BOOTP and DHCP

Shivkumar Kalyanaraman
Rensselaer Polytechnic Institute
shivkuma@ecse.rpi.edu
http://www.ecse.rpi.edu/Homepages/shivkuma

- **Bootstrapping (Diskless workstations)**
- **BOOTP**
- **Dynamic address allocation**
- **DHCP**
- **Ref: Chap 16, Doug Comer's TCP/IP book**

Bootstrapping

- **Computer loads a simple boot program. The boot program loads operating system.**
- **On diskless machine, the computer needs to know the network address of the o/s file**
- **It needs to know its *own IP address., subnet mask, IP address of default router, IP address of DNS server***
- **It only knows its h/w address.**

Configuration

- Different nodes have different parameters
- Configuration = Setting the parameters
- Key parameters for IP hosts:
  - IP Address
  - Default router address
  - Subnet mask
  - Name
  - DNS server addresses

Method 1: Reverse ARP (RARP)

- What is the IP addr of a hardware address?
- Need RARP server to respond.
- Once IP address is gotten, it does a "tftp" to get its boot image.
- Design of the RARP server complex unlike ARP which is integrated into TCP/IP host implementations
  - Needs to maintain table for multiple hosts (/etc/ethers)
  - Kernel does not process/parse files => RARP is a user process. But does not run over IP.

RARP (contd)

- RARP cannot use IP
- Needs unique Ethernet frame type (0x8035)
- Works through a filter like BPF or nit_if/nit_pf streams modules (fig: A.1, A.2)
- Multiple RARP servers needed for reliability, but unlike ARP where only one reply is sent, each RARP sends a unicast reply => additional traffic
- Possibility of collision between RARP replies
- RARP servers cannot be consolidated since RARP requests are broadcasts => router cannot forward (relay) RARP requests

Key RARP limitations

- RARP is a user process but works over link layer directly =>
  - RARP server system dependent
  - Needs to interface with link layer driver directly => separate filters and direct access to hardware needed
- Returns only IP address
  - Booting and configuration params not returned even though there is space in packet
  - Host needs ICMP and TFTP to complete booting
- Can't relay RARP requests to a central server.
  - Need RARP server per broadcast domain

Method 2: BOOTP

- Runs over UDP/IP! Issues w/ using UDP/IP:
  - IP software can broadcast (to 255.255.255.255) even if local IP address unknown => client broadcasts BOOTP request
  - Port number 67 for server and 68 for client (not an ephemeral port)
    - Delivers BOOTP reply to BOOTP client and not other UDP apps when reply is broadcast
    - Does not wake up other servers during broadcast reply

BOOTP (contd)

- BOOTP requests/replies sent w/ DF bit set.
- Server can send reply via broadcast or unicast:
  - For unicast reply, BOOTP server knows the IP address, but the link layer address is not in the ARP cache
  - Note that the server cannot send an ARP message because client does not know its IP address
  - Server can use ioctl(8) {or arp -s } to set the value of the cache => can do this only if it has permission

BOOTP features (contd)

- Else send broadcast reply
- Reply: IP Address, Boot Server IP address, Default Router, Boot file name, subnet mask
  - More information, but still only a single packet exchange
  - Client gets boot image using TFTP => booting still a 2-step process

BOOTP features (contd)

- Advantages of using UDP/IP:
  - Bootstrapping can occur across a router via a relaying mechanism
  - BOOTP uses checksum provided by UDP
- Multiple requests/replies
  - Process the first one
  - Client uses a transaction ID field to sort out replies
- Clients responsible for reliability =>
  - Uses timeout, retransmission & exponential backoff
  - Random initial timeout (betn 0 & 4s): simultaneous reboot after power restoration.

BOOTP Message Format

0 31b

Operation	H/W Type	H/W Length	Hops	
Transaction Identifier				
Seconds elapsed		Unused		
Client IP Address				
Your IP Address				
Server IP Address				
Router IP Address				
Client H/W address				16 B
Server Host Name				64 B
Bootfile Name				128 B
Vendor Specific Area				64 B

BOOTP Message (Cont)

- **Operation: 1 = Request, 2 = Reply**
- **H/w type: 1 = Ethernet**
- **H/w Address Length**
- **Hops: Initialized to zero. Incremented by BOOTP relays (routers)**

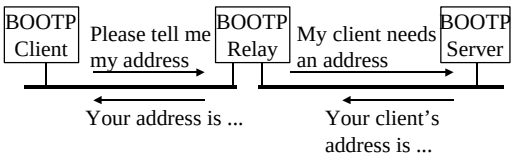

BOOTP Message (Cont)

- **Transaction ID: used to match responses with requests**
- **Seconds = Number of seconds since the client started to boot**
- **If a client knows its IP address, it places it in the Client IP address**
- **If server address/name fields are non-zero in the request, only the indicated host can answer the request**
- **Your IP Address: Clients IP address returned by the server**

BOOTP Message (Cont)

- **Boot File name: Generic name like "unix" in the request. Full name in response.**
- **Vendor specific area: Misnomer. Also used for general purpose info.**
  - **Magic cookie: First 4 octets = 99.130.83.99**
  - **Type-length-value: (Contd on next slide)**

Item	Code	Length
Padding	0	-
Subnet mask	1	4
Time of Day	2	4
End	255	-

Contents of Vendor-Specific Area

Item	Code	Length
Routers	3	4n
Time Server	4	4n
IEN116 Server	5	4n
Domain server	6	4n
Log server	7	4n
Quote server	8	4n
LPR servers	9	4n
Impress servers	10	4n
RLP Server	11	4n
Host name	12	4n
Boot size	13	2
Reserved	128-254	-

Method 3: DHCP

- **BOOTP limitation: cannot dynamically assign IP address**
- **Dynamic Host Configuration Protocol (DHCP)**
  - **BOOTP + Dynamic allocation of IP addresses => compatible with BOOTP. No new fields in header.**
  - **Addresses can be leased for a period. Reallocated to the same or other nodes after lease expiry.**
- **Non-mobile computers can get a permanent address.**

DHCP Message Format

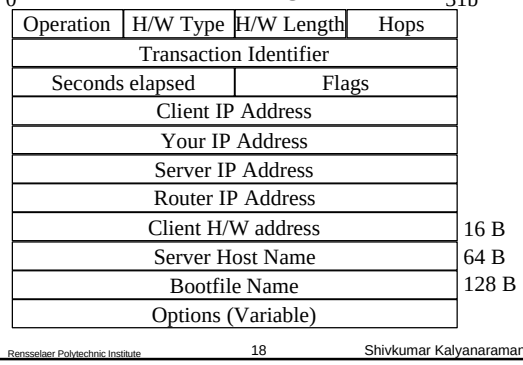

DHCP Message Format

- Slightly modified version of BOOTP message ⇒ A DHCP server can be programmed to answer BOOTP requests
- BOOTP's Unused field renamed to Flags
- Only one bit of 16-bit Flags has been defined
  - Left-most flag bit =1 ⇒ Servers, please reply using IP broadcast address
  - Servers by default send hardware unicast response
- Vendor specific field renamed to options
  - Size increased to 312 bytes (from 64 bytes)
  - Option type 53 specifies the "type of the message"

DHCP (contd)

Type	Meaning
1	DHCP Discover
2	DHCP Offer
3	DHCP Request
4	DHCP Decline
5	DHCP Ack
6	DHCP Nack
7	DHCP Release

- "Option overload"
  - Server Host name and boot file name when unused for their original purpose could be used to code more options

DHCP State Diagram

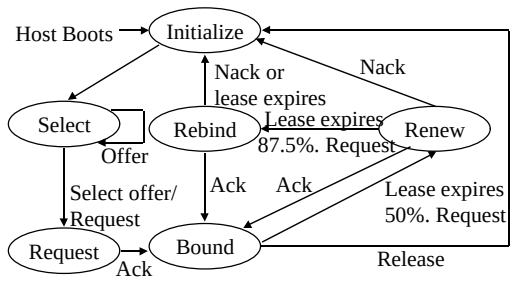

DHCP states

- Boots => INITIALIZE state
- DHCPDISCOVER: broadcast request to servers => SELECT state
- DHCPOFFER (from server) => remain in SELECT
- DHCPREQUEST => select one of the offers and notify server (goto REQUEST state) about the lease
- DHCPACK => server Oks request to lease => go to the BOUND state
- Renewal: after 50% of lease go to RENEW state
- Rebind: after 87.5% of time, if server has not responded, try again and go to REBIND.
- If server NACKs or lease expires, or client sends DCHPRELEASE, go to INITIALIZE, else come back to BOUND state

DHCP: Current Issues

- Interaction with DNS
- Should the *names* also be dynamically leased?
- Should the names be registered on DNS?
  - How to work with a directory service (given a fixed name, find a temporary IP address)?
- Currently there are no protocols for dynamic DNS updates.

Summary

- RARP allows finding an IP address
- BOOTP allows default router, subnet mask, DNS
- DHCP allows dynamic allocation
- DHCP is backward compatible with BOOTP

Initialization: RFCs

- **[RFC1533] S. Alexander, R. Droms, "DHCP Options and BOOTP Vendor Extensions", 10/08/1993, 30 pages.**
- **[RFC1534] R. Droms, "Interoperation Between DHCP and BOOTP", 10/08/1993, 4 pages.**
- **[RFC1541] R. Droms, "**Dynamic Host Configuration Protocol**", 10/27/1993, 39 pages.**
- **[RFC1542] W. Wimer, "Clarifications and Extensions for the Bootstrap Protocol", 10/27/1993, 23 pages.**
- **[RFC0951] W. Croft, J. Gilmore, "**Bootstrap Protocol**", 09/01/1985, 12 pages. (Updated by RFC1532, RFC1395, RFC1497)**
- **[RFC0906] R. Finlayson, "Bootstrap loading using TFTP", 06/01/1984, 4 pages.**