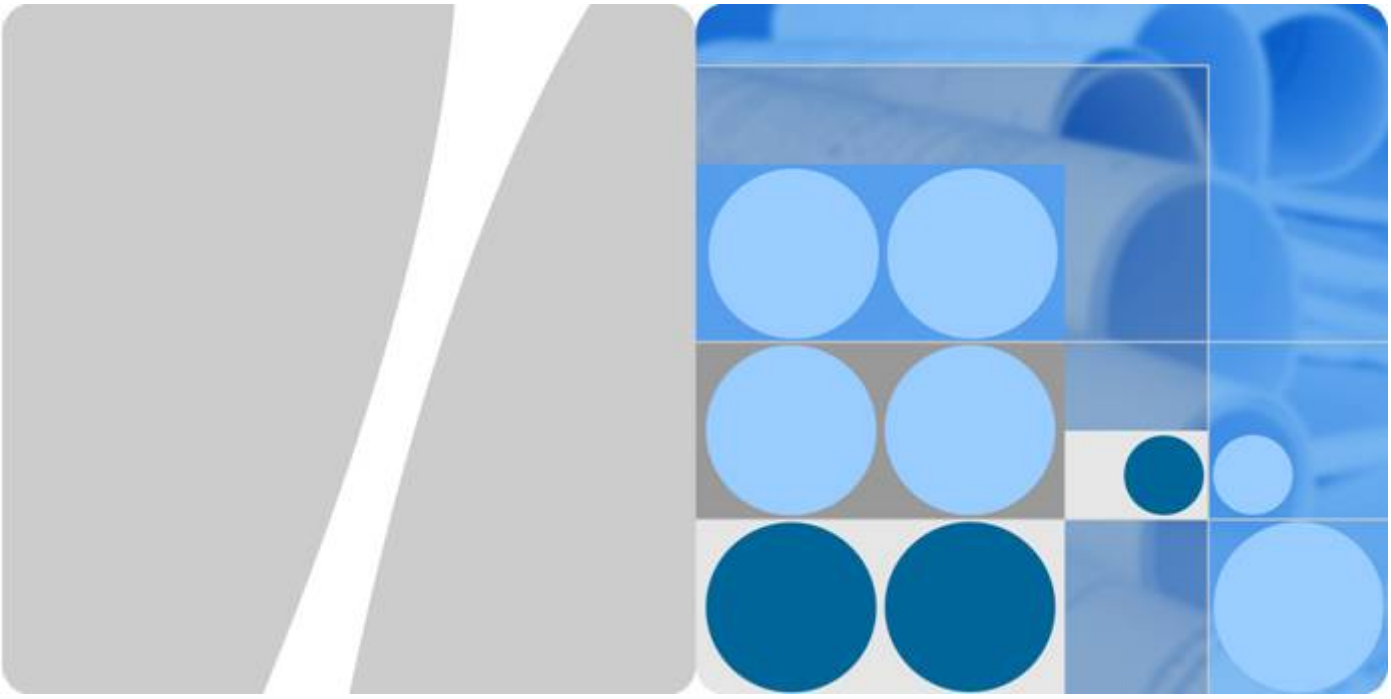

Quidway S6700 Series Ethernet Switches

V100R006C00

Feature Description - Security

Issue 01

Date 2011-07-15

HUAWEI TECHNOLOGIES CO., LTD.

HUAWEI

Trademarks and Permissions


and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute the warranty of any kind, express or implied.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: http://www.huawei.com

Email: support@huawei.com

About This Document

Intended Audience

This document describes the security feature in terms of its overview, principle, and applications.

This document together with other types of document helps intended readers get a deep understanding of the security feature.

This document is intended for:

- Network planning engineers
- Commissioning engineers
- Data configuration engineers
- System maintenance engineers

Symbol Conventions

The symbols that may be found in this document are defined as follows.

Symbol	Description
DANGER	Indicates a hazard with a high level of risk, which if not avoided, will result in death or serious injury.
WARNING	Indicates a hazard with a medium or low level of risk, which if not avoided, could result in minor or moderate injury.
CAUTION	Indicates a potentially hazardous situation, which if not avoided, could result in equipment damage, data loss, performance degradation, or unexpected results.
TIP	Indicates a tip that may help you solve a problem or save time.
NOTE	Provides additional information to emphasize or supplement important points of the main text.

Command Conventions

The command conventions that may be found in this document are defined as follows.

Convention	Description
Boldface	The keywords of a command line are in **boldface**.
Italic	Command arguments are in *italics*.
[ ]	Items (keywords or arguments) in brackets [ ] are optional.
{ x | y | ... }	Optional items are grouped in braces and separated by vertical bars. One item is selected.
[x | y | ...]	Optional items are grouped in brackets and separated by vertical bars. One item is selected or no item is selected.
{ x | y | ... }*	Optional items are grouped in braces and separated by vertical bars. A minimum of one item or a maximum of all items can be selected.
[x | y | ...]*	Optional items are grouped in brackets and separated by vertical bars. Several items or no item can be selected.
&<1-n>	The parameter before the & sign can be repeated 1 to n times.
#	A line starting with the # sign is comments.

Change History

Updates between document issues are cumulative. Therefore, the latest document issue contains all updates made in previous issues.

Changes in Issue 01 (2011-07-15)

Initial commercial release.

Contents

1 AAA and User Management

About This Chapter

1.1 Introduction

Definition

Authentication, Authorization, and Accounting (AAA) is a technology used for user authentication, authorization, and accounting.

- Authentication: checks whether a user can access the network.
- Authorization: authorizes a user to use specific services.
- Accounting: records the utilization of network resources.

The S6700 supports the following features:

Authentication:

- None authentication. To ensure the system security, the S6700 does not allow the administrators to log in without authentication.
- Local authentication
- RADIUS authentication
- HWTACACS authentication
- Combination of multiple authentication modes. If no response is received in the current authentication mode, another authentication mode is used.
- Extensible Authentication Protocol (EAP) termination
- EAP relay
- PAP/CHAP authentication, including the authentication through plain-text password and the authentication through the cipher-text password
- Authentication for the switching of user levels
- HWTACACS authentication for the switching of user levels
- Local authentication for the switching of user levels
- None authentication for the switching of user levels
- Super authentication for the switching of user levels

Authorization:

- None authorization
- Local authorization
- HWTACACS authorization
- if-authenticated authorization. If a user passes the authentication and the authentication mode is not none, the user is authorized; otherwise, the user cannot be authorized.
- Combination of multiple authorization modes. If no response is received in the current authorization mode, another authorization mode is used.
- Command line authorization for users
- Command line authorization through HWTACACS
- Command line authorization in local mode
- RADIUS COA

- RADIUS DM

Accounting:

- None accounting
- RADIUS accounting
- HWTACACS accounting
- Time-based accounting
- Sending start-accounting packets
- Sending stop-accounting packets
- Interim accounting

Purpose

AAA provides authentication, authorization, and accounting.

1.2 References

The following table lists the references of this document.

Document	Description	Remarks
RFC 2093	Generic AAA Architecture	
RFC 2094	AAA Authorization Framework	
RFC 2095	AAA Authorization Application Examples	
RFC 2096	AAA Authorization Requirements	
RFC 2058	Remote Authentication Dial In User Service (RADIUS)	
RFC 2059	RADIUS Accounting	
RFC 2138	Remote Authentication Dial In User Service (RADIUS)	
RFC 2139	RADIUS Accounting	
RFC 2809	Implementation of L2TP Compulsory Tunneling via RADIUS	
RFC 2865	Remote Authentication Dial In User Service (RADIUS)	
RFC 2866	RADIUS Accounting	
RFC 2867	RADIUS Accounting Modifications for Tunnel Protocol Support	
RFC 2868	RADIUS Attributes for Tunnel Protocol Support	

Document	Description	Remarks
RFC 2869	RADIUS Extensions	
RFC 0927	TACACS user identification Telnet option	
RFC 1492	An Access Control Protocol, Sometimes Called TACACS	

1.3 Principles

1.3.1 AAA

Basic Structure

AAA adopts the Client/Server model. This model has a good extensibility and is convenient for concentrated management of user information, as shown in **Figure 1-1**.

Figure 1-1 Schematic diagram of AAA basic structure

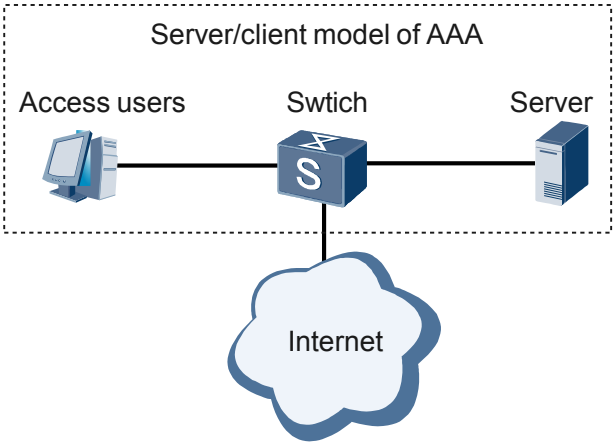

Authentication

AAA supports the following authentication modes:

- Non-authentication: completely trusts users and does not check their validity. It is rarely used.

  NOTE

  Non-authentication method of AAA is invalid for administrators.

- Local authentication: configures user information, including the user name, password, and attributes of local users, on a Network Access Server (NAS). Local authentication features

fast processing and low operation cost. The major limitation of local authentication is that the hardware restricts the capacity of information storage.

- Remote authentication: configures user information, including the user name, password and attributes, on the authentication server. AAA can remotely authenticate users through the Remote Authentication Dial In User Service (RADIUS) protocol or the HuaWei Terminal Access Controller Access Control System (HWTACACS) protocol. As the client, the NAS communicates with the RADIUS or HWTACACS server.

In the process of upgrading user levels, if several authentication modes are used in an authentication scheme, these authentication modes take effect in the sequence of their configurations.

- In the sequence of remote authentication and then local authentication:

  If a login account is available locally but unavailable on the remote server, the authentication mode is not changed from remote authentication to local authentication. Instead, it is considered that this account fails the authentication.

  The authentication mode is changed from remote authentication to local authentication only when the remote server has no response.
- The authentication mode **none** must be the last one adopted.

Authorization

AAA supports the following authorization modes:

- Non-authorization: does not authorize users.
- Local authorization: authorizes users according to related attributes of the local user accounts configured on the NAS.
- HWTACACS authorization: authorizes users through the TACACS server.
- if-authenticated authorization: authorizes users after the users pass the authentication in either local or remote authentication mode.
- RADIUS authorization: authorizes users after they pass RADIUS authentication. Authentication and authorization through the RADIUS protocol are bound together, so RADIUS cannot be used to perform only authorization.

If multiple authorization modes are configured to an authorization scheme, execute the modes based on the configuration sequence. The authorization mode **none** must be the last one adopted.

Accounting

AAA supports the following accounting modes:

- Non-accounting: provides free services for users.
- Remote accounting: supports remote accounting through the RADIUS server or the HWTACACS server.

1.3.2 RADIUS

RADIUS is one of commonly-used protocols to implement AAA. The RADIUS protocol was initially used for managing a large number of scattered users who use serial interfaces and modems. Now this protocol is widely applied to the NAS system.

To obtain the right to access certain networks or to use certain network resources, a user needs to set up a connection with the NAS over a network (such as the telephone network). In this case, the NAS authenticates the user and coordinates the connection.

After this authentication, the NAS sends authentication and accounting information of the user to the RADIUS server. RADIUS defines how the NAS and RADIUS server send user information, accounting information, and authentication and accounting results to each other. The RADIUS server receives the user's connection request, authenticates the user, and then sends the authentication result and configurations required by the user to the NAS.

The authentication information between the NAS and the RADIUS server is transmitted with a key. This helps protect the user password against theft on an insecure network.

Protocol Implementation

Using the User Datagram Protocol (UDP) as the transport protocol, RADIUS features a high real-time performance. Owing to the retransmission mechanism and standby server mechanism, RADIUS is of high reliability.

RADIUS is easy to implement and is applicable to the multithreading structure of the server for a large number of users.

As the RADIUS client, the NAS performs the following functions:

- Supports the standard RADIUS protocol and its extensions, including RFC 2865 and RFC 2866.
- Supports Huawei-developed private functions
- Supports active detection of the RADIUS server: After receiving an AAA authentication or accounting message, the RADIUS client starts the server detection if the status of the current server is Down. The RADIUS client then transforms the message into a packet that functions as the server-probe packet, and sends the packet to the server. If the client receives a response packet from the RADIUS server, the client considers the server as available.
- Caches the Accounting-stop packets locally and retransmits them: If the number of retransmission failures exceeds the value configured, packets are saved to the buffer queue. The system periodically scans the queue, extracts the packets and then sends them to the specific server and enables the timer to wait. If the transmission fails or no response packet is received from the server within the timeout period, the packets are placed back to the buffer queue again.

Supports auto switchover of the RADIUS server: If the current server does not work or the number of retransmission events exceeds the maximum number configured, another server in the server group replaces the current server to transmit packets.

Authentication and Accounting

The RADIUS server sets up a unique database to store user names and passwords that are required by authentication.

Figure 1-2 shows the main process of RADIUS messages transmission between the server and the client.

- When logging in to a network device such as a switch
- After the RADIUS client (an NAS server) on this network receives the user name and password, it sends an authentication request to the RADIUS server.

- If the request is valid, the server completes the authentication and sends the required authorization information back to the client. If the request is invalid, the server sends the authorization failure information back to the client.

User password information is encrypted before being transmitted between the RADIUS client and RADIUS server to avoid theft on an insecure network.

Figure 1-2 Message Exchange Between the RADIUS Client and the RADIUS Server

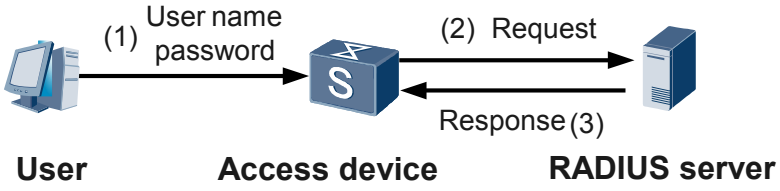

Message exchange for accounting is similar to message exchange for authentication or authorization.

1.3.3 HWTACACS

HWTACACS is a security protocol whose functions are enhanced on the basis of TACACS defined in RFC 1492. Similar to RADIUS, HWTACACS adopts the Client/Server model to communicate with the HWTACACS server, thus implementing AAA for various users, including Point-to-Point Protocol (PPP) users, Virtual Private Dial Network (VPDN) users, and login users.

Principle for HWTACACS

- HWTACACS authorizes users through command lines.
- When a user logs on to the switch through Telnet or SSH, set the command line authorization mode to HWTACACS for the users with the certain level if the user needs to be authorized with command lines. Then each input command should pass through HWTACACS authorization. The command can be run only after the command authorization is passed. Otherwise, the HWTACACS server displays a message to inform the user that the command authorization fails and the command cannot be run.
- Local authentication is a backup of command-line authorization. If the HWTACACS server fails (Down, unreachable or timing out for response) and the command-line authorization cannot be performed, local authentication functions instead of command-line authorization.
- If the switch does not receive any authorization response from the HWTACACS server within the timeout period set by the user, it is considered that the command authorization times out, and thus the command cannot be run.
- The user can configure policies in the case that the server gives no response or command-line authorization fails when the local user is configured. There are two polices: keep the user online and cut off users when the number of authorization failures exceeds the threshold.

NOTE

The preceding policies are applicable in the case that the HWTACACS server fails or no local user is configured. The following two cases cannot trigger the preceding policies:

- The server works normally but the executed command fails to pass the HWTACACS authorization.
- The server does not work and the command-line authorization fails. Local authorization is then used. The authorization fails because the level of the input command is higher than the locally configured level.

Figure 1-3 shows the process of command-line authorization supported by HWTACACS.

Figure 1-3 Process of command-line authorization supported by HWTACACS

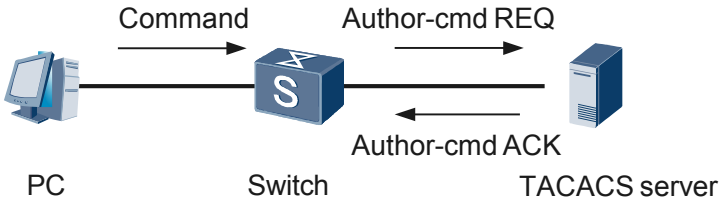

- HWTACACS authorizes user upgrade level.
- After logging on to the switch through Telnet or SSH, the user can upgrade or degrade its own level by running the **super** command. The switch must authenticate the user password.
- **Figure 1-4** shows the process of upgrading the HWTACACS user level. The switch sends the user password to the HWTACACS server for authentication. If the password passes the authentication, the user level can be upgraded. The modified level takes effect only during the current login.

Figure 1-4 Process of upgrading the HWTACACS user level

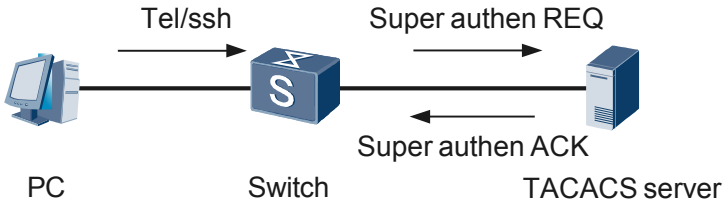

- If the switch does not receive the authentication result in the timeout period specified by the users, the authentication fails and the user level cannot be upgraded.
- Alternatively, you can upgrade the user level by running the **super** command and verify the upgrade with the super password. In this way, if the HWTACACS server fails to upgrade the user level due to its fault (Down, unreachable, or timing out for response), the upgrade of the user level can be switched to local processing.

NOTE

When the switch authenticates the user during user level upgrade, passwords of users at different levels can be different. When the HWTACACS server authenticates the user during user level upgrade, passwords of users at each different levels should be the same.

- The user changes the password on the HWTACACS server.
- When the password change policy is enabled on the HWTACACS, the user can change the password before or after the password expires.

- Public network users are authenticated on private networks and thus the HWTACACS server can manage public network devices.

Message Exchange Process of Implementing AAA Through HWTACACS

The process of transmitting HWTACACS messages is similar to that of transmitting RADIUS messages. The difference is that the HWTACACS server replies an authentication acknowledgement packet rather than the user authority after the user passes authentication. The user authority is returned only after the authorization process is complete.

Comparisons Between HWTACACS and RADIUS

Compared with RADIUS, HWTACACS is more reliable in transmission and encryption and is more suitable for security control. **Table 1-1** shows the differences between HWTACACS and RADIUS.

Table 1-1 Comparisons between HWTACACS and RADIUS

HWTACACS	RADIUS
Uses the Transmission Control Protocol (TCP) to provide reliable transmission.	Uses UDP.
Encrypts the main structure of a packet except the standard HWTACACS header.	Encrypts only the password field in the authentication packet.
Separates authorization from authentication.	Performs authentication together with authorization.
Is suitable for security control.	Is suitable for accounting.
Authorizes users to use the commands to configure a switch.	Does not support authorization with configuration commands.

1.3.4 Domain-based User Management

In current AAA implementations, users belong to different domains. By default, the domain to which a user belongs depends on the string following at sign (@) in the user name. For example, the user named user@hua belongs to the domain named hua. If there is no at sign (@) in a user name, the user belongs to the default domain.

Access User Management

A switch can manage users based on their domains. The default authorization, RADIUS/HWTACACS template, and authentication and accounting schemes can be configured in a domain.

Authentication, authorization and accounting modes should be pre-configured in the AAA view. To perform AAA for access users, you need to apply the authentication schemes, authorization schemes, and accounting schemes in the domain view.

The default authentication scheme, authorization scheme and accounting scheme adopts local authentication, local authorization, and non-accounting. If no authentication scheme,

authorization scheme or accounting scheme are applied to a new domain, the default authentication scheme and accounting scheme are adopted. In addition, to use the RADIUS or HWTACACS schemes for a user, you must pre-configure the RADIUS or HWTACACS server template in the system view and then apply it in the view of the domain to which the user belongs.

When a domain and the users in the domain are configured with the same attribute, the user-based configuration is preferred over the domain-based configuration.

The authorization configured in a domain has a lower priority than the authorization delivered by an AAA server. That is, the authorization delivered by an AAA server is preferred. When the AAA server does not have or does not support the authorization, the authorization configured in a domain takes effect. In this manner, you can increase services flexibly by means of domain management, regardless of the authorization by the AAA server.

2 NAC

About This Chapter

2.1 Introduction to NAC

Definition

Network access control (NAC) is a security access framework. The idea of NAC is the concept that the security is end-to-end. You need to start considering the security on the terminals of users but not just traditionally start considering the security on network devices.

Purpose

The security control is performed on the user access.

2.2 References

The following table lists the references of this document.

Document	Description	Remarks
RFC3748	Extensible Authentication Protocol (EAP)	-
Portal 2.0	Portal protocol standard for Huawei broadband products (V2.01)	-

2.3 Principles

2.3.1 Basic Structure of NAC

Figure 2-1 Basic structure of NAC

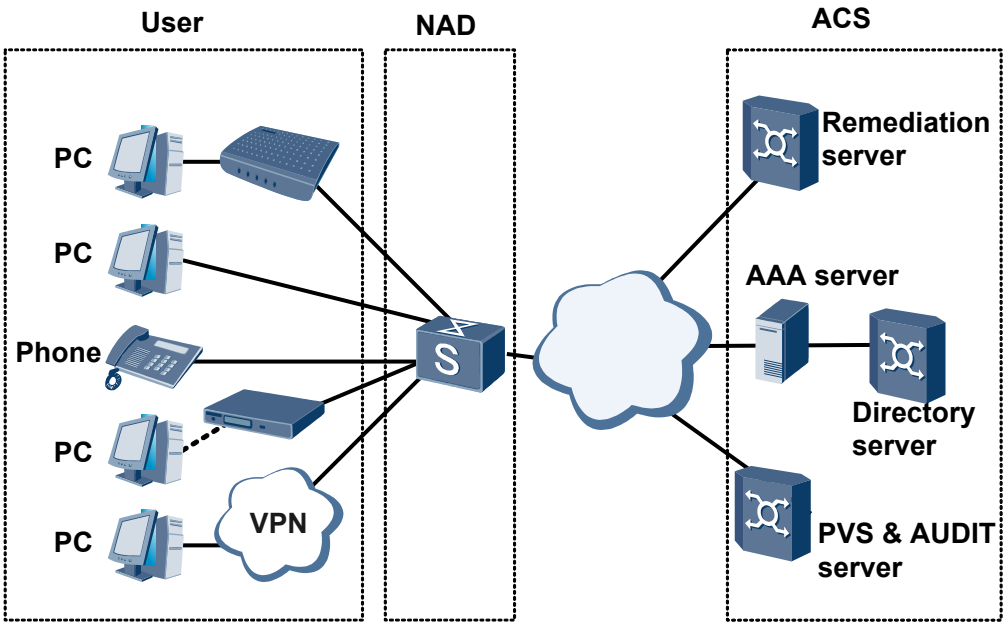

As shown in **Figure 2-1**, NAC is a solution to solve the security access control of the network. NAC includes the following main parts:

- User: indicates access users who need to be authenticated. If 802.1X authentication is adopted, users need to install the client software.
- NAD: indicates a network access device (NAD), such as a router and an S6700. An NAD authenticates and authorizes access users. In general, an NAD needs to cooperate with the AAA server. In this manner, the access of illegal terminals is prevented and the threats of insecure terminals are reduced; unauthorized access of legal terminals is prevented and core resources are protected.
- ACS: indicates access control servers (ACSs). ACSs are used to check the robustness of the terminal security and manage policies; ACSs are used to manage user behaviors, audit violations, strengthen the behavior audit, and prevent malicious terminal damages.

2.3.2 Basic Concepts of NAC

802.1x authentication

802.1x authentication is also called Extensible Authentication Protocol Over Ethernet (EAPOE) authentication and is mainly used to solve the problem of the access authentication of local area network (LAN) users.

The Institute of Electrical and Electronics Engineers (IEEE) 802.1x standard, 802.1x in brief, is an interface-based network access control protocol. Interface-based network access control indicates the authentication and control implemented for access devices on an interface of a LAN access control device. User devices connected to the interface can access the resources on the LAN only after passing the authentication.

The 802.1x protocol is concerned about only the status of an access interface. When a legal user accesses an interface by using the account and password, the interface is enabled; when an illegal

user accesses an interface or no user accesses an interface, the interface is disabled. The authentication result is about the change of the interface status but is not involved with the IP address negotiation and assignment that need to be considered in common authentication technologies. 802.1x authentication is the most simplified implementation solution among various authentication technologies.

802.1x supports interface-based authentication and MAC-based authentication.

- Interface-based authentication: When interface-based authentication is adopted, all the other access users can use network resources and do not need to be authenticated, as long as the first user on an interface passes the authentication. After the first user gets offline, other users cannot use network resources.
- MAC-based authentication: When MAC-based authentication is adopted, all access users on an interface need to be independently authenticated.

802.1x supports the following authentication modes:

- EAP termination authentication: A network access device terminates the EAP packets of users, parses user names and passwords, encrypts the passwords, and then sends them to the AAA server for authentication.
- EAP transparent transmission authentication: EAP transparent transmission authentication is also called EAP relay authentication. A network access device directly encapsulates authentication information about 802.1x users and EAP packets to the attribute fields in RADIUS packets and sends them to the RADIUS server. The network access device does not need to convert EAP packets to standard RADIUS packets and send the standard RADIUS packets to the RADIUS server to complete the authentication.

802.1x supports the following interface control modes:

- Automatic identification: Initially, an interface is in the unauthorized state and allows only the receiving and sending of EAPOL packets; an interface does not allow a user to access network resources. After the user passes the 802.1x authentication, the interface switches to the authorized state and the user can access network resources.
- Forcible authorization: An interface is always in the authorized state and users can access network resources without authentication.
- Forcible unauthorization: An interface is always in the unauthorized state and users cannot access network resources.

MAC address authentication

MAC address authentication is an authentication method that controls the network access authority of a user based on a MAC address and the user does not need to install any client software. The user name and password are the MAC address of the user or the fixed user name configured on a device. After a device first detects the new MAC address of the user by sending DHCP or ARP messages, the device starts to authenticate the user.

MAC bypass authentication

A network access device firstly triggers a user to use 802.1x authentication. If the user does not perform 802.1x authentication for a long time, the MAC address of the user is used as authentication information and is used as the user name and password to be sent to the AAA server for authentication.

Web authentication

Web authentication is also called Portal authentication. The basic principle is that when a user opens a browser for the first time and enters any Website address, the user is forcibly redirected to the authentication page of the Web server and the user can access network resources only after passing the authentication. Unauthenticated users can access only some specified site servers. Web authentication uses the Portal protocol to finish the authentication process after a user name and password are entered on a Web page.

The Portal protocol is mainly used in the message exchange between Web servers and other devices. The Portal protocol is based on the client/server model and uses the User Datagram Protocol (UDP) as the transmission protocol. During Web authentication, the Web authentication server communicates with the S6700 acting as a client through the Portal protocol. After obtaining the user name and password entered by the user on the authentication page, the Web authentication server transmits them to the S6700 through the Portal protocol.

Guest VLAN

After the guest VLAN function is enabled, the S6700 adds the interface to which a user belongs to a guest VLAN, if the user authentication fails. Then, users on the guest VLAN do not need to be authenticated when accessing the resources of the guest VLAN but need to be authenticated when accessing external resources. In this manner, the requirement for allowing unauthenticated users to access specific resources is met.

The guest VLAN function is used in MAC address authentication, MAC address bypass authentication, or 802.1x authentication.

2.3.3 802.1X Authentication

802.1X authentication is also called EAPOE authentication and is mainly used to solve the problem of the access authentication of LAN users.

EAP is not an authentication mechanism but a common architecture. EAP is used to transmit actual authentication protocols. The advantage of EAP is that when a new authentication protocol is developed, the basic EAP mechanism does not need to be changed. Currently, there are more than 20 types of EAP protocols.

The following figure shows the EAP protocol stack and the supported authentication protocols:

Figure 2-2 EAP protocol framework

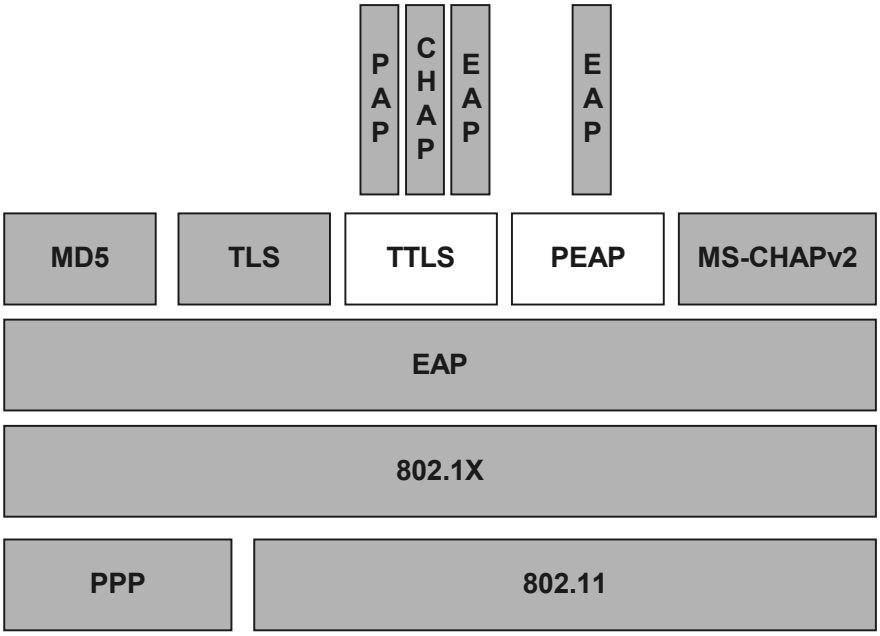

EAP Termination Authentication

Figure 2-3 EAP termination authentication

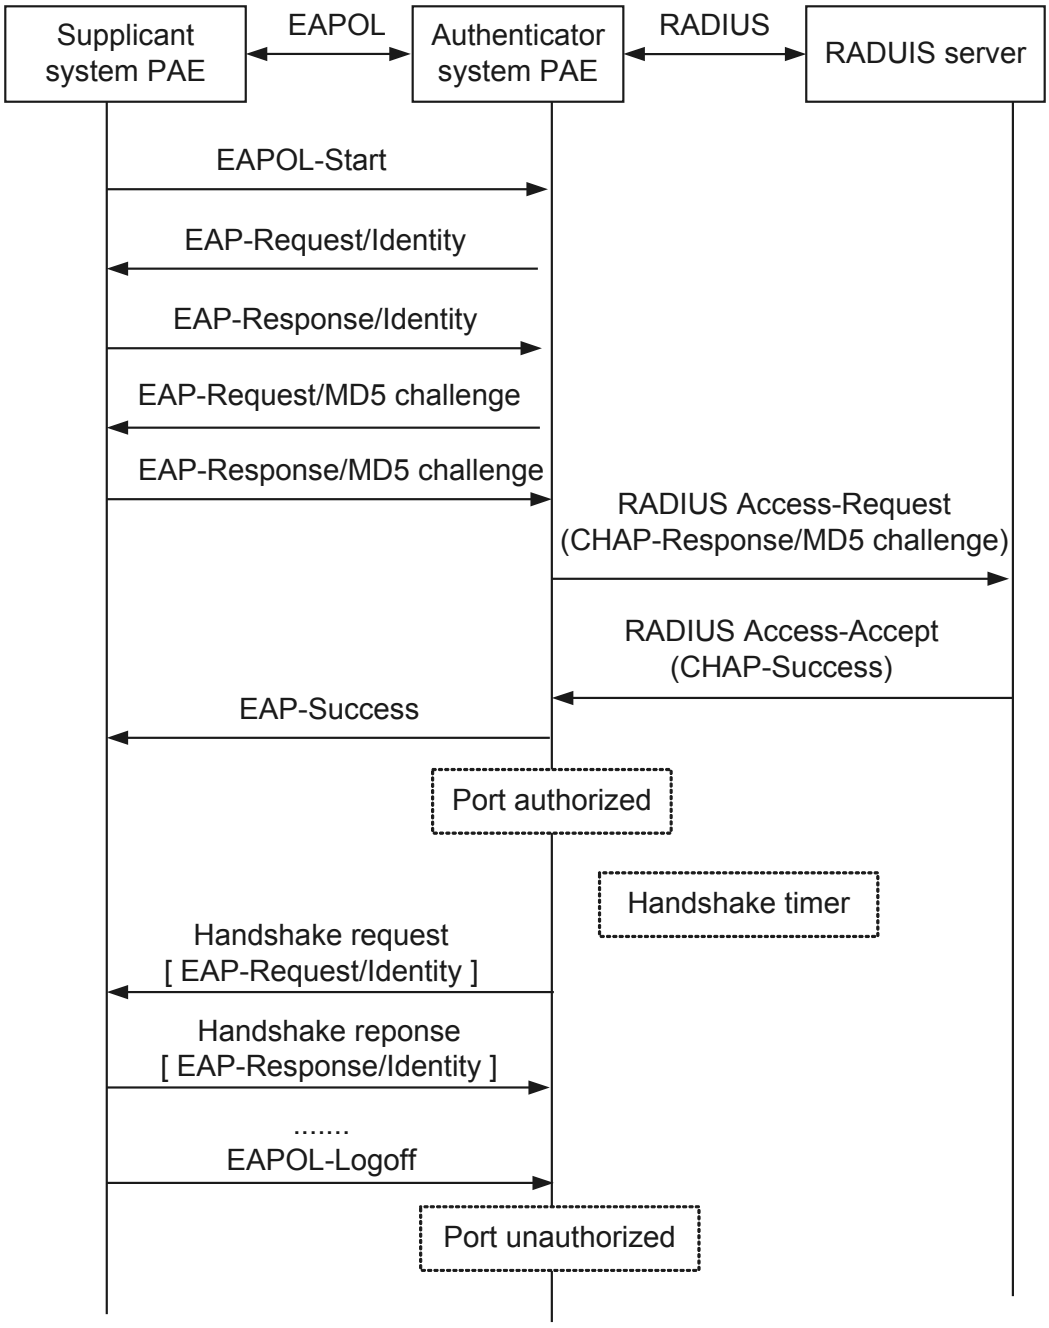

The process of EAP termination authentication is as follows (A device indicates a network access device):

1. The EAP client sends an EAP-Start packet to a device.
2. The device sends an EAP-Request/Identity packet to the EAP client.

3. The EAP client responds with an EAP-Response/Identity packet carrying user name information.
4. The device sends an EAP-Request/MD5-Challenge packet to the EAP client.
5. The EAP client responds with an EAP-Response/MD5-Challenge packet and the device obtains password information about the EAP client.
6. The device carries user account information for authentication in the AAA system.
7. After passing the authentication, the device notifies the EAP client of the authentication success and the interface is enabled.
8. The device detects whether the EAP client remains online according to EAP detection.

EAP Transparent Transmission Authentication

Figure 2-4 EAP transparent transmission authentication

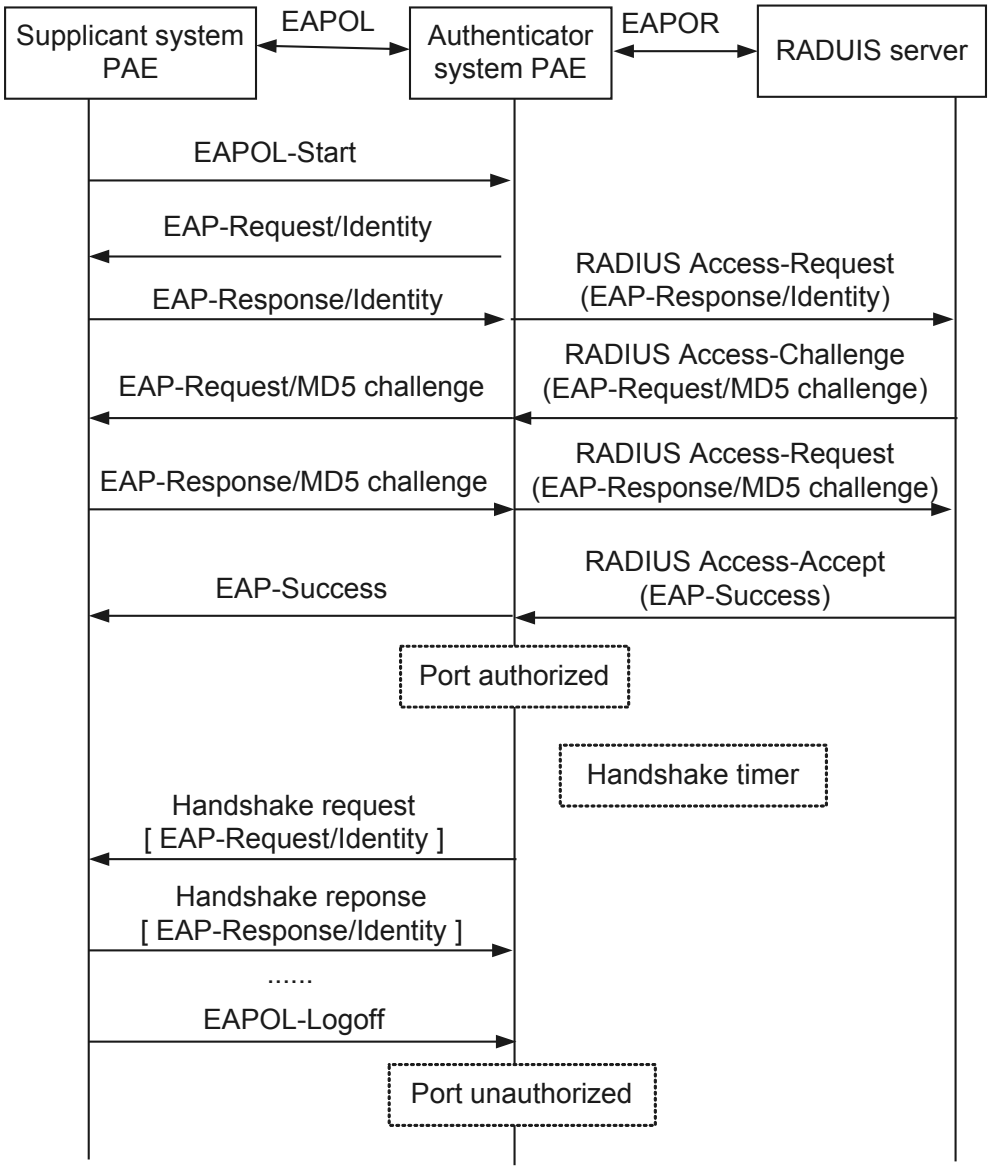

The process of EAP transparent transmission authentication is as follows (A device indicates a network access device):

1. The EAP client sends an EAP-Start packet to a device.
2. The device sends an EAP-Request/Identity packet to the EAP client.
3. The EAP client responds with an EAP-Response/Identity packet and the device transparently transmits the packet to the RADIUS server.
4. After receiving an RADIUS challenge packet, the device sends an EAP-Request/MD5-Challenge packet to the EAP client.
5. The EAP client responds with an EAP-Response/MD5-Challenge packet and the device transparently transmits the packet to the RADIUS server.
6. After passing the authentication, the device notifies the EAP client of the authentication success and the interface is enabled.
7. During the login of the EAP client, the device detects whether the EAP client remains online according to EAP handshake packets.

2.3.4 MAC Address Authentication

Switches need to support the access of terminals such as printers that do not support 802.1X authentication and temporary visitors. MAC address authentication and Web authentication need to be supported.

MAC Address Authentication

MAC address authentication is an authentication method that controls the network access authority of a user based on an interface and a MAC address. The user does not need to install any client software. The user name and password are the MAC address of the user device. After detecting the MAC address of the user for the first time, a network access device starts authenticating the user.

MAC Bypass Authentication

An access device firstly triggers a user to use 802.1X authentication. If the user does not perform 802.1X authentication for a long time, the MAC address of the user is used as authentication information and is used as the user name and password to be sent to the AAA server for authentication.

As shown in **Figure 2-5**, when the S6700 initiates EAP authentication requests for multiple times (times being configurable) and the terminal does not respond to the EAP authentication requests, the S6700 obtains the MAC address of the terminal and initiates authentication in the AAA server.

Figure 2-5 MAC bypass authentication

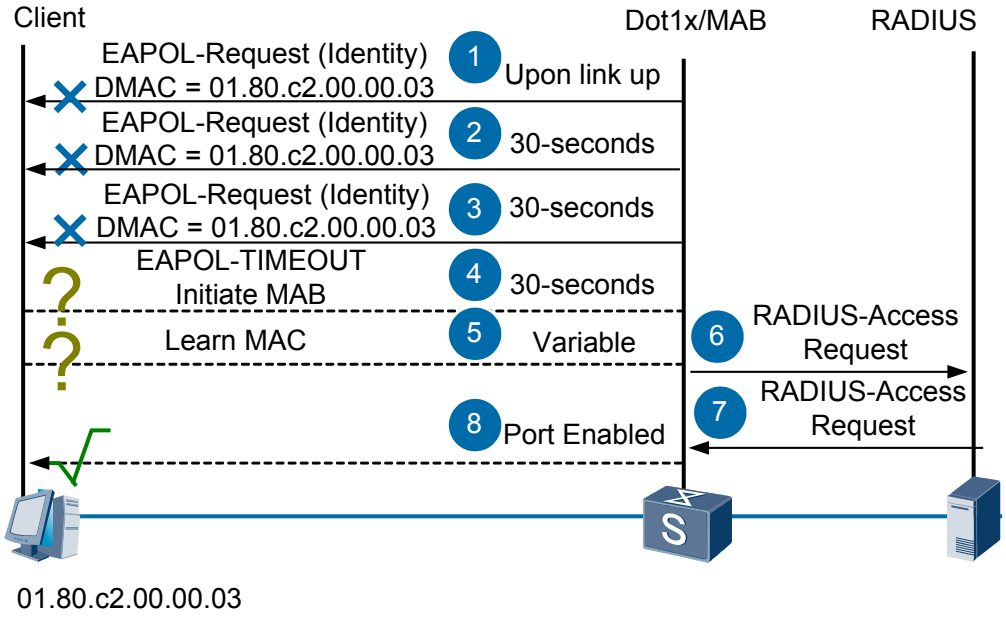

2.3.5 Web Authentication

Switches need to support the access of terminals such as printers that do not support 802.1X authentication and temporary visitors. MAC address authentication and Web authentication need to be supported.

Temporary users such as visitors do not support 802.1X terminals. After accessing any Website, a temporary user is redirected to the login Website. After the user name and password are entered, the S6700 submits them to the AAA server for authentication.

Figure 2-6 Web authentication

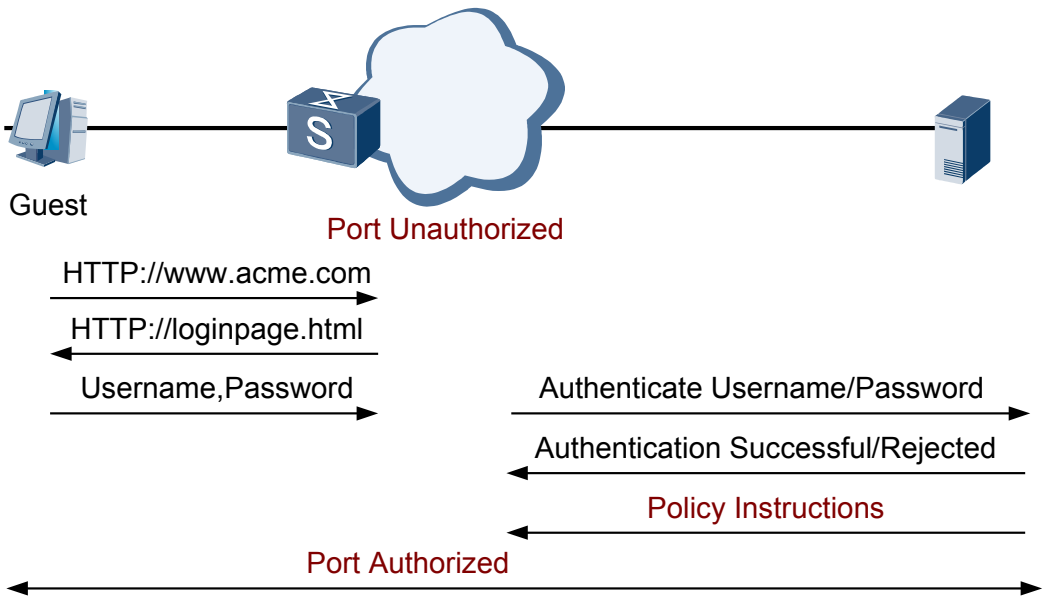

2.3.6 Guest VLAN

When 802.1x authentication, MAC address authentication, or MAC address bypass authentication is adopted, you can configure a guest VLAN. When a user terminal does not respond to the authentication request or the user does not pass the authentication, the user is added to the guest VLAN and can access only the resources in the guest VLAN. Users in the guest VLAN can access resources in the guest VLAN without authentication but must be authenticated when they access external resources. Thus certain resources are available for users who do not pass the authentication.

Figure 2-7 Guest VLAN

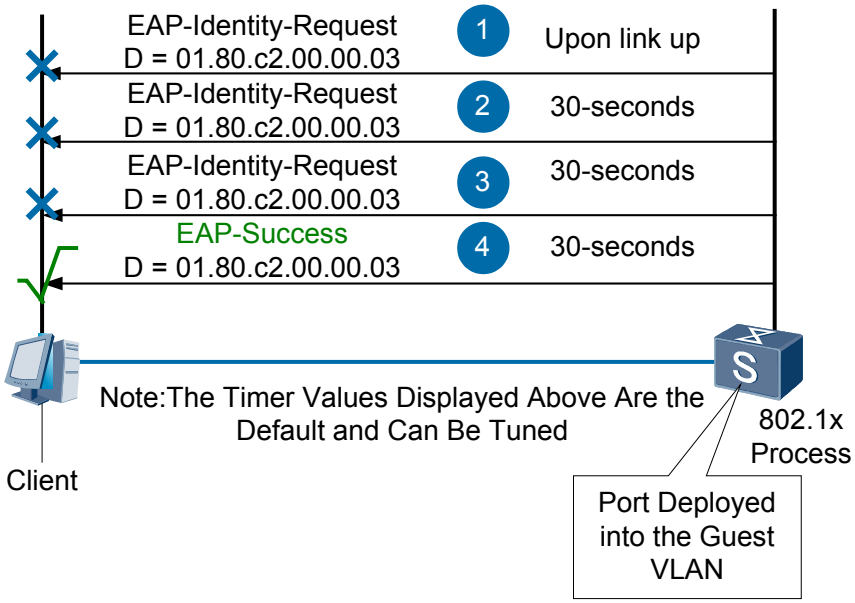

2.4 Applications

2.4.1 Web Authentication in NAC

As shown in **Figure 2-8**, the authentication terminal does not need to be configured with the terminal software. The Switch configured with Web authentication redirects users to the login page. The users need to enter user names and passwords. Then the Switch sends the user names and passwords to the access control server (ACS) for authentication. Before being authenticated, users can access only the resources in the separated area.

After users are authenticated, an HTTP link is established between the authentication terminal and the ACS. The ACS checks security of the authentication terminal. If the authentication terminal is authenticated, the user can access the common information area or the core information area according to the authority of the users.

Figure 2-8 Web authentication

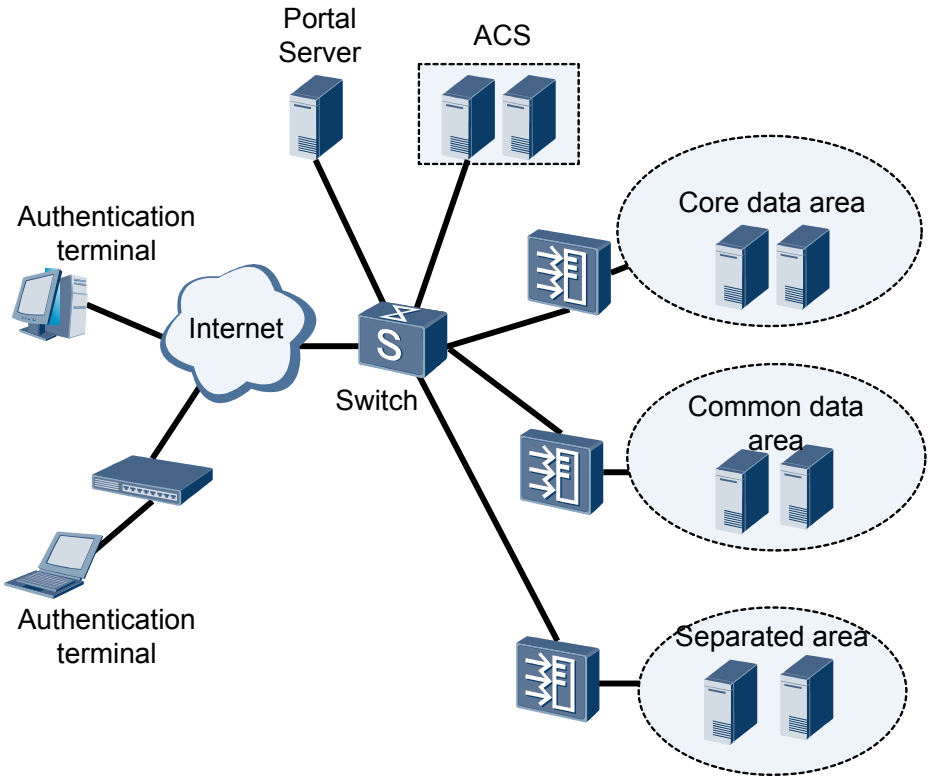

In **Figure 2-8**, the authentication terminal does not need to be installed with any terminal software and can access servers in the separated area. The network must support the Portal protocol because Web authentication uses the Portal protocol to complete the authentication process.

2.4.2 802.1x Authentication in NAC

As shown in **Figure 2-9**, after the authentication terminal has the terminal software installed, it sends an authentication request to the Switch. The Switch exchanges information with the authentication terminal and sends user information to the ACS. After authentication is successful, the Switch enables the interface connected to the authentication terminal and allows the authentication terminal to access the network.

Figure 2-9 802.1x Authentication

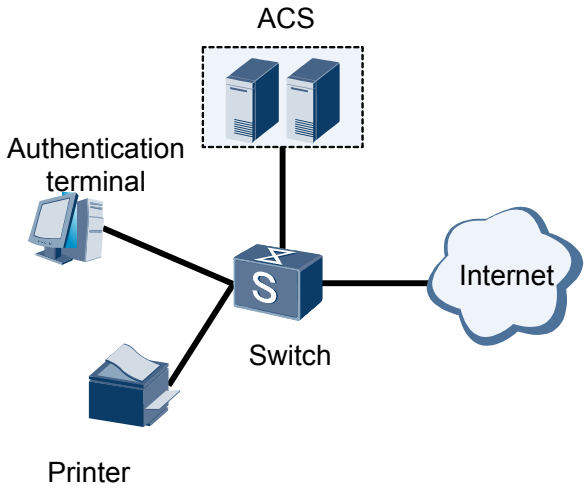

802.1x authentication is relevant to the change of the interface status and is the most simplified implementation because common authentication technologies involve IP address negotiation and assignment. 802.1x authentication, however, requires that the terminal software be installed.

You can configure MAC address bypass authentication for the terminals that do not require high security.

3 MFF

About This Chapter

3.1 Introduction to MFF

Definition

MAC-Forced Forwarding (MFF) is a method for Layer-2 separation and Layer-3 interworking between customer hosts in the same broadcast domain. MFF provides Ethernet Access Nodes (EANs) with traffic rate limitation and prohibits traffic from being directed from one customer interface to another. Access between customer hosts can be implemented only after Layer-3 forwarding is performed. In this manner, traffic policing and accounting can be performed to all customers.

Purpose

An access network solution needs to ensure Layer-2 separation between customer hosts and enable traffic to be forwarded through gateways. This is to implement traffic policing and accounting and ensure network security. Assuming that the connection interface of each customer host is a separate IP interface, an EAN implements Layer-2 separation through Layer-3 forwarding.

Point-to-Point Protocol over Ethernet (PPPoE) and Virtual Local Area Network (VLAN) per-customer can address the problem of Layer-2 separation. PPPoE, however, cannot take advantage of Ethernet broadcast domains when being employed in multicast applications because premature traffic replication wastes a large amount of bandwidth. Whereas, VLAN per-customer is restricted by the total number of VLANs. Although VLAN stacking is not limited by the total number of VLANs, it increases the complexity of network planning.

MFF can take advantage of Ethernet broadcast domains without wasting IP addresses or being restricted by the total number of VLANs.

Figure 3-1 MFF application scenario

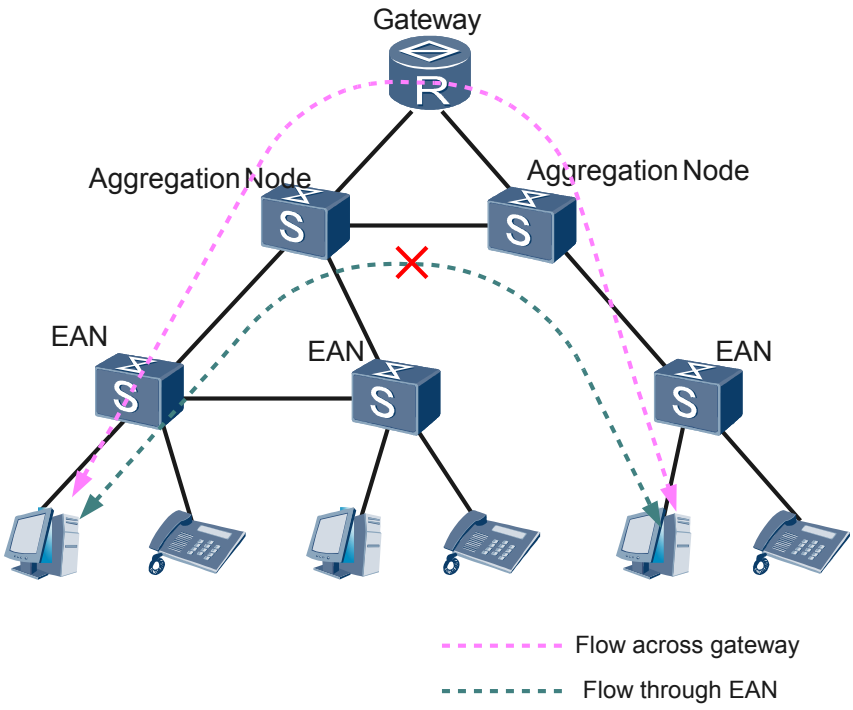

Figure 3-1 shows a scenario where an access network service provider offers Internet access, Internet Protocol Television (IPTV), and voice services. Commonly, the access network service provider needs to perform unified management and accounting of services in this scenario. MFF enables traffic from a customer premise to pass through gateways before being directed to other customers through Layer-3 forwarding. This implements Layer-2 separation and avoids bandwidth waste in multicast applications.

Benefits

Benefits Brought to Service Providers

- Implementation of Layer-2 separation and Layer-3 interworking on an Ethernet access network saves bandwidth, simplifies network planning, and prevents malicious attacks from occupying bandwidth.

Benefits Brought to Customers

- MFF improves network security and stability, and defends customers against malicious attacks.

3.2 References

For more information about MFF, refer to the following document.

Document No.	Description	Remarks
RFC4562	MAC-Forced Forwarding	None

3.3 Technical Description

The major aspects of an MFF solution are as follows:

- Obtaining the IP and MAC addresses of the Access Routers (ARs)
- Responding to ARP requests
- Limiting the traffic rate
- Restricting access to application servers

The functions of responding to ARP requests and limiting the traffic rate can be implemented only after the IP and MAC addresses of the ARs are learnt. Restriction of access to application servers, which is a special setting for responding to ARP requests, ensures the customer hosts to access corresponding application servers.

3.3.1 Obtaining the IP Addresses and MAC Addresses of the Gateway

The method in which an EAN learns IP addresses of gateway varies with the method of obtaining an IP address. An EAN can learn IP addresses of gateway by using the following methods:

- If the gateway IP address is configured statically, configure this IP address statically on the EAN. After the EAN determines the gateway IP address, if the EAN does not learn the

gateway MAC address, the EAN discards the ARP request packet. Then the EAN constructs an ARP request packet with the user IP address and MAC address to request the gateway MAC address. After receiving the ARP reply packet, the EAN learns the gateway MAC address.

- On a network where IP addresses are allocated dynamically using DHCP, an EAN listens on DHCP ACK packets from network interfaces and obtains the IP addresses of the gateway. DHCP option 121 or option 3 of the DHCP ACK packets contains the IP addresses of the gateway that users are authorized to access. Then the EAN parses the field and learns the IP addresses of the gateway. The EAN constructs an ARP packet with the users' IP and MAC addresses and sends an ARP request to each AR. Then the EAN learns the MAC addresses of the gateway from the ARP Reply packets.

The methods of obtaining IP addresses and MAC addresses of gateway are as follows:

- Assigning IP addresses to the gateway statically

  MFF supports static assignment of IP addresses.

  On a network where IP addresses are manually assigned to gateway, manually configure the gateway that can be accessed by user hosts. After you assign an IP address to an AR, the MFF module identifies online users by capturing the ARP packets from the user side. Each ARP packet containing a new source IP address triggers the MFF module to record information about a new user until the number of users exceeds the upper threshold.

  MFF supports a maximum of 1024 static and dynamic users. When the number of users exceeds the upper threshold, the EAN discards packets from new users and stops responding to ARP requests from the gateway.

  After the MFF module learns MAC addresses of the gateway, the EAN responds to ARP request packets from the user side. The EAN uses the MAC addresses of the gateway to control ARP learning of user hosts.

- Dynamically assigning gateway using DHCP snooping

  MFF supports dynamic assignment of gateway that can be accessed by user hosts through DHCP snooping.

  When DHCP snooping and MFF are configured in a VLAN simultaneously, the EAN focuses on interaction between user hosts and DHCP servers. When the DHCP server authorizes a user to access the network, a DHCP binding entry is generated. At the same time, the EAN records information about the authorized user. The EAN parses DHCP option 121 or 3 carried in the DHCP ACK packet and obtains the list of accessible gateway.

  The EAN obtains IP and MAC addresses of users according to obtained information about the gateway and encapsulates the information into an APR request packet. Then the EAN sends the ARP request packet to the network interface to detect MAC addresses of the gateway.

  If the user host is authorized to access multiple gateway, the EAN uses the MAC address of the first AR to respond to the ARP request packet when capturing an ARP request that will be received by a non-AR device. When capturing an ARP request packet that be received by an AR, the EAN responds to the request packet using the MAC address of the AR.

  When a new AR is configured, an entry is used. The gateway with the same IP address in the same VLAN can share an entry. When the number of entries exceeds the upper threshold, no more AR can be configured in the VLAN. If the number of new gateway in the list of accessible gateway exceeds the number of available entries, no AR can be configured. If all the accessible gateway have not been configured and each entry is unavailable, no new AR can be configured and the IP address of the user host is not recorded by the EAN.

- MFF supports both static and dynamic configuration of gateway in the same VLAN.

  The user host obtaining an IP address dynamically through DHCP detects packets through DHCP snooping and instructs the MFF module to record information and form an entry. Each user host can access any AR in the VLAN, regardless of whether the AR is statically or dynamically configured.

- Detecting gateway MAC addresses

  MFF supports detection of gateway MAC addresses. After the detection function is enabled, the EAN scans recorded information about the gateway every 30 seconds. If a gateway fails to learn a MAC address, the EAN constructs an ARP request packet by using information about a user and sends the packet to the network interface. Then the EAN obtains the MAC address of the gateway from a ARP Reply packet sent by the gateway. If the MAC address changes, the EAN updates information about the gateway immediately.

  NOTE

  If no user is recorded in the VLAN, the EAN does not send ARP request packets until the user goes online.

3.3.2 Responding to ARP Requests

After capturing an ARP request from a customer host, the EAN replies to the request with an ARP reply with the MAC address of the AR as the source address. In this manner, the IP addresses recorded in the ARP tables of the customer hosts correspond to the MAC addresses of the gateway, and all the packets from the customer hosts are destined for the gateway when being directed through Layer-2 forwarding. In this manner, traffic policing and accounting can be performed and network security is improved.

The EAN can respond to network-side ARP request packets. The EAN uses the user IP address and MAC address to respond to ARP request packets sent by the gateway and server.

NOTE

If the source IP address of the ARP packets that the EAN receives differs from what specified in the MFF table, the EAN discards the ARP packets as invalid ones.

3.3.3 Limiting the Traffic Rate

MFF provides filtering for upstream traffic and rate limitation for broadcast traffic.

Thanks to the function of responding to ARP requests, the customer hosts do not learn the MAC addresses of other hosts. Certain malicious users, however, may send a great amount of traffic which is not destined for the MAC addresses of the ARs. This wastes a great deal of bandwidth. Therefore, the EAN restricts upstream traffic from customer hosts. On a VLAN enabled with MFF, if the destination MAC address of upstream traffic does not pertain to the specified AR, the EAN discards the packets to avoid bandwidth waste.

Similarly, most of the broadcast and multicast traffic from customers is stopped by the EAN, and is not forwarded to the network side or other user interfaces. ARP, DHCP, Internet Group Management Protocol (IGMP), and EAPOL packets are not stopped by the EAN.

- ARP packets are replied to or discarded after being intercepted by the MFF module
- DHCP packets are processed in the original method of the DHCP snooping module
- IGMP packets are transparently transmitted to the network side to ensure smooth operation of IGMP
- EAPOL packets are processed in the original flow specified in 802.1x

All traffic from the network side is not restricted by EANs. Therefore, you can discard unknown multicast packets to filter out unknown downstream multicast traffic. IGMP packets from network interfaces can be forwarded to the customer side normally.

By controlling the broadcast traffic from the user side, MFF reduces the possibility of bandwidth waste.

3.3.4 Restricting Access to Application Servers

You can specify the IP addresses of application servers on an EAN to set a list of accessible application servers. If the source IP address of the ARP request packet is the same as the IP address of one of the application servers, the EAN replies to the AR with the requested MAC address. In this manner, the customer can access all the set application servers.

NOTE

You can set a list of accessible application servers in the VLAN view. Then, all the VLAN users can access the listed application servers.

3.3.5 Discarding IPv6 Packets

If IPv6 packets are broadcast in a VLAN, users can learn about MAC addresses of each other, and the MFF user isolation function becomes invalid. To solve the problem, the S6700 discards IPv6 packets.

3.3.6 Transparent Transmission of ARP Probe Packets

On a network enabled with MFF, if the gateway performs accounting for users based on online duration, the gateway must know whether a user is online at a specified moment. By default, the S6700 responds to user requests for the gateway. The S6700 always responds for the gateway until the IP address lease expires and the binding entries ages out. The gateway detects that users are online even if the users have gone offline. The S6700 transparently transmits ARP probe packets from the gateway. If the user status changes, the MFF-enabled S6700 sends the updated user information to the gateway.

3.4 Applications

MFF isolates clients at Layer 2 and connects clients at Layer 3 in a broadcast domain. MFF captures Address Resolution Protocol (ARP) request packets and sends ARP response packets with the gateway MAC address through proxy ARP. All traffic from users is forwarded to the gateway so that the gateway monitors the traffic and prevents attacks. The network security is ensured.

ARP anti-spoofing can be configured on the S6700 to prevent unauthorized users from modifying ARP entries and ensure that authorized user can access the Internet.

Figure 3-2 MFF isolates users at Layer 2

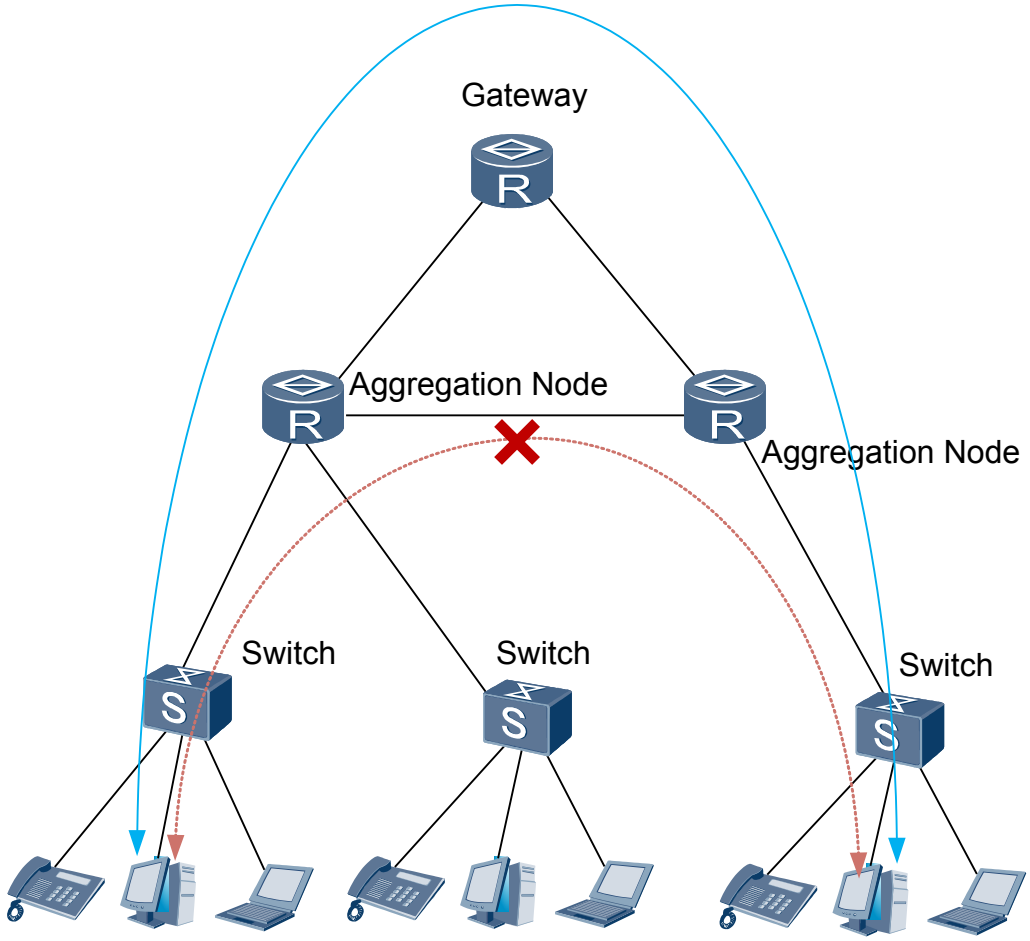

As shown in **Figure 3-2**, user traffic is sent to the gateway, but not the switch. Users are isolated at Layer 2.

4 ACL

About This Chapter

4.1 Introduction to the ACL

Definition

An Access Control List (ACL) uses the 5-tuple to match data packets, and permits, denies, or perform other actions on the matching packets.

The 5-tuple (source address, source port, destination address, destination port, and protocol number) identifies a type of IP data packets.

Figure 4-1 IP data packet

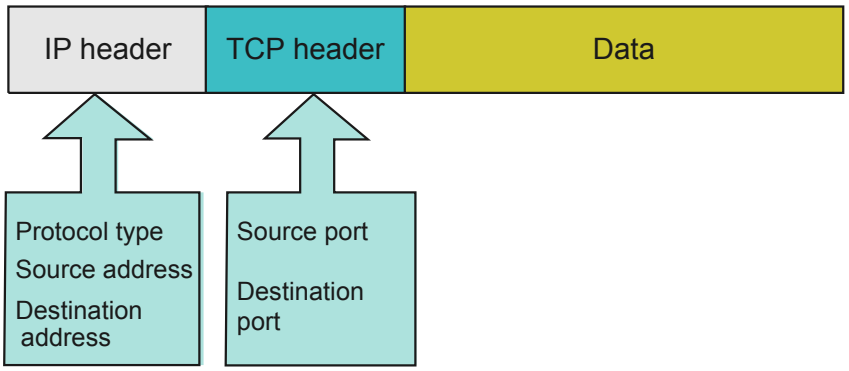

For example, you can set a rule in the ACL to prevent any user terminal from logging in to a device using Telnet, or to allow every user terminal to send emails to a device using the Simple Mail Transfer Protocol (SMTP).

Table 4-1 shows the classification of ACLs.

Table 4-1 Classification of ACL

Classification Rule	Type	Application Scenario	Description
Support for IPv4 or IPv6	ACL4	ACL4s are applicable to IPv4.	ACL4 and ACL6 commands are different.
	ACL6	ACL6s are applicable to IPv6.	ACL4 and ACL6 commands are different.
Naming mode	Numbered ACL	A numbered ACL is identified by a unique number. You can reference the ACL by specifying the ACL number.	-

Classification Rule	Type	Application Scenario	Description
	Named ACL	A named ACL is identified by a character string. You can reference the ACL by specifying the ACL name. Compared with numbered ACLs, named ACLs are easy to identify and remember. The S6700 provides flexible ACL command modes.	You can configure numbered ACLs and named ACLs simultaneously. The names of IPv4 ACLs and IPv6 ACLs are unique, whereas an IPv4 ACL can use the same name as an IPv6 ACL.
Functions	Basic ACL	A basic ACL defines rules based on the source IP addresses, fragment flag, or time range of packets.	The number of a basic ACL ranges from 2000 to 2999.
	Advanced ACL	An advanced ACL can define rules based on the following parameters including the source IP addresses of packets, destination IP addresses, IP priorities, ToSs, DSCP values, IP protocol types, ICMP types, TCP source/destination ports, and UDP source/destination ports. Compared with basic ACLs, advanced ACLs define more accurate, diversified, and flexible rules.	The number of an advanced ACL ranges from 3000 to 3999.
	Layer 2 ACL	A Layer 2 ACL defines rules based on the information in Ethernet frame headers of packets, such as the source MAC address, destination MAC address, and Ethernet frame protocol number.	The number of a Layer 2 ACL ranges from 4000 to 4999.

Purpose

ACLs filter specific data packets by defining a series of rules, and identify the objects that need to be filtered. After identifying specific objects, the ACLs permit or reject the data packets based on the preset policy.

The S6700 selects data packets by using a serials of rules defined in ACLs.

ACL rules are applied to the following modules:

- Policy-based routing (PBR)

- Route filtering
- QoS
- Modular QoS Command Line (MQC)
- Device security

The S6700 can classify Layer 2 and Layer 3 packets by using ACLs. To classify the combination of Layer 2 and Layer 3 packets, use MQC and ACLs together.

4.2 References

The following table lists the references of this document.

Document	Description	Remarks
RFC 4314	Defines several new access control rights and clarifies which rights are required for different IMAP (Internet Message Access Protocol) commands.	-

4.3 Principles

4.3.1 Principles

An ACL manages all rules configured by users and provides rule matching methods. Service packets are permitted or rejected according to matching rules.

ACL Rule Management

As a group of rules, each ACL can store multiple rules. When the number of configured ACL groups or rules exceeds the maximum, the system displays a configuration failure message. A rule is identified by a rule ID, which is set by a user or generated by the system according to the rule step. All rules in an ACL are arranged in ascending order of rule IDs. There is a step between rule IDs. For example, if the ACL step is set to 5, rules are numbered 5, 10, 15, and so forth. You can add new rules based on the step.

ACL Rule Matching

When a packet reaches a device, the search engine retrieves information from the packet to constitute the key value and matches the key value with rules in the rule group. When a matching rule is found, the system stops the matching. The system then processes the packet according to the action defined in the rule. If the permit action is defined, the system forwards the packet. If the deny action is defined, the system discards the packet. If no rule matches the packet, the system does not process the packet.

4.3.2 Matching Order of ACL Rules

An ACL can consist of multiple **deny** and **permit** statements. Each statement describes a rule. These rules may overlap, that is, one rule contains another rule but the two rules are not

completely the same. The matching order of ACL rules determines the priorities of ACL rules to be matched with a packet.

The S6700 supports two matching orders, that is, configuration order and automatic order.

Configuration Order

By default, ACL rules are matched according to their configuration order. The rule that is configured first is matched first.

Automatic Order

The automatic order follows the depth-first principle.

The depth-first principle means that the statement with smallest range of hosts is placed before the others. This can be implemented by checking the address wildcard. A smaller wildcard indicates a narrower host range.

For example, 129.102.1.1 0.0.0.0 specifies a host with the IP address being 129.102.1.1, and 129.102.1.1 0.0.0.255 specifies a network segment ranging from 129.102.1.1 to 129.102.1.255. Therefore, the former rule is placed before the latter because it specifies a narrower range. The detailed standards are as follows:

- Basic ACL4 rules
  1. The rule that defines the smallest source IP address range is matched first. A greater number of 0 bits in the wildcard mask indicates a smaller source IP address range.
  2. The rule that is configured first is matched first if the source IP address ranges are the same.

  NOTE

  A wildcard mask is also called inverse mask and is in dotted decimal notation. In a binary wildcard mask, the 1 bits represent the host ID and the 0 bits represent the network ID. For example, Class C subnet 192.168.1.0 corresponds to the subnet mask 255.255.255.0. The wildcard mask is 0.0.0.255.
- Advanced ACL4 rules
  1. The rule that defines the protocol type is matched first.
  2. The rule that defines the smallest source IP address range is matched first if the protocol types are the same. A greater number of 0 bits in the wildcard mask indicates a smaller source IP address range.
  3. The rule that defines the smallest destination IP address range is matched first if the protocol types and source IP address ranges are the same. A greater number of 0 bits in the wildcard mask indicates a smaller destination IP address range.
  4. The rule that defines the smallest Layer 4 port number (TCP/UDP port number) range is matched first if the protocol types, source IP address ranges, and destination IP address ranges are the same.
  5. The rule that is configured first is matched first if the preceding ranges are the same.
- Layer 2 ACL rules:
  1. The rule that defines the smallest source MAC address range is matched first. A greater number of 1 bits in the mask indicates a smaller source MAC address range.
  2. The rule that defines the smallest destination MAC address range is matched first if the source MAC address ranges are the same. A greater number of 1 bits in the mask indicates a smaller destination MAC address range.

3. The rule that is configured first is matched first if the source and destination MAC address ranges are the same.

4.3.3 Setting the Step for an ACL

Setting the Step for an ACL Rule Group

You can run the **step** command to set the step for an ACL rule group. The step means the difference between the IDs that are automatically assigned to rules in the ACL rule group. For example, if the step is set to 5, the rule IDs are multiples of 5 (beginning with 5), such as 5, 10, and 15. By default, the step of an ACL rule group is 5.

The rule IDs in an ACL rule group realign automatically when the step changes. For example, if the original rule IDs are 5, 10, 15, and 20, the rule IDs become 2, 4, 6, and 8 after you run the **step 2** command to set the step to 2.

If rule IDs are not evenly distributed, the distribution of the rule IDs becomes even after you run the **step** command. For example, if the current step is 5 and the rule IDs are 1, 3, 10, and 12, the rule IDs become 2, 4, 6, and 8 after you run the **step 2** command to set the step to 2.

NOTE

If the current step is 2 and rule IDs are 1, 3, 10, and 12, the rule IDs remain unchanged after you run the **step 2** command. To change the rule IDs to 2, 4, 6, and 8, you need to run the **undo step** command to change the rule IDs to 5, 10, 15, and 20, and then run the **step 2** command to change the rule IDs to 2, 4, 6, and 8.

Restoring the Default Step

You can run the **undo step** command to restore the default step and realign rule IDs.

The **undo step** command can be used to realign ACL rule IDs based on the default step. For example, ACL rule group 1 contains four rules with IDs being 1, 3, 5, and 7, and the step is 2. After the **undo step** command is run, the rule IDs become 5, 10, 15, and 20 and the step becomes 5.

Function of the Step

The step can be used to change the number of supported ACL rules if the number of ACLs supported by a device is fixed. A greater step indicates a smaller number of supported ACL rules in an ACL.

4.3.4 ACL Supporting Fragmented Packets

Traditional packet filtering matches only the first fragmented IP packet (the first fragment) and allows all the subsequent fragmented packets to pass through. This brings risks to the network because attackers may construct the subsequent fragmented packets.

Packet filtering supported by the S6700 can filter all fragmented packets based on Layer 3 information contained in the packets.

In an ACL rule, the keyword **fragment** is used to identify that the rule is valid for only non-first fragmented packets. For the first fragmented packet and non-fragmented packets, this rule is invalid. Configuration rule entries that do not contain this keyword are valid for all packets.

4.3.5 Time Range of an ACL

A time range specifies a period of time. In practice, some ACL rules are required to be valid during a certain period and invalid out of the period. That is, the ACL rules are used to filter packets based on the time range. For example, if staff of a company are forbidden to browse entertainment websites in business hours and are allowed to visit entertainment websites in after-hours by using a specified device, you can define the time range for an ACL. You can set one or more time ranges, and run the **rule** command to apply the time ranges. Then packets can be filtered based on the set time ranges.

4.3.6 Differences Between an ACL4 and an ACL6

The implementation principles of an ACL4 and an ACL6 are mostly the same. **Table 4-2** shows the differences between an ACL4 and an ACL6.

Table 4-2 Differences between an ACL4 and an ACL6

ACL4	ACL6
It supports the step.	It does not support the step.
It supports IPinIP and basic protocols such as ICMP, IGMP, OSPF, TCP, and UDP.	Besides the basic protocols, it supports IPv6-AH and IPv6-ESP used for filtering packets.

4.3.7 User defined ACLs

Basic ACLs, advanced ACLs, and Layer 2 ACLs cannot match Layer 4 to Layer 7 information. User defined ACLs obtain the contents of packets to generate the matching rule according to the defined offset position and offset value of the packets. In this manner, matching rules of packets can be defined flexibly.

NOTE

The user defined ACLs are applicable to only incoming traffic.

The number of a user defined ACLs ranges from 5000 to 5999.

The user defined ACLs support the offset beginning from the Layer 2 header, Layer 3 header of IPv4 and IPv6, and Layer 4 header of packets. Up to eight matching ranges can be configured; each matching range can match up to four bytes; the maximum matching length is 32 bytes.

4.4 Applications

4.4.1 Application of ACLs in Route Filtering

ACLs can be applied to various dynamic routing protocols to filter the routing information advertised and received by the dynamic routing protocols.

For example, on a campus network, users can visit the Internet by using the Switch. Some users are forbidden to access the Internet and some servers such as educational management systems reject external access to ensure information security. In this case, you can define ACL rules on the Switch connected to the Internet to filter packets.

Figure 4-2 Application of ACLs in route filtering

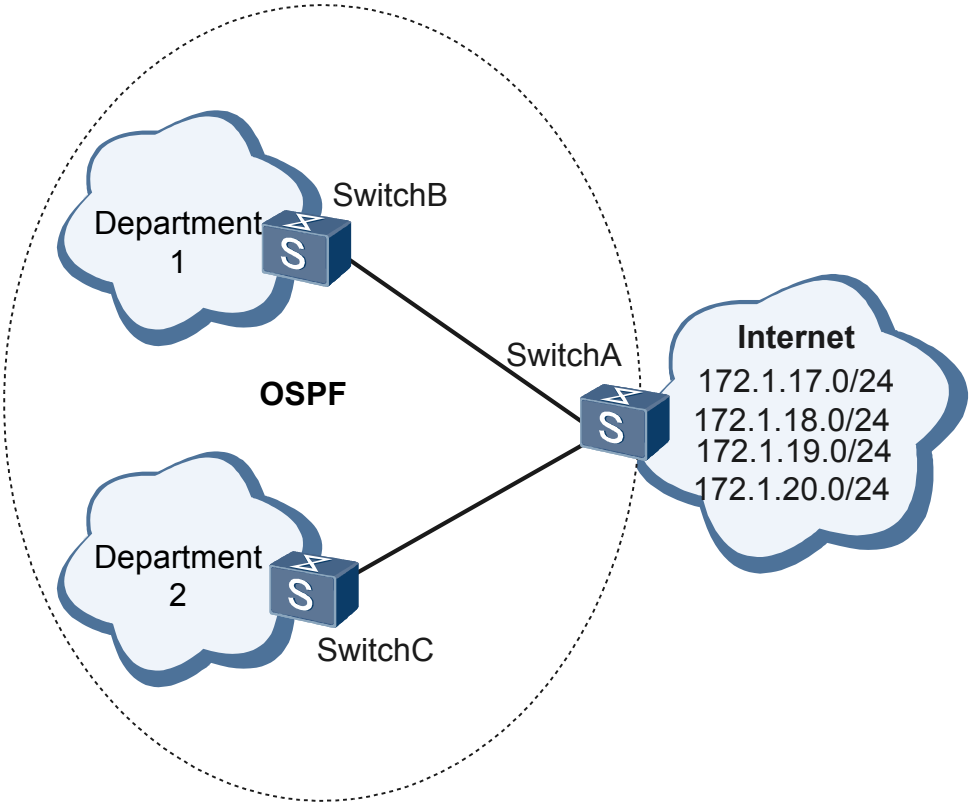

As shown in **Figure 4-2**, SwitchA connects the intranet running OSPF to the Internet. ACLs are defined on SwitchA and applied to OSPF to control route advertisement and receiving. In **Figure 4-2**:

- Switch A provides routes 172.1.17.0/24, 172.1.18.0/24, and 172.1.19.0/24 for Switch B.
- Switch C accepts only route 172.1.18.0/24.

4.4.2 Application of ACLs in QoS

As shown in **Figure 4-3**, NetworkA, NetworkB, and NetworkC connect to NetworkD by using SwitchA. They have different requirements for voice, video, and data services. For example, NetworkA has high requirements for video services. To ensure the quality of video services on Network A, configure an ACL on SwitchA to identify all the packets sent from NetworkA, and then bind the ACL to the traffic policy. In this way, all the packets sent from NetworkA are processed by SwitchA before being forwarded. The quality of video services on Network A is thus ensured. Packets from other networks are forwarded without QoS guarantee because no ACL is matched.

Figure 4-3 Application of ACLs in QoS

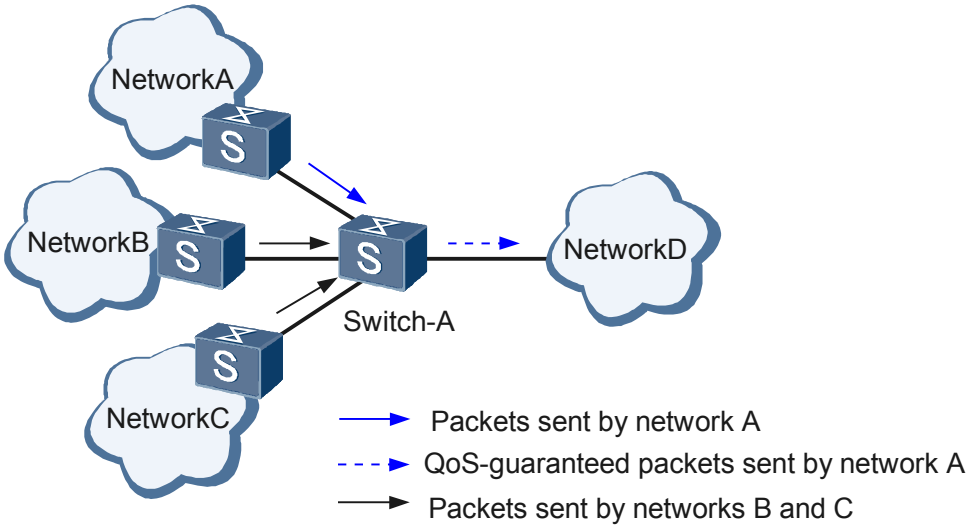

4.4.3 Application of ACLs in the Firewall

On an intranet, different departments have different rights to access the Internet and take different security measures. Firewall policies are also different. To improve security, apply an ACL to the firewall to filter packets.

Figure 4-4 Application of ACLs in the firewall

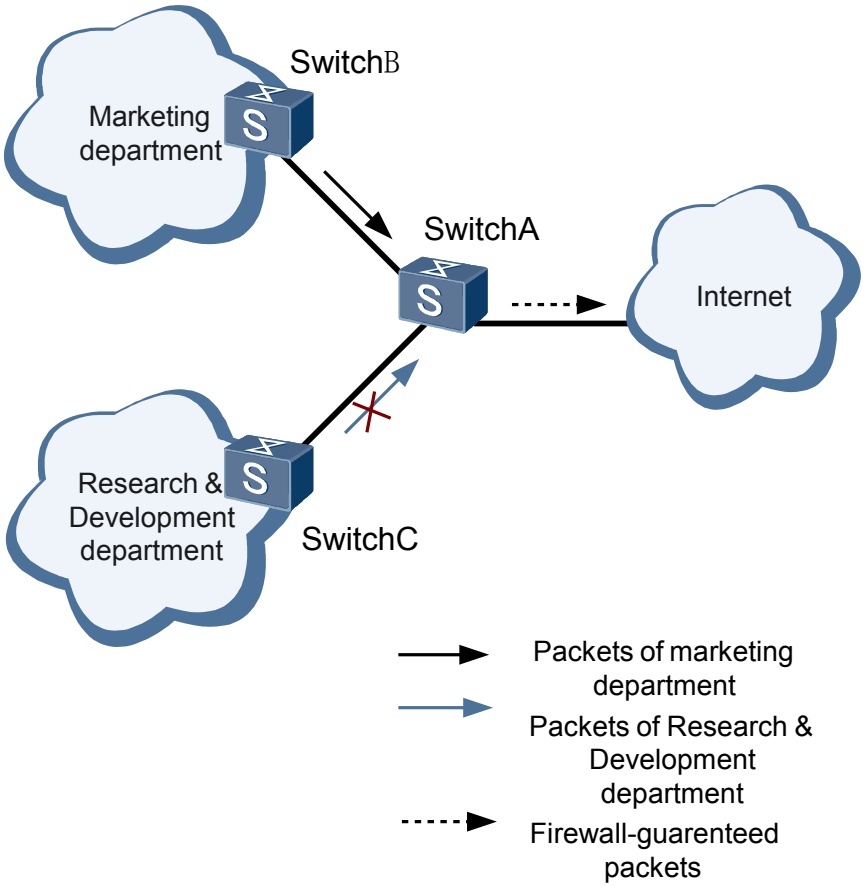

As shown in **Figure 4-4**, staff in the Marketing department are allowed to access the Internet with security ensured and staff in the Research & Development department are forbidden to access the Internet. You can deploy an ACL and configure the firewall on the Switch to meet this requirement.

4.4.4 Application of ACLs in IPSec

Communicating parties encrypt data and authenticate the data source by using IPSec on an IP network to ensure confidentiality, integrity, authenticity, and anti-replay. An IPSec tunnel is established on the two devices connecting two networks, whereas users on LANs have different requirements for security. In this case, an ACL can be configured on the LAN egress device to filter packets entering the IPSec tunnel. Services packets allowed by the ACL are protected and services packets rejected by the ACL are not protected.

Figure 4-5 Application of ACLs in IPSec

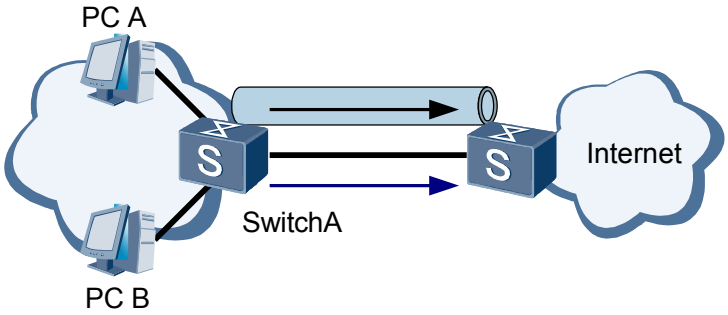

As shown in **Figure 4-5**, an IPSec tunnel is established between SwitchA and SwitchB. An ACL is configured on SwitchA to permit all the packets from PC A to pass through. The ACL is bound to the IPSec policy. In this way, all the packets from PC A are encrypted by SwitchA before being forwarded. All the packets from PC B are forwarded directly without encryption because no ACL is matched.

5 IP Address Anti-spoofing

About This Chapter

5.1 Introduction to IP Address Anti-spoofing

Definition

IP address anti-spoofing includes IP source guard and Unicast Reverse Path Forwarding (URPF). IP source guard checks for IP packets against the binding table. If user IP addresses are dynamically allocated by DHCP, a dynamic binding table is generated after DHCP snooping is enabled. If user IP addresses are configured statically, a static binding table is required. Before the S6700 forwards an IP packet, it compares the source IP address, source MAC address, interface number, and VLAN ID in the IP packet with entries in the binding table. If an entry is matched, the S6700 considers the IP packet as a valid packet and forwards it. Otherwise, the S6700 considers the IP packet as an attack packet and discards it.

URPF prevents network attacks based on source address spoofing.

When the S6700 receives a packet, it searches for a route to the destination address of the packet. If a route is found, the S6700 forwards the packet. Otherwise, the S6700 discards the packet. After URPF is configured, the S6700 obtains the source address and inbound interface of the packet. The S6700 takes the source address as the destination address to retrieve the corresponding outbound interface in the FIB and compares the retrieved interface with the inbound interface. If they mismatch, the S6700 considers the source address as a spoofing address and discards the packet. This allows URPF to effectively protect a device against malicious attacks by blocking packets from bogus source addresses.

Purpose

Some attackers on networks aim at source IP addresses. They use spoofed IP addresses to access network resources, steal users' information, or block authorized users from accessing networks.

Network attacks by using source address spoofing often occur on the Internet. URPF can prevent source IP address spoofing attacks.

Benefits

- IP source guard and URPF prevent source IP address spoofing attacks and reduce maintenance costs.
- IP source guard and URPF improve network security and stability and defend against source IP address spoofing attacks.

5.2 References

None.

5.3 Principles

5.3.1 IP Source Guard

IP source guard checks for IP packets against the binding table. Before the Switch forwards an IP packet, it compares the source IP address, source MAC address, interface number, and VLAN ID in the IP packet with entries in the binding table. If a binding entry is matched, the Switch considers the IP packet as a valid packet and forwards the IP packet. Otherwise, the Switch considers the IP packet as an attack packet and discards the IP packet.

As shown in **Figure 5-1**, user A goes online by using DHCP. SwitchA generates a binding table based on the DHCP ACK message. The binding table includes the user source IP address, source MAC address, interface number, and VLAN ID. When SwitchA receives an IP packet from a user, it matches the IP packet against the binding table. If a binding entry is matched, SwitchA forwards the IP packet. Otherwise, SwitchA discards the IP packet. This allows IP packets sent by authorized users to pass through and bogus IP packets sent by attackers to be discarded.

Figure 5-1 IP source guard

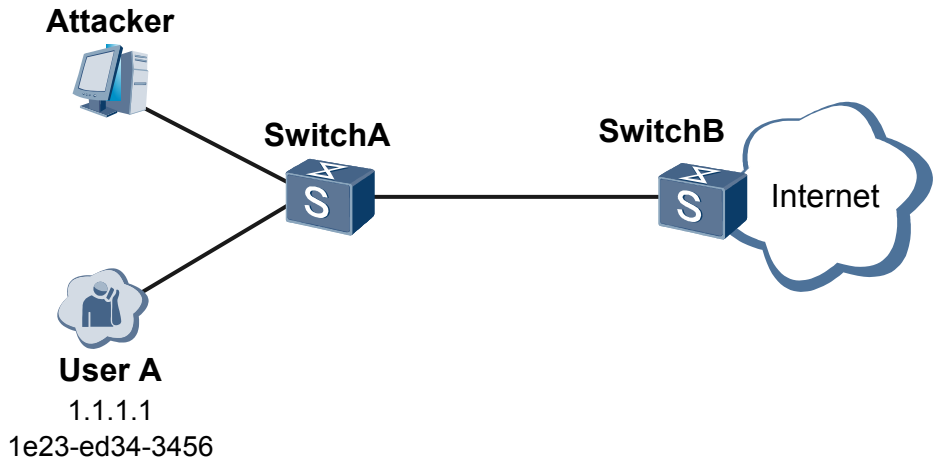

5.3.2 URPF

Before a packet is forwarded, URPF obtains the source address and inbound interface of the packet. URPF then compares the retrieved interface with the inbound interface. If they mismatch, URPF considers the source address as a spoofing address and discards the packet. This allows URPF to effectively protect a device against malicious attacks by blocking packets from bogus source addresses.

Figure 5-2 URPF

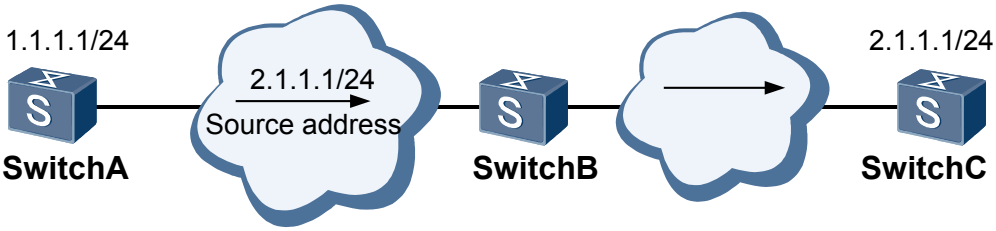

As shown in **Figure 5-2**, a bogus packet with source IP address 2.1.1.1 is sent from SwitchA to SwitchB. After receiving the bogus packet, SwitchB sends a response packet to the actual destination device SwitchC at 2.1.1.1. As a result, SwitchB and SwitchC are both attacked.

If URPF is enabled on SwitchB, when SwitchB receives the bogus packet with source IP address 2.1.1.1, URPF discards the packet because the source IP address 2.1.1.1 mismatches the incoming interface.

5.4 Applications

5.4.1 IP Source Guard

IP source guard can be configured on the S6700 that is directly connected to user devices or is used as an egress gateway. IP source guard enables the S6700 to check IP packets to defend against IP address spoofing attacks.

As shown in **Figure 5-3**, the S6700 functions as an egress gateway for an enterprise. Interface1 and Interface2 on the S6700 are added to VLAN 10 and VLAN 20. Host A belongs to VLAN 20 and host B belongs to VLAN 10. In this networking, IP packet checking against the binding table is configured on Interface1 and Interface2. The S6700 then matches the source IP address, source MAC address, and VLAN ID in IP packets with entries in the binding table. By doing this, the S6700 prevents users in a VLAN from attacking users in another VLAN. Even if host B is an attacker and communicates with the server by sending a bogus packet with the IP address of host A, the S6700 discards this bogus packet and forwards IP packets from host A.

Figure 5-3 IP source guard

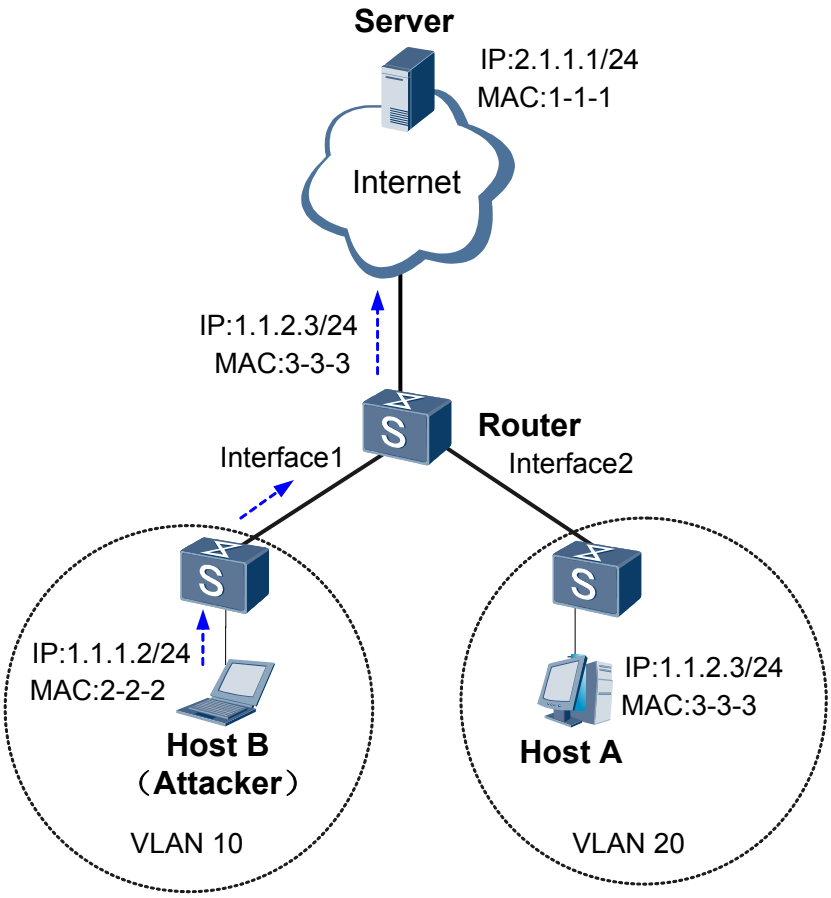

5.4.2 URPF

On a complex network, the routes recorded on the local end and remote end may be different. A URPF enabled device on this network may discard the packets transmitted along the correct path, but forward the invalid packets.

The S6700 supports the following URPF modes to solve the preceding problem:

- URPF strict check
- URPF loose check

URPF Strict Check

If route symmetry is ensured, using URPF strict check is recommended. In this mode, packets can pass the check only when the forwarding table has a corresponding routing entry and the inbound interface of the packets matches the outbound interface in the routing entry.

If there is only one path between two network edge S6700s, route symmetry is ensured. URPF strict check can be used to ensure network security.

As shown in **Figure 5-4**

Assume that PC A in AS1 generates a packet with the bogus source IP address 2.2.2.2 and sends the packet to the server in AS3. After AR_C receives this packet, it checks the inbound interface, and determines that the packet with the source address 2.2.2.2 must enter AR_C through Interface1 but not Interface2. AR_C then considers the packet as a bogus packet and discards it.

The packet sent from AS2 to the server is forwarded after passing URPF check.

Figure 5-4 URPF strict check application

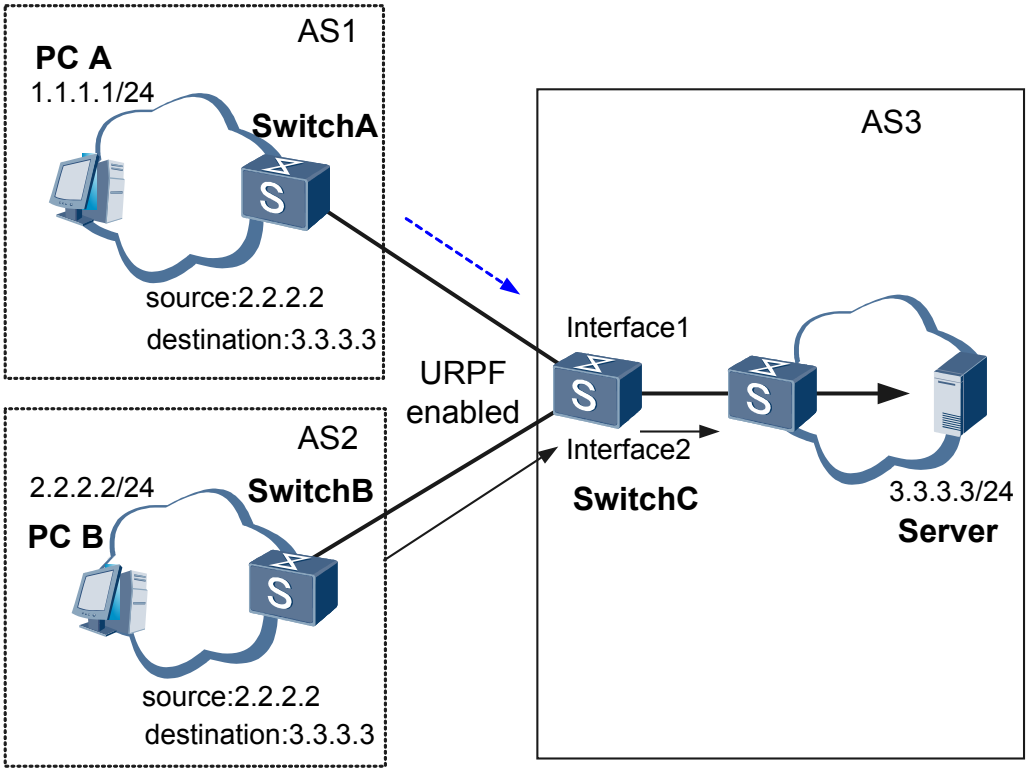

URPF Loose Check

In URPF loose check mode, packets pass URPF check and are forwarded only if there is the route of the source IP address, regardless of whether the outgoing interface of a default route is the same as the incoming interface of the packets. If there are multiple connections between two network edge Switchs, route symmetry cannot be ensured. In this case, use URPF loose check to ensure network security.

There are two situations when multiples connections are set up between two network edge Switchs: single-homed client connected to a single ISP and multi-homed client connected to several ISPs.

- Single-homed client connected to a single ISP

Figure 5-5 URPF enabled on a single-homed client connected to a single ISP

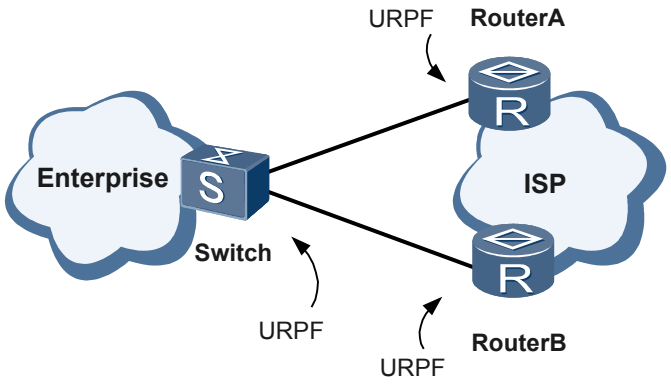

As shown in **Figure 5-5**, multiple connections are set up between Enterprise and ISP to ensure reliability. Route symmetry between Enterprise and the two ISP Switchs cannot be ensured. You must use URPF loose check.

- Multi-homed client connected to several ISPs

Figure 5-6 URPF enabled on a multi-homed client connected to several ISPs

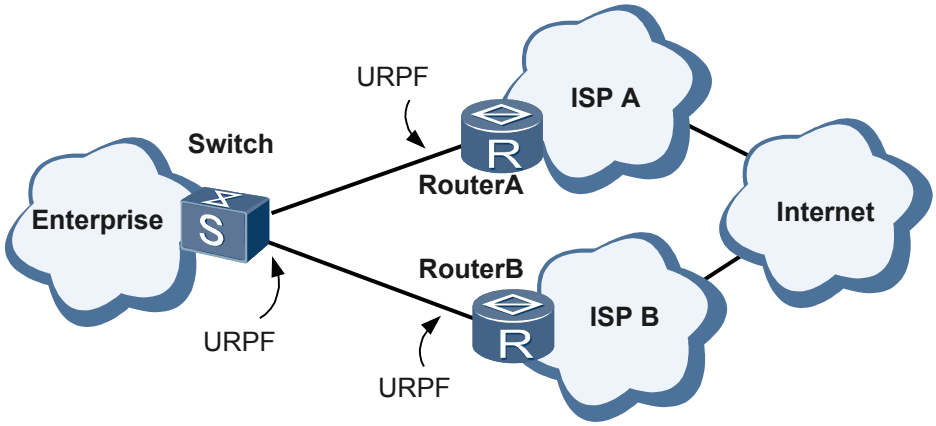

As shown in **Figure 5-6**, Enterprise is connected to several ISPs. Route symmetry between Enterprise and the two ISP Switchs cannot be ensured. You must use URPF loose check.

URPF applied in the scenario where a client is connected to multiple ISPs has the following features:

- If certain special packets are required to pass URPF check in any case, you can specify the source address in an Access Control List (ACL).
- The S6700s connected to users may only have a default route to the ISP. Therefore, matching the default routing entry needs to be supported.

6 ARP Security

About This Chapter

6.1 Introduction to ARP Security

Definition

Address Resolution Protocol (ARP) security prevents ARP protocol attacks and ARP-based network scanning attacks by using strict ARP learning, Dynamic ARP Inspection (DAI), ARP Anti-spoofing and rate limit on ARP packets.

Purpose

ARP security defends against ARP-oriented attacks.

There are a lot of ARP attack types:

- ARP attacks may aim at user hosts or the S6700.
- Attacks can be initiated by using virus or unauthorized software.
- Depending on attack impact, ARP attacks are classified into address spoofing attack and Denial of Service (DoS) attack.
  - Address spoofing attack
    - The attacker sends incorrect MAC addresses to the gateway. The gateway updates the ARP entries. As a result, user hosts cannot go online.
    - The attacker sends an incorrect ARP reply to a user host. After obtaining the incorrect gateway address, the user host cannot go online.
  - DOS attack
    - The attacker sends a lot of bogus ARP request and reply packets to the device. The ARP table of the device overflows and the device cannot cache valid ARP entries. As a result, the device cannot forward valid packets.
    - The attacker sends a lot of bogus ARP request and reply packets to the device or triggers ARP Miss packets on the device. The device will be busy processing these ARP packets, and cannot process valid service packets.

      A typical scenario where ARP Miss packets are triggered is as follows: An attack uses tools to scan the devices on the local network segment or other network segments, the S6700 searches for the corresponding ARP entries before responding to the attacker. The MAC addresses corresponding to the destination IP addresses of the packets do not exist; therefore, the ARP module of the S6700 sends ARP Miss packets to the upper-layer software, requesting the upper-layer software to send ARP request packets to obtain destination MAC addresses of the packets. If the attacker sends a lot of scanning packets, a lot of ARP Miss packets will be generated.

  ARP anti-spoofing can prevent unauthorized users; however, the ARP DoS attacks have greater impact on networks.

Benefits

- ARP security defends against ARP-oriented attacks and reduces maintenance costs.
- ARP security improves network security and stability, and protects users against ARP-oriented attacks.

6.2 References

The following table lists the references of this document.

Document	Description
RFC826	Ethernet Address Resolution Protocol
RFC903	Reverse Address Resolution Protocol
RFC1027	Using ARP to Implement Transparent Subnet Gateways
RFC1042	Standard for the Transmission of IP Datagrams over IEEE 802 Networks

6.3 Principles

You can configure different ARP security features to defend against different attacks.

Table 6-1 ARP security features against different attacks

Attack Type	ARP Security Feature
ARP flood attack	Source address-based or interface-based ARP suppression Source address-based ARP Miss packet suppression
ARP spoofing attack	Strict ARP learning ARP anti-spoofing Dynamic ARP Inspection (DAI) ARP gateway anti-collision

6.3.1 ARP Packet Suppression

If the S6700 receives a large number of ARP packets, it is busy in learning ARP entries and responding to ARP packets and cannot process other services. ARP packet suppression can be configured to protect the CPU. ARP packet suppression is classified into source IP address-based ARP packet suppression and ARP packet rate limit. ARP packet rate limit can be configured globally, on an interface or a VLAN.

If a user sends a large number of ARP packets in a short period of time, an attack occurs. When the S6700 detects the attack, it limits the rate of ARP packets from the user to protect the CPU so that other services can be processed by the CPU.

When ARP-oriented attacks occur on the S6700, a global rate limit for ARP packets can be configured to limit the rate of ARP packets sent to the SRU so that the CPU is protected and other services can be processed by the CPU.

When ARP-oriented attacks occur on an interface, a rate limit for ARP packets can be configured on the interface so that the CPU is protected and other services can be processed by the CPU. Other interfaces can still learn ARP entries.

When ARP-oriented attacks occur on an interface, VLAN-based ARP packet rate suppression can be configured so that the CPU is protected and other services can be processed by the CPU. Other interfaces in the same VLAN can still learn ARP entries.

6.3.2 ARP Miss Packet Suppression

If a host sends a large number of IP packets with unreachable destination IP addresses to attack a device, ARP Miss packets are generated. In addition, the following situations occur:

- The device sends a large number of ARP Request packets to the destination network segment, which increases the load of the destination network segment.
- The device resolves destination IP addresses continuously, which increases the burden of the CPU.

The rate limit for ARP Miss packets can be configured globally, on an interface, or a VLAN.

If a large number of ARP Miss packets are triggered in a short period of time, an attack occurs. When the S6700 detects the attack, it limits the rate of ARP Miss packets with a specified source IP address to protect the CPU so that other services can be processed by the CPU.

If a large number of ARP Miss messages are triggered in a short period of time, a rate limit for ARP Miss packets can be configured to limit the rate of ARP Miss packets sent to the SRU so that the CPU is protected and other services can be processed by the CPU.

6.3.3 Gratuitous ARP Packet Discarding

A gratuitous ARP packet is a packet in which the destination IP address is the sender IP address. Gratuitous ARP packets are used to:

- Check duplicate IP addresses: A host should not receive an ARP Reply packet after sending an ARP Request packet with the destination address being its own IP address. If an ARP Reply packet is received, another host is assigned the same IP address.
- Advertise a new MAC address. If the MAC address of a host changes because its network adapter is replaced, the host sends a gratuitous ARP packet to notify all hosts of the change before ARP entries are aged out.

Any user can send gratuitous ARP packets. If a gratuitous packet uses the IP address of another device, the ARP entry corresponding to the user is modified incorrectly. ARP spoofing attacks occur. A large number of gratuitous ARP packets may cause CPU overload and affect other services.

When the S6700 learns ARP entries by using ARP Request and Reply packets, it does not learn information in gratuitous ARP packets to prevent ARP spoofing attacks.

6.3.4 Strict ARP Learning

Strict ARP learning indicates that the S6700 learns only the ARP entries for ARP Reply packets in response to the ARP Request packets sent by itself. By doing this, the S6700 can deny the ARP Request and Reply packets sent by some attackers.

Figure 6-1 Strict ARP learning

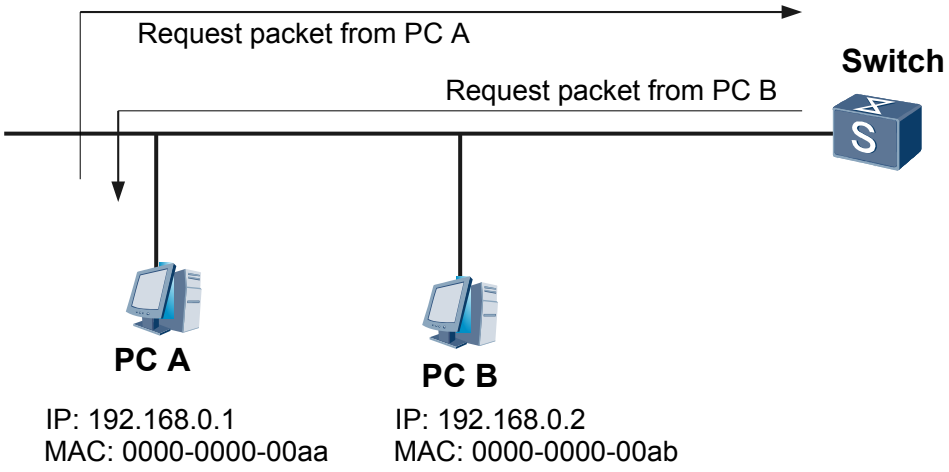

As shown in **Figure 6-1**, PC A sends an ARP Request packet to the S6700. Generally, the S6700 sends an ARP Reply packet to PC A and updates the ARP entry by adding the MAC address of PC A to the corresponding ARP entry. If strict ARP learning is configured, the S6700 sends an ARP Reply packet to PC A but does not add the MAC address of PC A to the corresponding ARP entry. When the received ARP Request packet does not match the original ARP entry, the S6700 sends an ARP Request packet to PC A again. After the S6700 receives an ARP Reply packet from PC A, it updates the ARP entry by adding the MAC address of PC A to the corresponding ARP entry.

6.3.5 ARP Anti-spoofing

When the S6700 receives the first ARP packet from a user, it adds an ARP entry to the ARP table. If the S6700 receives an ARP packet from the user again, it updates the ARP entry based on the latest ARP packet, including the MAC address, interface number, VLAN ID, and aging time. If attackers send incorrect ARP packets to update ARP entries of authorized users, ARP entries on the S6700 are incorrect and packets from authorized users cannot be forwarded correctly.

To defend against such attacks, configure the ARP Anti-spoofing to prohibit other users from updating an ARP entry or some information in the ARP entry after the S6700 learns the ARP entry for the first time.

The S6700 provides the following anti-spoofing modes:

- fixed-all: When the S6700 receives an ARP attack packet, the S6700 discards the packet, if the MAC address, interface number, and VLAN ID in the received ARP packet do not match any ARP entry in the ARP table.

- fixed-mac: When the S6700 receives an ARP attack packet, it updates the interface number and VLAN ID except for the MAC address. This mode is applied to port switching scenarios.
- send-ack: The S6700 does not modify the MAC address, VLAN ID, and interface number in a received ARP packet immediately when receiving it. Instead, the S6700 sends a unicast request packet to the user corresponding to this MAC address in the original ARP table.

6.3.6 DAI

Dynamic ARP Inspection (DAI) runs on Layer 2 interfaces and defends against ARP-oriented attacks by using DHCP snooping. When the S6700 receives an ARP packet, it compares the source IP address, source MAC address, interface number, and VLAN ID in the ARP packet with entries in the binding table. If the packet matches a binding entry, the S6700 considers the ARP packet as a valid packet and forwards it. Otherwise, the S6700 considers the ARP packet as an attack packet and discards it.

6.3.7 ARP Gateway Anti-collision

When the S6700 functions as a gateway, if an attacker sends ARP packets with the source IP address as the gateway address, ARP entries on all the downstream devices connected to the S6700 are modified incorrectly. As a result, all the traffic is sent to the attacker and the attacker intercepts user information.

After ARP gateway anti-collision is enabled, the S6700 generates ARP anti-collision entries and discards the packets with the same source MAC address in the Ethernet header in a period of time. This can prevent ARP packets with the bogus gateway address from being broadcast in a VLAN. The S6700 can send gratuitous ARP packets. The function of sending a correct gratuitous ARP packet can be enabled. The gratuitous ARP packet is broadcast to all the users so that incorrect ARP entries are corrected and user security is ensured.

6.4 Applications

6.4.1 ARP Anti-spoofing

On an intranet, the S6700 functions as a user gateway. After an authorized user goes online, the S6700 generates an ARP entry mapping the authorized user. If another user sends ARP spoofing packets, the ARP entry mapping the authorized user is modified on the S6700. As a result, packets from the authorized user cannot be forwarded.

To solve the problem, configure ARP anti-spoofing on the S6700 to prevent unauthorized users from modifying ARP entries so that authorized user can access the Internet.

Figure 6-2 ARP anti-spoofing

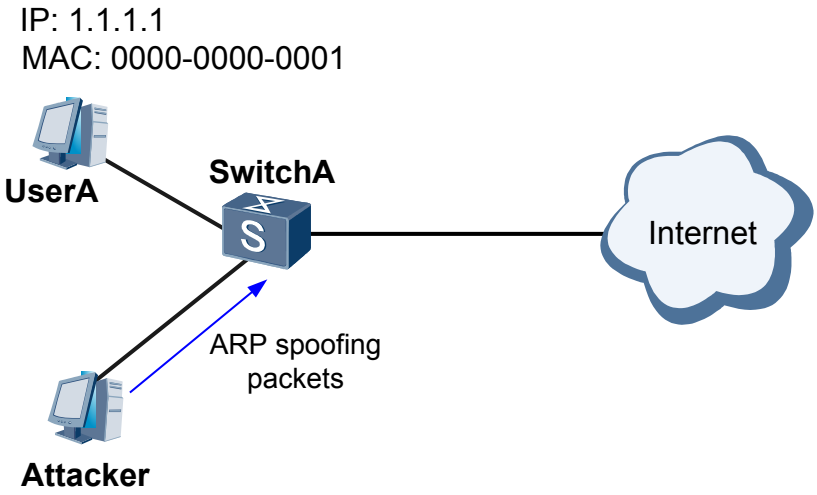

As shown in **Figure 6-2**, after a user goes online, an ARP entry mapping the user is generated on the S6700 and the user can access the Internet. If attackers send ARP spoofing packets, the ARP entry is not modified because ARP anti-spoofing is configured.

6.4.2 DAI

When the S6700 functions as a Layer 2 device, Dynamic ARP Inspection (DAI) can be configured on the S6700 to check validity of ARP packets. The S6700 forwards valid packets and discards invalid packets to protect users or gateways against ARP spoofing attacks.

Figure 6-3 DAI

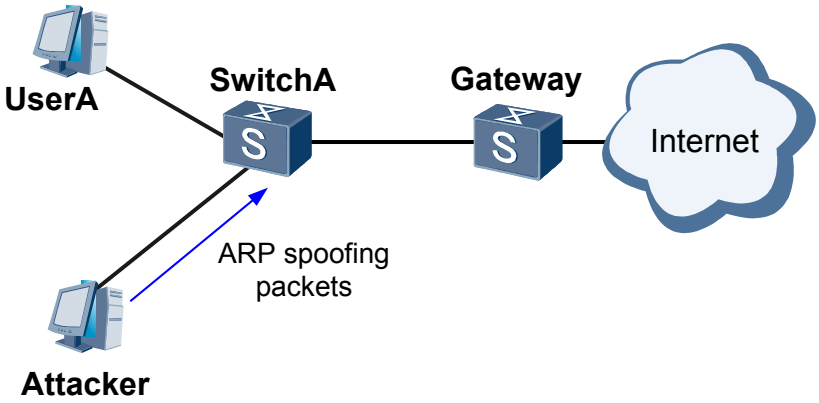

As shown in **Figure 6-3**, the S6700 functions as a Layer 2 device and users go online by using the Dynamic Host Configuration Protocol (DHCP). After users go online, the S6700 generates a binding table based on the source IP address, source MAC address, interface number, and VLAN ID. When a user sends an ARP packet, the S6700 matches information in the ARP packet against the binding table. If the ARP packet matches a binding entry, the S6700 forwards it. Otherwise, the S6700 discards it. ARP packets sent from authorized users can pass through

because there are matching entries, whereas bogus ARP packets sent from attackers are discarded because they match no binding entry.

7 DHCP Snooping

About This Chapter

7.1 Introduction to DHCP Snooping

This section describes the functions of DHCP snooping.

Definition

DHCP snooping is a security feature of DHCP. The S6700 creates and maintains the DHCP snooping binding table to filter out untrusted DHCP information that is sent from untrusted zones. The DHCP snooping binding table contains the MAC address, IP address, lease, VLAN ID, interface number of each user in an untrusted zone.

When DHCP snooping is enabled on an S6700, the S6700 listens on DHCP packets and records the IP addresses and MAC addresses in the received DHCP Request packets or Ack messages. A physical interface can be configured as a trusted interface or an untrusted interface. A trusted interface can forward received DHCP Reply packets, whereas an untrusted interface discards the received DHCP relay packets. By using DHCP snooping, the S6700 can prevent bogus DHCP servers and ensure that clients obtain IP addresses from valid DHCP servers.

Purpose

DHCP snooping prevents the following attacks:

- Bogus DHCP server attack
- Man-in-the-middle attack and IP/MAC spoofing attack
- Denial of Service (DoS) attack
- DoS attack by changing the value of the Client Hardware Address (CHADDR)

Benefits

DHCP snooping ensures that:

- Clients obtain IP addresses from valid DHCP servers.
- The IP addresses and MAC addresses of DHCP clients are recorded, and the binding entries can be used by other features.

7.2 References

For more information about DHCP, refer to the following documents.

Document No.	Description
RFC 3046	DHCP Relay Agent Information Option
RFC 2132	DHCP Options and BOOTP Vendor Extensions

7.3 Principles

7.3.1 Concepts

Option 82 Field

- Format of DHCP Messages with the Option 82 Field

  The Options field in a DHCP message is used to carry other control information and parameters that are not defined in certain protocols. **Figure 7-1** shows the format of the Options field. Option 82 refers to the option field whose code is 82. The S6700 identifies the path for sending DHCP Request messages by setting the Option 82 field.

  Figure 7-1 Format of the Option 82 field

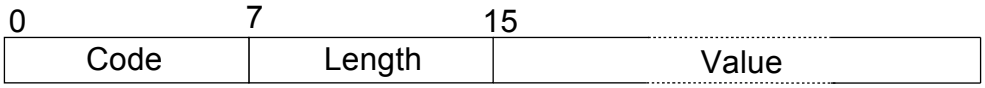

  The Options field consists of the Code, Length, and Value. The description of each field is as follows:

Field	Length	Description
Code	1 byte	Indicates the subsequent contents.
Length	1 byte	Indicates the length of the subsequent contents.
Value	Its value is specified by the Length field.	Indicates the information contents.

  When a DHCP Request message sent from a DHCP client passes through the S6700, the S6700 adds the Option 82 field to this Request message. On receiving the DHCP Request message with the Option 82 field, the DHCP server returns a DHCP Reply message containing the same Option 82 field to the S6700. The S6700 can then determine the interface to which the DHCP Reply message is sent based on the Option 82 field.

  As shown in **Figure 7-2**, the Code field in Option 82 is 82; the Length field indicates the total number of bytes in the Agent Information field; the iN field indicates sub-options of the Agent Information field. Each sub-option is a SubOpt/Length/Value tuple. In **Figure 7-3**, the SubOpt field indicates the sub-option number and the Length field identifies the number of bytes only in the sub-option value field. In the Option 82 field, at least one sub-option must be defined and the sub-option can be defined as null, and the minimum length of the Option 82 field is 2.

  The initially assigned sub-options are as follows:

  - 1: agent circuit ID sub-option
  - 2: agent remote ID sub-option

  A DHCP server uses the agent circuit ID sub-option for IP address and other parameter assignment policies.

In addition to sub-option 1, the S6700 also supports sub-option 9 for showing information on the circuit ID added by Huawei devices.

Functions of sub-option 9 are as follows:

- If the Option 82 field in a DHCP Reply message forwarded by an interface contains Sub-option 9 and this option contains the Huawei Device Identifier field, Huawei device can parse the Option 82 field and obtain interface information successfully. It then removes the Huawei Device Identifier field from sub-option 9 before forwarding the Reply message.
- When the S6700 receives a Reply message, the S6700 determines whether the Option 82 field contains sub-option 9 if the Reply message carries the Option 82 field. If the Option 82 field contains sub-option 9, a binding table based on sub-option 9 is generated. If the Option 82 field does not contain sub-option 9, a binding table based on sub-option 1 is generated.

Figure 7-2 Format of DHCP messages with the Option 82 field

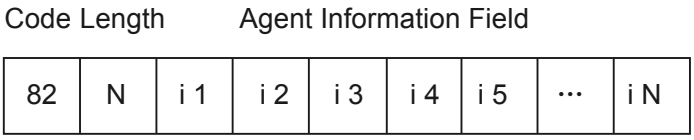

Figure 7-3 Suboptions of the Option 82 field

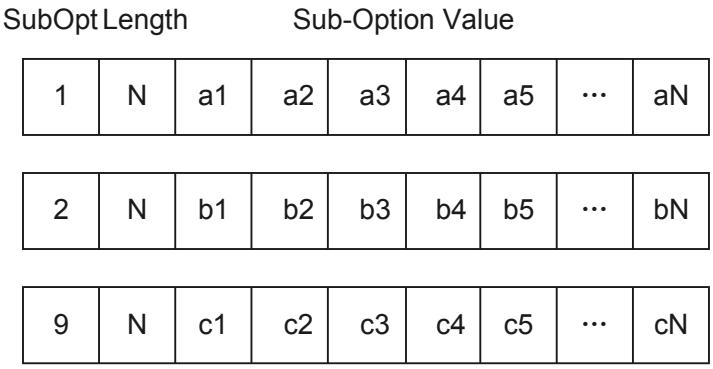

The Option 82 field can be used on the S6700 at Layer 2 or Layer 3. When the Option 82 field is used on the S6700 at Layer 3, the DHCP server performs IP address assignment policies or other policies. When the Option 82 field is used on the S6700 at Layer 2, the S6700 can determine the interface to which the DHCP Reply messages are sent and generate binding entries of IP addresses and MAC addresses by analyzing the Option 82 field.

- Option 82 Field Appended by the S6700 at Layer 2

  As shown in **Figure 7-4**, the client is connected to the S6700 and then the DHCP relay agent or the DHCP server through a Layer 2 network.

  To listen to DHCP messages on the S6700, enable DHCP snooping globally on the S6700 and append the Option 82 field to DHCP Discover messages. Then, the DHCP server adds the Option 82 field to DHCP Offer messages. The S6700 determines the interface to which DHCP Offer messages are sent by analyzing the Option 82 field and generates DHCP

snooping binding entries of IP addresses and MAC addresses. The S6700 removes the Option 82 field from DHCP Offer messages before forwarding them to the client.

NOTE

To append the Option 82 field to DHCP messages on the S6700 at Layer 2, you need to first enable DHCP snooping on the S6700.

Figure 7-4 Option 82 field appended by the S6700 at Layer 2

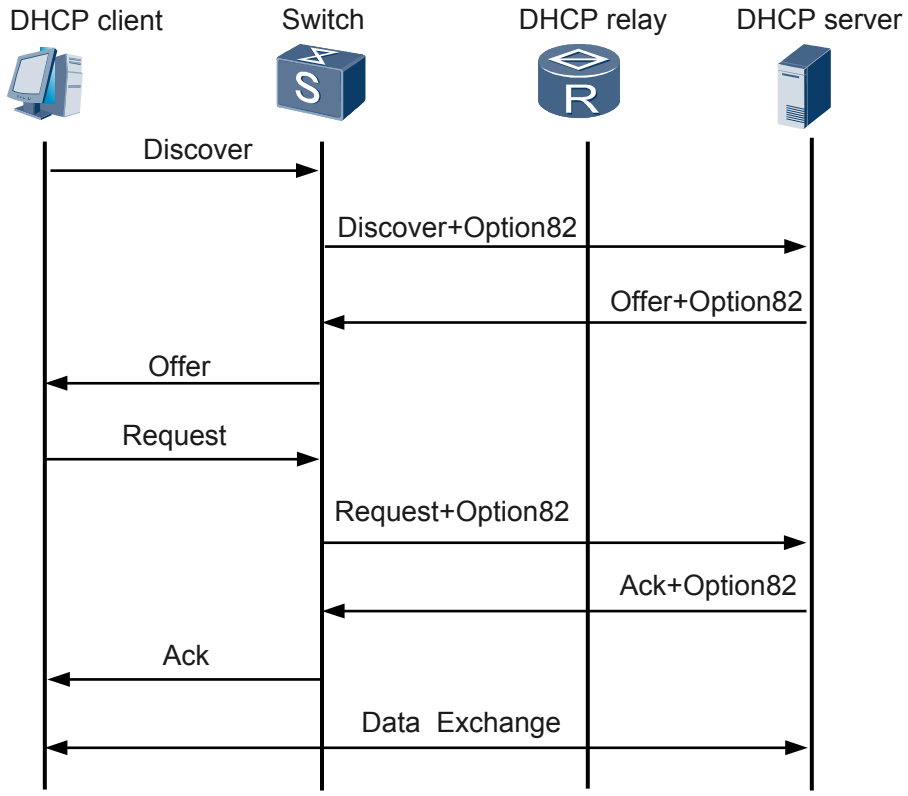

- Option 82 Field Appended by the S6700 at Layer 3

  With regard to the Layer 3 mode, the S6700 functions as the DHCP relay agent.

  As shown in **Figure 7-5**, after Option 82 is enabled on the S6700 that functions as the DHCP relay agent, the S6700 appends the Option 82 field to the DHCP Discover message and DHCP Request message. The DHCP server then performs IP address assignment policies and other policies based on the Option 82 field.

  The DHCP Reply messages returned by the DHCP server also carry the Option 82 field. Upon receiving the DHCP Reply messages, the S6700 removes the Option 82 field before forwarding them to the client.

Figure 7-5 Option 82 field appended by the S6700 at Layer 3

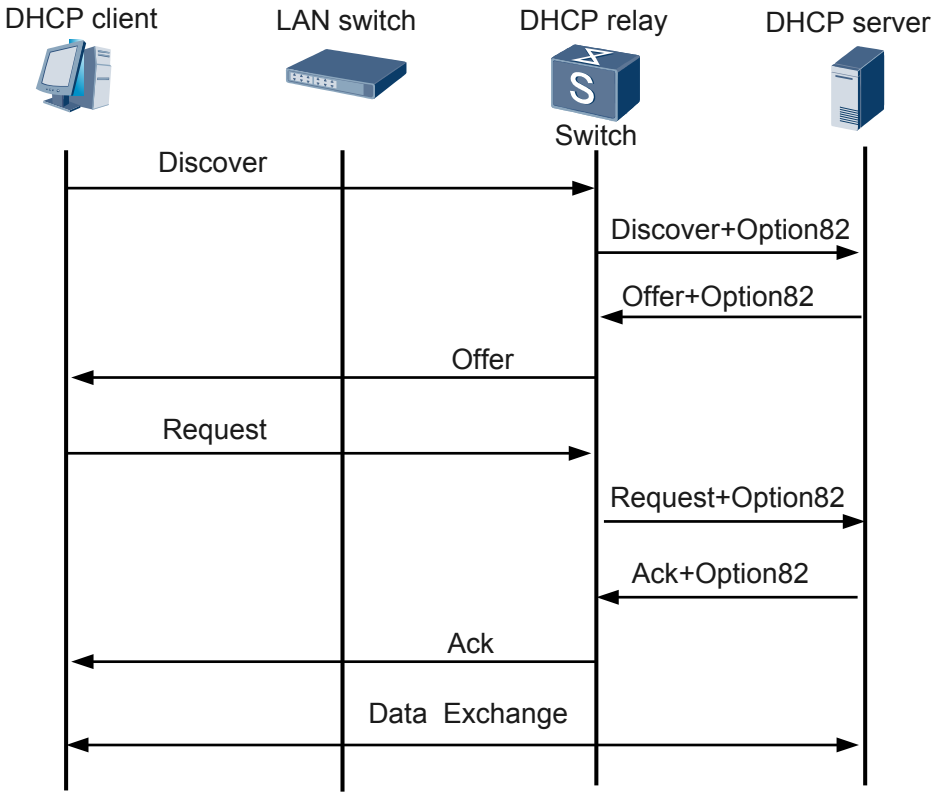

- Implementation of Option 82

  After the Option 82 function is enabled, the S6700 checks whether the DHCP Request message sent by a client contains the Option 82 field.

  - If the DHCP Request message contains the Option 82 field:

    The S6700 checks the appending of the Option 82 field, including two modes: **Insert** and **Rebuild**

    - If the current interface is configured with the **Rebuild** mode, it indicates that this interface does not trust the Option 82 field contained in the received message and must modify sub-option 1 contained in the Option 82 field.
    - If the current interface is configured with the **Insert** mode, it indicates that this interface trusts the Option 82 field contained in the received message and does not need to modify sub-option 1 contained in the Option 82 field. The S6700 then must check whether the Option 82 field contains sub-option 9. If not, the interface adds sub-option 9. If the Option 82 field contains sub-option 9, the S6700 checks whether this option contains the Huawei Device Identifier field. If not, the S6700 adds the Huawei Device Identifier field following the other manufacturer information field.

  - If the DHCP Request message does not contain the Option 82 field:

    The Huawei device adds the Option 82 field with sub-option 1 regardless of whether the Option 82 field is appended in **Insert** or **Rebuild** mode.

  During the forwarding of the DHCP Reply message, the S6700 first checks whether the Reply message contains sub-option 1 or sub-option 9 and the sub-option contains the Huawei Device Identifier field. If so, the S6700 can successfully parse the Option 82 field.

The S6700 then removes the Huawei Device Identifier field from sub-option 1 or sub-option 9, and then forwards the Reply message.

DHCP Snooping Binding Table

The DHCP snooping binding table contains dynamic binding entries and static binding entries.

- Dynamic DHCP snooping binding entries

  As shown in **Figure 7-6**, the S6700 generates DHCP snooping binding entries according to the DHCP ACK packets received on trusted interfaces.

  Figure 7-6 Generating dynamic DHCP snooping binding entries

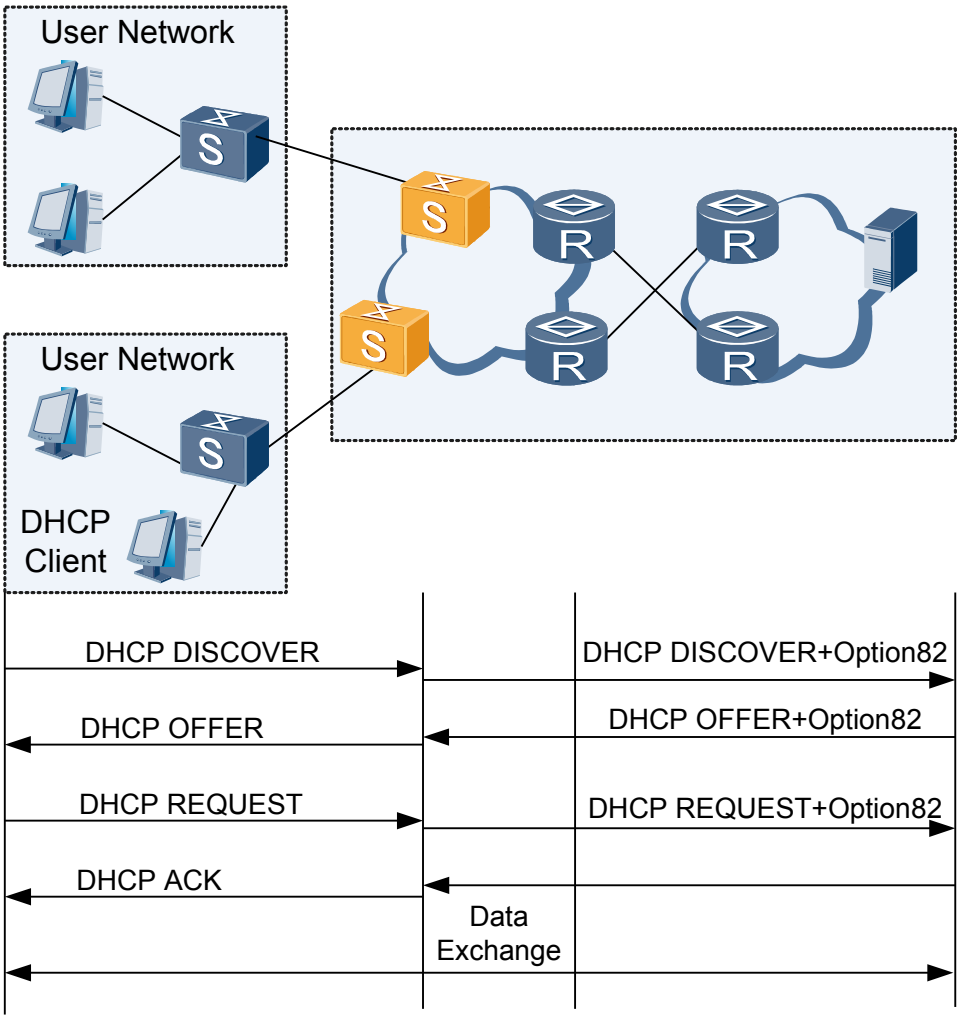

- Static DHCP snooping binding entries

  If a user accesses the network by using a static IP address, the S6700 discards the packets of the user because the IP address does not match any entry in the DHCP snooping binding table on the S6700. You can configure a static DHCP snooping binding entry by using commands. When configuring a static entry, you must learn about the IP address, MAC address, VLAN ID, and interface number of the user.

7.3.2 Bogus DHCP Server Attack

Principle of the Attack

DHCP Discover messages are sent in broadcast mode. The bogus DHCP server can thus listen to the Discover messages. The bogus DHCP server then replies incorrect messages such as the incorrect IP address of the gateway, incorrect DNS server, and incorrect IP address to the DHCP client. This causes the DoS.

Figure 7-7 DHCP client sending DHCP Discover messages

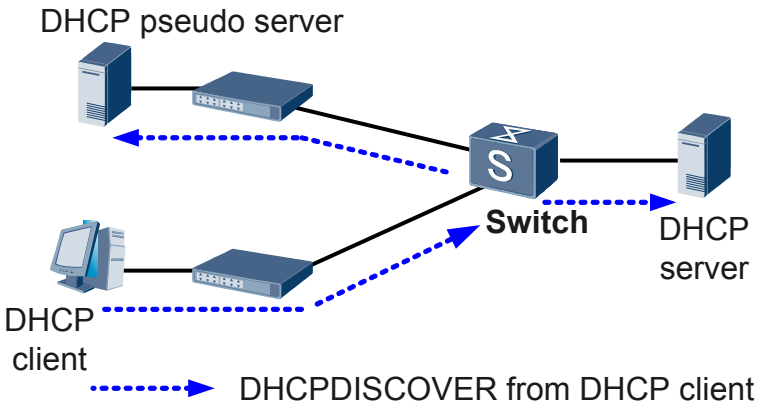

Figure 7-8 Bogus DHCP server attack

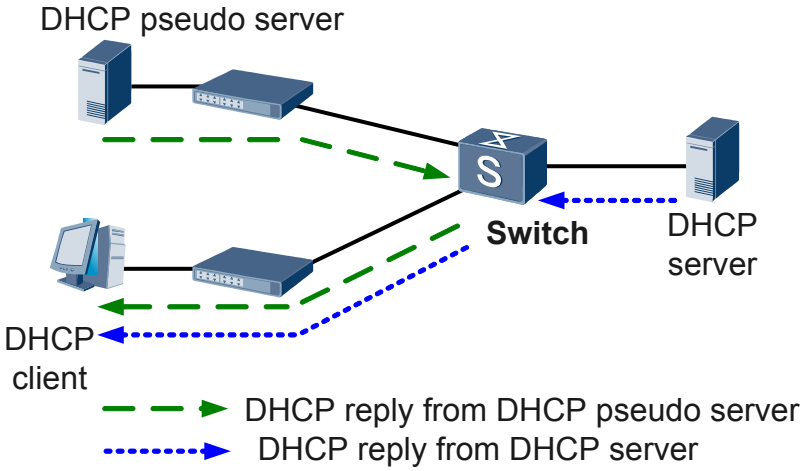

Solutions

To prevent the bogus DHCP server attack, use the trusted/untrusted operation mode of DHCP snooping.

You can configure a physical interface to be trusted or untrusted. DHCP Reply messages, including Offer, ACK, and NAK messages, received from an untrusted interface are directly discarded so that the bogus DHCP server attack can be prevented. See **Figure 7-9**.

Figure 7-9 Trusted/Untrusted operation mode of DHCP snooping

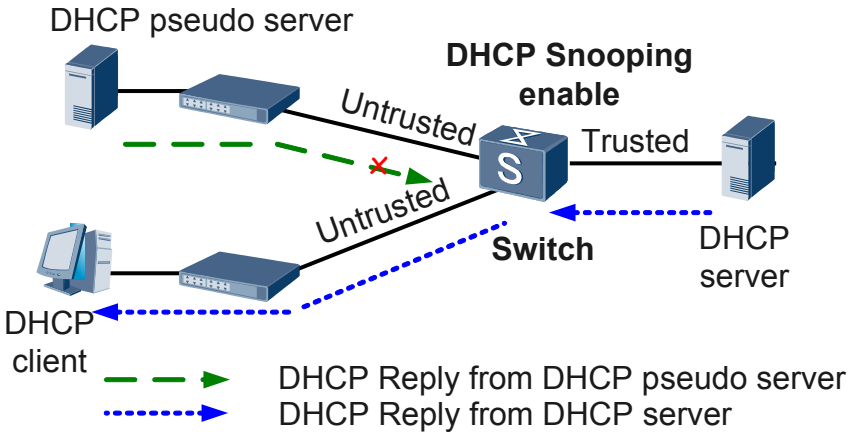

7.3.3 Middleman Attack and IP/MAC Spoofing Attack

Middleman Attack

A middleman sends a packet carrying its own MAC address and the IP address of the DHCP server to the client. The client then learns the IP and MAC addresses of the middleman and considers the middleman as the DHCP server. The packets sent from the client to the DHCP server pass the middleman. The middleman then sends a packet carrying its own MAC address and the IP address of the client. The DHCP server can learn the IP and MAC address of the middleman and consider the middleman as the client. The packets sent from the DHCP server to the client pass the middleman, as shown in **Figure 7-10**.

The middleman can implement data exchange between the DHCP server and the client. The DHCP server considers that all the packets are sent to or received from the DHCP client. In the same manner, the DHCP client considers that all the packets are sent to or received from the DHCP server. In fact, all the packets that are exchanged between the DHCP server and the client are bogus packets processed by the middleman.

Figure 7-10 Middleman attack

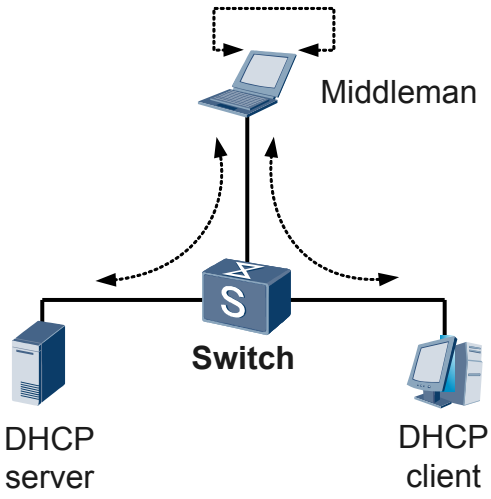

IP/MAC Spoofing Attack

The attacker sends a packet carrying valid IP and MAC addresses of a client to the DHCP server. The DHCP server misidentifies the attacker as a valid client and learns the IP and MAC addresses. The actual valid client, however, cannot access the service provided by the DHCP server, as shown in **Figure 7-11**.

Figure 7-11 IP/MAC spoofing attack

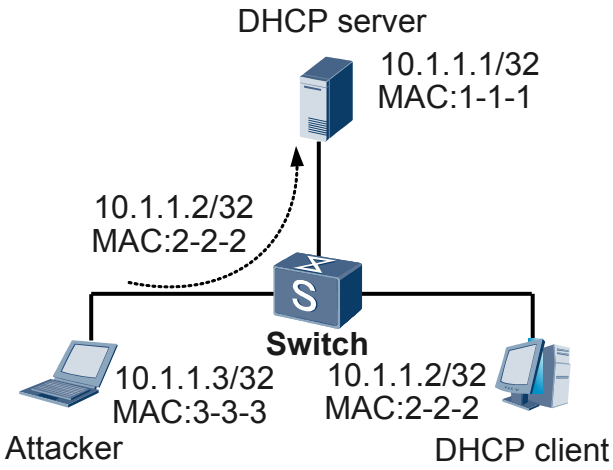

Solutions

To prevent the middleman attack and IP/MAC spoofing attack, use the DHCP snooping binding table.

By default, the Discard policy is used on the S6700. After receiving an ARP packet, an interface matches the source IP address and source MAC address of the ARP packet head with the entries in the DHCP snooping binding table. The packet is forwarded if they match and is discarded if they do not match, as shown in **Figure 7-12**.

- For the clients configured with static IP addresses, ARP packets or IP packets sent by them are discarded. This is because these clients do not obtain IP addresses by sending DHCP Request messages and no DHCP snooping binding entry exists. In this manner, these clients are prevented from accessing the network illegally. To allow the users with statically allocated IP addresses to access the network, you must configure a static DHCP snooping binding table.
- Similarly, for the clients who embezzle valid IP addresses of other clients, the packets from these clients are discarded because they do not obtain IP addresses by sending DHCP Request messages and hence the MAC address and interface information in the DHCP snooping binding table corresponding to the IP address are different from those of the embezzler. In this way, these clients are prevented from accessing the network illegally.

Figure 7-12 Binding table of IP addresses and MAC addresses

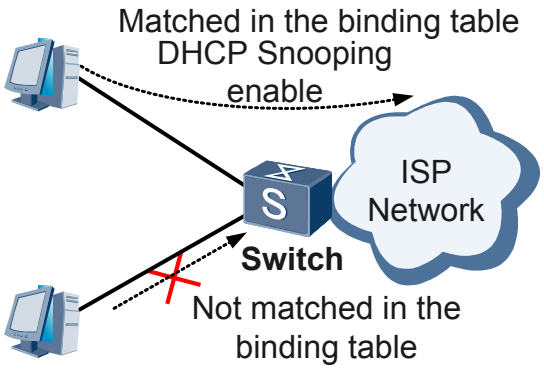

The entries in the DHCP snooping binding table are classified into the following types:

- Static entries configured through command lines. These entries can be deleted only through command lines.
- Dynamic entries automatically learnt through DHCP snooping. These entries are aged according to the lease.

The dynamic entries in the DHCP snooping binding table are automatically generated according to DHCP ACK messages from the DHCP server. The procedure for generating dynamic entries is different according to the layer where the S6700 is located:

- Layer 2

  If Option 82 is enabled, the S6700 can intercept the DHCP message and append the Option 82 field in the DHCP Request message. The DHCP server then returns the DHCP Reply message carrying the Option 82 field. The S6700 determines the interface to which the DHCP Reply message is sent by analyzing the Option 82 field and generates the DHCP snooping binding entries.

  If Option 82 is disabled, the S6700 identifies interface information according to the MAC address table.

- Layer 3

  For the untrusted interface, the S6700 obtains the IP address of the interface assigned by the DHCP server, the MAC address of the interface, and the interface that the messages pass by monitoring the DHCP Reply message. An IP and MAC binding entry of the untrusted interface is then generated. The binding entry has the same lease as the IP address of the client. When the lease expires or the client releases this IP address, the entry is automatically deleted.

If a user is disconnected abnormally after obtaining an IP address, the user cannot send DHCP Release messages to release the IP address. In this case, you can enable the association function between ARP and DHCP snooping. The system performs ARP detection for the IP addresses whose DHCP snooping entries expire and are not contained in ARP entries. If no user is detected within the specified number of detection times, the system deletes the binding relation in the DHCP binding table and notifies the DHCP server of releasing the IP address.

NOTE

The association function between ARP and DHCP snooping is used only when the S6700 functions as the DHCP relay agent.

7.3.4 DoS Attack by Changing the Value of the CHADDR Field

The attacker may change the CHADDR field carried in DHCP messages instead of the source MAC address in the frame header to apply for IP addresses continuously. The S6700, however, only checks the validity of packets based on the source MAC address in the frame header. The attack packets can still be forwarded normally. The MAC address limit cannot take effect in this manner.

NOTE

CHADDR is short for Client Hardware Address.

Figure 7-13 DoS attack by changing the value of the CHADDR field

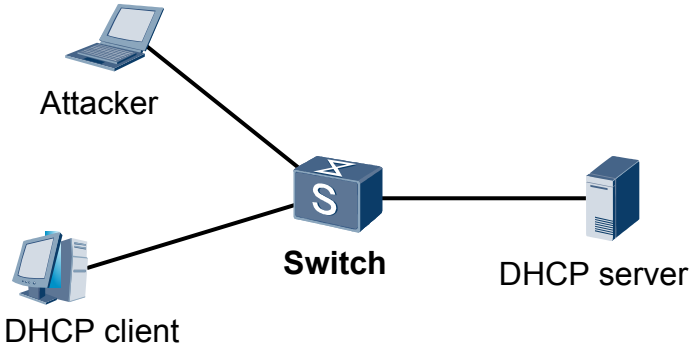

Figure 7-14 Format of DHCP messages

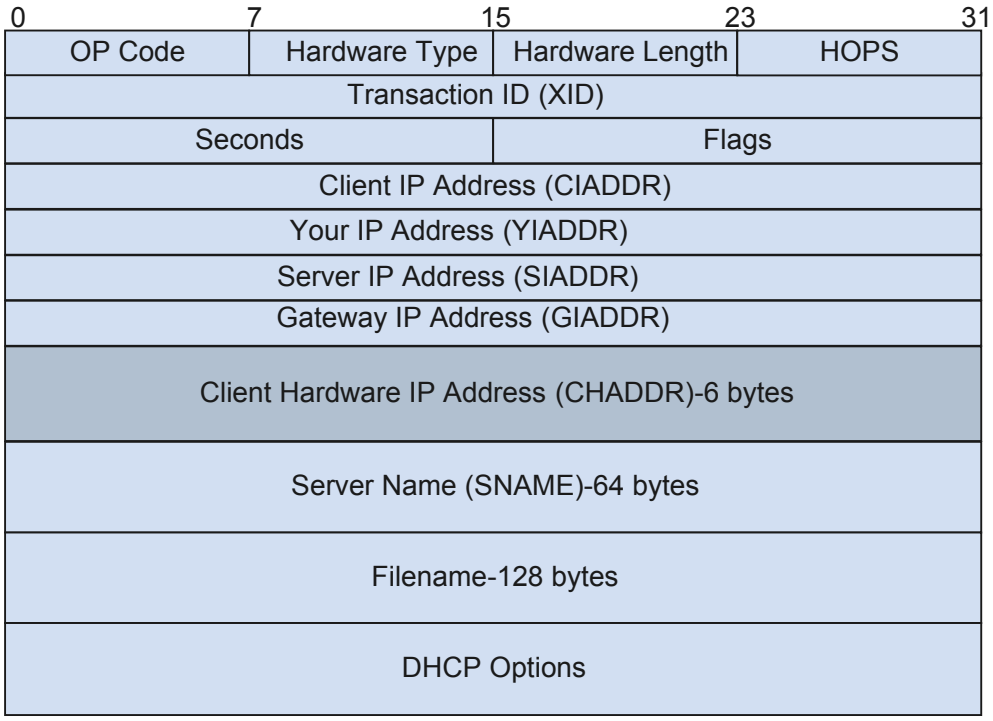

NOTE

For details on the format of DHCP messages, refer to the *Quidway S6700 Series Ethernet Switches Feature Description - IP Services*.

You can configure DHCP snooping on the S6700 to check the CHADDR field carried in a DHCP Request message. If the CHADDR field matches the source MAC address in the frame header, the message is forwarded. Otherwise, the message is discarded.

7.3.5 Man-in-the-Middle Attacks

The following attacks may occur on a network where DHCP is used:

- An attacker sends an ARP packet carrying its own MAC address and DHCP server IP address to the DHCP client. The DHCP client then learns the entry containing the DHCP server IP address and the attacker's MAC address, and considers the attacker as the DHCP server.
- An attacker sends an ARP packet carrying its own MAC address and DHCP client IP address to the DHCP server. The DHCP server then learns the IP address and the attacker's MAC address and considers the attacker as the DHCP client.

The attacker can learn information exchanged between the DHCP server and the DHCP client.

Such an attacker is called man-in-the-middle, as shown in **Figure 7-15**.

Figure 7-15 Man-in-the-middle attacks

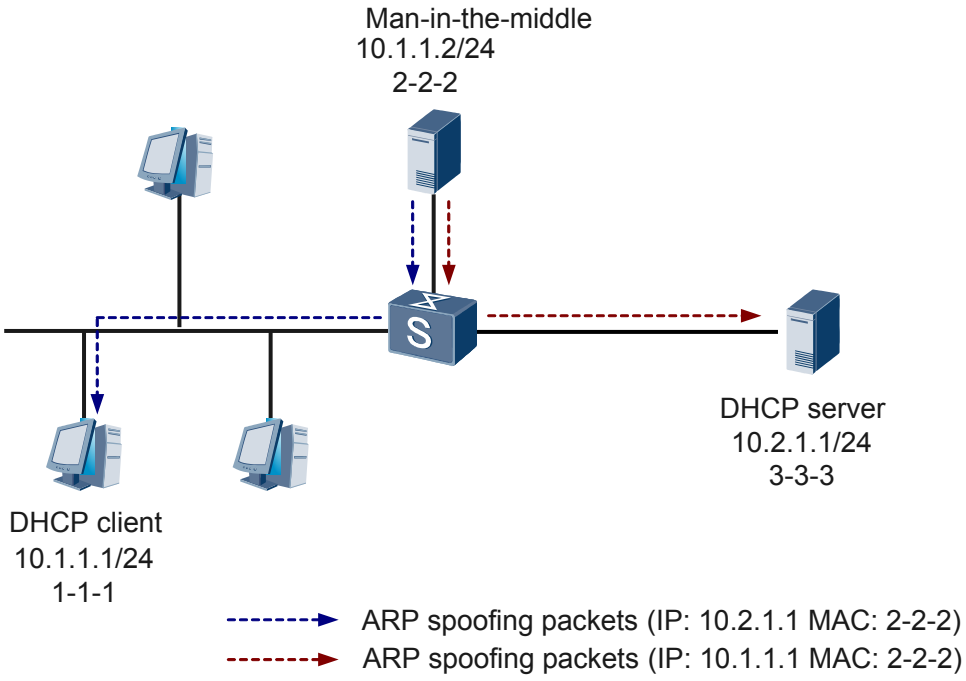

To defend against man-in-the-middle attacks, enable DHCP snooping and configure defense against ARP man-in-the-middle attacks.

After DHCP snooping is enabled, the S6700 creates and maintains a DHCP snooping binding table. The binding table contains the IP address, MAC address, VLAN ID, and interface

information. After defense against ARP man-in-the-middle attacks is configured, only the received ARP packets that match entries in the binding table are forwarded. Otherwise, packets are discarded. Attack packets are discarded because the IP address and MAC address in the ARP packet sent by the man-in-the-middle are different from the entry in the binding table, as shown in **Figure 7-16**.

Figure 7-16 Defense against man-in-the-middle attacks

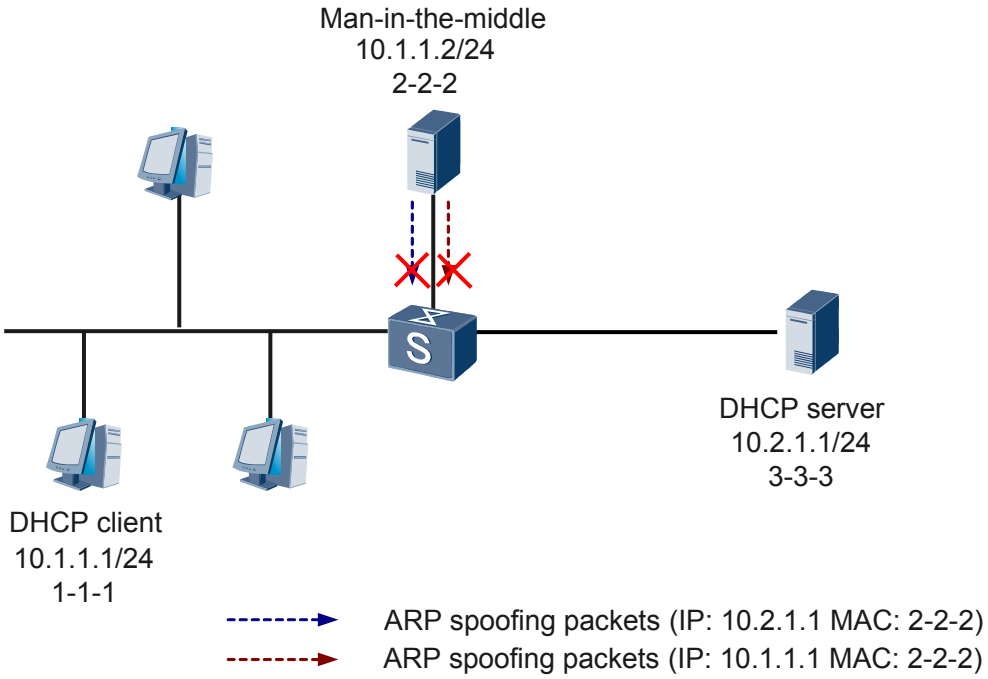

7.4 Applications

DHCP snooping can be applied on the S6700 at Layer 2 or Layer 3 as shown in **Figure 7-17** and **Figure 7-18**.

NOTE

The S6700 provides functions of Layer 2 and Layer 3 devices. The Layer 2 device mentioned later refers to the S6700 at Layer 2 and the Layer 3 device mentioned later refers to the S6700 at Layer 3.

Figure 7-17 Using DHCP snooping on the S6700 at Layer 2

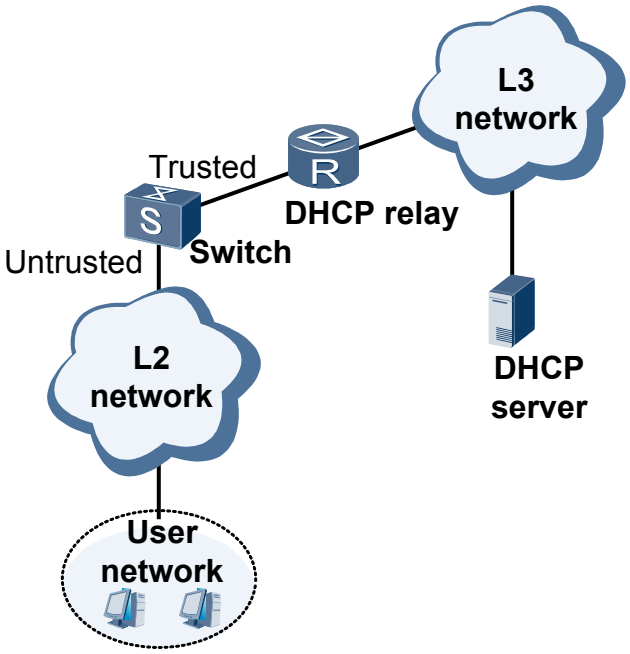

Figure 7-18 Using DHCP snooping on the S6700 at Layer 3

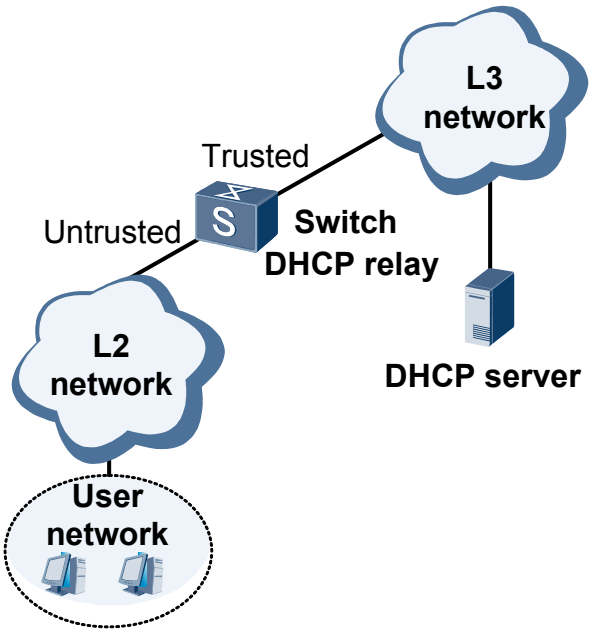

DHCP functions can be configured at both Layer 2 and Layer 3 to prevent DHCP attacks:

- Configure a static binding table so that users using static IP addresses can access the Internet.
- Configure the trusted interface to defend against attacks from bogus DHCP servers.

- Configure SwitchA to check the CHADDR field in DHCP Request messages to defend against attacks by sending bogus DHCP Request messages to extend IP address leases.
- Configure IP source guard to defend against IP/MAC spoofing attacks.
- Configure ARP packet checking to defend against ARP man-in-the-middle attacks.
- Configure SwitchA to check the MAC address in DHCP Request messages to defend against DoS attacks by sending DHCP Request messages with variable CHADDR values.

8 ND Snooping

About This Chapter

8.1 Introduction of DHCP Snooping

Definition

Neighbor discovery (ND) is a group of messages and processes that identify relationships between neighboring nodes in IPv6 network. IPv6 ND corresponds to a combination of the Address Resolution Protocol (ARP), ICMP router discovery, and ICMP Redirect of IPv4. ND snooping provides the following functions:

- Detecting address conflicts
- Determining neighbor reachability
- Resolving the neighboring node address
- Configuring the host address

Purpose

ND snooping prevents the attacks of the bogus ND server.

You can configure ND snooping on the S6700, configure the network-side interface as the trusted interface, and configure user-side interfaces as untrusted interfaces. The RA messages received from untrusted interfaces are forwarded directly.

Benefits

ND snooping is a security feature of the S6700. It ensures that the S6700 obtains interaction messages from an authorized ND server.

8.2 References

None.

8.3 Principles

8.3.1 ND Snooping Message Types

IPv6 ND provides the following types of ICMPv6 messages:

- Router Solicitation (RS): After startup, a host sends an RS message to a device, and waits for the device to respond with a Router Advertisement (RA) message.
- Router Advertisement (RA): A device periodically advertises RA messages that contain prefixes and flag bits.
- Neighbor Solicitation (NS): Through NS messages, an IPv6 node obtains the link-layer address of its neighbor, checks whether the neighbor is reachable, and performs duplicate address detection.

- Neighbor Advertisement (NA): After receiving an NS message, an IPv6 node responds with an NA message. In addition, the IPv6 node actively sends NA messages when the link layer changes.
- Redirect: When finding that the inbound interface and outbound interface of a packet are the same, a device can send Redirect messages to instruct the host that sends the packet to choose a better next hop.

8.3.2 Implementation

The ND snooping technology is a security feature of ND. By capturing and analyzing the preceding types of messages, it establishes and maintains the prefix management table and ND dynamic binding table. The prefix management table contains information about the prefix and the prefix lease. The ND dynamic binding table contains information about IPv6 addresses, MAC addresses, interfaces, and VLAN IDs.

By maintaining the prefix management table and ND dynamic binding table, the device enabled with ND snooping allows authorized users to access the network and prevents unauthorized users from attacking network devices and authorized users.

8.4 Applications

When being deployed on a Layer 2 network, the S6700 is located between the ND server (usually a router) and the user network. To prevent unauthorized users from forging the ND server, you can configure interfaces as trusted or untrusted interfaces on the S6700.

By maintaining the prefix management table and ND dynamic binding table, the S6700 enabled with ND snooping allows authorized users to access the network and prevents unauthorized users from attacking network devices and authorized users.

Figure 8-1 shows ND snooping applied to the S6700.

Figure 8-1 ND snooping enabled on the S6700 of the Layer 2 network

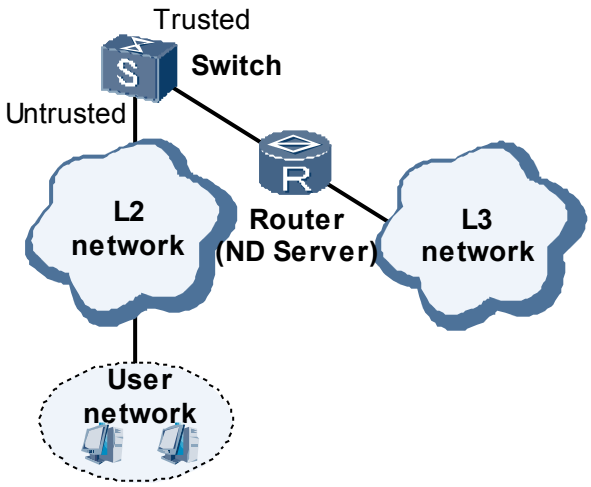

9 Local Attack Defense

About This Chapter

9.1 Introduction to Local Attack Defense

Definition

The S6700 supports CPU attack defense and attack source tracing.

When an attack occurs, the local attack defense function ensures non-stop service transmission.

CPU attack defense has the following features:

- Rate limit(CAR)

  CAR features include blacklist, CPCAR, deny and rate limit functions. A blacklist is a set of unauthorized users. The S6700 adds users with the specific characteristic into the blacklist by using Access Control List (ACL) rules and discards the packets on CPU.. CPCAR limits the rate of protocol packets and the deny function directly discards the packets of a certain protocol. The rate limit function uniformly limits the rate of packets sent to the CPU to protect the CPU.

- Application layer protocol association

  The ALP function on the S6700 protects session-based application layer data, such as data of File Transfer Protocol (FTP) sessions, Border Gateway Protocol (BGP) sessions and open shortest path first (OSPF) sessions. ALP ensures non-stop service transmission when attacks occur.

  When the S6700 detects setup of an FTP session, a BGP session or a OSPF session, ALP is enabled to protect the session. The packets matching characteristics of the session are sent at a high rate. Reliability and stability of session-related services are ensured.

Attack source tracing

The attack source tracing function analyzes packets sent to the CPU and collects statistics on the packets, and automatically identifies attack sources. Then an alarm is reported to the administrator to suppress packets from the attack sources and ensure network security.

The S6700 supports the whitelist for attack source tracing and discards the packets sent from the attack source.

Purpose

As the Internet technology and network scale develop quickly and various network applications emerge, enterprises are using more network applications. How to protect confidential data and resources in an open network environment is becoming a concern.

Therefore, protecting the CPU is a necessary and important factor for processing services and system responses. A large number of packets including valid packets and malicious attack packets on a network must be processed by devices' CPUs. The malicious attack packets affect services and may even cause a system breakdown. In addition, excess valid packets can also lead to high CPU usage, CPU performance deterioration, and service interruption.

The CPU attack defense and source tracing functions protect the S6700 against attacks. When an attack occurs, these functions ensure service transmission and minimize the impact of the attack on network services.

Benefits

CPU attack defense and attack source tracing ensure network security and stable running of devices.

9.2 References

None.

9.3 Principles

9.3.1 Working Mechanism of CPU Attack Defense

Token Bucket

The token bucket technology limits the rate of incoming packets on an interface by limiting the speed at which tokens are placed in the token bucket. If there are enough tokens in the bucket, the rate of packets is within the rate limit and packets are directly sent out. If the tokens in the token bucket are insufficient, excess packets are buffered and are not discarded.

The number of tokens in the token bucket increases at a fixed rate until there are enough tokens in the token bucket. The system periodically takes out the packets in the buffer, and then sends them out. Each time the system sends buffered packets, the system compares the number of buffered packets with that of tokens in the token bucket and then deletes tokens of the same number as sent packets, until the number of tokens in the token bucket is too small to send packets.

Rate Limit for Packets Sent to the CPU

CPU attack defense provides hierarchical device protection: rate limit based on protocols, rate limit based on queues, and rate limit for all the packets. If all the preceding rate limits are set, the smallest rate limit takes effect.

Figure 9-1 Overview of Rate Limit for Packets Sent to the CPU

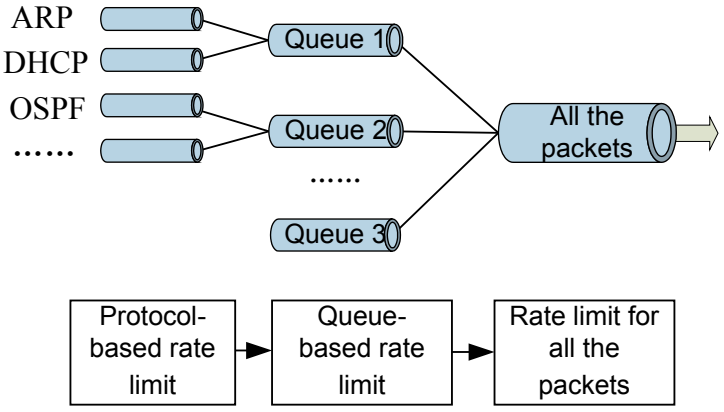

After the rate limit for protocol packets is set, each protocol has certain bandwidth so that protocol packets can be processed. In addition, protocol packets are processed independently.

After the rate limit for protocol packets is set, a queue is specified for each type of protocols. For example, a queue is allocated to management protocols such as Telnet and SSH and a queue is allocated to routing protocols. Queues are scheduled based on weights or priorities. Services with the highest priority are processed first.

After the rate limit for all the packets is set, the number of packets sent to the CPU is limited and more protocol packets can be processed. This function cannot protect the CPU when the CPU exception occurs.

If no independent network is deployed on a management interface or attacks occur on the independent network, the CPU is busy processing attack packets and fails to process other services. Consequently, the CPU usage becomes high, or even exceptions occur. To ensure processing capabilities and security of the switch, configure the rate limit on the management interface.

NOTE

After the rate limit is configured on the management interface, the switch may fail to be managed when severe attacks occur. Users cannot log in from the management interface. Remove viruses on the host or re-plan the networking.

Rate Limit Associated with Application-Layer Protocols

After the OSPF, BGP, or FTP connection is set up, the rate limit as shown in **Figure 9-1** is invalid. Rate limit is based on the application-layer protocols.

9.3.2 Attack Source Tracing

The attack source tracing function protects the CPU against Denial of Service (DoS) attacks. The S6700 enabled with attack source tracing analyzes packets sent to the CPU, collects statistics on the packets, and specifies a threshold for the packets. Excess packets are considered to be attack packets. The S6700 finds the source user address or source interface of the attack by analyzing the attack packets and generates logs or alarms. Accordingly, the network administrator can take measures to defend against the attacks.

As shown in **Figure 9-2**, attack source tracing involves the following processes:

- Parsing packets
- Analyzing traffic
- Identifying an attack source
- Sending logs or alarms to the network administrator and Attack Punishment

Figure 9-2 Attack source tracing processes

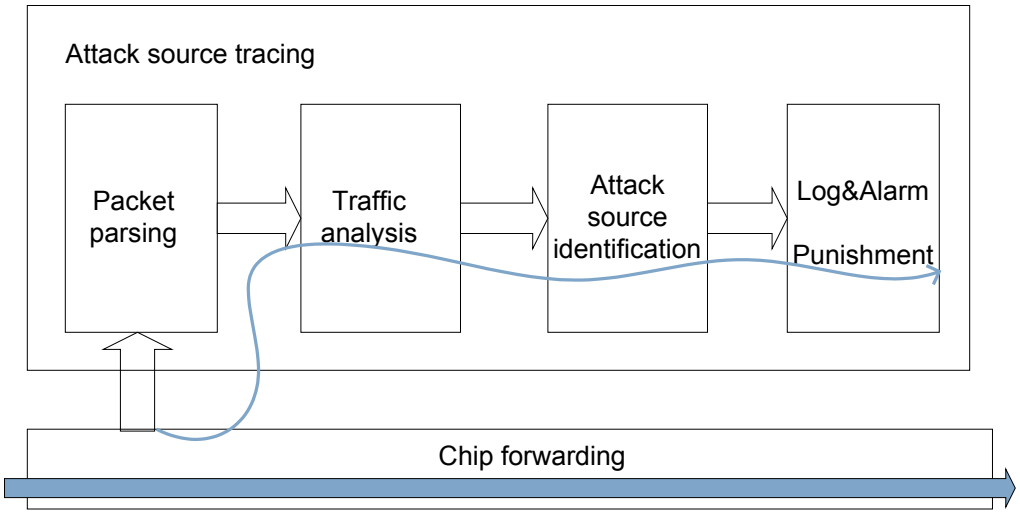

Figure 9-2 contains four phases:

1. Analyze packets based on users and ports. Users are identified by MAC addresses; ports are identified by physical port numbers and VLAN IDs (including inner VLAN IDs).
2. Count the number of received packets based on protocols and MAC addresses (or port information).
3. When the number of packets exceeds the threshold, the system considers that an attack occurs.
4. When detecting an attack, the system reports a log and a trap, or carries out Punishment. The current Punishment action is Deny.

9.4 Applications

9.4.1 CPU Attack Defense

CPU attack defense protects the S6700 against attacks and ensures non-stop service transmission when attacks occur.

Figure 9-3 CPU attack defense

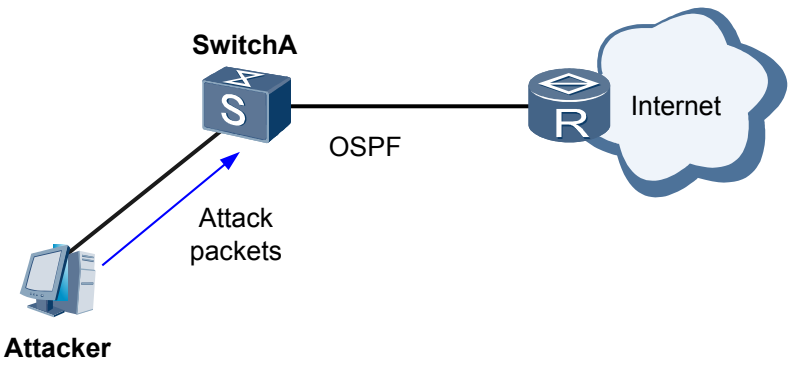

In this section, OSPF is used as an example. As shown in **Figure 9-3**, an attacker sends a large number of attack packets, which occupy too many CPU resources. As a result, OSPF packets cannot be exchanged and OSPF neighbor relationships cannot be set up.

The S6700 supports protocol-based rate limit, queue rate limit, and rate limit for all packets. These rate limit methods limit the attack packet rate within an allowed range so that OSPF packets can be properly processed. Therefore, authorized users can access networks.

After the OSPF connection is set up, the protocol-based rate limit does not take effect. Rate limit is performed based on the application-layer protocols.

9.4.2 Attack Source Tracing

Attack source tracing protects the CPU against attacks by sending a large number of protocol packets.

Figure 9-4 Attack source tracing

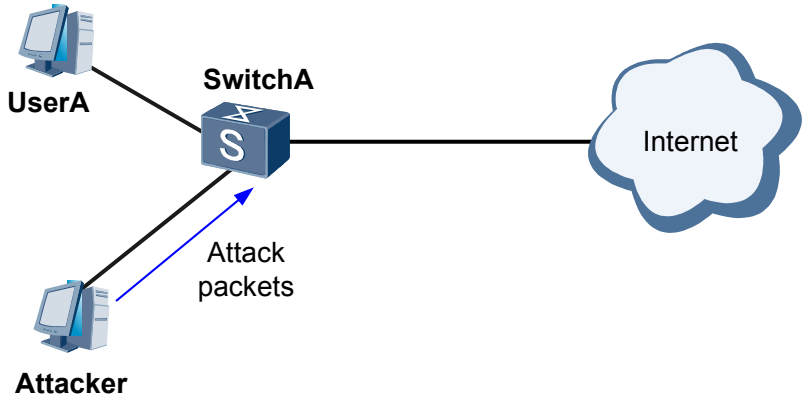

As shown in **Figure 9-4**, an attacker sends a large number of attack packets. As a result, the CPU is busy processing attack packets and fails to process other services. After attack source tracing is enabled, the S6700 can locate the MAC address of the attacker by collecting the statistics on the attack packets and analyzing them. Then, the S6700 generates an alarm and sends it to the network administrator. The network administrator can add the attacker to the blacklist by logging in to the S6700. This enables the S6700 to discard packets from the attacker to protect the CPU.

10 Traffic Suppression

About This Chapter

10.1 Introduction to Traffic Suppression

Definition

The S6700 support suppression of broadcast packets, unknown multicast packets, and unknown unicast packets received by Ethernet interfaces. When the number of broadcast packets, unknown multicast packets, or unknown unicast packets exceeds the configured threshold, the system discards the excessive packets. This reduces the traffic to an acceptable range and ensures the normal transmission of network services.

Broadcast traffic, unknown multicast traffic, and unknown unicast traffic need to be broadcast. If one or more interfaces do not need to broadcast the traffic, configure the S6700 to block these interfaces.

Purpose

When a Layer 2 Ethernet interface on the S6700 receives broadcast packets, multicast packets, or unknown unicast packets it forwards these packets to other Layer 2 Ethernet interfaces in the same VLAN. If there are a large number of such packets, the interface bandwidth is used and performance of the S6700 is diminished. To resolve the problem, suppress these packets to ensure that the rate of these packets stays within the desired range.

Benefits

Traffic suppression reduces broadcast storms and ensures high forwarding performance of the S6700.

10.2 References

Document No.	Description
IEEE 802.1d	Media Access Control (MAC) Bridges: specifies an architecture and protocol for the interconnection of IEEE802 LANs below the MAC service boundary.

10.3 Traffic Suppression

When forwarding unknown unicast, broadcast, or multicast packets, a switch broadcasts the packets to all its interfaces. If there are a large number of such packets, a lot of bandwidth resources are consumed and the security performance of the switch is degraded. Therefore, you need to control such packets and ensure that traffic is controllable on switches.

As shown in **Figure 10-1**, Switch A is directly connected to Switch B; when processing packets received from Switch A, XGE 0/0/2 on Switch B needs to check the types of the packets. If the packets are unknown unicast packets, multicast, or unknown broadcast packets, XGE 0/0/2 controls the traffic of the packets.

Figure 10-1 Networking diagram of storm control

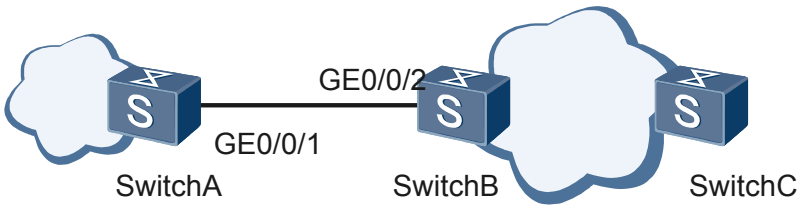

By default, the maximum percentage of broadcast packets, multicast packets, and unknown unicast packets that an Ethernet interface allows to pass through is 100% respectively. That is, the broadcast packets, unknown multicast packets, and unknown unicast packets are not suppressed on an Ethernet interface by default. The S6700 supports the limit on packets sent by an interface based on the percentage. The formula used to calculate the rate of traffic allowed by a device is: Rate of traffic allowed to pass through = Percentage of suppressed traffic x Rate set on an interface. The S6700 supports the setting of the maximum number of broadcast packets, multicast packets or unknown unicast packets allowed to pass through an interface per second.

If traffic suppression based on the inbound interface is configured, block the unknown unicast, multicast, and broadcast traffic on the outbound interface so that traffic is not broadcast and security of users or network devices connected to the outbound interface is ensured.

10.4 Applications

As shown in **Figure 10-2**, SwitchA is connected to a Layer 2 network and to a Layer 3 SwitchB. To limit the number of broadcast, multicast, or unknown unicast packets forwarded on the Layer 2 network, configure traffic suppression on the Layer 2 Ethernet interface of SwitchA.

Figure 10-2 Traffic suppression network diagram