DHCP Technology White Paper

This document provides generic technical information, some of which might not be applicable to your products.

The information in this document is subject to change without notice.

Contents

Overview

Technical background

To send and receive data on the Internet, a host must have an IP address and information such as the IP address of a DNS server and router. Dynamic Host Configuration Protocol (DHCP) is a client/server protocol widely used for dynamic assignment of IP addresses and other configuration data to endpoints.

DHCP was developed based on Bootstrap Protocol (BOOTP).

BOOTP is a client/server protocol designed for diskless clients to obtain an IP address and such information as the DNS server IP address, boot image filename, and gateway IP address from a server. It is suitable only for small-sized, stable networks on which each host gets a permanent network connection and the network configuration is seldom changed. The network administrator only needs to configure a BOOTP configuration file that contains a set of parameters for each BOOTP client. This configuration file might remain unchanged for weeks.

DHCP enhances BOOTP to provide dynamic assignment of IP addresses and complex configuration data on networks that are complex or prone to change. It dynamically issues configuration data, and assigns and withdraws IP addresses for reuse as endpoints connect to and disconnect from the network. This helps conserve IP resources on networks and is especially important for networks that have mobile endpoints such as laptops and wireless clients.

Benefits of the H3C DHCP implementation

The H3C implementation of DHCP delivers the following benefits:

- Full DHCP capabilities, including DHCP client, DHCP relay agent, and DHCP server capabilities.
- Strong support for services and flexible deployments.
- Excellent serviceability and configurability.
- Interoperability with devices of mainstream vendors, including Windows and Linux servers.
- Easy management and low deployment costs.

DHCP implementation

Concepts

- **DHCP server**—A DHCP server assigns IP addresses and other network configuration parameters to DHCP clients.
- **DHCP client**—A DHCP client obtains an IP address and other network configuration parameters from a DHCP server. It is the initiator of an IP address allocation process.
- **DHCP relay agent**—A DHCP relay agent forwards DHCP messages between DHCP servers and clients that are located on different subnets.
- **DHCP snooping**—A DHCP snooping device snoops DHCPACK and DHCPREQUEST messages to record IP-to-MAC bindings for security purposes.

DHCP message format

Figure 1 shows the DHCP message format.

Figure 1 DHCP message format

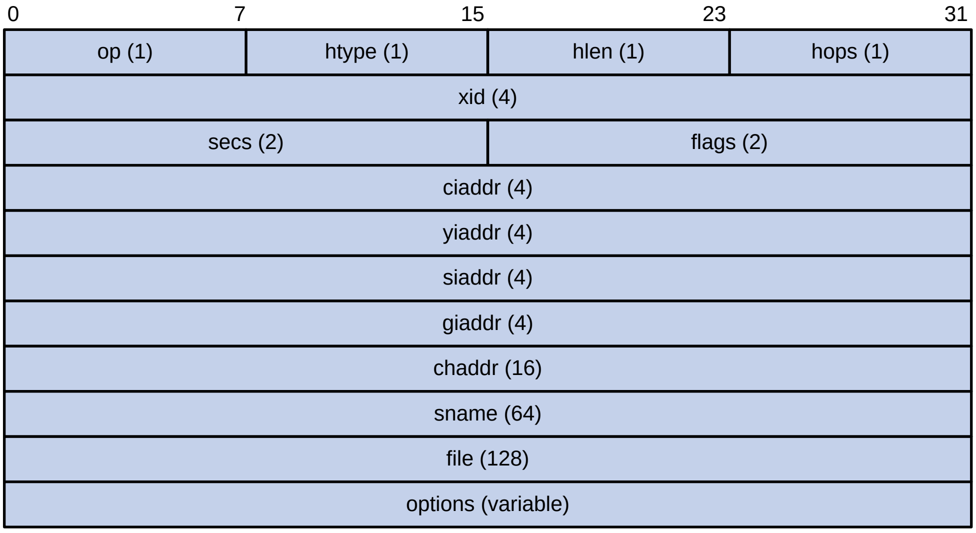

A DHCP message contains the following fields:

- **op**—General type of the message. A value of 1 indicates a request message and a value of 2 indicates a reply message.
- **htype**—Hardware address type.
- **hlen**—Hardware address length. This field is only applicable to Ethernet and has a fixed length of 6.
- **hops**—Number of relay agents a request message has traversed.
- **xid**—Transaction ID, a random number chosen by the client to identify an IP address allocation transaction.
- **secs**—Filled in by the client. It is the number of seconds that has elapsed since the client began the address acquisition or renewal process. This field is reserved and set to 0.
- **flags**—The leftmost bit is defined as the BROADCAST (B) flag. If this flag is set to 0, the DHCP server unicasts the reply. If this flag is set to 1, the DHCP server broadcasts the reply. The remaining bits of the flags field are reserved for future use.
- **ciaddr**—Client IP address. This field is filled in by the client if the client has an IP address that is valid and usable. If the client has not obtained a valid IP address, this field is set to zero.
- **yiaddr**—Your IP address. It is an IP address offered by the DHCP server to the DHCP client.
- **siaddr**—Server IP address, from which the client obtained configuration parameters.
- **giaddr**—Gateway IP address. It is the IP address of the first relay agent at which a request message arrived.
- **chaddr**—Client hardware address.
- **sname**—Server host name, from which the client obtained configuration parameters.
- **file**—Boot file (also called system software image) name and path information, assigned by the server to the client.
- **options**—Optional parameters field, which is variable in length. Optional parameters include the message type, lease duration, domain name server IP address, and WINS IP address.

DHCP server operating mechanisms

A DHCP server is typically placed on the same subnet as its clients, because some DHCP messages are broadcast. If you use a DHCP server to provide services for clients on different subnets, you must deploy a DHCP relay agent between the server and its clients.

The generic DHCP operating mechanisms are the same, whether or not a DHCP relay agent is used. The following information describes these mechanisms by using a DHCP server located on the same subnet as its clients.

Address allocation

DHCP supports the following address allocation mechanisms:

- **Static (or manual) allocation**—The network administrator assigns an IP address to a client, such as a WWW server, and DHCP conveys the assigned address to the client.
- **Automatic allocation**—DHCP assigns a permanent IP address to a client.
- **Dynamic allocation**—DHCP assigns an IP address to a client for a limited period of time, which is called a lease.

The network administrator can select these mechanisms for DHCP clients as needed.

Dynamic IP address allocation process

The following are phases for a DHCP client to obtain a valid IP address from a DHCP server:

1. **Discovery**—The client discovers available DHCP servers.
2. **Offering**—Each available DHCP server offers an IP address and other configuration parameters to the client.
3. **Selection**—The client selects an IP address assigned by one of the DHCP servers.
4. **Acknowledgment**—The DHCP server confirms the offered IP address.

Figure 2 Dynamic IP address allocation process

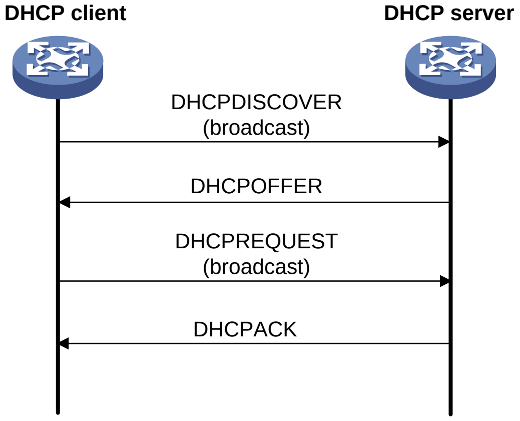

Discovery phase

Typically, a DHCP client does not have IP addresses of DHCP servers when it starts up. At startup, the DHCP client broadcasts a DHCPDISCOVER message to discover available DHCP servers.

All available DHCP servers send replies to the client upon receipt of the DHCPDISCOVER message.

Offering phase

Upon receipt of the DHCPDISCOVER message, a DHCP server offers an IP address lease and other configuration data to the client in a DHCPOFFER message. The configuration data might include information such as the gateway IP address and DNS server address.

The DHCP server maintains assignable IP addresses and other configuration information in address pools.

A DHCP server uses the following process to offer an IP address to a client:

1. Selects an IP address for the client in the following order:
    a. The IP address manually bound to the client’s MAC address or ID.
    b. An IP address that was ever assigned to the client.
    c. The IP address designated by the client in Option 50 in the DHCPDISCOVER message.
    d. The first assignable IP address found in the DHCP address pool selected from all available pools.
    e. An IP address whose lease has expired.
    f. An IP address that caused an IP conflict.

    If no IP address is assignable, the server does not respond to the discovery message.

    If the server finds an assignable IP address, the server proceeds to verify that the IP address is not used by any other device.
2. Sends ICMP echo requests to ping the selected IP address until it receives a response or the maximum number of ping attempts is reached.
    - If the server receives a response to a ping attempt, the server marks the IP address as a conflicting IP address and continues to select a new IP address for the client.
    - If the server does not receive a response to any one of its ping attempts, the server determines that the IP address is available for use and offers the IP address to the client.

Selection phase

If multiple DHCP servers send DHCPOFFER messages to the client, the client typically accepts the first received offer. Then, the client sends a formal request for the offered IP address in a DHCPREQUEST broadcast message. This message contains the IP address of the server that supplies the accepted parameters in Option 54 (the server ID option). The client sends this information to notify DHCP servers that it will use the IP address offered by the DHCP server specified in Option 54.

Acknowledgment phase

Upon receipt of the DHCPREQUEST message from the DHCP client, the DHCP server looks up its lease records based on the MAC address in the message. If a lease record exists for the MAC address, the server returns a DHCPACK message to confirm that the client can use the IP address. If no matching lease record is found or if the IP address cannot be assigned, the DHCP server returns a DHCPNAK message to deny the IP address allocation. In this case, the client needs to broadcast a DHCPDISCOVER message again to request a new IP address.

Upon receipt of the DHCPACK message, the DHCP client broadcasts a gratuitous ARP packet to check for IP conflict. If the client does not receive a reply within the specified time, the client uses this IP address. If the client receives a reply, the client sends a DHCPDECLINE message to the server to request a new IP address.

Reuse of previously allocated IP addresses

A DHCP client does not send a DHCPDISCOVER message when it accesses the network again. Instead, it directly broadcasts a DHCPREQUEST message to request the IP address previously allocated to it. This IP address is sent in Option 50.

Upon receipt of the message, the DHCP server determines whether the client can use the IP address, as follows:

- If the requested IP address is available, the DHCP server returns a DHCPACK message, as shown in Figure 3.

  Figure 3 Messages exchanged when the requested IP address is available

  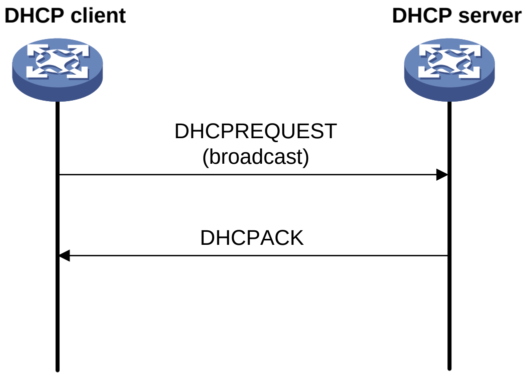

- If the requested IP address is not available, for example, because the IP address is already assigned to another DHCP client, the DHCP server returns a DHCPNAK message. Then, the client needs to send a DHCPDISCOVER message to request a new IP address, as shown in Figure 4.

  Figure 4 Messages exchanged when the requested IP address is not available

  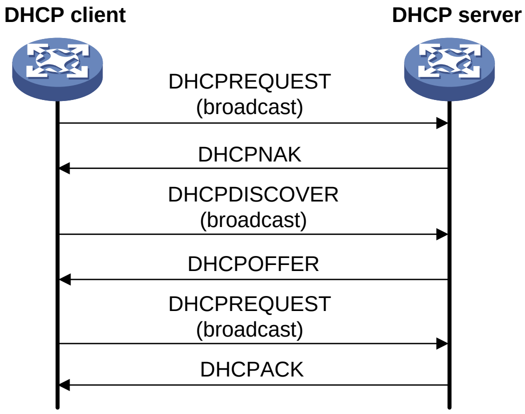

IP address lease renewal

The IP address dynamically allocated by a DHCP server to a client has a lease. The DHCP server reclaims the IP address when the IP lease expires. To extend the IP address lease, the client must renew the lease before it expires.

The following is the typical IP address lease renewal process:

1. When half the lease duration (T1) elapses, the DHCP client sends a DHCPREQUEST unicast message to renew the lease.
2. Upon receipt of the message, the DHCP server renews the lease and returns a DHCPACK unicast message if the IP address is still available. If the IP address is not available, the DHCP server returns a DHCPNAK unicast message.
3. The client acts depending on the reply state:
   - If the client has not received a reply when 87.5% of the lease duration (T2) elapses, the client broadcasts the DHCPREQUEST message again for lease renewal.

- If the client receives a DHCPACK message, the client continues to use the IP address.
- If the client receives a DHCPNAK message, the client initiates a process to obtain a new IP address.

Figure 5 shows this process.

Figure 5 IP address lease renewal process

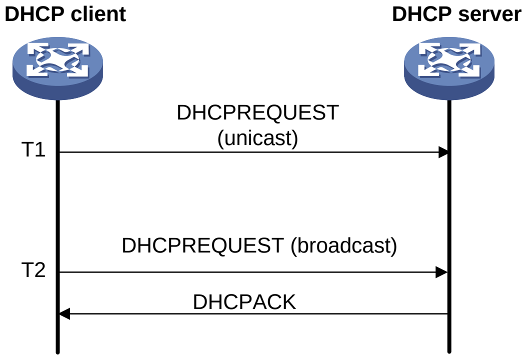

Release of an IP address

If a DHCP client does not need its IP address any longer, it sends a DHCPRELEASE message to the DHCP server. However, the DHCP server will retain the configuration information for the client so the client can reuse this information when it requests an IP address again.

Acquisition of additional configuration information

To obtain more configuration information than an IP address from the DHCP server, the DHCP client sends a DHCPINFORM message with Option 55. In this option, the client specifies the requested configuration parameters.

Upon receipt of the DHCPINFORM message, the DHCP server returns a DHCPACK message to assign the requested network parameters to the client.

Dynamic subnet allocation on a CUPS network

A control plane and user plane separation (CUPS) network decouples the control plane (CP) from the data plane, or user plane (UP). On a CUPS network, the CP provides control services such as user identification, and address allocation and management. The UP provides forwarding services.

A CP uses a control channel (for example, an OpenFlow connection) to issue table entries to a UP and uses a protocol channel (a VXLAN tunnel) to exchange protocol packets with the UP.

On a CUPS network, a CP acts as a DHCP server to allocate subnets from the primary network segment in an address pool to UPs.

As shown in Figure 6, the CP device connects to UP devices through VXLAN tunnels and acts as the DHCP server. The CP subnets the primary network segment depending on the configured subnet mask length.

When a UP receives a DHCPDISCOVER message from a client, the UP forwards the message to the CP. Then, the CP performs the following operations:

1. Examines whether a subnet has been assigned to the UP and whether the subnet has an assignable IP address for the client.
    - If a subnet has been assigned and that subnet has assignable IP addresses, the CP selects an IP address from the subnet for the client.

- If no subnet has been assigned to the UP or no assignable addresses are available, the CP goes to the next step.

2. Allocates a subnet to the UP, and then selects an IP address from the subnet for the client. The CP also issues the subnet route to the UP. After the routing table converges, users from external networks can access endpoints attached to the UP.
3. When all users on the subnet go offline, the CP reclaims the idle subnet to conserve IP resources.

Figure 6 Dynamic subnet allocation

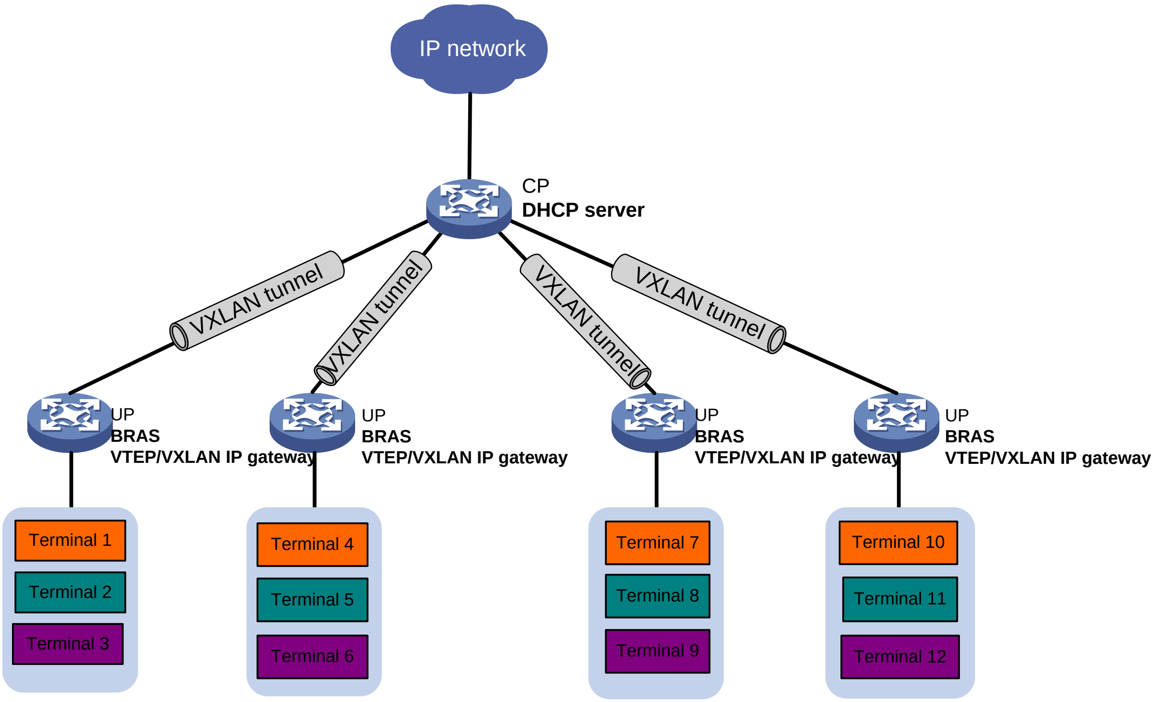

ODAP

The On-Demand Address Pool (ODAP) feature makes full use of address resources through dynamic subnet allocation. It is applicable to large-scale networks that have a large number of users.

An ODAP network contains an ODAP server for overall IP subnet resource management and one or multiple ODAP clients for subnet-specific IP resource management.

Figure 7 ODAP network

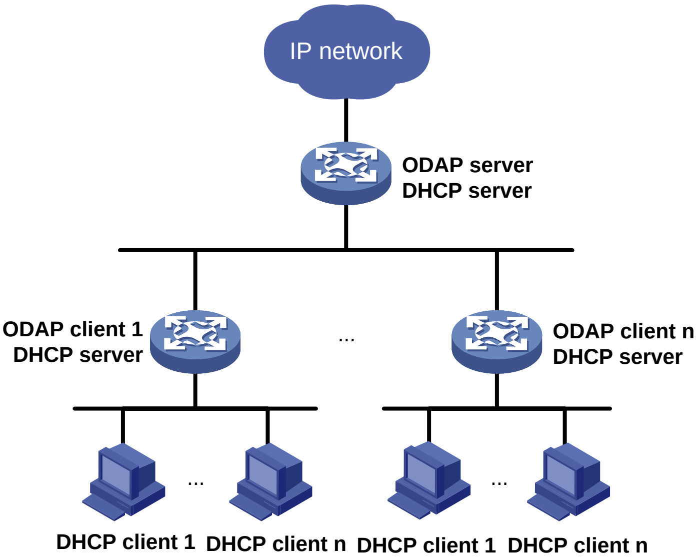

As shown in Figure 7, the ODAP server acts as a DHCP server to assign subnets to ODAP clients. Each IP pool on the ODAP server contains a primary network segment. This network segment can be divided into multiple subnets for allocation depending on the configured subnet mask length.

Each ODAP client acts as a DHCP client to obtain subnets from the ODAP server and acts as a DHCP server to allocate IP addresses to users.

The following is the generic address allocation process on an ODAP network:

1. Upon receipt of a DHCPDISCOVER message from an end-user client, the ODAP client examines whether it has an assignable IP address for the client.
   - If a subnet has been assigned, and an assignable IP address is available, the ODAP client selects an IP address from the subnet for the end-user client.
   - If no subnets or assignable addresses are available, the ODAP client goes to the next step.
2. The ODAP client requests a subnet from the ODAP server. Subnet acquisition uses the same process as IP address acquisition except that the messages exchanged between ODAP server and ODAP client contain Option 220 and Option 221.
   - In the request message, Option 220 contains the requested subnet information, and Option 221 contains information about the VPN instance in which the ODAP client resides.
   - In the reply sent by the ODAP server, Option 220 contains the allocated subnet information, and Option 221 contains information about the VPN instance in which the ODAP server resides.
3. The ODAP client uses the obtained subnet to allocate IP addresses to the end-user DHCP client.

Unaware of ODAP client errors or disconnections, an ODAP server cannot reclaim allocated subnets when such an event occurs. To avoid the waste of subnet resources, you can force the ODAP server to renew its allocated subnet information.

The following is the generic renewal process:

1. The ODAP server sends a DHCPFORCERENEW message to all ODAP clients to obtain information about the subnets that they have obtained.
2. Each ODAP client sends back their own subnet information in a DHCPREQUEST message.
3. The ODAP server collects subnet statistics from all received DHCP REQUEST messages and reclaims subnets that do not have statistics.

IP pool group

As the name implies, an IP pool group contains multiple IP pools. These IP pools can be the same type or different types. For example, an IP pool group can contain common IP pools, local BAS IP pools, and remote BAS IP pools.

The IP pool group feature addresses the following requirements in an AAA environment, especially on a hybrid non-CUPS and CUPS network:

- The AAA server selects different DHCP servers for users depending on their location.
- The AAA server acts as the DHCP server for address allocation, and also as a relay agent to forward DHCP messages exchanged between DHCP clients and DHCP servers.
- On a hybrid non-CUPS and CUPS network, the AAA server must use both common IP pools and subnet IP pools to assign IP addresses and subnets.

The following information describes the generic DHCP dynamic address allocation process on a CUPS network for which an IP pool group is used:

1. When the device receives a DHCP discovery or request message from a client attached to a user plane device, the device performs the following operations:
    - If the client is already online with at an IP address, the device renews the lease on that IP address for the client.
    - If no IP address was ever assigned to the client, but a static binding exists for the client in an IP pool, the server assigns the IP address in the static binding to the client.
    - If neither of the above situations exists, the address allocation goes to the next step.
2. The device examines the IP pool group for common IP pools or local BAS IP pools bound to the UP address or remote interface of the UP device.
    - If a common IP pool or a local BAS IP pool is available, the device assigns the client an IP address in the bound IP pool.
    - If none of the IP pools in the pool group is available, the address allocation goes to the next step.
3. The device examines the IP pools in the IP pool group in the following order until an IP address or a DHCP server is found for the client or until the allocation attempt fails:
    a. Common IP pools or local BAS IP pools bound to the UP device.
    b. Common IP pools or local BAS IP pools not bound to the UP device.
    c. Remote BAS IP pools bound to the UP device.
    d. Remote BAS IP pools not bound to the UP device.

    If an address is available from a common IP pool or a local BAS IP pool, the server offers that address to the client.

    If a remote BAS IP pool is used, the device forwards the request from the client to the DHCP server configured in the remote BAS IP pool.

    If none of these IP pools are available for address allocation, the address allocation process fails.

NOTE:

You can bind an IP pool to a UP device by binding that pool to the UP address or to the remote interface on the CP for that UP.

Restrictions

If there are multiple DHCP servers on a network, one DHCP server does not know which IP addresses have been assigned by the other DHCP servers.

DHCP clients on one subnet cannot communicate with a DHCP server on another subnet without a DHCP relay agent.

DHCP relay agent operating mechanism

To use one DHCP server for centralized assignment of IP addresses and configuration data to clients on different subnets, you must deploy a DHCP relay agent.

Figure 8 DHCP relay agent working process

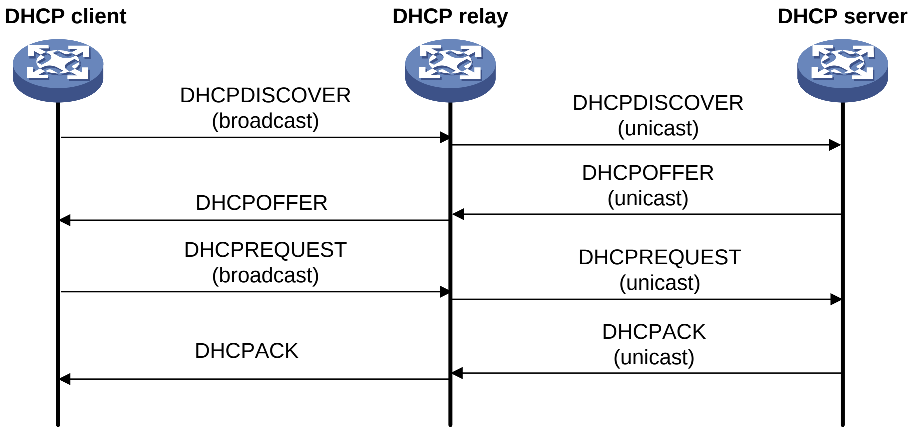

As shown in Figure 8, a DHCP relay agent forwards requests and replies between a DHCP client and a DHCP server.

The following is the IP assignment process that involves a DHCP relay agent:

1. Upon receipt of a DHCPDISCOVER or DHCPREQUEST message from a client, the DHCP relay agent examines the hops field.

   This field contains the number of relay agents that the message has traversed.

   - If the value in the hops field exceeds the upper limit, the DHCP relay agent discards the DHCP request to avoid the risk of a loop.
   - If the value does not exceed the limit, the DHCP relay agent goes to the next step.
2. The DHCP relay agent performs the following operations and then forwards the message:
   a. Examines the **giaddr** field.
      - If the **giaddr** field is 0, the DHCP relay agent fills in this field with the IP address of the interface that received the request. If the incoming interface has multiple IP addresses, it selects the primary IP address.
      - If the **giaddr** field is not 0, the relay agent leaves it intact.

   b. Increases the value in the hops field by one.

   c. Sets the TTL value in the request message to its default TTL value, instead of decreasing the value by one.

   d. Changes the destination IP address in the DHCP request message to the IP address of the DHCP server or the next DHCP relay agent.
3. Upon receipt of the message, the DHCP server returns an IP address and other configuration parameters to the relay agent. The IP address and configuration parameters are selected based on the **giaddr** field.
4. Upon receipt of the reply from the DHCP server, the relay agent performs the following operations:

   a. Examines the **giaddr** field in the reply to determine whether the reply is intended for a client attached to one of its interfaces.

If the IP address in the **giaddr** field does not belong to one of its local interfaces, the relay agent discards the reply.

b. Examines the broadcast flag in the reply.

- If this flag is set to 1, the DHCP relay agent broadcasts the reply to DHCP clients.
- If this flag is not set to 1, the DHCP relay agent unicasts the reply to the DHCP client. The destination IP address is in the **yiaddr** field and the link layer address is in the **chaddr** field.

DHCP snooping operating mechanism

Basic DHCP snooping features

DHCP snooping is a security feature that operates on Layer 2 network devices. It listens to DHCPREQUEST and DHCPACK messages exchanged between DHCP clients and servers to record clients' MAC and IP addresses for security purposes.

DHCP Snooping trusted/untrusted ports

To prevent DHCP clients from obtaining IP addresses and configuration data from an unauthorized DHCP server on the network, configure DHCP snooping trusted ports and untrusted ports.

- **Trusted**—Configure a port as a trusted port if it is connected to an authorized DHCP server, directly or indirectly. A trusted port forwards DHCP messages normally to guarantee that DHCP clients can obtain valid IP addresses.
- **Untrusted**—Configure a port as an untrusted port if it is connected to an unauthorized DHCP server. The device discards the DHCPACK, DHCPNAK, and DHCPOFFER messages received on an untrusted port to prevent DHCP clients from receiving invalid IP addresses.

DHCP extensions

DHCP security features

H3C provides various DHCP security features to protect DHCP clients, servers, and relay agents against DHCP attacks and protect the network against unauthorized access.

The following information describes some of the most widely used DHCP security features.

Maintaining DHCP relay entries

The DHCP relay agent maintains a binding table that records information about all authorized clients.

The DHCP relay agent can dynamically record IP-to-MAC bindings (also called DHCP relay entries) for clients in that binding table after they obtain IP addresses.

Aging of DHCP relay entries

A DHCP relay agent forwards the DHCPRELEASE messages from DHCP clients to DHCP servers without processing them. It might retain a relay entry for a client for a period of time after the DHCP server releases the MAC and IP binding for that client. To resolve this issue, periodic refresh of dynamic relay entries was introduced.

With this feature enabled, the DHCP relay agent periodically sends a DHCPREQUEST message to the DHCP server for each relay entry. Each request message contains the following information:

- The IP address in a relay entry as the desirable IP address.
- The MAC address of the DHCP relay interface.

Then, the relay agent maintains the relay entries depending on what it receives from the DHCP server.

- If the server returns a DHCPACK message or does not return any message within the required interval, the DHCP relay agent removes the relay entry. The relay agent also sends a DHCPRELEASE message to release the IP address if it receives a DHCPACK message.
- If the server returns a DHCPNAK message, the relay agent retains the relay entry, knowing that the lease on the IP address still exists.

DHCP flood attack protection

DHCP flood attack protection limits the incoming rate of DHCP messages on interfaces on a per-source MAC basis. This feature is configurable on DHCP servers and DHCP relay agents.

When an interface enabled with this feature receives a DHCP message from a client (MAC address), the device creates a DHCP flood attack protection entry in check state. If the number of DHCP messages from that MAC address reaches the upper limit during the detection period, the device determines that the client is launching a DHCP flood attack. The device changes the DHCP flood attack protection entry to the restrain state and starts to discard the DHCP messages sent from that client. When the aging timer for the entry expires, the device examines the drop rate of the DHCP messages sent from the MAC address.

- If the drop rate is lower than the DHCP flood attack threshold, the DHCP relay agent deletes the entry. If a DHCP message from that MAC address arrives after the deletion, the device creates a new flood attack protection entry.
- If the drop rate is equal to or higher than the DHCP flood attack threshold, the device restarts the aging timer for the entry.

Interface-based DHCP attack suppression

Interface-based DHCP attack suppression protects interfaces against DHCP flood attacks by limiting the rate of incoming DHCP messages on interfaces. This feature is configurable on DHCP servers and DHCP relay agents.

When an interface enabled with this feature receives a DHCP message, the device creates a DHCP attack suppression entry in check state for the interface. If the incoming DHCP message rate on the interface reaches the threshold, the device determines that a DHCP attack occurs on the interface. Then, it changes the suppression entry to the restrain state. To protect the CPU against DHCP attacks, the server limits the rate of incoming DHCP messages on the interface before the aging timer for the suppression entry expires.

When the aging timer for the DHCP attack suppression entry expires, the device examines the incoming DHCP message rate on the interface.

- If the incoming DHCP message rate is below the suppression threshold, the device deletes the entry. If a DHCP message arrives at that interface after the deletion, the device creates a new attack suppression entry for that interface.
- If the incoming DHCP message rate reaches or exceeds the suppression threshold, the device restarts the aging timer for the entry.

DHCP starvation attack protection

A DHCP starvation attack occurs when an attacker constantly sends forged DHCP requests that use different MAC addresses in the **chaddr** field to a DHCP server. This attack exhausts the IP address resources of the DHCP server so legitimate DHCP clients cannot obtain IP addresses. In the worst scenario, the DHCP server might fail to work because of exhaustion of system resources.

If the forged DHCP messages are encapsulated in frames that use different source MAC addresses, you can use the following method to mitigate DHCP starvation:

- Set the MAC learning limit on a port.
- Disable the port from forwarding unknown frames when the MAC learning limit is reached.

This method has the following drawbacks:

- While it prevents the attacker from obtaining too many IP addresses, it also prevents some of the legitimate clients attached to that port from obtaining IP addresses.
- This method is not applicable if the forged DHCP requests are encapsulated with the same source MAC address.

To prevent DHCP starvation attacks without affecting legitimate DHCP clients, enable the DHCP server and relay agent to perform MAC address check on DHCP requests.

When a DHCP request arrives, the DHCP server or relay agent compares the client MAC address in its **chaddr** field with the source MAC address in its frame header for inconsistency.

- If the MAC addresses are the same, the device determines that the request is legitimate and continues to process the request.
- If the MAC addresses are different, the device determines that the request is illegitimate and drops the request.

DHCP server proxy

To protect DHCP servers against attacks by malicious users, you can set up a DHCP relay agent as a DHCP server proxy between DHCP servers and clients.

As a DHCP server proxy, a DHCP relay agent always changes the server's IP address to the relay interface's IP address before forwarding the responses from DHCP servers to clients. Then, the DHCP clients take the DHCP relay agent as the DHCP server for subsequent DHCP transactions.

DHCP Option 82

DHCP Option 82, also known as the DHCP Relay Agent Information Option, is defined in RFC 3046 for DHCP servers to assign parameters based on location of clients for security purposes. It enables a DHCP snooping device or relay agent to insert location information about clients in messages sent to the DHCP servers. Then, the DHCP servers can assign IP addresses and other parameters to the clients in a VLAN or switching domain depending on their location.

Option 82 has two sub-options: sub-option 1 (Circuit ID) and sub-option 2 (Remote ID). The circuit ID identifies the interface or VLAN on the device on which the request was received. The remote ID identifies the remote host.

Figure 9 Option 82 application scenario

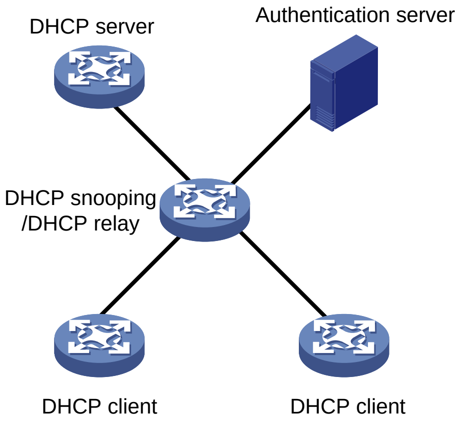

As shown in Figure 9, Option 82 works as follows:

1. A DHCP client sends a DHCPDISCOVER message to obtain an IP address.

   Before a DHCP client obtains an IP address, a DHCP snooping device or DHCP relay agent does not forward any traffic from that client except for DHCP messages.

2. Upon receipt of the message, the DHCP snooping device or relay agent adds the client location information in Option 82 and forwards the message to the DHCP server.
3. The DHCP server selects an IP address for the client based on the client's location and returns a DHCPOFFER message with the Option 82 field. The Option 82 field in the DHCPOFFER message contains the same information as that added by the snooping device or relay agent.
4. The DHCP snooping device or DHCP relay agent forwards the DHCPOFFER message to the client, with the Option 82 field removed.
5. After receiving the DHCPREQUEST from the client, the DHCP snooping device or relay agent adds the client location information in Option 82 and forwards the DHCPREQUEST to the DHCP server.
6. The DHCP server replies with a DHCPACK message, with Option 82 intact.
7. The DHCP snooping device or relay agent forwards the DHCPACK message to the client, with Option 82 removed.

DHCP Option 82 not only allows for granular IP address management and enhanced security. You can also use DHCP Option 82 in conjunction with PBR to apply different forwarding rules to traffic from different IP addresses.

Autoconfiguration

The autoconfiguration feature allows a device to obtain a configuration file from a remote file server when starting up without any configuration file.

The following is the generic autoconfiguration process:

1. The starting device sets an interface (such as the VLAN-interface 1 or the first Ethernet interface) as the DHCP client to request configuration parameters from the specified DHCP server.

   These startup configuration parameters include the IP address and name of a TFTP file server, and a bootfile name.
2. After getting the startup configuration parameters, the DHCP client sends a TFTP request to obtain the configuration file from the specified TFTP server for system initialization.
3. If the client fails to obtain startup configuration parameters, it starts up without loading any configuration file.

Application scenarios

Single-subnet DHCP

As shown in Figure 10, the DHCP clients dynamically obtain IP addresses and other configuration parameters from the DHCP server on the same subnet.

Make sure the network ranges in the address pool and the IP address of the client-facing interface on the DHCP server are on the same network segment.

Figure 10 DHCP server deployed on the same subnet as DHCP clients

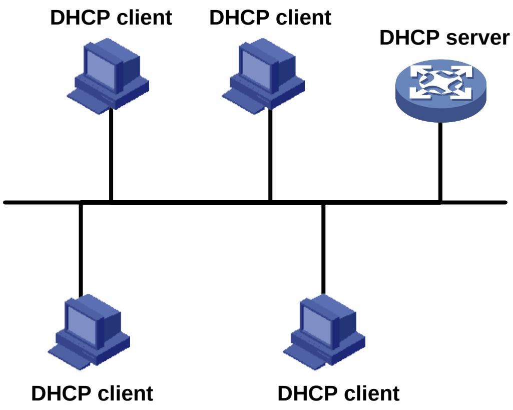

Multi-subnet DHCP

As shown in Figure 11, deploy a DHCP server to provide services for clients on different subnets. Enable a DHCP relay agent on each clients-attached interface to forward messages exchanged between the clients and the server. These interfaces are called relay interfaces.

To ensure that the DHCP relay agent and the DHCP server can send unicast packets to each other, configure a route between them.

Make sure the IP address of each relay interface is on the same network segment as one of the IP address pool on the DHCP server. This configuration ensures that the clients attached to them can obtain an IP address that is on the same subnet as their gateway for successful communication on the network.

Figure 11 Multi-subnet DHCP

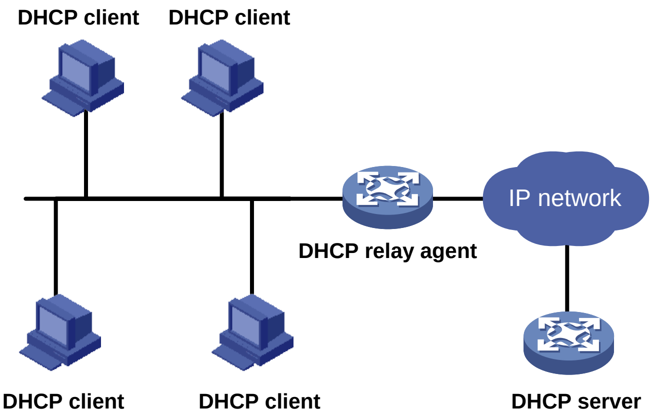

DHCP snooping

As shown in Figure 12, configure DHCP snooping on the Layer 2 switch between the DHCP clients and the DHCP server. After receiving a DHCP request from a client, the DHCP snooping device adds Option 82 to the request message, so that the DHCP server can assign an IP address based on the location of the DHCP client.

By default, all ports on a DHCP snooping device are untrusted ports. You need to configure the port connected to the DHCP server as a trusted port for forwarding DHCP response messages to the DHCP clients.

Figure 12 Application of Option 82

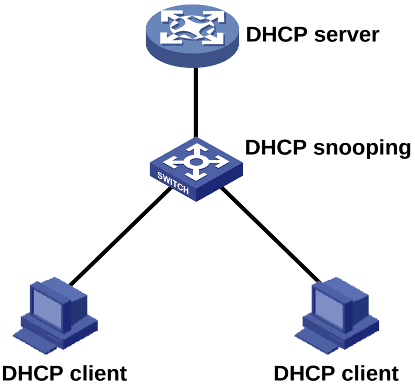

Autoconfiguration for automated device onboarding

When a DHCP client starts up without a configuration file, it obtains an IP address and a TFTP server address from the DHCP server, and then obtains a configuration file from the TFTP server.

Figure 13 Autoconfiguration application

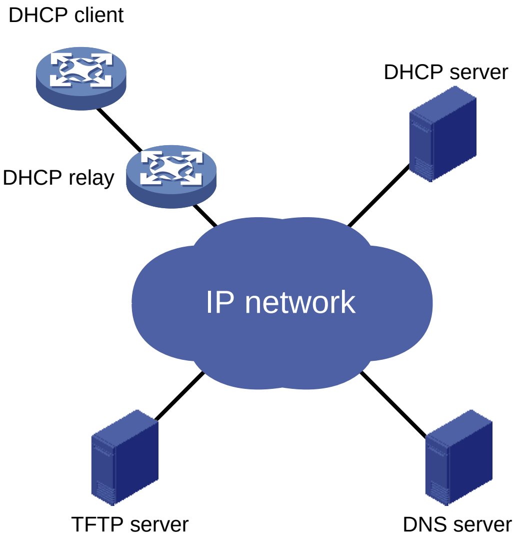

DHCP server, relay, and snooping deployment

As shown in Figure 14, deploy DHCP servers and a DHCP relay agent to provide services for DHCP clients on different subnets. Configure DHCP snooping on the client-attached Layer 2 switches to enhance security.

Figure 14 DHCP comprehensive configuration

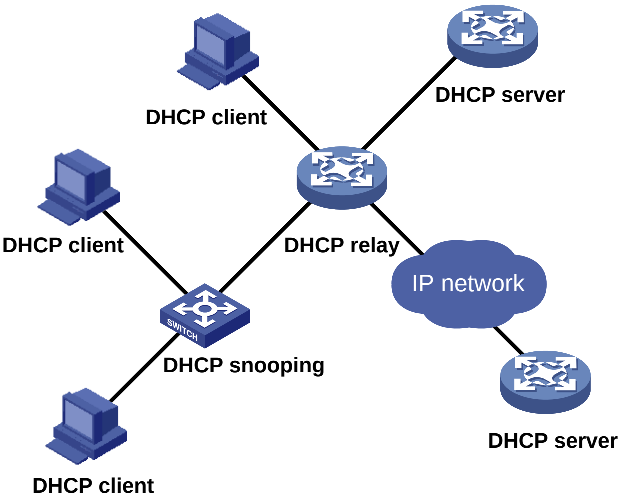

References

- RFC 951, *BOOTSTRAP PROTOCOL(BOOTP)*
- RFC 1497, *BOOTP Vendor Information Extensions*
- RFC 2131, *Dynamic Host Configuration Protocol(DHCP)*
- RFC 2132, *DHCP Options and BOOTP Vendor Extensions*
- RFC 3046, *DHCP Relay Agent Information Option*