original *security* zone.

See Using Tunnel Zones, on page 14 for a brief walkthrough of a tunnel zone implementation, and a discussion of the implications of rezoning without explicitly handling rezoned traffic.

Configurations Supporting Tunnel Zone Constraints

Only access control rules support tunnel zone constraints.

No other configurations support tunnel zone constraints. For example, you cannot use QoS to rate limit a plaintext tunnel as a whole; you can only rate limit its individual encapsulated sessions.

Using Tunnel Zones

This example procedure summarizes how you might rezone GRE tunnels for further analysis, using tunnel zones. You can adapt the concepts described in this example to other scenarios where you need to tailor traffic inspection to connections encapsulated in plaintext, passthrough tunnels.

Consider a situation where your organization's internal traffic flows through the Trusted security zone. The Trusted security zone represents a set of interfaces across multiple managed devices deployed in various locations. Your organization's security policy requires that you allow internal traffic after deep inspection for exploits and malware.

Internal traffic sometimes includes plaintext, passthrough, GRE tunnels between particular endpoints. Because the traffic profile of this encapsulated traffic is different from your "normal" interoffice activity—perhaps it is known and benign—you can limit inspection of certain encapsulated connections while still complying with your security policy.

In this example, after you deploy configuration changes:

- Plaintext, passthrough, GRE-encapsulated tunnels detected in the Trusted zone have their individual encapsulated connections evaluated by one set of intrusion and file policies.
- All other traffic in the Trusted zone is evaluated with a different set of intrusion and file policies.

You accomplish this task by *rezoning* GRE tunnels. Rezoning ensures that access control associates GRE-encapsulated connections with a custom *tunnel* zone, rather than their original Trusted *security* zone. Rezoning is required due to the way access control handles encapsulated traffic; see Passthrough Tunnels and Access Control, on page 5 and Tunnel Zones and Prefiltering, on page 14.

Procedure

Step 1 Configure custom intrusion and file policies that tailor deep inspection to encapsulated traffic, and another set of intrusion and file policies tailored to nonencapsulated traffic.

Step 2 Configure custom prefiltering to rezone GRE tunnels flowing through the Trusted security zone.

Create a custom prefilter policy and associate it with access control. In that custom prefilter policy, create a tunnel rule (in this example, **`GRE_tunnel_rezone`**) and a corresponding tunnel zone (**`GRE_tunnel`**). For more information, see Configure Prefiltering, on page 8.

Table 1: GRE_tunnel_rezone Tunnel Rule

Rule Component	Description
Interface object condition	Match internal-only tunnels by using the Trusted security zone as both a Source Interface Object and Destination Interface Object constraint.
Tunnel endpoint condition	Specify the source and destination endpoints for the GRE tunnels used in your organization. Tunnel rules are bidirectional by default. If you do not change the **Match tunnels from...** option, it does not matter which endpoints you specify as source and which as destination.
Encapsulation condition	Match GRE traffic.
Assign Tunnel Zone	Create the **`GRE_tunnel`** tunnel zone, and assign it to tunnels that match the rule.
Action	Analyze (with the rest of access control).

Step 3 Configure access control to handle connections in rezoned tunnels.

In the access control policy deployed to your managed devices, configure a rule (in this example, **`GRE_inspection`**) that handles the traffic you rezoned. For more information, see Create and Edit Access Control Rules.

Table 2: GRE_inspection Access Control Rule

Rule Component	Description
Security zone condition	Match rezoned tunnels by using the GRE_tunnel security zone as a Source Zone constraint.

Rule Component	Description
Action	Allow, with deep inspection enabled. Choose the file and intrusion policies tailored to inspect encapsulated internal traffic.

Caution If you skip this step, the rezoned connections may match **any** access control rule not constrained by security zone. If the rezoned connections do not match any access control rules, they are handled by the access control policy default action. Make sure this is your intent.

Step 4 Configure access control to handle nonencapsulated connections flowing through the Trusted security zone.

In the same access control policy, configure a rule (in this example, **`internal_default_inspection`**) that handles non-rezoned traffic in the Trusted security zone.

Table 3: internal_default_inspection Access Control Rule

Rule Component	Description
Security zone condition	Match non-rezoned internal-only traffic by using the Trusted security zone as both a Source Zone and Destination Zone constraint.
Action	Allow, with deep inspection enabled. Choose the file and intrusion policies tailored to inspect nonencapsulated internal traffic.

Step 5 Evaluate the position of the new access control rules relative to preexisting rules. Change rule order if necessary.

If you place the two new access control rules next to each other, it does not matter which you place first. Because you rezoned GRE tunnels, the two rules cannot preempt each other.

Step 6 Save all changed configurations.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes.

Creating Tunnel Zones

The following procedure explains how to create a tunnel zone in the object manager. You can also create zones when editing a tunnel rule.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Chose **Tunnel Zone** from the list of object types.

Step 3 Click **Add Tunnel Zone**.

Step 4 Enter a **Name** and, optionally, a **Description**.

Step 5 Click **Save**.

What to do next

- Assign tunnel zones to plaintext, passthrough tunnels as part of custom prefiltering; see Configure Prefiltering, on page 8.

Moving Prefilter Rules to an Access Control Policy

You can move prefilter rules from a prefilter policy to the associated access control policy.

Before you begin

Note the following conditions before you proceed:

- Only prefilter rules can be moved to an access control policy. Tunnel rules cannot be moved.
- The prefilter rules can be moved only to the associated access control policy.
- The prefilter rules with configured interface groups cannot be moved.
- The **Action** parameter in the prefilter rule is changed to a suitable action in the access control rule when moved. To know what each action in the prefilter rule maps to, see the following table:

Action in the prefilter rule	Action in the access control rule
Analyze	Allow
Block	Block
Fastpath	Trust

- Similarly, based on the action configured in the prefilter rule, the logging configuration is set to an appropriate setting after the rule is moved, as mentioned in the following table.

Action in the prefilter rule	Enabled Logging configurations in the access control rule
Analyze	None of the log settings are enabled.
Block	• Log at Beginning of Connection • Event Viewer • Syslog Server • SNMP Trap

Action in the prefilter rule	Enabled Logging configurations in the access control rule
Fastpath	• Log at Beginning of Connection • Log at End of Connection • Event Viewer • Syslog Server • SNMP Trap

- The comments in the prefilter rule configuration are lost after moving the rule. However, a new comment is added in the moved rule mentioning the source prefilter policy.
- While moving rules from the source policy, if another user modifies those rules, the management center displays a message. You may continue with the process after refreshing the page.

Procedure

Step 1 In the prefilter policy editor, select the rules that you want to move with a left-click on your mouse.

Tip To select multiple rules, use the Ctrl (Control) key on your keyboard.

Step 2 Right-click the selected rules and choose **Move to another policy**.

Step 3 Select the destination access control policy from the **Access Policy** drop-down list.

Step 4 From the **Place Rules** drop-down list, choose where you want to position the moved rules:

- To position as the last set of rules in the **Default** section, choose **At the bottom (within the Default section)**.
- To position as the first set of rules in the **Mandatory** section, choose **At the top (within the Mandatory section)**.

Step 5 Click **Move**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes.

Prefilter Policy Hit Counts

Hit count indicates the number of times a policy rule has triggered for a matching connection.

For complete information on viewing prefilter policy hit counts, see Viewing Rule Hit Counts.

Large Flow Offloads

On Firepower 4100/9300 chassis, certain traffic that you configure to be fastpathed by a prefilter policy is handled by the hardware (specifically, in the NIC), not by your threat defense software. Offloading these connection flows results in higher throughput and lower latency, especially for data-intensive applications such as large file transfers. This feature is especially useful for data centers. This is called *static flow offload*.

In addition, by default, threat defense devices offload flows based on other criteria, including trust. This is called *dynamic flow offload*.

Offloaded flows continue to receive limited stateful inspection, such as basic TCP flag and option checking. The system can selectively escalate packets to the firewall system for further processing if necessary.

Examples of applications that can benefit from offloading large flows are:

- High Performance Computing (HPC) Research sites, where the threat defense device is deployed between storage and high compute stations. When one research site backs up using FTP file transfer or file sync over NFS, the large amount of data traffic affects all connections. Offloading FTP file transfer and file sync over NFS reduces the impact on other traffic.
- High Frequency Trading (HFT), where the threat defense device is deployed between workstations and the Exchange, mainly for compliance purposes. Security is usually not a concern, but latency is a major concern.

The following flows can be offloaded:

- (Static flow offload only.) Connections that are fastpathed by the prefilter policy.
- Standard or 802.1Q tagged Ethernet frames only.
- (Dynamic flow offload only):
  - Inspected flows that the inspection engine decides no longer need inspection. These flows include:
    - Flows handled by access control rules that apply the Trust action and are based on security zone, source and destination network and port matching only.
    - TLS/SSL flows that are not selected for decryption using an SSL policy.
    - Flows that are trusted by the Intelligent Application Bypass (IAB) policy either explicitly or due to exceeding flow bypass thresholds.
    - Flows that match file or intrusion policies that result in trusting the flow.
    - Any allowed flow that no longer needs to be inspected.
  - The following IPS preprocessor inspected flows:
    - SSH and SMTP.
    - FTP preprocessor secondary connections.
    - Session Initiation Protocol (SIP) preprocessor secondary connections.
  - Intrusion rules that use keywords (also referred to as *options*)

Important For details, exceptions, and limitations to the above, see Flow Offload Limitations, on page 20.

Use Static Flow Offload

To offload eligible traffic to hardware, create a prefilter policy rule that applies the **Fastpath** action. Use prefilter rules for TCP/UDP, and tunnel rules for GRE.

(Not recommended.) To disable static flow offload and as a by-product, dynamic flow-offload, use FlexConfig to run the **no flow-offload enable** command. For information about this command, see the *Cisco ASA Series Command Reference*, available from https://www.cisco.com/c/en/us/support/security/adaptive-security-appliance-asa-software/products-command-reference-list.html.

Use Dynamic Flow Offload

Dynamic flow offload is enabled by default.

To disable dynamic offload:

```
> configure flow-offload dynamic whitelist disable
```

To re-enable dynamic offload:

```
> configure flow-offload dynamic whitelist enable
```

Note that dynamic offload occurs only if static flow offload is enabled, regardless of whether prefiltering is configured.

Flow Offload Limitations

Not all flows can be offloaded. Even after offload, a flow can be removed from being offloaded under certain conditions. Following are some of the limitations:

Device Limitations

The feature is supported on the following devices:

- Firepower 4100/9300 running FXOS 1.1.3 or higher.

Flows that cannot be offloaded

The following types of flows cannot be offloaded.

- Any flows that do not use IPv4 addressing, such as IPv6 addressing.
- Flows for any protocol other than TCP, UDP, and GRE.

Note PPTP GRE connections cannot be offloaded.

- Flows on interfaces configured in passive, inline, or inline tap mode. Routed and switched interfaces are the only types supported.
- Flows that require inspection by Snort or other inspection engines. In some cases, such as FTP, the secondary data channel can be offloaded although the control channel cannot be offloaded.

- IPsec and TLS/DTLS VPN connections that terminate on the device.
- Flows that require encryption or decryption. For example, connections decrypted due to an SSL policy.
- Multicast flows in routed mode. They are supported in transparent mode if there are only two member interfaces in a bridge group.
- TCP Intercept flows.
- TCP state bypass flows. You cannot configure flow offload and TCP state bypass on the same traffic.
- Flows tagged with security groups.
- Reverse flows that are forwarded from a different cluster node, in case of asymmetric flows in a cluster.
- Centralized flows in a cluster, if the flow owner is not the control unit.
- Flows that include IP options cannot be dynamically offloaded.

Additional Limitations

- Flow offload and Dead Connection Detection (DCD) are not compatible. Do not configure DCD on connections that can be offloaded.
- If more than one flow that matches flow offload conditions are queued to be offloaded at the same time to the same location on the hardware, only the first flow is offloaded. The other flows are processed normally. This is called a *collision*. Use the **show flow-offload flow** command in the CLI to display statistics for this situation.
- Dynamic flow offload disables all TCP normalizer checks.
- Although offloaded flows pass through FXOS interfaces, statistics for these flows do not appear on the logical device interface. Thus, logical device interface counters and packet rates do not reflect offloaded flows.

Conditions for reversing offload

After a flow is offloaded, packets within the flow are returned to the threat defense for further processing if they meet the following conditions:

- They include TCP options other than Timestamp.
- They are fragmented.
- They are subject to Equal-Cost Multi-Path (ECMP) routing, and ingress packets move from one interface to another.

History for Prefiltering

Feature	Minimum Management Center	Minimum Threat Defense	Details
Moving prefilter rules to an access control policy	6.7	Any	You can move prefilter rules from a prefilter policy to the associated access control policy. New/modified pages: In the prefilter policy page, the right-click menu for the selected rules provides a new **Move to another policy** option. Supported platforms: management center
Time-based rules	6.6	Any	Ability to apply prefilter and tunnel rules depending on the date and time, as determined by the time zone of the threat defense device. See description in History for Access Control Rules.
View Object Details from prefilter rule page	6.6	Any	Feature introduced: Option to view details for an object or object group when viewing prefilter rules. New options: Right-clicking a value in any of the following columns in the prefilter rule list offers an option to view object details: Source Networks, Destination Networks, Source Port, Destination Port, and VLAN Tag. Supported platforms: Secure Firewall Management Center