

Day in the Life of a Packet

PAN-OS Packet Flow Sequence PAN-OS 6.1

Contents

SECTI	ON 1: OVERVIEW	3		
SECTI	ON 2: INGRESS STAGE	3		
2.1	PACKET PARSING	5		
2.2	TUNNEL DECAPSULATION	6		
2.3	IP DEFRAGMENTATION	6		
SECTI	ON 3: FIREWALL SESSION LOOKUP	6		
3.1	FIREWALL SESSION SETUP	7		
3.2. 3.3. 3.4. 3.5. 3.6. 3.7. 3.8. 3.9.	ZONE PROTECTION CHECKS TCP STATE CHECK FORWARDING SETUP NAT POLICY LOOKUP USER-ID SECURITY POLICY LOOKUP DOS PROTECTION POLICY LOOKUP SESSION ALLOCATION	77889999		
SECTI		D		
4.1. 4.2.	SECURITY PROCESSING	1 1		
SECTI	ON 5: APPLICATION IDENTIFICATION (APP-ID)1	1		
SECTI	SECTION 6: CONTENT INSPECTION12			
SECTION 7: FORWARDING/EGRESS12				
SECTION 8: SUMMARY14				

Section 1: Overview

This document describes the packet handling sequence inside of PAN-OS devices. The ingress and forwarding/egress stages handle network functions and make packet-forwarding decisions on a per-packet basis. The remaining stages are session-based security modules highlighted by App-ID and Content-ID. This decoupling offers stateful security functions at the application layer, and the resiliency of per-packet forwarding and flexibility of deployment topologies.

Section 2: Ingress Stage

The ingress stage receives packets from the network interface, parses those packets, and then determines whether a given packet is subject to further inspection. If the packet is subject to further inspection, the firewall continues with a session lookup and the packet enters the security processing stage. Otherwise, the firewall forwards the packet to the egress stage. Section 3 summarizes cases when the firewall forwards packets without inspection, depending on the packet type and the operational mode of the interface.

Note: During packet processing, the firewall may discard a packet because of a protocol violation. In certain cases, due to firewall attack prevention features, it discards packets without configurable options. Section 2.1 enumerates such cases when the firewall discards packets at this stage.



Revision A

[4]

©2015, Palo Alto Networks, Inc.

2.1 Packet Parsing

Packet parsing starts with the Ethernet (Layer-2) header of the packet received from the wire.

The ingress port, 802.1q tag, and destination MAC address are used as keys to lookup the ingress logical interface. If the interface is not found, the packet is discarded. The hardware interface counter "receive error" and global counter "flow_rcv_dot1q_tag_err" are incremented.

Next, the IP header is parsed (Layer-3).

IPv4: The firewall will discard the packet for any one of the following reasons:

- Mismatch of Ethernet type and IP version
- Truncated IP header
- IP protocol number 0
- TTL zero
- Land attack
- Ping of death
- Martian IP address
- IP checksum errors

IPv6: The firewall will discard the packet for any one of the following reasons:

- Mismatch of Ethernet type and IP version
- Truncated IPv6 header
- Truncated IP packet (IP payload buffer length less than IP payload field)
- JumboGram extension (RFC 2675)
- Truncated extension header

Next, the Layer-4 (TCP/UDP) header is parsed, if applicable.

TCP: The firewall will discard the packet for any one of the following reasons:

- TCP header is truncated,
- Data-offset field is less than 5
- Checksum error
- Port is zero

• Invalid combination of TCP flags

UDP: The firewall will discard the packet for any one of the following reasons:

- UDP header truncated
- UDP payload truncated (not IP fragment and UDP buffer length less than UDP length field)
- Checksum error

2.2Tunnel Decapsulation

The firewall performs decapsulation/decryption at the parsing stage. After parsing the packet, if the firewall determines that it matches a tunnel, i.e. IPSec, SSL-VPN with SSL transport, then it performs the following sequence:

- The firewall decapsulates the packet first and discards it if errors exist.
- The tunnel interface associated with the tunnel is assigned to the packet as its new ingress interface and then the packet is fed back through the parsing process, starting with the packet header defined by the tunnel type. Currently, the supported tunnel types are IP layer tunneling, thus packet parsing (for a tunneled packet) starts with the IP header.

2.3IP Defragmentation

The firewall parses IP fragments, reassembles using the defragmentation process, and then feeds the packet back to the parser starting with the IP header. At this stage, a fragment may be discarded due to tear-drop attack (overlapping fragments), fragmentation errors, or if the firewall hits system limits on buffered fragments (hits the max packet threshold).

Section 3: Firewall Session Lookup

A packet is subject to firewall processing depending on the packet type and the interface mode. The following table summarizes the packet processing behavior for a given interface operation mode and packet type:

Deskat Temp	Interface operational modes				
Раскет Туре	Layer-3	Layer-2	Virtual-Wire	Тар	
IPv4 unicast	inspect & forward	inspect & forward	inspect & forward	inspect & drop	
IPv4 Multicast (224.0.0.1-239.255.255.255)	inspect & forward	inspect & forward(flood)	forward, but inspect only if multicast firewalling is on	inspect & drop	
IP broadcast (255.255.255.255)	drop	forward only (flood)	forward, but inspect only if multicast firewalling is on	drop	
IP local broadcast	drop	forward only (flood)	forward, but inspect only if multicast firewalling is on	drop	

IPv6	inspect and forward if enabled	forward, but inspect only if IPv6 firewalling is on (default)	forward, but inspect only if IPv6 firewalling is on (default)	drop, but inspect only if IPv6 firewalling is on (default)
Non-IP	process if applicable, not forward	forward only	forward only	drop

If the packet is subject to firewall inspection, it performs a flow lookup on the packet. A firewall session consists of two unidirectional flows, each uniquely identified. In PAN-OS's implementation, the firewall identifies the flow using a 6-tuple key.

- Source and destination addresses: IP addresses from the IP packet.
- **Source and destination ports:** Port numbers from TCP/UDP protocol headers. For non-TCP/UDP, different protocol fields are used (e.g. for ICMP the ICMP identifier and sequence numbers are used, for IPSec terminating on device the Security Parameter Index (SPI) is used, and for unknown, a constant reserved value is used to skip Layer-4 match).
- **Protocol:** The IP protocol number from the IP header is used to derive the flow key.
- **Security zone:** This field is derived from the ingress interface at which a packet arrives.

The firewall stores active flows in the flow lookup table. When a packet is determined to be eligible for firewall inspection, the firewall extracts the 6-tuple flow key from the packet and then performs a flow lookup to match the packet with an existing flow. Each flow has a client and server component, where the client is the sender of the first packet of the session from firewall's perspective, and the server is the receiver of this first packet.

Note: The distinction of client and server is from the firewall's point of view and may or may not be the same from the end hosts' point of view.

Based on the above definition of client and server, there will be a client-to-server (C2S) and server-to- client (S2C) flow, where all client-to-server packets should contain the same key as that of the C2S flow, and so on for the S2C flow.

3.1 Firewall Session Setup

The firewall performs the following steps to set up a firewall session:

3.2. Zone Protection Checks

After the packet arrives on a firewall interface, the ingress interface information is used to determine the ingress zone. If any zone protection profiles exist for that zone, the packet is subject to evaluation based on the profile configuration.

3.3. TCP State Check

If the first packet in a session is a TCP packet and it does not have the SYN bit set, the firewall discards it (default).

If SYN flood settings are configured in the zone protection profile and action is set to SYN Cookies, then TCP SYN cookie is triggered if the number of SYN matches the activate threshold. SYN cookie implementation functions as follows:

- The seed to encode the cookie is generated via random number generator each time the data plane boots up.
- If an ACK packet received from the client does not match cookie encoding, it treats the packet as non-SYN packet.
- A session that passes SYN cookie's process is subject to TCP sequence number translation because the firewall acted as a proxy for TCP 3-way handshake.

If the SYN Flood protection action is set to Random Early Drop (RED) instead, which is the default, then the firewall simply drops any SYN messages that are received after hitting the threshold. SYN Cookies is preferred when you want to permit more legitimate traffic to pass through while being able to distinguish SYN flood packets and drop those instead. RED, on the other hand, will drop SYN packets randomly and can impact legitimate traffic equally.

Note: You can configure the firewall to allow the first TCP packet, even if it does not have SYN bit set. Altering the default behavior and allowing non-SYN TCP packets through poses a security risk by opening up the Firewall to malicious packets not part of a valid TCP connection sequence. Although this is not a recommended setting, it might be required for scenarios with asymmetric flows. You should configure the firewall to reject TCP non-SYN when SYN cookies are enabled.

3.4. **Forwarding Setup**

This stage determines the packet-forwarding path. Packet forwarding depends on the configuration of the interface. The following table summarizes the packet-forwarding behavior:

Interface Mode	Forwarding action
Тар	Egress interface/zone is the same as the ingress interface/zone from a policy perspective. The firewall discards the packet.
Virtual Wire	Egress interface is the peer interface configured in the virtual wire
Layer-2	Egress interface for the destination MAC is retrieved from the MAC table. If the information is not present, the frame is flooded to all interfaces in the associated VLAN broadcast domain, except for the ingress interface.
Layer-3	The firewall uses the route lookup table to determine the next hop, or discards the packet if there is no match.

3.5. **NAT Policy Lookup**

This is applicable only in Layer-3 or Virtual Wire mode. At this stage, the ingress and egress zone information is available. The firewall evaluates NAT rules for the original packet.

For destination NAT, the firewall performs a second route lookup for the • translated address to determine the egress interface/zone.

• For source NAT, the firewall evaluates the NAT rule for source IP allocation. If the allocation check fails, the firewall discards the packet.

3.6. User-ID

The firewall uses the IP address of the packet to query the User-IP mapping table (maintained per VSYS). The corresponding user information is fetched. The firewall next takes this user information to query the user-group mapping table and fetches the group mapping associated with this user (it returns all groups the user belongs to).

There is a chance that user information is not available at this point. In that case, if captive portal policy is setup, the firewall will attempt to find out the user information via captive portal authentication (discussed in Section 4).

3.7. Security Policy Lookup

At this stage, the ingress and egress zone information is available. The firewall uses application *ANY* to perform the lookup and check for a rule match. In case of a rule match, if the policy action is set to 'deny', the firewall drops the packet. The firewall denies the traffic if there is no security rule match. The firewall permits intra-zone traffic by default. You can modify this default behavior for intra-zone and inter-zone traffic from the security policies rulebase.

Note: The firewall applies security rules to the contents of the original packet, even if there are NAT rules configured.

3.8. DoS Protection Policy Lookup

Next, the firewall checks the DoS (Denial of Service) protection policy for traffic thresholds based on the DoS protection profile.

If the DoS protection policy action is set to "Protect", the firewall checks the specified thresholds and if there is a match (DoS attack detected), it discards the packet.

If the policy action is either allow or deny, the action takes precedence regardless of threshold limits set in the DoS profile.

3.9. Session Allocation

The firewall allocates a new session entry from the free pool after all of the above steps are successfully completed. Session allocation failure may occur at this point due to resource constraints:

- VSYS session maximum reached, or
- The firewall allocates all available sessions.

After the session allocation is successful:

• The firewall fills session content with flow keys extracted from the packet and the forwarding/policy results.

- Session state changes from INIT (pre-allocation) to OPENING (post-allocation).
- If the application has not been identified, the session timeout values are set to default value of the transport protocol. You can configure these global timeout values from the Firewall's device settings. Application specific timeout values override the global settings, and will be the effective timeout values for the session once application is identified.

After setup, session installation takes place:

- Firewall queries the flow lookup table to see if a match exists for the flow keys matching the session. If a flow lookup match is found (session with same tuple already exists), then this session instance is discarded as session already exists, else
- Session is added to the flow lookup table for both C2S and S2C flows and firewall changes the session's state from OPENING to ACTIVE.

The firewall then sends the packet into Session Fast Path phase for security processing.

Section 4: Firewall Session Fast Path

A packet that matches an existing session will enter the fast path. This stage starts with Layer-2 to Layer-4 firewall processing:

- If the session is in discard state, then the firewall discards the packet. The firewall can mark a session as being in the discard state due to a policy action change to deny, or threat detection.
- If the session is active, refresh session timeout.
- If the packet is a TCP FIN/RST, the session TCP half closed timer is started if this is the first FIN packet received (half closed session) or the TCP Time Wait timer is started if this is the second FIN packet or RST packet. The session is closed as soon as either of these timers expire.
- If NAT is applicable, translate the L3/L4 header as applicable.

If an application uses TCP as the transport, the firewall processes it by the TCP reassembly module before it sends the data stream into the security-processing module. The TCP reassembly module will also perform window check, buffer out-of-order data while skipping TCP retransmission. The firewall drops the packets if there is a reassembly error or if it receives too many out-of-order fragments, resulting in the reassembly buffers filling up.

4.1. Security Processing

A packet matching an existing session is subject to further processing (application identification and/or content inspection) if packet has TCP/UDP data (payload), or it is a non-TCP/UDP packet.

If the firewall does not detect the session application, it performs an App-ID lookup. If App-ID lookup is non-conclusive, the content inspection module runs known protocol decoder checks and heuristics to help identify the application.

If the firewall detects the application, the session is subject to content inspection if any of the following apply:

- Application Layer Gateway (ALG) is involved.
- Application is tunneled application.
- Security rule has security profile associated.

The Application Identification (App-ID) and Content Inspection stages are discussed in detail in later sections (Section 5 and 6).

4.2. Captive Portal

If the user information was not available for the source IP address extracted from the packet, and the packet is destined to TCP/80, the firewall performs a captive portal rule lookup to see if the packet is subject to captive portal authentication. If captive portal is applicable, the packet is redirected to the captive portal daemon.

Note: Since captive portal is applicable to http traffic and also supports a URL category based policy lookup, this can be kicked in only after the TCP handshake is completed and the http host headers are available in the session exchange.

Section 5: Application Identification (App-ID)

The firewall first performs an application-override policy lookup to see if there is a rule match. If there is, the application is known and content inspection is skipped for this session.

If there is no application-override rule, then application signatures are used to identify the application. The firewall uses protocol decoding in the content inspection stage to determine if an application changes from one application to another.

After the firewall identifies the session application, access control, content inspection, traffic management and logging will be setup as configured.

 Security policy lookup: The identified application as well as IP/port/protocol/zone/user/URL category in the session is used as key to find rule match.

- If the security policy has logging enabled at session start, the firewall generates a traffic log, each time the App-ID changes throughout the life of the session.
- If security policy action is set to allow and it has associated profile and/or application is subject to content inspection, then it passes all content through Content-ID.
- If security policy action is set to allow, the firewall performs a QoS policy lookup and assigns a QoS class based on the matching policy.
- If security policy action is set to allow and the application is SSL or SSH, perform a decryption policy lookup and set up proxy contexts if there is a matching decryption rule.

Section 6: Content Inspection

The firewall performs content Inspection, if applicable, where protocol decoders' decode the flow and the firewall parses and identifies known tunneling applications (those that routinely carry other applications like web-browsing).

If the identified application changes due to this, the firewall consults the security policies once again to determine if the session should be permitted to continue.

If the application does not change, the firewall inspects the content as per all the security profiles attached to the original matching rule. If it results in threat detection, then the corresponding security profile action is taken.

The firewall forwards the packet to the forwarding stage if one of the conditions hold true:

- If inspection results in a 'detection' and security profile action is set to allow, or
- Content inspection returns no 'detection'.

The firewall then re-encrypts the packet before entering the forwarding stage, if applicable (SSL forward proxy decryption and SSH decryption).

Section 7: Forwarding/Egress

The firewall identifies a forwarding domain for the packet, based on the forwarding setup (discussed earlier).

The firewall performs QoS shaping as applicable in the egress process. Also, based on the MTU of the egress interface and the fragment bit settings on the packet, the firewall carries out fragmentation if needed.

[12]

If the egress interface is a tunnel interface, then IPSec/SSL-VPN tunnel encryption is performed and packet forwarding is reevaluated.

Finally the packet is transmitted out of the physical egress interface.



Section 8: Summary

Palo Alto Networks next-generation firewalls use a unique Single Pass Parallel Processing (SP3) Architecture–which enables high-throughput, low-latency network security, all while incorporating unprecedented features and technology. Palo Alto Networks solves the performance problems that plague today's security infrastructure with the SP3 architecture, which combines two complementary components-Single Pass software, Parallel Processing hardware. The result is an excellent mix of raw throughput, transaction processing, and network security that today's high performance networks require.