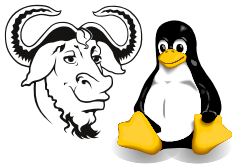

DHCP and `tcpdump`

1 Background

The format of DHCP packets was established with RFC 951 for the *bootstrap protocol*, or *bootp*. DHCP was made to be backwardly compatible with the bootp protocol so that the infrastructure of bootp relay agents on routers would not need to be replaced. The DHCP extensions to bootp are bootp *options*. Figure 1 shows the arrangement of the fields in the DHCP message, while table 1 on the next page gives a description of each one.

See *The DHCP Handbook* (Second Edition), Ralph Droms and Ted Lemon, Sams, October 2002. A copy is available in the library Reserved Collection for a one day loan period, call number TK 5105.585 .D766 2003. There is also a second copy on the shelves. Also RFC 2131 and 2132 are quite well written, and are quite easy to read.

0 7	8 15	16 23	24 31
op	htype	hlen	hops
xid			
secs		flags	
ciaddr			
yiaddr			
siaddr			
giaddr			
chaddr (16 bytes)			
sname (64 bytes)			
file (128 bytes)			
options (variable size)			

Figure 1: The fields in the fixed-format section of a DHCP message.

2 `tcpdump` and DHCP

The manual page for the current version of `tcpdump` (version 3.7.1; an RPM is available from our server) unfortunately does not explain the detail of all the fields in the DHCP protocol. To understand them all, it is necessary to look at the source. Here is my summary after reading `~/RPM/BUILD/tcpdump-3.7.1/tcpdump-3.7.1/print-bootp.c`. I have put a copy of this source file on the subject site.

Table 2 on page 3 shows how `tcpdump` displays the DHCP fields listed in table 1 on the following page. Table 3 on page 4 shows how `tcpdump` shows the DHCP options. Note that many of these are essential for DHCP, for example, the DHCP message type, which is optional only for the old bootp protocol.

Field	Description
`op`	Message operation code: 1 in message from client, 2 in message from server
`htype`	Link-layer address type from RFC 1700. For Ethernet, `htype` is 1.
`hlen`	Link-layer address length, in bytes. (number of bytes in `chaddr` field)
`hops`	Number of relay agents that have forwarded this message.
`xid`	*Transaction identifier*; used by clients to match responses from servers with previously transmitted requests.
`secs`	Number of seconds since client began DHCP transaction
`flags`	Least significant bit is set to 1 to indicate messages to client must be broadcast
`ciaddr`	Client's IP address, set by client after reaches `BOUND` state (i.e., address is valid)
`yiaddr`	Client's IP address, set by server to inform client of its address ("your" IP address)
`siaddr`	IP address of the next server for the client to use (i.e., for the client to download an operating system kernel using `tftp`)
`giaddr`	Relay agent (or "gateway") IP address: relay agent fills this in with the address of the interface through which it received the DHCP message
`chaddr`	Client's link layer address (i.e., on our LAN, the Ethernet address)
`sname`	Name of the next server for client to use in the configuration process
`file`	filename the client should request from the next server (i.e., an operating system kernel, or kickstart file)

Table 1: DHCP Message fields; see figure 1 on the preceding page for the arrangement of these fields in a DHCP message.

Some other information will be provided by `tcpdump` that is not directly concerned with DHCP: for example, a packet with the IP don't fragment flag is marked with a trailing `(DF)`.

Field	Format in tcpdump	Short Description
htype	htype-#⟨*length*⟩	length of link-layer address, bytes
hops	hops:⟨*hops*⟩	number of relay agents
xid	xid:0x⟨*32-bit hex ID*⟩	transaction ID
secs	secs:⟨*seconds*⟩	seconds since session started
flags	flags:0x⟨*hex digits*⟩	LSb is broadcast flag
ciaddr	C:⟨*IP address*⟩	Client's IP address
yiaddr	Y:⟨*IP address*⟩	'your' IP address (bootp client)
siaddr	S:⟨*IP address*⟩	Server's IP address
giaddr	G:⟨*IP address*⟩	Gateway's IP address
chaddr	ether ⟨*MAC address*⟩	Ethernet address
sname	sname "⟨*servername*⟩"	name of next server
file	file "⟨*filename*⟩"	file name to download

Table 2: How tcpdump represents some of the fixed DHCP fields. See table 1 on the preceding page for more details of each field.

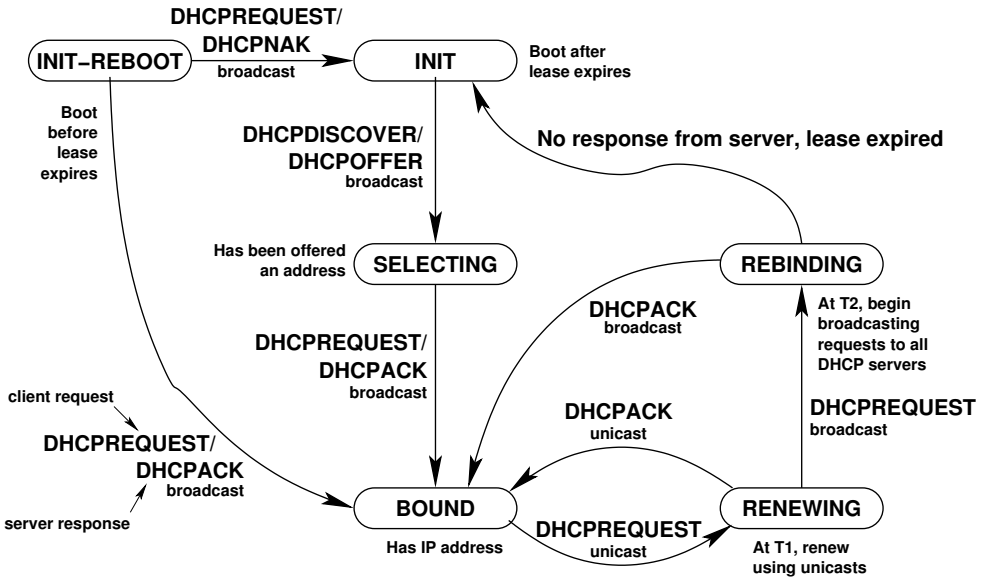

Figure 2: A state diagram showing states of a DHCP client. Note that T is the lease time, $T1 = \frac{T}{2}$, $T2 = \frac{7T}{8}$. See also table 4 on page 5 from the DHCP RFC 2131 (available in full at `/home/nfs/ietf/rcf/rfc2131.txt`), which sumarises DHCP messages.

Format in tcpdump	**Short Description**
SM:⟨*dotted quad IP*⟩	Subnet mask (as an IP address)
CID:⟨*client ID*⟩	Client ID; may be an Ethernet address, or an identifier string provided by client. Examples: CID:"cisco-0008.e3aa.3ac0-VL1"[len 25] and one with an Ethernet client ID: CID:[ether]00:08:02:40:4e:c5
SID:⟨*name or IP*⟩	Server ID
DG:⟨*name or IP*⟩	Default gateway, IP address
NTP:⟨*name or IP*⟩	Network Time Protocol server, IP address
NS:⟨*server*⟩,...	Name servers, IP addresses
HN:"⟨*host name*⟩"	Host name
DN:"⟨*domain name*⟩"	Domain name
VC:"⟨*class*⟩"	Vendor Class (variable length ASCII string). Some examples: VC:"Linux 2.4.18-3 i686", VC:"Linux 2.4.18-6mdk i686", VC:"MSFT 98", VC:"MSFT 5.0", VC:"Hewlett-Packard JetDirect"
PR:⟨*option*⟩+⟨*option*⟩...	Parameter Request—for the parameters that are listed in the request
WNS:⟨*name or IP*⟩,...	WINS (NETBIOS) name server, IP address
WNT	NETBIOS node
WSC	NETBIOS scope, ASCII string
RD	Perform Router Discovery, binary value
SR	Static Route, a list of IP address pairs: address of destination, address of router. But useless in CIDR
VO	Vendor Options — period-separated decimal bytes (variable length)
MSZ:⟨*integer*⟩	Maximum Message size (16 bit short integer)
FQDN	Fully-qualified domain name; a request from client to server to use a particular FQDN. Server only responds to this, and does not send unless requested by client. Format is: first byte is flags, used to indicate state of negotiation. Actual name begins at the fourth byte.
LT:⟨*seconds*⟩	Lease time
RN:⟨*seconds*⟩	Renewal time ($T1$)
RB:⟨*seconds*⟩	Rebinding time ($T2$)

Table 3: How tcpdump represents various DHCP options.

Message		Use	tcpdump
DHCPDISCOVER	—	Client broadcast to locate available servers.	DHCP:DISCOVER
DHCPOFFER	—	Server to client in response to DHCPDISCOVER with offer of configuration parameters.	DHCP:OFFER
DHCPREQUEST	—	Client message to servers either (a) requesting offered parameters from one server and implicitly declining offers from all others, (b) confirming correctness of previously allocated address after, e.g., system reboot, or (c) extending the lease on a particular network address.	DHCP:REQUEST
DHCPACK	—	Server to client with configuration parameters, including committed network address.	DHCP:ACK
DHCPNAK	—	Server to client indicating client's notion of network address is incorrect (e.g., client has moved to new subnet) or client's lease as expired	DHCP:NACK
DHCPDECLINE	—	Client to server indicating network address is already in use.	DHCP:DECLINE
DHCPRELEASE	—	Client to server relinquishing network address and cancelling remaining lease.	DHCP:RELEASE
DHCPINFORM	—	Client to server, asking only for local configuration parameters; client already has externally configured network address.	DHCP:INFORM

Table 4: DHCP Messages: this is "table 2" from RFC 2131; the RFC is available in full from ictlab at /home/nfs/ietf/rcf/rfc2131.txt.