DHCP

This lecture describes the Dynamic Host Configuration Protocol, a protocol used by hosts to request information they can use to configure their network interfaces.
After this lecture you should be able to: explain: (a) the need for dynamic assignment of IP addresses, (b) the purpose of address leases, and (c) the reason for avoiding unnecessary IP addresses changes; state the DHCP message that would be transmitted in response to another DHCP message, in case of imminent lease expiry or when leaving a network; describe how a client and server can verify that an address is not in use; and encode an arbitrary option and value as a type-length-value sequence.

Introduction

IP addresses were originally allocated manually to specific hosts. This becomes impractical when hosts are constantly being added and removed from a network.

DHCP (Dynamic Host Configuration Protocol) is a protocol that allows hosts to configure themselves. A host (the "client") requests configuration information from a server and uses that information to configure its network interfaces.

Exercise 1: What configuration information might be required to configure a network interface?

DHCP is an extension to an earlier, simpler, protocol called BOOTP (bootstrap protocol) and uses the same packet format. It is backwards compatible in the sense that a DHCP server can respond to BOOTP requests. DHCP extends the earlier BOOTP protocol in that it can "lease" addresses for a limited time (e.g. 4 hours) and allows a wider range of configuration parameters to the specified.

Although DHCP will work across subnets by making use of BOOTP "relay" hosts, DHCP servers are often located in the router that acts as the gateway for a network.

Dynamic Address Allocation

Although DHCP servers can be configured to always supply the same IP address to a host, most often addresses are allocated out of a "pool" of the unallocated addresses in a subnet (IP network). This is called dynamic address allocation.

Dynamic assignment of IP addresses helps conserve addresses. In many cases hosts may be connected to a network intermittently (e.g. at an airport or Internet cafe) and the DHCP server would soon run out of IP addresses if it permanently assigned one IP address to each host. Instead, addresses are "leased" to hosts for a fixed amount of time and must be renewed before the lease expires. If the lease is not renewed then the server can return the address to the pool for possible re-assignment to another host.

Dynamic allocation of addresses also simplifies the use of networking equipment because it can be done without user interaction. For example, unplugging an Ethernet cable and plugging it back in can alert the machine that it may have been moved to a different network. It can then request an IP address suitable for the new network.

The ease of dynamic address allocation also makes it practical to allow a host to roam from one network to another. When a host detects it has moved from one physical network to another (e.g. between Wireless LANs) it can request a new address.

On the other hand, network connections are often long-lived (file downloads, persistent connections to mail servers, etc). If the client's IP address changes then these connections will be broken. It is therefore desirable that addresses not change more often than necessary. The DHCP protocol ensures that addresses do not change more often than necessary by (1) allowing a client to request assignment or renewal of a specific address, (2) requiring the server to remember the address assigned to each client and reuse it when possible.

A client attempts to renew its lease before the lease time runs out, typically about half-way through the lease.

Unique IP Addresses

Of course, each host should have a unique IP address. DHCP servers and clients take steps to ensure that the same IP address is not assigned to two hosts at the same time.

A client can check to see if another machine on its subnet has the same IP address by issuing an ARP request for that address. If it receives a response with a MAC address that is not its own then it knows that another host is still using that IP address.

A server can also check to see if an IP address is in use by pinging the IP address before allocating it. If a response is received then the address is still in use.

DHCP Packet Format

The DHCP protocol is implemented by exchanging packets between the client and server over the UDP (User Datagram Protocol) transport protocol. Packets for the server are sent to port 67 while packets to the client are sent on port 68.

Figure 1, from RFC 2131, shows the format of the DHCP protocol messages. The DHCP packet has a fixed initial header followed by a variable-length portion. The fixed-length header consists of the following fields:

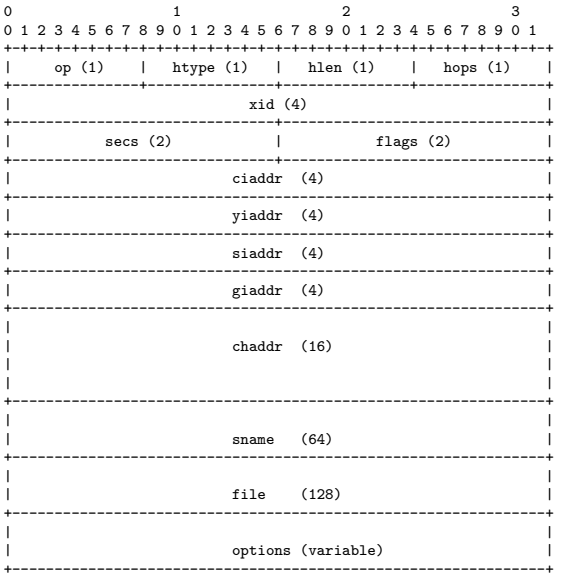

Figure 1: Format of a DHCP message

- op: BOOTP protocol operation (1=request, 2=reply)
- xid: exchange ID; a random number used to match responses with requests
- secs: seconds since request/renewal process began
- flags: a bit is used to indicate that all messages must be broadcast
- ciaddr: client's IP address (set by client if renewing)
- yiaddr: "your" IP address (sent by server)
- siaddr: server IP address (typically DHCP server)
- giaddr: gateway IP address (zero unless a relay is being used)
- chaddr: client hardware address (in case server allocates based on MAC address)
- sname/file: optional server and boot file name (often used for additional options)
- options: additional configuration information (including a magic cookie at the start, the type of DHCP message and an end marker)

  For example:

  (i) option 52 is used to indicate that the sname and/or file fields contain additional options.

  (ii) option 53 is used to indicate the type of DHCP message (e.g. 1 = DHCPDISCOVER, 2=DHCPOFFER, etc)

  (iii) option 43 is used for vendor-specific information (e.g. Microsoft). These options are also encoded in TLV form and are terminated with a 255 (End) code.

TLV Fields

DHCP allows different configuration values to be provided to the host. The intent is that everything required for a host to become active on the internet should be provided. A set of standard configuration

items is defined in the RFC 1533. Additional vendor-specific values can also be sent.

Many protocols encode information in the TLV (Type, Length, Value) format. This consists of a two fixed-length values (Type and Length) and one variable-length value. The first value (Type) is a number that defines the type of information (IP address, netmask, etc). The second value (Length) defines the length of the value field (typically, the number of bytes). The third field consists of the number of bytes specified in the second field and contain the value of the information element.

The advantage of TLV encoding is that a sequence of TLV fields can be parsed without interpreting each field. This allows new types to be added to the protocol without having to change existing implementations ("backwards compatibility").

The size of the type and length fields are set by the expected number of types and the maximum length of the value field,. DHCP TLV encoding uses 1-byte Type and and 1-byte length fields.

To save space, DHCP uses a different format for a few of the most commonly used Types ("tags"). These include Pad (Tag 0) and End (Tag 255).

Exercise 2: What is the maximum possible number of DHCP options? What is the maximum length of a DHCP option value?

Exercise 3: The Host Name DHCP option is encoded in TLV format using a Type of 0x0c and a value of "1234". How many bytes are required to transmit the host name? What are the values of the bytes?

DHCP Protocol

Figure 3 below is taken from from RFC 2131 and shows the sequence of messages exchanged between a DHCP client and server.

At the start of the process the client does not have an IP address so it uses a source address of 0. Since the client does not know the address(es) of the DHCP servers it sends request (and discover) packets to the broadcast IP address (255.255.255.255). In some cases the responses can also be sent to the broadcast address.

1. The client sends DHCP Discover message to find DHCP servers. If it already has an IP address it can include it in the message.

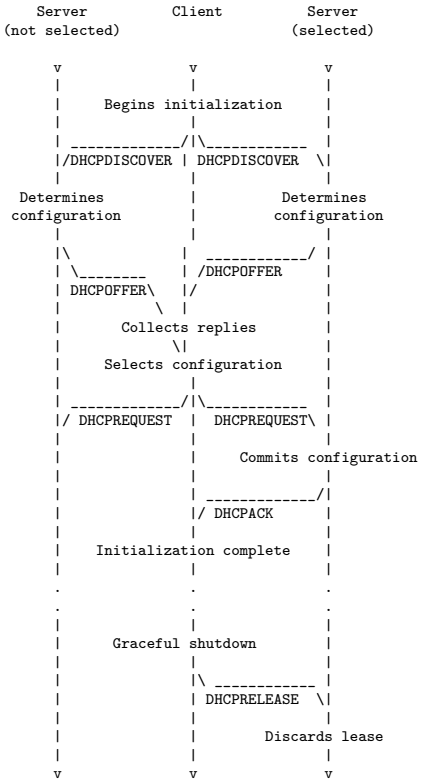

Figure 3: Timeline diagram of messages exchanged between DHCP client and servers when allocating a new network address

2. The server(s) respond(s) with DHCP Offer which include the suggested address.

3. The client chooses one of the offers and responds with a DHCP Request for a specific address.

4. The server acknowledges with a DHCP ACKnowledge message.

If the client no longer needs the address it can return it to the pool with a DHCP Release message.

Exercise 4: Why can there be more than one response to a single DHCPDISCOVER packet? According to the diagram above, what is the response to a DHCPREQUEST packet?

If the client discovers an address is already in use it can decline an offer by sending a DHCP Decline message. Similarly a server can deny a client request by sending DHCP NACK packet.

Lease times are measured relative to message transmission (in seconds as 32 bit values) so they can be quite long. If the lease time is set to the maximum value the lease is considered permanent.

Exercise 5: What is the longest non-permanent lease?