Attack Prevention Technology White Paper

Keywords: Attack prevention, denial of service

Abstract: This document introduces the common network attacks and the corresponding prevention measures, and describes the features and network applications of the H3C firewall attack prevention technology.

Acronyms:

Acronym	Full spelling
DMZ	De-Militarized Zone
DDoS	Distributed Denial of Service
DoS	Denial of Service

Table of Contents

1 Overview

Attack prevention is one of the important functions of firewalls. With this function, a firewall can determine whether received packets are attack packets according to the packet contents and behaviors and, if detecting an attack, take corresponding measures to protect internal network hosts and devices against the attack.

The attack prevention function allows a firewall to detect network attacks including denial of service (DoS), scanning and snooping, and malformed packet attacks. The attack prevention measures include blacklist filtering, packet attack characteristics identification, abnormal traffic detection, and intrusion detection statistics.

1.1 Background

With the popularity of network technologies, network attacks are more likely to happen. In addition, various viruses are full of networks, facilitating network attacks.

At present, the common network attacks fall into the following three types:

- DoS attack

A DoS attack sends a large number of data packets to a target system, so that the system cannot receive requests from legitimate clients or is suspended and cannot work normally. DoS attacks mainly fall into two types: SYN flood and Fraggle. Different from other types of attacks, DoS attacks do not search for ingresses of the target network but prevent legitimate clients from accessing network resources by interfering with the normal operation of the target network.

- Scanning and snooping attack

A scanning and snooping attack uses ping-type programs (including ICMP ping and TCP ping) to identify active hosts on a network and to further locate potential targets, and uses TCP and UDP port scanning to detect the target operating systems and enabled service types. In this way, the attacker can get a general idea of the service types provided by a target system and the potential security bugs, preparing for a further intrusion of the target system.

- Malformed packet attack

A Malformed packet attack sends defective packets, such as overlapping IP

fragments and packets with illegal TCP flags, to a target system, so that the system crashes when it processes such packets. There are mainly two types of malformed packet attacks: Ping of Death and Teardrop.

Of the above attacks, DoS attacks are the most common. This type of attack is easy to implement and the attack effect is obvious. A successful DoS attack can result in a sharp performance degradation of the target server, making the clients unable to access the server normally. It also makes a service-providing company lose reputation, which is a serious and long-term damage.

A firewall must have an attack prevention technology that can protect the network against various common network attacks, ensuring that the network can work normally under attacks.

1.2 Benefits

Through the packet pattern identification technology, attack prevention can identify dozens of attack types, such as large ICMP packet attack, malformed TCP packet attack, and Tracert probe packet attack.

For a DoS or DDoS attack exploiting a protocol supported by a target server, attack prevention can distinguish attack traffic from normal traffic by using an abnormality detection algorithm based on behavior patterns. Once it detects attack traffic, it drops the attack traffic, protecting the server against DoS attacks. The attack prevention function can dig out the following attacks: SYN flood, ICMP flood, and UDP flood.

2 Attack Prevention Implementation

2.1 ICMP Redirect Attack

1. Introduction

An ICMP redirect attack sends ICMP redirect messages to hosts on a subnet to request the hosts to change their routing tables. Normally, a device sends ICMP redirect messages only to hosts on the local subnet. An attacker may send fake ICMP redirect messages to hosts on another subnet to change their routing tables, interfering with the normal forwarding of IP packets.

2. Prevention measures

Detect whether received packets are ICMP redirect messages. If yes, drop or forward the packets and log the event, depending on the configuration.

2.2 ICMP Unreachable Attack

1. Introduction

Different systems process ICMP unreachable packets (type 3) differently. Upon receiving a network unreachable ICMP packet (code 0) or host unreachable ICMP packet (code 1), some systems conclude that the destination is unreachable and drops all subsequent packets destined for the destination. An ICMP unreachable attack exploits this mechanism to cut off the connection between the destination system and a host.

2. Prevention measures

Check whether received packets are ICMP unreachable packets. If yes, drop or forward the packets and log the event, depending on the configuration.

2.3 Address Scanning Attack

1. Introduction

An address scanning attacker uses a ping-type program to send ping packets to a target network. All responding hosts are considered active and may become attack targets. The attacker can also use TCP/UDP packets, such as TCP ping, to connect to the target address and determine whether a system is open by checking whether there is a response.

2. Prevention measures

Check received ICMP, TCP, and UDP packets and collect statistics on how many destination IP addresses a source IP address is sending such packets to. If the number reaches the pre-defined threshold in a period, drop such packets received later from the source IP address, log the event, and add the source IP address to the blacklist, depending on the configuration.

2.4 Port Scanning Attack

1. Introduction

Port scanning attackers usually use some scanning tools to initiate connections to a series of TCP/UDP ports on target hosts, and determine whether the hosts are using the ports to provide services. When using TCP packets to scan ports, such an attacker sends connection request (TCP SYN) packets to the target host. The target host responds with TCP ACK packets if the requested TCP ports are open, or TCP RST packets if the requested TCP ports are not open. The attacker can determine whether the requested services are enabled on the target host or not by analyzing the responses of the target host. When using UDP packets to scan ports, the attacker sends connection request packets (UDP packets) to the target host. The target host responds with ICMP packets if the requested UDP ports are open, or ICMP unreachable packets if the requested UDP ports are not open. The attacker can determine whether the requested services are enabled on the target host or not by analyzing the responses of the target host. After determining which ports are open on the target host, the attacker usually further attacks the target host using the ports.

2. Prevention measures

Check received TCP and UDP packets and collect statistics on how many destination ports a source IP address is sending such packets to. If the number reaches the pre-defined threshold in a period, drop such packets received later from the source IP address, log the event, and add the source IP address to the blacklist, depending on the configuration.

2.5 IP Source Route Attack

1. Introduction

The Source Route option in IP packets is usually used in network path troubleshooting and temporary transmission of some special services. Packets carrying the Source Route option ignore the forwarding entries of devices along the transmission path during the forwarding process. For example, if you want an IP packet to pass through routers R1, R2, and R3, you can specify the IP addresses of

the interfaces on the three routers in the Source Route option of the packet. In this case, the IP packet will pass through R1, R2, and R3 in turn, regardless of the routing tables of the routers. During the transmission of an IP packet carrying the Source Route option, the source address and destination address are always changing. An attacker may forge some legal IP addresses by configuring the Source Route option, thus mingling in the target network.

2. Prevention measures

Check whether received packets carry the Source Route option. If yes, drop or forward the packets and log the event, depending on the configuration.

2.6 IP Route Record Attack

1. Introduction

In IP routing technology, a Route Record option is added to an IP packet to record the path the packet passes from source to destination, that is, the routers that have processes the packet. The IP Route Record option is usually used in network path troubleshooting. However, it may also be used by an attacker to get the network structure by analyzing the path information carried this option.

2. Prevention measures

Check whether received packets carry the Route Record option. If yes, drop or forward the packets and log the event, depending on the configuration.

2.7 Tracert Attack

1. Introduction

The Tracert program usually sends UDP packets with a large destination port number and an increasing TTL (starting from 1). Some Tracert programs may send ICMP ping packets. The TTL of a packet is decreased by 1 when the packet passes each router. The protocol prescribes that if a router receives a packet with a TTL of 0, the router must sends an ICMP time exceeded message back to the source IP address of the packet. A Tracert attacker may run the Tracert program and analyzes the source

IP addresses in returned ICMP error messages to obtain the paths to destinations and figure out the network topology.

2. Prevention measures

Check whether received ICMP messages are ICMP time exceeded messages (type 11) or destination port unreachable messages (type 3, code 3). If yes, drop or forward the packets and log the event, depending on the configuration.

2.8 Land Attack

1. Introduction

A land attack uses the TCP three-way handshake function. During such an attack, the attacker forges large amounts of TCP SYN packets with both the source address and destination address being the IP address of the target, causing the target to send SYN ACK messages to itself. After the target receives the SYN ACK messages, it sends ACK messages to itself and creates TCP connections, which will be kept until they time out. In this way, the attacker may deplete the resources of the target. Different operating systems respond differently to land attacks. For instance, UNIX hosts will crash down while Windows NT hosts will slow down.

2. Prevention measures

Check the source address and destination address of each received IP packet. If the two addresses are the same or the source address is a loopback address (127.0.0.1), drop or forward the packet and log the event, depending on the configuration.

2.9 Smurf Attack

1. Introduction

A simple Smurf attack sends ICMP echo request messages with the destination address being a broadcast address on the target network. When all hosts on the target network reply to the requests, network congestion occurs. An advanced Smurf attack uses the address of a target host or network as the source address of the ICMP echo request messages, and keeps sending such ICMP echo request

messages to the target, making the target crash down.

2. Prevention measures

Check whether the destination address of a received ICMP echo request message is a subnet broadcast address or network address. If yes, drop or forward the packets and log the event, depending on the configuration.

2.10 Fraggle Attack

1. Introduction

A Fraggle attack is similar to a Smurf attack. The difference is that it uses UDP messages instead of ICMP messages. A Fraggle attacker sends to a subnet broadcast address large amounts of UDP packets with the source address being the target network address or target host address and the destination port number being 7 (Echo service) or 19 (Chargen service). All hosts in the subnet that are enabled with Echo or Chargen service will send reply messages to the target network or host, resulting in network congestion in the target network or crash of the target host.

2. Prevention measures

Check whether the destination port number of a received UDP packet is 7 or 19. If yes, drop or forward the packets and log the event, depending on the configuration.

2.11 WinNuke Attack

1. Introduction

A WinNuke attack sends Out-of-Band (OOB) data packets to the NetBIOS port (139) of a target running a Windows system. Such a packet has the Urgent flag bit in its header set. The Urgent Pointer fields of WinNuke attack packets are overlapped, resulting in NetBIOS fragment overlaps. This causes the target host that has established TCP connections with other hosts to crash when it processes these NetBIOS fragments.

2. Prevention measures

Check received UDP packets. If the destination port number is 139 and the TCP Urgent flag is set, drop or forward the packets and log the event, depending on the configuration.

2.12 SYN Flood Attack

1. Introduction

An SYN Flood attack sends a great quantity of forged SYN packets to a target server, using a forged or nonexistent address as the source address. After receiving the SYN packets, the server replies with SYN ACK packets. As the source address of the attack packets is unreachable, the server can never receive the expected ACK packets, resulting in large amounts of half-open connections. In this way, the attack exhausts the network resources, making the server unable to service normal clients.

2. Prevention measures

Use the firewall as a relay between the server and its clients. When a client initiates a connection to the server, the firewall does not forward the SYN packet to the server; it sends an SYN ACK to the client on behalf of the client. Only when it receives the expected ACK packet from the client, does it allow a connection to be established between the client and the server.

2.13 ICMP Flood Attack

1. Introduction

An ICMP flood attack sends a large number of ICMP messages to the target in a short time by, for example, using the ping program, causing the target too busy to process normal network data packets.

2. Prevention measures

Check the rate of ICMP packets destined for the specific destination address using the intelligent traffic detection technology. If the rate exceeds the specified upper threshold, treat all such subsequent ICMP packets as attack packets: drop or forward

the packets and log the event, depending on the configuration. When the packet rate drops below the lower threshold, forward subsequent ICMP packets normally.

2.14 UDP Flood Attack

1. Introduction

Similar to ICMP flood attack, UDP flood attack sends a large number of UDP messages to the target in a short time, so that the target gets too busy to transmit the normal network data packets.

2. Prevention measures

Check the rate of the UDP packets destined for the specific destination address using the intelligent traffic detection technology. If the rate exceeds the specified upper threshold, treat the subsequent UDP packet attack packets: drop or forward the packets and log the event, depending on the configuration. When the packet rate drops below the lower threshold, forward subsequent UDP packets normally.

3 Technical Characteristics of the H3C Implementation

The H3C attack prevention implementation features the following:

1. Support for security zone-based configuration

Traditionally, a firewall is used between an internal network and the external network, and firewall/router policies are configured on the inbound and outbound interfaces. Nowadays, the internal network/external network/DMZ mode is more likely used in firewall networking environments and firewalls tend to provide more ports. As a result, configuring and maintaining the firewalls using the traditional configuration mode becomes a heavy burden. What is more, the massive configuration tasks mean a higher security risk.

Some traditional firewalls support global configuration besides the interface-based configuration. However, a firewall under global configuration uses the same security policy on all its ports and therefore cannot satisfy the diverse requirements of different security zones connected.

The H3C attack prevention function supports security zone-based configuration. All attack detection policies are configured on security zones. Therefore, the policy configuration of a firewall is simple and flexible. This not only reduces the configuration workload, but also satisfies the requirements of differentiated policy configuration for different security zones.

2. Abundant alarm log information

The H3C implementation of attack prevention provides abundant alarm log information, which can be used by third party software. The log function and audit function allow you to monitor attacks in real time, query and analyze attack history records, facilitating the trace of attack events.

3. Flexible attack prevention measures

H3C provides flexible measures to prevent attacks of the previously described types. You can choose to log the attacks, drop the attack packets, add attack sources to the blacklist (applicable to scanning attacks), enable TCP proxy function (applicable to TCP SYN flood attacks), or instruct the server to release the oldest half-open connections (applicable to TCP SYN flood attacks).

4 Application Scenario

4.1 SYN Flood Attack Prevention

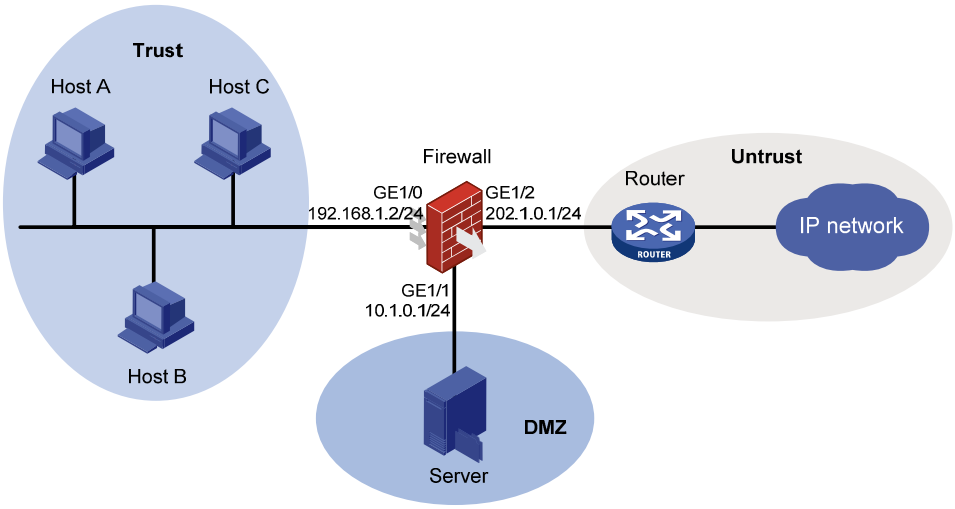

Figure 1 Network diagram for SYN flood attack prevention configuration

In the SYN flood attack prevention scenario illustrated in Figure 1 , the internal network is in the Trust zone, the internal server is in the DMZ zone, and the external network is in the Untrust zone.

Configure a security policy on the H3C firewall to perform SYN flood attack detection for the server in the DMZ zone. Based on the actual traffic of the server, configure the maximum connection rate and the maximum number of half-open connections allowed on the server. If the server is under an SYN flood attack, the firewall will output a SYN flood attack log, and you can choose to use the TCP proxy to process subsequent connection requests to the server in the DMZ zone, ensuring that TCP connection requests reaching the server are normal requests.