



Personal Stateful Firewall Overview

This chapter provides an overview of the Personal Stateful Firewall In-line Service.

This chapter covers the following topics:

- [Firewall Overview, on page 1](#)
- [Supported Features, on page 2](#)
- [How Personal Stateful Firewall Works, on page 13](#)
- [Understanding Rules with Stateful Inspection, on page 16](#)

Firewall Overview

The Personal Stateful Firewall is an in-line service feature that inspects subscriber traffic and performs IP session-based access control of individual subscriber sessions to protect the subscribers from malicious security attacks.

The Personal Stateful Firewall in-line service works in conjunction with the following products:

- GGSN
- HA
- PDSN
- P-GW

The Personal Stateful Firewall supports stateless and stateful inspection and filtering based on the configuration.

In stateless inspection, the firewall inspects a packet to determine the 5-tuple—source and destination IP addresses and ports, and protocol—information contained in the packet. This static information is then compared against configurable rules to determine whether to allow or drop the packet. In stateless inspection the firewall examines each packet individually, it is unaware of the packets that have passed through before it, and has no way of knowing if any given packet is part of an existing connection, is trying to establish a new connection, or is a rogue packet.

In stateful inspection, the firewall not only inspects packets up through the application layer / layer 7 determining a packet's header information and data content, but also monitors and keeps track of the connection's state. For all active connections traversing the firewall, the state information, which may include IP addresses and ports involved, the sequence numbers and acknowledgement numbers of the packets traversing the connection, TCP packet flags, etc. is maintained in a state table. Filtering decisions are based not only on rules but also on the connection state established by prior packets on that connection. This enables to prevent a variety of DoS, DDoS, and other security violations. Once a connection is torn down, or is timed out, its entry in the

state table is discarded. For more information see the [Connection State and State Table in Personal Stateful Firewall, on page 17](#) section.

The Enhanced Charging Service (ECS) / Active Charging Service (ACS) in-line service is the primary vehicle that performs packet inspection and charging. For more information on ECS, see the *Enhanced Charging Service Administration Guide*.

Qualified Platforms

PSF is a StarOS in-line service application that runs on Cisco ASR 5500 and virtualized platforms. For additional platform information, refer to the appropriate *System Administration Guide* and/or contact your Cisco account representative.

License Requirements

The Personal Stateful Firewall is a licensed Cisco feature. A separate feature license may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

Supported Features

The Personal Stateful Firewall supports the following features:

Protection against Denial-of-Service Attacks

Denial-of-Service (DoS) and Distributed Denial-of-Service (DDoS) attacks can deprive network resources/services unavailable to its intended users.

DoS attacks can result in:

- A host consuming excessive resources — memory, disk space, CPU time, etc. — eventually leading to a system crash or providing very sluggish response.
- Flooding of the network to the extent that no valid traffic is able to reach the intended destination.
- Confusing target TCP/IP stack on destination hosts by sending crafted, malformed packets eventually resulting in system crash.

In this release, malformed check is enhanced for IPv6 and ICMPv6 packets. Port scan and Flooding attacks are also enhanced to support IPv6. Protection against other L4 attacks are similar to IPv4. The Attacking server feature is also enhanced to store IPv6 servers.

DoS attacks can destroy data in affected mobile nodes. Stateful Firewall is designed to defend subscribers and prevent the abuse of network bandwidth from DoS attacks originating from both the Internet and the internal network.

Types of Denial-of-Service Attacks

Personal Stateful Firewall can detect the following DoS attacks.

The DoS attacks are listed based on the protocol layer that they work on.

- IP-based Attacks:
 - Land attacks

- Jolt attacks — Detected in downlink and uplink direction
- Teardrop attacks — Detected in downlink and uplink direction
- Invalid IP option length — Detected in downlink and uplink direction
- IP-unaligned-timestamp attack — Detected only in downlink direction
- Short IP header length
- IP checksum errors
- IP reassembly failure (downlink)
- IP reassembly failure (uplink)
- Source router — Detected only in downlink direction
- IPv6 header checks
- Source-IP based flood attack — Detected in downlink and uplink directions

- TCP-based Attacks:
 - Data packets received after RST/FIN
 - Invalid SEQ number received with RST
 - Data without connection established
 - Invalid TCP connection requests
 - Invalid TCP pre-connection requests
 - Invalid ACK value (cookie enabled)
 - Invalid TCP packet length
 - Short TCP header length
 - TCP checksum errors
 - SEQ/ACK out-of-range
 - TCP null scan attacks
 - Post connection SYN
 - No TCP flags set
 - All TCP flags set
 - Invalid TCP packets
 - Flows closed by RST before 3-Way handshake
 - Flows timed-out in SYN_RCVD1 state
 - Flows timed-out in SYN_RCVD2 state
 - TCP-SYN flood attacks — Detected in downlink and uplink direction
 - FTP bounce attack — Detected only in downlink direction
 - MIME flood attacks — Detected in downlink and uplink direction
 - Exceeding reset message threshold — Detected in downlink and uplink direction
 - Source port zero
 - WinNuke attack — Detected in downlink and uplink direction
 - TCP-window-containment — Detected only in downlink direction

- UDP-based Attacks:
 - Invalid UDP echo response
 - Invalid UDP packet length
 - UDP checksum errors
 - Short UDP header length
 - UDP flood attack — Detected in downlink and uplink direction

- ICMP-based Attacks:
 - Invalid ICMP response
 - ICMP reply error
 - Invalid ICMP type packet
 - ICMP error message replay attacks
 - ICMP packets with duplicate sequence number
 - Short ICMP header length
 - Invalid ICMP packet length
 - ICMP flood attack — Detected in downlink and uplink direction
 - Ping of death attacks — Detected in downlink and uplink direction
 - ICMP checksum errors
 - ICMP packets with destination unreachable message
 - ICMP echo packets with ID zero

- Other DoS Attacks:
 - Port scan attacks — Detected in downlink and uplink direction

Various header integrity checks are performed for IPv6 to ensure the integrity of an IPv6 packet. IPv6 packets with unknown extension headers will not be dropped by Firewall; such packets will be allowed by Firewall. Firewall performs the following header checks:

- Limiting extension headers
- Hop-by-hop Options filtering
- Destination Options filtering
- Router Header filtering
- Fragment Header filtering

Source-IP based Flood Attack Detection

The Source-IP based flood attack detection feature is implemented to limit the number of new flows originating from a source IP per unit time to various destinations in both uplink and downlink directions. Limiting the number of connections from a single source IP to different destinations will be applied only per SMGR instance. CLI support is added for enabling and disabling uplink/downlink IP Sweep protection. Statistics support is added to show details of packets dropped due to IP Sweep protection and the list of internet hosts doing a flood attack.

An Internet host can flood mobiles by sending packets while firewall downlink action is configured as “permit”. At present there is no mechanism to prevent flood attacks from a host going to multiple destinations.

Protection against Port Scanning

Port scanning is a technique used to determine the states of TCP/UDP ports on a network host, and to map out hosts on a network. Essentially, a port scan consists of sending a message to each port on the host, one at a time. The kind of response received indicates whether the port is used, and can therefore be probed further for weakness. This way hackers find potential weaknesses that can be exploited.

Port scans are mainly classified into two types:

- **Horizontal scan:** In the case of Horizontal scanning, a port scanner probes multiple destination addresses (i.e. hosts) for the same port for profiling active hosts.

- **Vertical scan:** In the case of Vertical scanning, a port scanner probes a set of ports on the same machine to find out which services are running on the machine.

Vertical scans get detected easily when compared to horizontal scans. Stateful Firewall provides protection against port scanning by implementing port scan detection algorithms. Port scan attacks are detected in both the downlink and uplink direction—traffic from external network towards mobile subscribers. Horizontal scan detection is limited to a particular session manager in the downlink direction. Detection in the uplink direction will be supported in a future release.

Limitations:

- The maximum number of requests analyzed per source need to be limited due to memory constraints and this may result in false negatives.

Application-level Gateway Support

A stateful firewall while ensuring that only legitimate connections are allowed, also maintains the state of an allowed connection. Some network applications require additional connections to be opened up in either direction and information regarding such connections is sent in the application payload. For these applications to work properly, a stateful firewall must inspect, analyze, and parse these application payloads to get the additional connection information, and open partial connections/pinholes in the firewall to allow the connections.

To parse application payloads, firewall employs ALGs. ALGs also check for application-level attacks. Stateful Firewall provides ALG functionality for the following protocols. ALG support for Simple Mail Transfer Protocol (SMTP) and HTTP is ECS functionality.

Table 1: Protocol Support for ALG Type

ALG Type	Support in IPv4 Firewall	Support in IPv6 Firewall
File Transfer Protocol (FTP)	Yes	Yes
H323	Yes	No
Point-to-Point Tunneling Protocol (PPTP)	Yes	Yes
Real Time Streaming Protocol (RTSP)	Yes	Yes
Session Initiation Protocol (SIP TCP)	Yes	Yes
Session Initiation Protocol (SIP UDP)	Yes	Yes
Trivial File Transfer Protocol (TFTP)	Yes	Yes

PPTP ALG Support

PPTP exchanges IP or port specific information over its control connection and that information will be used to transfer the data over tunnel. If a PPTP client resides behind NAT and uses private IP to communicate with the outside world, it is possible that the information exchange over PPTP control flow consists of private IPs. So NAT translates the private IP specific information to public IP (NATed IP) for good communication. To achieve this, PPTP ALG is supported.

To establish a GRE session, PPTP exchanges call IDs from both peers to form a unique triple value, that is, client IP, server IP and Call ID. For Many-to-One NAT, PPTP analyzer is implemented to analyze the PPTP Control Flow traffic. It can be configured to send all the PPTP Control Flow packets to PPTP analyzer. PPTP analyzer analyzes the packet and allocates a new unique Call ID. Packet payload will be modified for the new

Call ID and the binding between the two Call IDs will be maintained. Similarly, the PPTP first packet will be NAT-ed, Call ID translated and sent to the PPTP Server. This Call ID translation happens for all the downlink packets after the first packet. For GRE Data Tunnel Flow translation, it can be configured to send all the GRE downlink packets to PPTP analyzer. PPTP analyzer then analyzes the GRE header and translates the GRE Call ID if a Call ID binding exists.



Important

In 9.0 and earlier releases, GRE flow was detected and identified as a 3-tuple. In later releases, GRE flow is identified as a 5-tuple as part of the PPTP analyzer implementation. For the PPTP traffic to work after upgrade, a PPTP analyzer need to be configured in the rulebase or add a downlink allow rule for GRE traffic in Firewall-and-NAT policy.

TFTP ALG Support

Trivial File Transfer Protocol (TFTP) ALG enables Firewall or NAT enabled users to seamlessly use applications using TFTP Protocol. TFTP ALG feature analyzes the TFTP packets and selectively allows the downlink data flow by creating pin holes. This feature also ensures NAT/PAT IP/Port translation for NAT enabled users.

TFTP ALG analyzes the packets for basic TFTP signatures. A TFTP analyzer is implemented for this purpose. A routing rule is created for routing the packets to TFTP analyzer. Potential TFTP packets are parsed and information like query type and mode are stored. After confirming that the packet is TFTP, a dynamic route is created for MS IP, MS Port, Server IP and Protocol. When the data flow starts, dynamic route is matched and data is sent to the TFTP analyzer. For NAT enabled calls, same Client port used for the control connection will be used for Data flow.

Stateful Packet Inspection and Filtering Support

As described in the Overview section, stateful packet inspection and filtering uses Layer-4 information as well as the application-level commands up to Layer-7 to provide good definition of the individual connection states to defend from malicious security attacks.

Personal Stateful Firewall overcomes the disadvantages of static packet filters by disallowing any incoming packets that have the TCP SYN flag set (which means a host is trying to initiate a new connection). If configured, stateful packet filtering allows only packets for new connections initiated from internal hosts to external hosts and disallows packets for new connections initiated from external hosts to internal hosts. TCP packets with SYN flag set and destination port zero will be dropped like TCP packets with source port zero are dropped.

TCP stateful processing is enhanced for processing IPv6 packets. The functionality is similar to IPv4 packets.

Stateless Packet Inspection and Filtering Support

Stateful Firewall service can be configured for stateless processing. In stateless processing, packets are inspected and processed individually.

Stateless processing is only applicable for TCP and ICMP protocols. By nature UDP is a stateless protocol without any kind of acking or request and reply mechanism at transport level.

When TCP FSM is disabled, flows can start with any kind of packet and need not respect the TCP FSM. Such flows are marked as dummy (equivalent to flows established during flow recovery timer running). For these

flows only packet header check is done; there will be no FSM checks, sequence number validations, or port scan checks done.

When ICMP FSM is disabled, ICMP reply without corresponding requests, ICMP error message without inner packet data session, and duplicate ICMP requests are allowed by firewall.

Host Pool, IMSI Pool, and Port Map Support

This section describes the Host Pool, IMSI Pool, and Port Map features that can be used while configuring access ruledefs.

Host Pool Support

Host pools allow operators to group a set of host or IP addresses that share similar characteristics together. Access rule definitions (ruledefs) can be configured with host pools. Up to 10 sets of IP addresses can be configured in each host pool. Host pools are configured in the ACS Host Pool Configuration Mode.

Host pools are enhanced to support IPv6 addresses and address ranges. It can also be a combination of IPv4 and IPv6 addresses.

IMSI Pool Support

IMSI pools allow the operator to group a set of International Mobile Station Identifier (IMSI) numbers together. Up to 10 sets of IMSI numbers can be configured in each IMSI pool. IMSI pools are configured in the ACS IMSI Pool Configuration Mode.

Port Map Support

Port maps allow the operator to group a set of port numbers together. Access ruledefs can be configured with port maps. Up to 10 sets of ports can be configured in each port map. Port maps are configured in the ACS Port Map Configuration Mode.

The Personal Stateful Firewall uses standard application ports to trigger ALG functionality. The operator can modify the existing set to remove/add new port numbers.

Port Control Protocol Support

The Port Control Protocol (PCP) feature provides a mechanism to control how incoming packets are forwarded by upstream devices such as Network Address Translation IPv4/IPv4 (NAT44) and IPv4 firewall devices, and to reduce application keepalive traffic.



Important

The PCP feature is customer specific. Contact your Cisco account representative for more information.



Important


PCP is a licensed Cisco feature. Contact your Cisco account representative for more information. A separate feature license may be required. Contact your Cisco account representative for detailed information on specific licensing requirements. For information on installing and verifying licenses, refer to the *Managing License Keys* section of the *Software Management Operations* chapter in the *System Administration Guide*.

The PCP server is supported on ASR 5500 chassis running in-line services such as NAT44 and Stateful Firewall(s) individually or in collocated configurations. PCP supports the following functions:

- A host to control how incoming packets are forwarded by upstream devices such as Network Address Translation (NAT44) and Stateful Firewall (IPv4).
- A host to reduce the application keepalive messages.
- A host to operate a server for a long duration (e.g. webcam) or a short duration (e.g. while playing a game or on a phone call) when behind a NAT device, including a CGN operated by an Internet service provider.
- Applications to create mappings from an external IP address and port to an internal (target) IP address and port. These mappings are required for successful inbound communications destined to machines located behind a NAT or Firewall.
- Applications to create mappings in NAT and Firewall, and reducing the incentive to deploy ALGs in NAT and Firewalls.

The following figure shows IPv4 Firewall and PCP Server on ASR 5500.

Figure 1: IPv4 Firewall and PCP Server



The PCP service has to be associated with a PCP server IP address. The PCP server IP address is picked from the destination context associated with the subscriber. Only, if such an IP address is available and the status is up, the PCP service will listen to PCP requests on that IP address. The PCP service will be bound only to an IPv4 address and listens on UDP port (5351 (default port) or can be configured).

In case of system failure, the PCP service recovers along with subscriber's PCP enabled status. In case of stand-alone recovery and ICSR, only the subscriber PCP enabled status will be check-pointed.

PCP supports interworking with the following existing NAT ALGs:

- FTP
- RTSP
- SIP

Bulk Statistics Support for the Port Control Protocol Feature

Bulk statistics reporting for the PCP feature is supported.

For the PCP feature the following bulk statistics are available in the ECS schema:

- total-pcp-svc-req
- total-pcp-svc-rsp

- total-pcp-svc-unknown-rsp
- total-pcp-svc-invalid-rsp
- total-pcp-svc-map-req
- total-pcp-svc-map-valid-req
- total-pcp-svc-map-invalid-req
- total-pcp-svc-map-rsp
- total-pcp-svc-map-rsp-success
- total-pcp-svc-peer-rsp-error
- total-pcp-svc-peer-req
- total-pcp-svc-peer-valid-req
- total-pcp-svc-peer-invalid-req
- total-pcp-svc-peer-rsp
- total-pcp-svc-peer-rsp-success
- total-pcp-svc-peer-rsp-error
- total-pcp-svc-announce-req
- total-pcp-svc-announce-vaild-req
- total-pcp-svc-announce-invaild-req
- total-pcp-svc-announce-rsp
- total-pcp-svc-announce-rsp-success
- total-pcp-svc-announce-rsp-error
- total-pcp-svc-subscribers
- current-pcp-svc-subscribers

Flow Recovery Support

Stateful Firewall supports call recovery during session failover. Flows associated with the calls are recovered.

A recovery-timeout parameter is configurable for uplink and downlink directions. If the value is set to zero, firewall flow recovery is disabled. If the value is non-zero, then firewall will be bypassed for packets from MS/Internet until the time configured (uplink/downlink). Once the manager recovers, the recovery-timeout timer is started. During this time:

- If any ongoing traffic arrives from the subscriber and no association is found, and flow recovery is enabled, basic checks like header processing, attacks, etc. are done (stateful checks of packet is not done), and if all is okay, an association is created and the packet is allowed to pass through.
- If any ongoing traffic arrives from the Internet to MS and no association is found, and flow recovery is not enabled, it is dropped. No RESET is sent. Else, basic checks like header processing, flooding attack check are done (stateful checks are not done), and if all is okay, an association is created and the packet is allowed to pass through.
- In case flow recovered from ongoing traffic arrives from Internet to MS, and MS sends a NACK, the Unwanted Traffic Suppression feature is triggered, i.e. upon repeatedly receiving NACK from MS for a 5-tuple, further traffic to the 5-tuple is blocked for some duration and not sent to MS.
- If any new traffic (3-way handshake) comes, whether it is a new flow or a new flow due to pin-hole, based on the direction of packet and flow-recovery is enabled, basic checks like header processing, attacks, etc. are done (stateful checks are not done) and if all is okay, an association is created and the packet is allowed to pass through.

For any traffic coming after the recovery-timeout:

- If any ongoing traffic arrives, it is allowed only if an association was created earlier. Else, it is dropped and reset is sent.
- If any new traffic (3-way handshake) arrives, the usual Stateful Firewall processing is done.

If the recovery-timeout value is set to zero, Stateful Firewall flow recovery is not done.

Stateful Firewall now supports IPv6 flows recovery similar to IPv4 flows.

ICSR Support for Dynamic Firewall Access Rules

ICSR recovery is supported for dynamic Firewall access rules. Firewall access rules can be enabled either statically or dynamically. After ICSR switchover, the Firewall access rules can be dynamically activated or deactivated from the Gx server. This feature currently works only for default bearers and not for dedicated bearers.

The following attributes are used to activate or deactivate Firewall access rules from the Gx server:

- Charging-Rule-Install
- Charging-Rule-Remove

Earlier, Firewall access rules enabled dynamically by PCRF were checkpointed only for standalone recovery. ICSR checkpointing was not done for dynamically enabled access rules. After ICSR switchover, the dynamically enabled Firewall rule will be enabled.

SNMP Thresholding Support

Personal Stateful Firewall allows to configure thresholds to receive notifications for various events that are happening in the system. Whenever a measured value crosses the specified threshold value at the given time, an alarm is generated. And, whenever a measured value falls below the specified threshold clear value at the given time, a clear alarm is generated. The following events are supported for generating and clearing alarms:

- Dos-Attacks: When the number of DoS attacks crosses a given value, a threshold is raised, and it is cleared when the number of DoS attacks falls below a value in a given period of time.
- Drop-Packets: When the number of dropped packets crosses a given value, a threshold is raised, and it is cleared when the number of dropped packets falls below a value in a given period of time.
- Deny-Rule: When the number of Deny Rules cross a given value, a threshold is raised, and it is cleared when the number of Deny Rules falls below a value in a given period of time.
- No-Rule: When the number of No Rules cross a given value, a threshold is raised, and it is cleared when the number of No Rules falls below a value in a given period of time.

Logging Support

Stateful Firewall supports logging of various messages on screen if logging is enabled for firewall. These logs provide detailed messages at various levels, like critical, error, warning, and debug. All the logs displaying IP addresses are enhanced to display IPv6 addresses.

Logging is also supported at rule level, when enabled through rule a message will be logging whenever a packet hits the rule. This can be turned on/off in a rule.

These logs are also sent to a syslog server if configured in the system.

Enhanced Syslog Reporting

Feature Summary and Revision History

Summary Data

Applicable Product(s) or Functional Area	All Products
Applicable Platform(s)	<ul style="list-style-type: none"> • ASR 5500 • VPC - Di • VPC - Si
Feature Default	Disabled - Configuration Required
Related Changes in This Release	Not Applicable
Related Documentation	<ul style="list-style-type: none"> • <i>ASR 5500 System Administration Guide</i> • <i>NAT Administration Guide</i> • <i>PSF Administration Guide</i>

Revision History

Revision Details	Release
With this feature, the port information of the dropped packet is included in the logs.	21.3.1
First introduced.	Pre 21.2

Feature Description

Firewall and NAT attack logs provide information on the source IP address, destination IP address, protocol, or attack type for any packet dropped due to an attack. Prior to this release, when an attack happened, the logs did not carry any information about the ports.

With this feature, the port information of the dropped packet is included in the logs. The port information such as source port and destination port are important while configuring access rules to allow or block certain traffic.

Following are some important points to be considered:

- Typically, attack logs are at INFO/DEBUG level. At this level, there are too many logs generated even for normal traffic. Hence, to segregate the attack logs, the attack logs are moved to the WARNING level.
- Firewall and NAT attack logs are moved to WARNING level from Info/Debug level. The source port and destination port are logged as part of Firewall/NAT attack logs.
- Both IPv4 and IPv6 traffic is supported.

- The source port and destination port are valid for TCP/UDP protocols. However, for other protocols, the ports are logged as zero.

Previous Behavior: Earlier, the attack logs did not carry any port information and the logs were of the type Info/Debug.

New Behavior: With this feature, firewall and NAT attack log levels have been changed to WARNING from INFO/DEBUG for event IDs 96188, 96995, 96186, 96185, 96159, and 96203. Source port and destination port information are now displayed in the attack logs.

Impact on Customer: The attack logs are per packet logs seen at WARNING level. If you enable WARNING and above logs for Firewall (NAT) facility, and when there is an attack, log rate is very high.

Limitations

Following are the limitations of this feature:

- The attack logs are per packet logs and when an attack is in progress, log generation rate is very high.
- Under extreme attack conditions, evlogd CPU might go up.
- Event IDs for the attack logs:
 - firewall 96188 warning
 - firewall 96995 warning
 - firewall 96186 warning
 - firewall 96203 warning
 - firewall 96159 warning
 - firewall 96185 warning
- When there are too many logs generated under attack conditions, the following event IDs must be disabled:
 - eventid 96188: Disables Firewall Attack log generation
 - eventid 96186: Disables Port Scan Attack log generation
 - eventid 96995: Disables NAT Attack log generation
 - eventid 96203: Disables logging for TCP reset message threshold breach
 - eventid 96159: Disables logging for packets denied by rule
 - eventid 96185: Disables logging for ICMP unreachable message threshold breach

Configuring Logging Event ID

When there are too many logs generated under attack conditions, use the following command to disable the event IDs:

- To disable firewall attack log generation:


```
[local]asr5500(config)# logging disable eventid 96188
```
- To disable port scan log generation:

```
[local]asr5500(config)# logging disable eventid 96186
```

- To disable NAT attack log generation:

```
[local]asr5500(config)# logging disable eventid 96995
```

- To disable logging for TCP reset message threshold breach:

```
[local]asr5500(config)# logging disable eventid 96203
```

- To disable logging for packets denied by rule:

```
[local]asr5500(config)# logging disable eventid 96159
```

- To disable logging for ICMP unreachable message threshold breach:

```
[local]asr5500(config)# logging disable eventid 96185
```

How Personal Stateful Firewall Works

This section describes how Personal Stateful Firewall works.



Important

In release 8.x, Stateful Firewall for CDMA and early UMTS releases used rulebase-based configurations, whereas later UMTS releases used policy-based configurations. In release 9.0, Stateful Firewall for UMTS and CDMA releases, both use policy-based configurations. For more information, please contact your local service representative.

Firewall-and-NAT policies are configured in the Firewall-and-NAT Policy Configuration Mode. Each policy contains a set of access ruledefs and the firewall configurations. Multiple such policies can be configured, however, only one policy is applied to a subscriber at any point of time. A Firewall-and-NAT policy defined with no Firewall or NAT enabled will not drop the call but will disable the configured Firewall or NAT.

The policy used for a subscriber can be changed either from the CLI, or by dynamic update of policy name in Diameter and RADIUS messages.

The Firewall-and-NAT policy to be used for a subscriber can be configured in:

- ACS Rulebase: The default Firewall-and-NAT policy configured in the ACS rulebase has the least priority. If there is no policy configured in the APN/subscriber template, and/or no policy to use is received from the AAA/OCS, only then the default policy configured in the ACS rulebase is used.
- APN/Subscriber Template: The Firewall-and-NAT policy configured in the APN/subscriber template overrides the default policy configured in the ACS rulebase. To use the default policy configured in the ACS rulebase, in the APN/subscriber configuration, the command to use the default rulebase policy must be configured.
- AAA/OCS: The Firewall-and-NAT policy to be used can come from the AAA server or the OCS. If the policy comes from the AAA/OCS, it will override the policy configured in the APN/subscriber template and/or the ACS rulebase.



Important

The Firewall-and-NAT policy received from the AAA and OCS have the same priority. Whichever comes latest, either from AAA/OCS, is applied.

The Firewall-and-NAT policy to use can be received from RADIUS during authentication.

Disabling Firewall Policy

Stateful Firewall processing is disabled for subscribers in the following cases.



Important

By default, Stateful Firewall processing for subscribers is disabled.

- If Stateful Firewall is explicitly disabled in the APN/subscriber template configuration.
- If the AAA/OCS sends the SN-Firewall-Policy AVP with the string “disable”, the locally configured firewall policy does not get applied.
- If the SN-Firewall-Policy AVP is received with the string “NULL”, the existing policy will continue.
- If the SN-Firewall-Policy AVP is received with a name that is not configured locally, the subscriber session is terminated.

Mid-session Firewall Policy Update

The Firewall-and-NAT policy can be updated mid-session provided firewall policy was enabled during call setup. Firewall-and-NAT policy can also be updated during mid-session rulebase update through Gx and Gy if the new rulebase has Firewall-and-NAT policy configured and the old Firewall-and-NAT policy is configured through old rulebase.



Important

When the SN-Firewall-Policy AVP contains “disable” during mid-session firewall policy change, there will be no action taken as the Firewall-and-NAT policy cannot be disabled dynamically. The policy currently applied will continue.



Important

When a Firewall-and-NAT policy is deleted, for all subscribers using the policy, Firewall processing is disabled, also ECS sessions for the subscribers are dropped. In case of session recovery, the calls are recovered but with Stateful Firewall disabled.


Firewall-and-NAT Checkpointing

The “Firewall-and-NAT Policy ID” can be used for checkpointing the Firewall-and-NAT Policy unlike “Firewall-and-NAT Policy-Name” used in previous releases. Backward compatibility has been provided to support checkpointing the Firewall-and-NAT Policy-Name by older versions of peer-chassis. If the older chassis sends Policy-name in checkpointing, then this has to be handled appropriately by the new chassis.

Packet Flow in Stateful Firewall Processing


The following figures illustrate packet flow in Stateful Firewall processing for a subscriber.

Figure 2: Stateful Firewall Processing



335214

Figure 3: Stateful Firewall Processing (Contd.)



335215

Understanding Rules with Stateful Inspection

This section describes terms used in the Personal Stateful Firewall context.

- **Access Ruledefs:** The Personal Stateful Firewall's stateful packet inspection feature allows operators to configure rule definitions (ruledefs) that take active session information into consideration to permit or deny incoming or outgoing packets.

An access ruledef contains the criteria for multiple actions that could be taken on packets matching the rules. These rules specify the protocols, source and destination hosts, source and destination ports, direction of traffic parameters for a subscriber session to allow or reject the traffic flow.

An access ruledef consists of the following fields:

- Ruledef name
- Source IP address
- Source port number — not required if the protocol is other than TCP or UDP
- Destination IP address
- Destination port number — not required if the protocol is other than TCP or UDP
- Transport protocol (TCP/UDP/ICMP/ICMPv6/AH/ESP)
- Direction of connection (Uplink/Downlink)
- Bearer (IMSI-pool and APN)
- Logging action (enable/disable)
- IP version - IPv4 or IPv6
- Client port number — TCP or UDP
- Server port number — TCP or UDP

An access ruledef can be added to multiple Firewall-and-NAT policies.

A combined maximum of 4096 rules (host pools + IMSI pools + port maps + charging ruledefs + firewall/access ruledefs + routing ruledefs) can be created in a system. Access ruledefs are different from ACS ruledefs.

Firewall access ruledefs support IPv6 addresses and parameters like IP version and ICMPv6 protocol. The existing rule lines "ip src-address" and "ip dst-address" are capable of accepting both IPv4 and IPv6 addresses.

Access ruledefs can be switched on/off from PCRF (Gx). Charging ruledef attributes can be used for this purpose. Access ruledefs that need to be switched on/off must be configured as "dynamic-only" or "static-and-dynamic" in Firewall-and-NAT policy. If configured as "dynamic-only", the rule will be disabled by default and can be switched on from PCRF. If configured as "static-and-dynamic", the rule will behave as "dynamic-only" for Gx enabled call and as static rule for non-Gx calls.

- **Firewall-and-NAT Policy:** Firewall policies can be created for individual subscribers, domains, or all callers within a referenced context. Each policy contains a set of access ruledefs with priorities defined for each rule and the firewall configurations. Firewall-and-NAT policies are configured in the Firewall-and-NAT Policy Configuration Mode.
- **Service Definition:** User-defined firewall service for defining Stateful Firewall policy for initiating an outgoing connection on a primary port and allowing opening of auxiliary ports for that association in the reverse direction.
- **Maximum Association:** The maximum number of Stateful Firewall associations for a subscriber.

Connection State and State Table in Personal Stateful Firewall

This section describes the state table and different connection states for transport and network protocols.

After packet inspection, the Personal Stateful Firewall stores session state and other information into a table. This state table contains entries of all the communication sessions of which the firewall subsystem is aware of. Every entry in this table holds a list of information that identifies the subscriber session it represents. Generally this information includes the source and destination IP address, flags, sequence, acknowledgement numbers, etc.

When a connection is permitted through the Personal Stateful Firewall enabled chassis, a state entry is created. If a session connection with same information (source address, source port, destination address, destination port, protocol) is requested the firewall subsystem compares the packet's information to the state table entry to determine the validity of session. If the packet is currently in a table entry, it allows it to pass, otherwise it is dropped.

Transport and Network Protocols and States

Transport protocols have their connection's state tracked in various ways. Many attributes, including IP address and port combination, sequence numbers, and flags are used to track the individual connection. The combination of this information is kept as a hash in the state table.

TCP Protocol and Connection State

TCP is considered as a stateful connection-oriented protocol that has well defined session connection states. TCP tracks the state of its connections with flags as defined for TCP protocol. The following table describes different TCP connection states.

Table 2: TCP Connection States

State Flag	Description
TCP (Establishing Connection)	
CLOSED	A “non-state” that exists before a connection actually begins.
LISTEN	The state a host is in waiting for a request to start a connection. This is the starting state of a TCP connection.
SYN-SENT	The time after a host has sent out a SYN packet and is waiting for the proper SYN-ACK reply.
SYN-RCVD	The state a host is in after receiving a SYN packet and replying with its SYN-ACK reply.
ESTABLISHED	The state a host is in after its necessary ACK packet has been received. The initiating host goes into this state after receiving a SYN-ACK.
TCP (Closing Connection)	
FIN-WAIT-1	The state a connection is in after it has sent an initial FIN packet asking for a graceful termination of the TCP connection.
CLOSE-WAIT	The state a host's connection is in after it receives an initial FIN and sends back an ACK to acknowledge the FIN.

State Flag	Description
FIN-WAIT-2	The connection state of the host that has received the ACK response to its initial FIN, as it waits for a final FIN from its connection peer.
LAST-ACK	The state of the host that just sent the second FIN needed to gracefully close the TCP connection back to the initiating host while it waits for an acknowledgement.
TIME-WAIT	The state of the initiating host that received the final FIN and has sent an ACK to close the connection and waiting for an acknowledgement of ACK from the connection peer. Note that the amount of time the TIME-STATE is defined to pause is equal to the twice of the Maximum Segment Lifetime (MSL), as defined for the TCP implementation.
CLOSING	A state that is employed when a connection uses the unexpected simultaneous close.

UDP Protocol and Connection State

UDP is a connection-less transport protocol. Due to its connection-less nature, tracking of its state is a more complicated process than TCP. The Personal Stateful Firewall tracks a UDP connection in a different manner than TCP. A UDP packet has no sequence number or flag field in it. The port numbers used in UDP packet flow change randomly for any given session connection. So the Personal Stateful Firewall keeps the status of IP addresses.

UDP traffic cannot correct communication issues on its own and it relies entirely on ICMP as its error handler. This method makes ICMP an important part of a UDP session for tracking its overall state.

UDP has no set method of connection teardown that announces the session's end. Because of the lack of a defined ending, the Personal Stateful Firewall clears a UDP session's state table entries after a preconfigured timeout value reached.

ICMP Protocol and Connection State

ICMP is also a connection-less network protocol. The ICMP protocol is often used to return error messages when a host or protocol cannot do so on its own. ICMP response-type messages are precipitated by requests using other protocols like TCP or UDP. This way of messaging and its connection-less and one-way communication make the tracking of its state a much more complicated process than UDP. The Personal Stateful Firewall tracks an ICMP connection based on IP address and request message type information in a state table.

Like UDP, the ICMP connection lacks a defined session ending process, the Personal Stateful Firewall clears a state table entry on a predetermined timeout.

Firewall now supports ICMP Traceroute to handle ICMP packets with type value 30 that were being dropped. ICMP packets with ICMP type value 30 are called ICMP Traceroute packets.

It is now possible to allow/deny the ICMP echo packets having identifier value zero. By default, these packets are allowed. This feature will be effective only if Firewall is enabled (Firewall or Firewall+NAT) for a call. For only NAT enabled calls, there is no change in the behavior. Configuration is available only if Firewall license is present.

Application-Level Traffic and States

The Personal Stateful Firewall uses Deep Packet Inspection (DPI) functionality to manage application-level traffic and its state. With the help of DPI functionality, the Personal Stateful Firewall inspects packets up to Layer-7. It takes application behaviors into account to verify that all session-related traffic is properly handled and then decides which traffic to allow into the network.

Different applications follow different rules for communication exchange so the Personal Stateful Firewall manages the different communication sessions with different rules through DPI functionality.

The Personal Stateful Firewall also provides inspection and filtering functionality on application content with DPI. Personal Stateful Firewall is responsible for performing many simultaneous functions and it detect, allow, or drop packets at the ingress point of the network.

HTTP Application and State

HTTP is the one of the main protocols used on the Internet today. It uses TCP as its transport protocol, and its session initialization follows the standard TCP connection method.

Due to the TCP flow, the HTTP allows an easier definition of the overall session's state. It uses a single established connection from the client to the server and all its requests are outbound and responses are inbound. The state of the connection matches with the TCP state tracking.

For content verification and validation on the HTTP application session, the Personal Stateful Firewall uses DPI functionality in the chassis.

PPTP Application and State

Point-to-Point Tunneling Protocol (PPTP) is one of the protocols widely used to achieve Virtual Private Networks (VPN). PPTP allows the Point-to-Point Protocol (PPP) to be tunneled through an IP network. PPTP uses an enhanced GRE (Generic Routing Encapsulation) to carry PPP packets.

PPTP protocol has 2 connection states - Control connection (TCP) and Data connection (GREv1). PPTP exchanges IP or port specific information over its control connection and that information will be used to transfer the data over tunnel. If a PPTP client resides behind NAT and uses private IP to communicate with the outside world, it is possible that the information exchange over PPTP control flow has private IPs.

TFTP Application and State

Trivial File Transfer Protocol (TFTP) is an application layer protocol which is used by File Transfer applications. TFTP uses UDP (User Datagram Protocol) as its transport protocol and has only basic functionalities. TFTP file operations include sending a file and receiving a file. TFTP supports different modes for File Transfer which are netascii, ascii, octet, and binary.

TFTP has two connection states - Control connection and Data connection that operate on UDP. Initially, TFTP starts the control flow (uses UDP Port 69) for communicating the type of file operation to be performed. The Client initiates the connection towards Server on port 69 (UDP). Server replies to the Client from a port other than 69 and data is transferred in this flow. Negative reply is sent using different error codes supported by TFTP.

File Transfer Protocol and State

FTP is an application to move files between systems across the network. This is a two way connection and uses TCP as its transport protocol.

Due to TCP flow, FTP allows an easier definition of the overall session's state. As it uses a single established connection from the client to the server, the state of the connection matches with the TCP state tracking.

Personal Stateful Firewall uses application-port mapping along with FTP application-level content verification and validation with DPI functionality in the chassis. It also supports Pinhole data structure and Initialization, wherein FTP ALG parses FTP Port command to identify the initiation and termination end points of future FTP DATA sessions. The source/destination IP and destination Port of FTP DATA session is stored.

When a new session is to be created for a call, a check is made to see if the source/destination IP and Destination Port of this new session matches with the values stored. Upon match, a new ACS data session is created.

This lookup in the pinhole list is made before port trigger check and stateful firewall ruledef match. If the look up returns a valid pinhole then a particular session is allowed. Whenever a new FTP data session is allowed because of a pinhole match the associated pinhole is deleted. Pinholes are also expired if the associated FTP Control session is deleted in, or when the subscriber call goes down.

