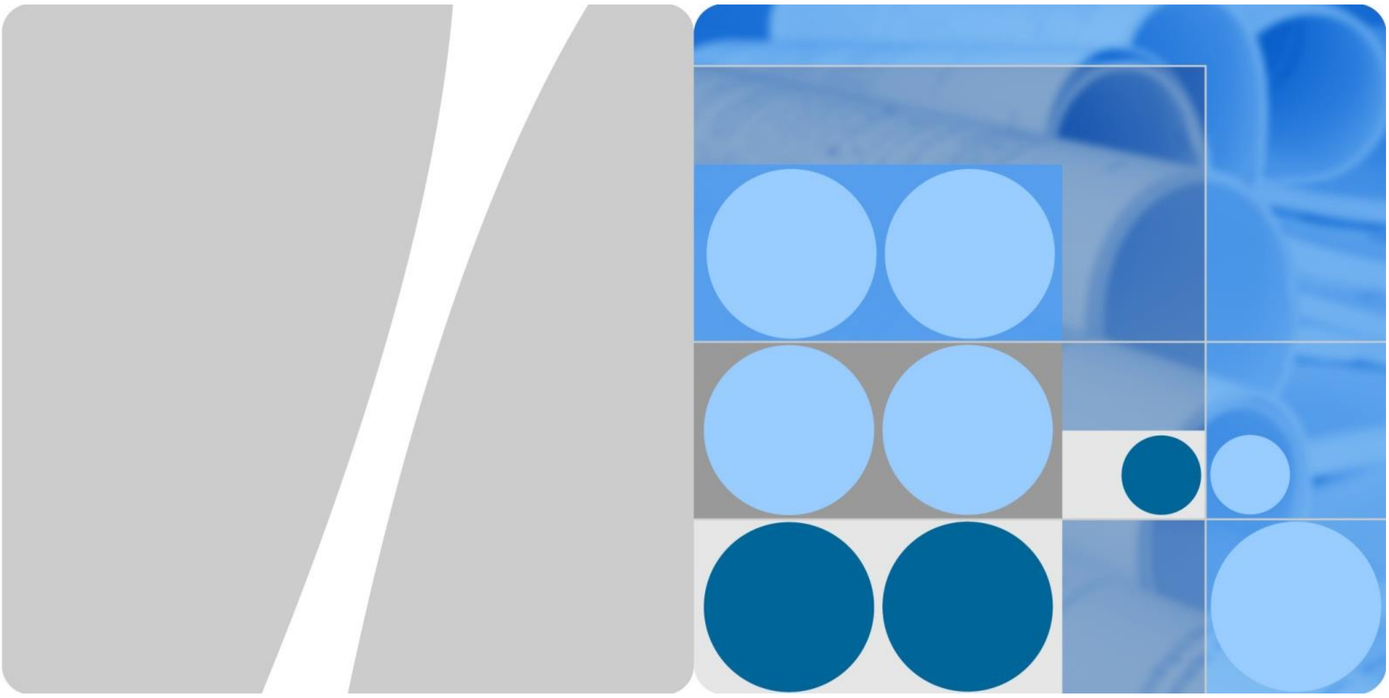

HUAWEI Data Center Network

Security Technology White Paper

Issue	**01**
Date	**2018-03-01**

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions

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 a 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

Contents

About This Document

Keywords

Cyberspace, network security, data center network

Abstract

Cyberspace development also brings security issues. This document describes basic network security capabilities of Huawei CloudEngine (CE) switches.

Acronyms and Abbreviations

Acronym/Abbreviation	Full Name
CP-CAR	Central Process CAR
GTSM	Generalized TTL Security Mechanism
URPF	Unicast Reverse Path Forwarding
VRP	Versatile Routing Platform
ACL	Access Control List
TACACS	Terminal Access Controller Access-Control System
AAA	Authentication, Authorization, Accounting

1 Data Center Network Security Overview

Explosive development of information and network technologies in the past decades brings a profound revolution to the society. This revolution promotes IT construction and interconnection around the world, gradual establishment of the cyberspace, and continuous social progress. However, malicious behaviors such as damage, theft, interference, espionage, and destruction also extend to the cyberspace, resulting in security issues in the cyberspace.

The security issues in the cyberspace are actually contradictions between attacks and defense. Attackers use vulnerabilities of resources such as computer systems or information in the systems to damage resource confidentiality, integrity, and availability. Resource owners need to identify, reduce, and eliminate resource vulnerabilities, reducing or eliminating security risks caused by attackers who use these vulnerabilities to attack resources. As the links of intermediate systems in the cyberspace, network devices are often prone to attacks. Because of interconnection in the cyberspace, security issues of network devices often cause adverse impact on the cyberspace. Security defense capabilities of network devices are the key to ensuring cyberspace security.

Huawei CE switches use the X.805 three-layer three-plane security isolation mechanism, as shown in Figure 1-1.

Figure 1-1 X.805 security architecture

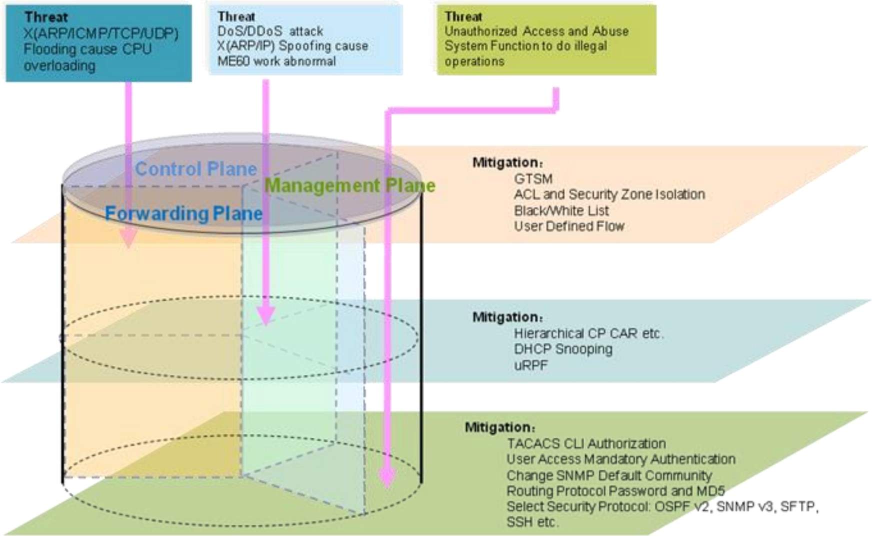

Network devices generally include three planes: management plane, control plane, and data plane. The control, management, and data planes are isolated from each other. If one plane is attacked, a CE switch can ensure normal running of the other two planes. Huawei CE switches provide security capabilities in the management, control, and data planes:

- **Management plane security**

  The management plane processes various operations and maintenance (O&M) activities of device administrators. The management plane's security capabilities protect O&M activities of device administrators. The management plane uses various security mechanisms to prevent potential security vulnerabilities of management protocols, and provides secure management channels to protect sensitive information during management.

- **Control plane security**

  The control plane maintains the running of various network protocols to control switching and routing of data flows. The control plane's security capabilities protect the running security of various control or signaling protocols to ensure the availability of the control plane and prevent leakage or abuse of network control information.

- **Data plane security**

  The data plane processes data flows entering the switch, and forwards data packets based on entries delivered by the control plane. With the data plane's security capabilities, the data plane filters or controls packets sent to the CPU to improve attack defense capabilities of the management or control plane and reduce device security risks. The filtering and control functions are complete on the data plane, and policy control is implemented on the CE switch.

This document describes a CE switch's basic security capabilities in the management, control, and data planes.

1.1 Security Capabilities of the Three Planes and Mitigated Attacks

The three planes (management, control, and data planes) of CE switches provide various security capabilities. The following table describes the security capabilities and mitigated attacks.

Table 1-1 Security capabilities and mitigated attacks

Security Capability	Description
Application layer association	Prevents many new session requests from occupying the bandwidth of successfully established sessions.
Defense against malformed packet attacks	Prevents the device buffer overflow and protocol stack breakdown caused by malformed packets from malicious attackers.
MD5 integrity check for routing protocols	Prevents incorrect routes caused by routing information intercepted and modified by man-in-the-middle.
ACL	Prevents unauthorized terminals from connecting

Security Capability	Description
	to the switch to perform unauthorized operations and access.
GTSM	Prevents network elements beyond the hop count allowed by the switch from communicating with the switch and defends against remote attacks by malicious users.
Attack source tracing and alarm	Records attack information when an attack occurs, facilitating source tracing and auditing.
CP-CAR	Limits the rate of packets sent to the CPU to prevent the CPU from being overloaded due to many requests, and prevents flood attacks.
Blacklist	Prevents malicious users from accessing the switch.
Whitelist	Ensures that authorized users preferentially access the switch.
Layer 2 loop detection	Prevents loops caused by misoperations, detects loops in real time, and reduces the impact of loops on the switch.
Security management center	Fast traces the attack source and locates faults through a unified management center when the switch is attacked.
URPF	Prevents DDoS attacks initiated by sending packets with bogus source IP addresses that are not on the network segment connected to the switch to attack remote network terminals.
Layer 2 ARP entry or MAC address entry limiting	Prevents attack with changed source addresses from exhausting entries on the switch and interrupting traffic forwarding.
DHCP snooping	Prevents attacks by packets with bogus addresses. If such attacks occur, other users cannot access the Internet.
System rights and account management	Prevents unauthorized users from accessing the system.

2 Network Security Threat Analysis

2.1 DoS Attack

Common network devices have powerful forwarding capabilities but limited processing capabilities in control and management planes. When attackers flood a network device with request messages in a denial of service (DoS) attack, the CPU of the network device fails to process messages in real time. As a result, normal service flows and internal processing flows are interrupted and services are denied.

DoS attacks are the greatest threat to network devices. When performing security hardening, consider defense against DoS attacks.

2.2 Information Disclosure

Unauthorized access to network devices poses a threat that can easily result in information disclosure. The following factors may lead to unauthorized access:

1. System configuration negligence: Network devices allow unauthorized login in some scenarios. After the live network is deployed, unauthorized login remains available due to negligence. Consequently, malicious users are able to gain access.
2. Negligent management: A configuration file template is usually used as a convenient way to deploy multiple network devices. If the administrator does not change the password for the administrator account on each network device, unauthorized access can occur.
3. Defects of IP network openness: Malicious users deploy sniffers and interception devices to intercept and parse IP packets in transmission, causing information disclosure.
4. Storage media information disclosure: Cards and storage devices of network devices are not encrypted when they are transferred to new locations.

2.3 Damage to Information Integrity

Because of IP network openness, packets may be maliciously tampered with by intermediate nodes or modified by man-in-the-middle attackers during transmission.

2.4 Unauthorized Access

A user can obtain the control or higher rights to a network device through unauthorized access. Possible causes of unauthorized access are as follows:

1. Network configuration vulnerabilities: Because the access control policy on the firewall is not configured properly, malicious users can crack into or force access to the system from a public network.
2. Unauthorized use of debugging tools provided by the system: Network devices allow access to internal system information to locate faults. Malicious users use diagnostic and debugging interfaces to obtain information.
3. Role-based management and control: The command control mechanism of a network device is based on user roles but not user accounts, so users may be able to gain access at an authority level higher than their own levels and run commands to read personal communication data or intercept system configurations.
4. No information isolation mechanism: The management information base (MIB) based on Simple Network Management Protocol (SNMP) does not have an information isolation mechanism. If users can access the MIB, they can traverse all MIB objects.

2.5 Identity Spoofing

Because of openness, IP networks do not have a powerful authentication and authorization mechanism for MAC and IP addresses, and ARP- or IP-based address spoofing attacks easily occur. As a result, a network device has to continuously update address entries required by the forwarding process and processes requests from spoofed addresses. Incorrect address entries may interrupt data forwarding, and insufficient entry learning capabilities make the network device vulnerable to DoS attacks.

2.6 Replay Attack

IP network openness makes it impossible for communication terminals to authenticate their peers at Layer 4 and lower layers. Hackers use this defect to initiate DoS attacks by repeatedly sending specific packets.

2.7 Computer Virus

A network device functions as a forwarding node and a network element (NE) that can be managed on the network. If computers on a network segment are infected with viruses, much spam traffic is generated, exhausting network bandwidth. In this case, a network device as an NE cannot obtain network resources. As a result, services are unavailable.

2.8 Human Error

Policies configured for service deployment during network construction are not deleted after service provisioning. The policies may be used by attackers to attack network devices.

During network reconstruction, engineers configure network devices incorrectly because of carelessness or lack of engineering skills, causing accidents. For example, incorrect connections of network cables cause loops, incorrect protocol configurations result in service interruptions, incorrect access control policies cause traffic to be blocked unexpectedly, and unnecessary access channels are activated.

Administrators carelessly share user name and password information with other people.

2.9 Physical Intrusion

Equipment room administrators often directly connect physical access devices to network devices. The network devices do not provide strong security defense capabilities. Attackers can obtain high-level permission through the physical connections, and access the network devices by avoiding door access control and intrusion monitoring systems.

3 Management Plane Security

Device management involves local and remote access management, and operation and maintenance. If device management security is damaged, attackers may log in to and control network devices to perform unauthorized actions. This affects the security and management of the entire network. Therefore, device management security is the key in network security. Establishing a strong security mechanism is important and critical for preventing unauthorized access and use. Device management security capabilities of CE switches involve access control, security management, file transfer, software integrity and sensitive information protection, and security audit.

3.1 Access Control

Network devices usually provide various access modes including local access through the console interface and remote access through Telnet and SSH. CE switches provide the following access control methods considering possible security risks of the access modes.

3.1.1 Authentication and Authorization

Physical interfaces, logical interfaces, and protocols including the console interface, Telnet, SSH, and SNMP that can offer system management provide access authentication to prevent unauthorized access. All these access modes support AAA authentication. Only the authenticated management users can access the device management interface. In addition, the hierarchical authorization mechanism is used for management users who pass authentication. The rights are classified into the visit level, monitoring level, configuration level, and management level in ascending order to reduce security risks caused by right allocation during device management.

3.1.2 Control for Enabling and Disabling Services

Some unnecessary access services may be enabled on some network devices by default, which poses potential security risks. CE switches provide control for enabling and disabling access services so that unnecessary access services can be disabled. For example, Telnet and SSH can be enabled or disabled on CE switches.

3.1.3 Changing Service Port Numbers

The default port numbers of some access services are known port numbers, and can be easily scanned and attacked. You can change these port numbers to private port numbers in the range of 1025 to 65535 to reduce the possibility of port scan and attack. The port numbers of Telnet, SSH, FTP (SFTP), and SNMP services can be changed.

3.1.4 Specifying Access Sources

CE switches can limit the access range to improve device security. For example, CE switches allow you to configure ACLs to permit specified IP addresses of access users or specify source interfaces so that the switches allow only packets from these interfaces to pass through.

3.1.5 Defense Against Brute Force Cracking

Dictionary attack is a common brute force cracking attack. When a dictionary attack occurs, you can take measures such as increasing password complexity or limiting login attempts.

Generally, the AAA server is responsible for access authentication. The AAA server usually requires users to use high-complexity passwords or limits login attempts. When the AAA server is unavailable, local authentication must provide the same features to prevent unauthorized users from logging in to the device.

CE switches have the minimum requirements for the local authentication password and password for upgrading the user level. The password must be a string of 8 to 128 case-sensitive characters and must contain at least two types of the following: uppercase letters, lowercase letters, digits, and special characters. The local user password cannot be the same as the user name or the user name in reverse order (the password is case-insensitive when being compared with the user name).

To limit the frequency of login attempts, you can configure the maximum number of authentication failures and reauthentication interval on CE switches to effectively delay an attack and increase the time required for the attack.

3.2 Security Management

The access control mechanism improves the device management security to a certain degree. However, some protocols cannot provide confidentiality and integrity protection for user and configuration information during communication management due to the lack of security mechanisms. CE switches provide two security management modes to further enhance the access security.

3.2.1 SSH

Secure Shell (SSH) provides secure remote connection services on an insecure network. When a user remotely logs in to a device through an insecure network, SSH ensures security with its encryption and authentication functions and protects the device against attacks such as IP address spoofing and plain-text password interception. SSH uses TCP to exchange data and establishes a secure channel over TCP.

3.2.2 SNMPv3

SNMP is used to manage network devices. Network administrators can use SNMP to obtain data from devices and configure devices. SNMP has three versions: v1, v2c, and v3. SNMPv1 and SNMPv2c encapsulate packets in UDP packets and transmit them in plain text, which easily causes user authentication and management information disclosure and poses potential security risks. CE switches support SNMPv3 with enhanced security. SNMPv3 transmits packets in cipher text and uses MD5/SHA and DES to authenticate and encrypt data, preventing information from being forged, tampered with, and leaked during transmission.

3.3 Software Integrity Protection

The software running on network devices undergoes many phases such as development, production, transmission, and delivery. If the software is tampered with or replaced by malicious personnel in a phase, or even is implanted with Trojan horses, viruses, or unauthorized programs and then installed on network devices on the live network, there are serious security risks. CE switches provide software integrity protection, and verify the integrity of software (software or patch package) during installation and upgrade, preventing the software from being tampered with. When the software is released, a digital signature is assigned to the software and the digital signature file containing the digital signature is incorporated into the software package. When an NMS distributes the software to an NE or the software is directly loaded on an NE, the digital signature is verified. The software is considered complete and valid only after the digital signature passes the verification. Otherwise, the software is considered invalid. When the software is loaded, MD5 authentication is performed online. If MD5 authentication is successful, the software is considered valid and can be loaded. Otherwise, the software is considered invalid.

3.4 Sensitive Information Protection

Network devices need to store some sensitive information such as user authentication information locally. The AAA server maintains and controls authentication information in a centralized manner, and provides security features to protect the sensitive information. If local authentication is used because the AAA server is unavailable, authentication information needs to be stored locally. CE switches encrypt and store the authentication information using an advanced encryption algorithm. In addition, CE switches can filter sensitive information in logs, protecting sensitive information.

3.5 Log Security

Logs record device information such as user operations and device running status, and are stored on devices as log files. Logs help network administrators monitor device running status and diagnose network faults. Log security involves:

- Log content security: Key sensitive information such as user password information is filtered and output in *** format.
- Log file security: Authentication in various access modes is used to ensure log file security. Log files can be viewed only by administrators. Administrators can log in to a device and run commands to view log files, or download log files from a remote device to the local PC and view the log files. Users can perform operations related to log files on a device only after they are authenticated and log in to the device successfully.

4 Control Plane Security

The control plane's security capabilities protect the running security of various control or signaling protocols to ensure the availability of the control plane and prevent leakage or abuse of network control information. CE switches use a mature network operating system and provide various types of network control protocols. This document only describes security capabilities of some basic network protocols, including TCP/IP security, routing service security, and switching service security.

4.1 TCP/IP Security

TCP/IP protocols are basic network protocols. Most routing protocols and application layer protocols use TCP/IP to transmit signaling packets. Attacks against TCP/IP affect the network stability. Due to inherent defects and flawed implementation of the TCP/IP protocol suite, increasing network attacks have a greater impact on TCP/IP networks. Attacks on network devices will make the network crash or become unavailable.

Attacks against TCP/IP networks are classified into flood attacks and malformed packet attacks.

- Flood attack: An attacker sends many data packets to attack the system. As a result, the system cannot receive requests from normal users or resources are exhausted. Flood attacks include SYN flood and Fraggle attacks.
- Malformed packet attack: An attacker sends malformed IP packets to the target system. As a result, the target system may break down when processing the malformed IP packets. Malformed packet attacks include Ping of Death and Teardrop attacks.

These two types of attacks are different from other types of attacks. Attackers do not search for access to the internal network, but prevent authorized users from accessing and using network devices or resources.

CE switches provide defense capabilities against these types of attacks, greatly enhancing system security.

4.1.1 Defense Against Malformed Packet Attacks

To prevent malformed packets from occupying too many resources and causing system crash or even network breakdown, CE switches directly discard received malformed packets.

- Flood attacks using packets without IP payloads: A CE switch considers IP packets with only IP headers and without any upper-layer data invalid and directly discards the packets.
- Attacks using IGMP null payload packets: If the length of an IGMP packet is less than 28 bytes, a CE switch considers the IGMP packet as a malformed packet and directly discards it.
- LAND attacks: A CE switch checks whether the source and destination addresses in a TCP SYN packet are the same. If they are the same, the CE switch considers the packet as a malformed packet and directly discards it.
- Smurf attacks: If the destination address in an ICMP echo request packet is a broadcast address or a subnet broadcast address, a CE switch considers the packet as a malformed packet and directly discards it.
- Attacks using packets with invalid TCP flag bits: A CE switch checks each flag bit of a TCP packet. If the URG, ACK, PSH, RST, SYN, and FIN flag bits are all 1 or 0, or both SYN and FIN flag bits are 1, the CE switch directly discards the packet.

4.1.2 Defense Against Fragment Attacks

The following describes several typical fragment attacks that may affect device availability and corresponding defense measures.

- Teardrop attacks: The offsets of fragmented packets may overlay. When the system reassembles fragmented packets, too many resources are occupied, causing network interruption. When processing Teardrop attack packets, a CE switch discards reassembled packets with overlapped offsets so that fragmented packets can be correctly reassembled.
- Attacks using packets with large offsets: An attacker sends fragmented packets with offset length greater than 65512 bytes. When the system reassembles fragmented packets, too many resources are occupied, causing network service interruption. When processing packets with large offsets, a CE switch checks the offset length and discards the packets if the offset length exceeds 65515 bytes.
- Attacks using repeated fragmented packets: An attacker sends repeated fragmented packets multiple times, including retransmitting the same fragment or sending different fragments with the same offset. As a result, the system fails to reassemble packets and the CPU usage is high. To defend against such attacks, a CE switch limits the rate of fragmented packets on LPUs, protecting the CPU. The Committed Access Rate (CAR) is configurable.

4.1.3 Defense Against Flood Attacks

Flood attacks include TCP SYN flood attacks, UDP flood attacks (including Fraggle attacks and UDP diagnosis port attacks), and ICMP flood attacks. CE switches limit the rate of TCP SYN flood attack packets and ICMP flood attack packets to prevent high CPU usage. The switches check port numbers in UDP packets and directly discard UDP packets with port numbers 7, 13, and 19 as UDP flood attack packets.

4.2 Routing Service Security

The routing service is one of the most important components on network devices and is vulnerable to network attacks. Therefore, necessary measures must be taken to ensure routing service security. The routing service is used to learn and advertise routing information.

Network devices often exchange routing protocol packets to learn and advertise routing information to find the optimal path for data forwarding. The routing service is mainly threatened by attacks on routing protocols, especially attacks on neighbor relationships and routing information.

Before exchanging routing information, two network devices need to establish a neighbor relationship using a routing protocol. Some attacks attempt to destroy or prevent neighbor relationship establishment to affect secure running of the routing service. For example, attackers can send malformed packets and reset TCP connections to terminate established peer sessions, or initiate DoS attacks to consume device resources such as memory and CPU resources to prevent neighbor relationship establishment. Some routing protocols provide automatic discovery mechanisms to simplify network device deployment. This mechanism assumes that peers are reliable, which makes it possible for a bogus device to establish a neighbor relationship with an authorized device and to inject invalid routes.

CE switches provide mechanisms such as neighbor authentication, Generalized TTL Security Mechanism (GTSM), and route filtering to reduce security risks of the routing service.

4.2.1 Neighbor Authentication

Most routing protocols support neighbor authentication. Neighbor authentication ensures that a CE switch receives only reliable routing information from trusted neighbors. A CE switch verifies authenticity of each neighbor and integrity of routing update messages. Each CE switch is initially configured with a shared key to authenticate each routing update message. Before sending a routing update message, each CE switch uses the predefined key to assign a digital signature for the routing update message and encapsulates the signature in the routing update message. The neighbor that receives the message authenticates the message to verify authenticity and integrity of the message. Routing protocols including BGP, IS-IS, OSPF, and RIP support neighbor authentication.

Neighbor authentication protects peers from session reset and session hijacking attacks from unauthorized routing peers. Neighbor authentication also protects routing information from invalid route injection, and prevents unauthorized routing peers from deleting or updating valid routes. However, neighbor authentication cannot prevent a peer that pretends to be an authorized router from injecting invalid routing information. Fortunately, this attack can be defended by route filtering, which will be described later.

Most routing protocols support two neighbor authentication modes: plain text authentication and MD5 authentication. In plain text authentication mode, the authentication key is transmitted in routing protocol messages in plain text. This mode cannot provide high security because routing protocol messages may be intercepted during transmission. In MD5 authentication mode, the authentication key is calculated using the MD5 algorithm, and then encapsulated in routing messages. Therefore, this mode is more secure than plain text authentication. MD5 authentication mode is recommended in actual deployment.

4.2.2 GTSM

If an attacker simulates a network device to send routing protocol packets to a device continuously, the device receives these packets, finds that the packets are sent to itself, and directly sends them to the routing protocol module in the control plane without checking the validity of these packets. As a result, the control plane is busy processing these packets, causing a high CPU usage.

GTSM is used to protect TCP/IP-based control protocols (such as routing protocols) against CPU-targeted attacks, such as CPU overload attacks. GTSM protects services above the IP layer by checking whether the time to live (TTL) value in the IP header is within a predefined range. It applies to the scenario where the neighbor relationship of a routing protocol is established between adjacent or neighboring devices and the TTL cannot be easily modified during packet forwarding. GTSM provides the following technical measures:

- For protocol peers that are directly connected, GTSM sets the TTL value of an outgoing protocol packet to 255. The forwarding plane of the GTSM-enabled peer directly discards protocol packets with a non-255 TTL value, protecting the control plane against attacks.
- For a multi-hop peer relationship, a proper TTL range, such as 251 to 255, can be defined. The forwarding plane of the peer directly discards protocol packets whose TTL value is out of the range, protecting the control plane against attacks.

4.2.3 Route Filtering

Route filtering is another mechanism that protects routing service security. Most routing protocols support route filtering to prevent specified routes from being advertised to a network. From the perspective of security, route filtering ensures that only valid routes are advertised.

Route filtering includes filtering of routing information exchanged between routing peers and filtering of routing information exchanged between different routing processes on the same routing device. Route filtering between routing peers controls route import between devices. Route filtering between routing processes controls route import between different routing protocols, generally between IGP and BGP.

4.3 Switching Service Security

The switching service is another important service provided by CE switches. The switching service security involves availability protection of Layer 2 networks. The Spanning Tree Protocol (STP), ARP, and DHCP are the most commonly used protocols on Layer 2 networks. The following sections describe security capabilities of these protocols on CE switches.

4.3.1 STP Security

On a Layer 2 switching network, packets are replicated and transmitted continuously when a loop occurs, causing a broadcast storm. The broadcast storm consumes all available bandwidth, making the network unavailable. STP is a Layer 2 management protocol. It blocks redundant links on the network to eliminate Layer 2 loops, and provides link backup. Devices running STP exchange BPDUs containing topology information to discover loops on the network and block some interfaces to prune the network into a loop-free tree network. STP prevents infinite looping of packets to ensure packet processing capabilities of devices.

Although STP effectively eliminates loops on Layer 2 networks, it does not provide any authentication or encryption method to protect BPDUs. As a result, it is vulnerable to attacks. Due to lack of authentication, an attacker can establish a session with an STP-enabled device and easily send bogus BPDUs to trigger network topology recalculation, causing network flapping and service interruption. In addition, BPDUs are not encrypted. Therefore, attackers can easily intercept non-encrypted BPDUs to obtain important topology information.

CE switches provide some protection mechanisms to reduce security risks against STP.

1. BPDU protection

   On a CE switch, the port that is directly connected to a user terminal such as a PC or a file server is configured as an edge port to ensure fast port status transition. Under normal circumstances, no BPDU is sent to edge ports. If the CE switch is attacked by bogus BPDUs, the CE switch automatically configures the ports that receive bogus BPDUs as non-edge ports and recalculates the spanning tree. As a result, network flapping occurs. MSTP provides BPDU protection to defend against such attacks. After BPDU protection is enabled, the CE switch shuts down the edge port that receives BPDUs, and informs the NMS of the event. The edge port that is shut down can only be manually enabled by the network administrator.

2. Root protection

   The root switch on a network may receive a BPDU with a higher priority due to incorrect configurations or network attacks. When this occurs, the root switch becomes a non-root switch, which causes incorrect changes of the network topology. Such changes may cause traffic to be switched from high-speed links to low-speed links, leading to network congestion. To address this issue, CE switches provide root protection. The root protection function protects the role of the root switch by retaining the role of the designated port. The port enabled with root protection remains as the designated port in all instances. When the port receives a BPDU with a higher priority, the port stops forwarding packets and enters the listening state, but does not turn into a non-designated port. If the port does not receive any BPDUs with a higher priority within a given period of time, it restores to the forwarding state.

3. Loop protection

   A CE switch maintains the status of the root port and blocked port by continuously receiving BPDUs from the upstream switch. If ports cannot receive BPDUs from the upstream switch due to link congestion or unidirectional link failures, the switch re-selects a root port. The original root port becomes a designated port, and the original blocked port changes to the forwarding state, which may cause loops on the network. The loop protection function prevents such network loops. After the loop protection function is enabled, the root port is blocked if it cannot receive BPDUs from the upstream switch. The blocked port remains in the blocked state and does not forward packets. This prevents loops on the network.

4. Defense against TC BPDU attacks

   After receiving TC BPDUs, a CE switch deletes the corresponding MAC address entries and ARP entries. If a malicious attacker sends bogus TC BPDUs to attack the CE switch, the CE switch receives a large number of TC BPDUs within a short period of time, and deletes its MAC entries and ARP entries frequently. As a result, the CE switch is heavily burdened, threatening the network stability. After defense against TC BPDU attacks is enabled, you can set the number of times TC BPDUs are processed within a given period of time in the MSTP process. (The default period is 2s, and the default number of times is 3.) If the number of TC BPDUs received by the MSTP process within the given period exceeds the specified threshold, TC BPDUs are processed only for the specified number of times in the MSTP process. After the timer expires, the remaining TC BPDUs are processed once in the MSTP process. In this way, the CE switch does not need to delete MAC entries and ARP entries frequently.

4.3.2 ARP Attack Defense

ARP is a basic Layer 2 protocol. It is easy to use but has no security mechanisms. Attackers often use ARP to attack network devices. Common ARP attacks are classified into ARP spoofing attacks and ARP flood attacks.

- ARP spoofing attacks

  An attacker sends forged ARP packets to devices (such as gateways or hosts) on a network. The devices then modify their ARP entries, leading to forwarding failures.

- ARP flood attacks

ARP flood attacks are also called DoS attacks. An attacker sends a large number of gratuitous ARP packets or forged ARP request packets in which destination IP addresses cannot be parsed to a device. The device creates a large number of invalid ARP entries and the ARP table becomes full quickly. As a result, the device cannot save valid ARP entries and fails to forward packets of authorized users.

CE switches provide various ARP security features and technologies to detect and defend against ARP attacks, ensuring the security of network devices and communication networks.

1. Dynamic ARP Inspection (DAI)

   When a CE switch functions as a Layer 2 access device, it broadcasts ARP request packets from a user in a VLAN. Attackers in the VLAN also receive the ARP request packets. When an attacker receives an ARP request packet destined for another device, the attacker can reply with a forged ARP reply packet and replaces the MAC address of the destination device with its own MAC address. The source device then incorrectly sends data packets of authorized users to the attacker, threatening the network communication security. DAI allows a CE switch to prevent the attack and is implemented as follows:

   - If DAI is enabled in a VLAN to which an untrusted interface belongs, the CE switch checks the validity of ARP packets (including ARP request packets, ARP reply packets, and gratuitous ARP packets) sent from the untrusted interface according to the DHCP snooping binding table. If corresponding DHCP snooping entries are found and the source IP address, source MAC address, port number, and VLAN ID in the ARP packets match those in the DHCP snooping binding entries, the CE switch allows the ARP packets to pass through. If the source IP address, source MAC address, port number, and VLAN ID in the ARP packets do not match DHCP snooping binding entries or no corresponding DHCP snooping entries are found, the CE switch considers the ARP packets as attack packets and directly discards them.
   - The CE switch does not check the validity of ARP packets sent from a trusted interface.

2. ARP packet validity check

   The MAC addresses in an ARP packet include the MAC addresses in the Ethernet header and the MAC addresses in the data field of the ARP packet. Normally, the MAC addresses in the Ethernet header are the same as those in the data field. An attacker can change the MAC addresses in the data field of ARP packets sent by authorized users and send the forged ARP packets to attack network devices. The source and destination MAC addresses in the forged ARP packets differ from those in the Ethernet header.

   CE switches provide ARP packet validity check. A CE switch compares the source and destination MAC addresses in the data field of a received ARP packet with those in the Ethernet header of the packet. If they are the same, the CE switch considers the ARP packet valid. If they are different, the CE switch considers the ARP packet invalid and directly discards it. The CE switch provides the following check modes:

   - Based on the source MAC address: The CE switch checks whether the source MAC address in the data field and that in the Ethernet header of an ARP packet are the same. If they are the same, the CE switch considers the ARP packet valid. If they are different, the CE switch considers the ARP packet invalid and directly discards it.
   - Based on the destination MAC address (only for ARP reply packets): The CE switch checks whether the destination MAC address in the data field of an ARP packet is all 1s. If the destination MAC address is all 1s, the CE switch considers the ARP packet invalid and directly discards it. The CE switch checks whether the destination MAC address in the data field and that in the Ethernet header are the same. If they are the same, the CE switch considers the ARP packet valid. If they are different, the CE switch considers the ARP packet invalid and directly discards it.

- Based on source and destination MAC addresses: For an ARP request packet, the CE switch only checks whether the source MAC address in the data field and that in the Ethernet header are the same. If they are the same, the CE switch considers the ARP packet valid. If they are different, the CE switch considers the ARP packet invalid and directly discards it. For an ARP reply packet, the CE switch checks whether the source MAC address in the data field and that in the Ethernet header are the same, checks whether the destination MAC address in the data field is all 1s, and checks whether the destination MAC address in the data field and that in the Ethernet header are the same. If the source MAC addresses are the same, the destination MAC address in the data field is all 1s, and the destination MAC addresses are the same, the CE switch considers the ARP packet valid; otherwise, the CE switch discards the packet.

3. ARP entry fixing

   When a device receives the first ARP packet from a user, it adds an ARP entry matching the user in the ARP table. If the device receives an ARP packet from the user again, it updates the ARP entry based on the ARP packet, including the MAC address, port number, VLAN ID, and aging time. If an attacker maliciously sends forged ARP packets to update the ARP entries of authorized users, the ARP entries on the device are incorrect and the authorized users cannot access the network properly.

   ARP entry fixing allows a CE switch not to update a learned ARP entry, to update some information in the entry, or to update the entry only upon acknowledgement. ARP entry fixing has three modes:

   - Fixed-all mode: If the MAC address, port number, and VLAN ID in a received ARP packet do not match those in the corresponding ARP entry, the CE switch directly discards the ARP packet.
   - Fixed-mac mode: When a CE switch receives an ARP packet, the CE switch is allowed to update the port number and VLAN ID in the corresponding ARP entry, but cannot update the MAC address in the entry.
   - Send-ack mode: When a CE switch receives an ARP packet, it does not directly update the corresponding ARP entry. Instead, the CE switch constructs an ARP request packet in which the destination IP address is the IP address in the ARP entry and matching the MAC address in the received ARP packet and broadcasts the ARP request packet. If the CE switch can receive an ARP reply packet, it updates the ARP entry according to the ARP reply packet. Otherwise, the ARP entry cannot be updated.

4. Strict ARP learning

   Generally, a network device learns received ARP request packets. ARP request packets are broadcast in a VLAN. Therefore, the network device can learn ARP request packets in the broadcast domain unconditionally. After strict ARP learning is configured, the network device only learns the ARP reply packets in response to the ARP request packets sent by itself, and does not process the ARP request packets sent by other devices or the ARP reply packets in response to the ARP request packets that are not sent by itself. In this way, the network device can reject most ARP request and reply packets sent by attackers.

5. ARP learning suppression

   In some cases, when ARP attacks occur on an interface connected to a network device, the network device receives a large number of ARP packets within a short period of time. As a result, ARP entries on the network device overflow and authorized users may fail to connect to the network. ARP learning suppression limits the number of ARP entries learned by each interface to prevent ARP entry overflow and ensure ARP entry security.

6. ARP packet rate limiting

   If a device receives a large number of ARP packets within a short period of time, the device is busy learning ARP entries and responding to ARP request packets, which consume many CPU resources and affect processing of other services. ARP packet rate limiting effectively prevents CPU resources of the device from being wasted in processing ARP packets to ensure normal processing of other services. After ARP packet rate limiting is configured, the device counts the number of ARP packets received within a specified period. If the number of ARP packets exceeds the configured threshold, the device does not process excess ARP packets. CE switches support ARP packet rate limiting based on the source MAC addresses, source IP addresses, destination IP addresses, and VLAN ID.

7. Rate limiting on ARP Miss messages

   If a device cannot find a matching ARP entry during packet forwarding, the device reports an ARP Miss message to the upper-layer software. After receiving the ARP Miss message, the upper-layer software generates a dynamic fake ARP entry, sends it to the device, and then sends an ARP request packet to request the MAC address of the destination host. After receiving the ARP reply packet, the upper-layer software sends the learned ARP entry to the device to replace the original fake entry so that traffic can be forwarded properly. If there are too many ARP Miss messages, many CPU resources of the device are used to process the ARP Miss messages, affecting processing of other services.

   Rate limiting on ARP Miss messages effectively prevents CPU resources of the device from being wasted in processing ARP Miss messages to ensure normal processing of other services. The device counts the number of ARP Miss messages received within a specified period. If the number of ARP Miss messages exceeds the configured threshold, the device does not process excess ARP Miss messages. CE switches support rate limiting on ARP Miss messages based on the source IP addresses and VLAN ID.

8. Gratuitous ARP packet discarding

   When a new device is connected to a network, the device broadcasts a gratuitous ARP packet to notify other devices on the network of its MAC address and check whether any device uses the same IP address as its own IP address in the broadcast domain. Any device can send gratuitous ARP packets and devices receive gratuitous ARP packets without identity authentication. Therefore, a large number of gratuitous ARP packets may be generated on a network and network devices are busy processing the gratuitous ARP packets, causing CPU overload and affecting processing of other services.

   After gratuitous ARP packet discarding is configured, a device directly discards received gratuitous ARP packets, reducing CPU resource consumption. This feature effectively prevents CPU resources of the device from being wasted in processing ARP packets to ensure normal processing of other services.

4.3.3 DHCP Snooping

DHCP is one of the basic protocols of Layer 2 networks. It provides a simple and efficient host configuration mechanism that is widely used on networks. DHCP faces the following security threats:

- Bogus DHCP server attack

  If a private DHCP server exists on a network, the private DHCP server interacts with DHCP clients that request IP addresses. As a result, the DHCP clients obtain incorrect IP addresses and network configuration parameters, and cannot go online.

- Man-in-the-middle attack and IP/MAC spoofing attack

An attacker uses forged packets to function as a man-in-the-middle role to communicate with the DHCP client or server, or uses spoofed IP and MAC addresses of authorized users to communicate with the DHCP server so that DHCP clients cannot obtain services.

- Attacks initiated by sending forged DHCP request packets to extend IP address leases

  If an attacker continuously sends forged DHCP request packets to extend IP address leases, some expired IP addresses cannot be reclaimed.

- Starvation attack

  An attacker may change the MAC address in the data frame header to continuously request IP addresses from the DHCP server. As a result, the addresses in the address pool on the server are exhausted and other users cannot obtain services.

- DoS attack by changing the CHADDR value

  An attacker may change the MAC address in the CHADDR field of DHCP packets to continuously request IP addresses from the DHCP server. As a result, the addresses in the address pool on the server are exhausted and other users cannot obtain services.

DHCP snooping is a DHCP security feature implemented by a CE switch to defend against the preceding potential threats. DHCP snooping enables the CE switch to establish and maintain a DHCP snooping binding table to filter out untrusted DHCP messages according to the table. A binding entry contains the MAC address, IP address, lease, binding type, VLAN ID, and port number. DHCP snooping acts as a firewall between DHCP clients and a DHCP server to effectively defend against the preceding common attacks as follows:

- **Defense against bogus DHCP server attacks**

  On a CE switch, network-side interfaces can be configured as trusted interfaces and user-side interface can be configured as untrusted interfaces. The CE switch discards DHCP reply packets received from untrusted interfaces. This prevents bogus DHCP server attacks and ensures that DHCP clients obtain IP addresses from the valid DHCP server.

- **Defense against man-in-the-middle attacks and IP/MAC spoofing attacks**

  A CE switch checks information such as the source IP address and source MAC address in IP packets or ARP packets against the DHCP snooping binding table. If matching entries exist, the CE switch forwards the packets. If no matching entry is found, the CE switch discards the packets.

- **Defense against attacks initiated by sending forged DHCP request packets to extend IP address leases**

  A CE switch checks the source IP address, source MAC address, VLAN ID, and port number in DHCP request packets against the DHCP snooping binding table. If matching entries exist, the CE switch forwards the DHCP request packets. If no matching entry is found, the CE switch discards the DHCP request packets.

- **Defense against starvation attacks**

  The maximum number of MAC addresses that can be learned on an interface of a CE switch is limited to prevent users from sending a large number of DHCP request packets by changing MAC addresses.

- **Defense against DoS attacks by changing the CHADDR value**

  A CE switch checks the CHADDR value in a DHCP request packet. If the CHADDR field value matches the source MAC address in the data frame header, the CE switch forwards the DHCP request packet. If the CHADDR field value does not match the source MAC address in the data frame header, the CE switch discards the DHCP request packet.

4.3.4 MFF

MAC-Forced Forwarding (MFF) isolates clients at Layer 2 and connects clients at Layer 3 in a broadcast domain.

Layer 2 isolation must be implemented for users on an Ethernet network because their services are different. These users need to communicate sometimes, so Layer 3 communication must be implemented for them. In traditional Ethernet networking, VLANs are assigned to implement Layer 2 isolation and Layer 3 communication among different clients. However, VLAN-based Layer 2 isolation has the following disadvantages:

- When many users need to be isolated at Layer 2, a large number of VLANs are required.
- To enable clients to communicate at Layer 3, you need to assign an IP network segment to each VLAN and configure an IP address for each VLANIF interface. Assigning too many VLANs wastes IP addresses.

MFF is introduced to solve the preceding problems. An MFF-enabled device intercepts ARP request packets and replies with ARP reply packets containing the gateway MAC address through the proxy ARP mechanism. In this way, all traffic (including the traffic in the same subnet) from users is forcibly sent to the gateway to implement Layer 2 isolation. If users who use different services need to communicate, the gateway can forward traffic exchanged between them to implement Layer 3 communication.

Because users communicate through the gateway, the gateway can monitor data traffic and prevent malicious attacks between users, which protect the network security.

5 Data Plane Security

In addition to security capabilities designed and implemented for security vulnerabilities and threats of various services, CE switches provide common network security policies to ensure that traffic entering the network matches specific network security policies. After these network security policies are enforced, CE switches discard abnormal packets, reducing device security risks. The data plane filters and controls packets sent to the CPU to implement these security policies. Therefore, the security capabilities of CE switches in this aspect are also called data plane security capabilities. The following sections describe common security mechanisms, including application layer association, URPF, IP source guard, CP-CAR, and traffic suppression and storm control.

5.1 Application Layer Association

Application layer association is the association between the protocol status (enabled or disabled) on the control plane and protocol packet sending of the forwarding engine on the data plane. A connection is established between the control plane and the data plane to ensure protocol status consistency. When a protocol is disabled, the system directly discards packets of this protocol to prevent attacks. When a protocol is enabled, the system limits the rate of the protocol packets sent to the CPU, ensuring that CPU resources are not exhausted and the network runs normally.

Application layer association supports application of whitelist and blacklist. A whitelist refers to a group of authorized users or users with high priorities, and a blacklist refers to a group of unauthorized users. When processing protocol packets, the data plane matches them with the whitelist first, and sends the packets matching the whitelist with high bandwidth and at a high rate. For common protocol packets, the data plane then checks whether the protocol packets are defined in application layer association and sends the packets matching an enabled protocol at a configurable rate. If the protocol of the protocol packets is disabled or the packets match the blacklist, the data plane directly discards the packets or sends them at a low rate.

5.2 URPF

Network attacks based on source address spoofing often occur on the Internet. A DoS attack occurs when there are a large number of such attack packets, greatly affecting the network security.

Unicast Reverse Path Forwarding (URPF) prevents network attacks based on source address spoofing. The reverse concept indicates a process opposite to normal route searching. When receiving a packet under normal circumstances, a CE switch obtains the destination IP address of the packet and then searches the forwarding table for the route to the destination. If the CE switch finds a matching route, it forwards the packet; otherwise, it discards the packet. Before a URPF-enabled CE switch sends or forwards a packet, the CE switch obtains the source address and inbound interface of the packet. The CE switch then uses the source address as the destination address to obtain the corresponding inbound interface and compares the obtained interface with the inbound interface. If they do not match, the CE switch considers the source address as a spoofed address and discards the packet. URPF can effectively protect the network against malicious attacks initiated by changing the source IP address of packets.

Route symmetry is required to ensure proper running of URPF. The data flows from a user to a host on the Internet and those from the host to the user must be transmitted along the same path between the user device and the Internet Service Provider (ISP) device. Otherwise, the URPF-enabled device discards some normal packets if the packets are received and sent from different interfaces.

5.3 IP Source Guard

As the network scale expands increasingly, more source IP address-based attacks occur. Some attackers use spoofed IP addresses to access and obtain network resources, intercept user information, or even block authorized users from accessing the network. IP source guard provides a mechanism to effectively defend against IP address spoofing attacks.

IP source guard checks IP packets against the binding table. Before a CE switch forwards an IP packet, it compares the source IP address, source MAC address, port number, and VLAN ID in the IP packet with entries in the binding table. If a matching entry is found, the CE switch considers the IP packet as a valid packet and forwards it. Otherwise, the CE switch considers the IP packet as an attack packet and discards it. For example, a user goes online through DHCP. A network device generates a binding entry for the user based on the DHCP ACK packet. The binding entry contains the user's source IP address, source MAC address, port number, and VLAN ID. When the network device receives an IP packet from a user, it checks whether the IP packet matches the binding entry of the user. If the IP packet matches the binding entry of the user, the network device considers the packet as a valid packet from an authorized user and allows the packet to pass through. If the IP packet is a forged IP packet from an attacker, the network device cannot find a matching entry in the binding table, and discards the packet.

5.4 CP-CAR

CP-CAR is used to set the rates of packets sent to the CPU. You can set the average rate, committed burst size (CBS), and priority for each type of packets. You can set different CAR rules for different types of packets to reduce their impact on each other and protect the CPU. CAR can also be used to set the total rate of packets sent to the CPU. When the total rate exceeds the upper limit, the system discards the packets to prevent CPU overload.

5.5 Traffic Suppression and Storm Control

Traffic suppression and storm control are security features used to prevent broadcast storms.

When a Layer 2 Ethernet interface on a switch receives broadcast, multicast, or unknown unicast packets, the switch forwards the packets to other Layer 2 Ethernet interfaces in the same VLAN because the switch cannot determine the outbound interface based on the destination MAC addresses of these packets. In this case, a broadcast storm may occur and the forwarding performance of the switch deteriorates.

Traffic suppression and storm control can control these packets and prevent broadcast storms. Traffic suppression limits the traffic using the configured threshold, and storm control blocks the traffic by shutting down interfaces.