ICMP, ARP & RARP, DHCP

- ARP is the address resolution protocol. It it a protocol by which IP addresses are mapped to corresponding MAC addresses in a LAN environment
- ICMP is the Internet Control Message Protocol which is and "internal" protocol IP uses for sending control and status information. ICMP message types include *unreachable, source quench, time exceeded, redirect, echo request & reply*
- RARP is useful for diskless workstations, needs an IP address,knows its Ethernet address.
- DHCP Dynamic Host Configuration Protocol (IP address allocations)

Internet Control Message Protocol

- The IP (Internet Protocol) relies on several other protocols to perform necessary control and routing functions:
- Control functions (ICMP)
- Multicast signaling (IGMP)
- Setting up routing tables (RIP, OSPF, BGP, PIM, …)

Introduction

- The **Internet Control Message Protocol (ICMP)** is ahelper protocol that supports IP with facility for
  - Error reporting
  - Simple queries
  - ICMP messages are encapsulated as IP datagrams:

Message Format

- **4 byte header:**
  - Type (1 byte): type of ICMP message
  - Code (1 byte): subtype of ICMP message
  - Checksum (2 bytes): similar to IP header checksum.
- Checksum is calculated over entire ICMP message
- If there is no additional data, there are 4 bytes set to zero.
- Each ICMP messages is at least 8 bytes long

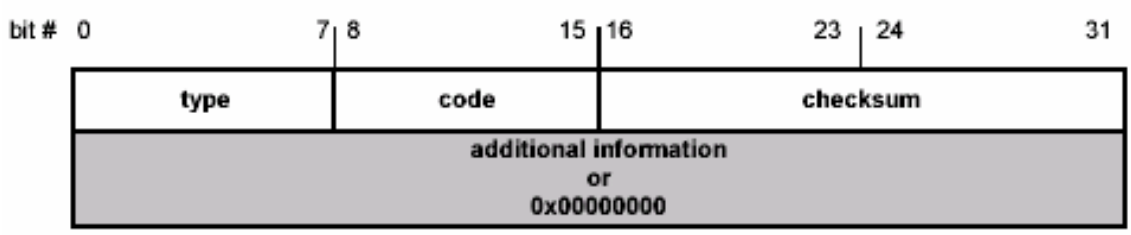

ICMP Query Message

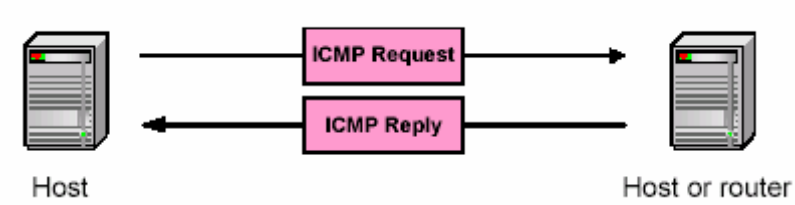

ICMP query:

- ICMP Request sent by host to a router or host
- ICMP Reply sent back to querying host

Example of ICMP Queries

Type/Code:	Description
8/0	Echo Request
0/0	Echo Reply
13/0	Timestamp Request
14/0	Timestamp Reply
10/0	Router Solicitation
9/0	Router Advertisement

The ping command uses Echo Request/ Echo Reply

Example Query: Echo Request and Reply

- Pings are handled directly by the kernel
- Each Ping is translated into an **ICMP Echo Request**
- The Pinged host responds with an **ICMP Echo Reply**

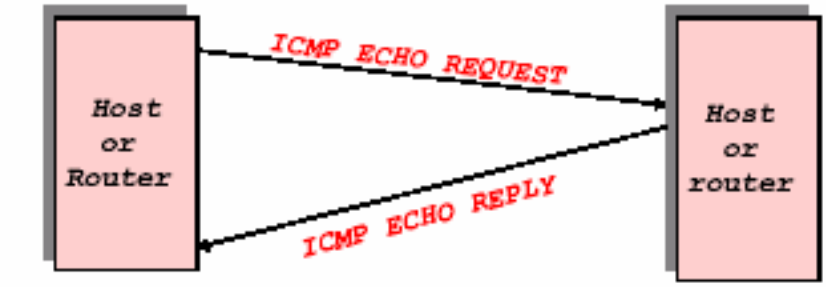

•The data portion of the request can be padded out to any size and is replicated in the reply

–Useful for testing MTU and/or fragmentation

–cf. “Ping of Death” DoS attack

Another Example: ICMP Timestamp

- A system (host or router) asks another system for the current time.
- Time is measured in milliseconds after midnight UTC (Universal Coordinated Time) of the current day
- Sender sends a request, receiver responds with reply

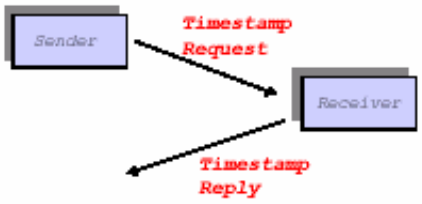

Type (= 17 or 18)	Code (=0)	Checksum
identifier		sequence number
32-bit sender timestamp		
32-bit receive timestamp		
32-bit transmit timestamp		

ICMP Error Message

- **ICMP error messages report error conditions**
- **Typically sent when a datagram is discarded**
- **Error message is often passed from ICMP to the application program**

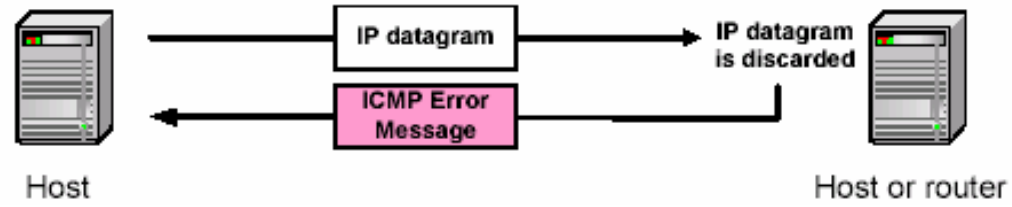

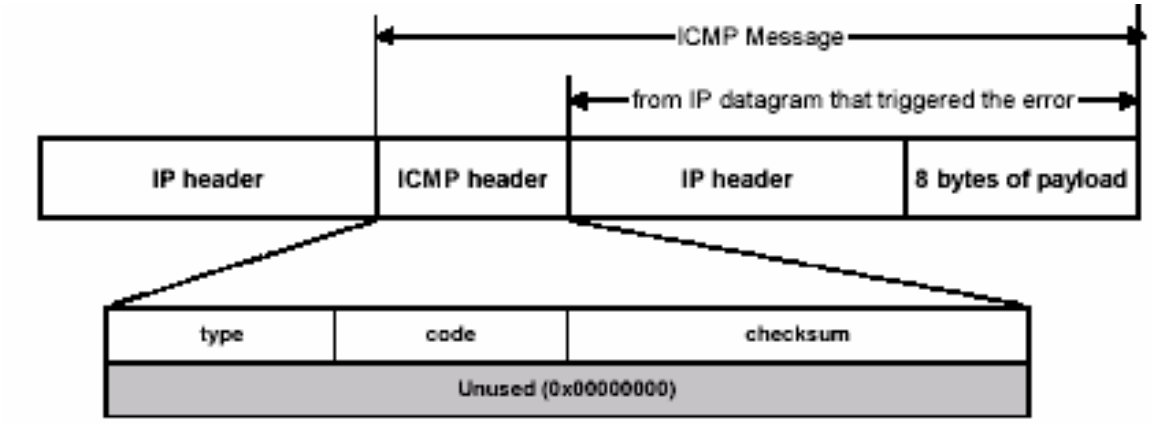

- ICMP error messages include the complete IP header and the first 8 bytes of the payload (typically: UDP, TCP)

Common ICMP Error Messages

Used by the “traceroute” utility to map the path through the IP network to a particular destination

Type	Code	Description	
3	0–15	Destination unreachable	Notification that an IP datagram could not be forwarded and was dropped. The code field contains an explanation.
5	0–3	Redirect	Informs about an alternative route for the datagram and should result in a routing table update. The code field explains the reason for the route change.
11	0, 1	Time exceeded	Sent when the TTL field has reached zero (Code 0) or when there is a timeout for the reassembly of segments (Code 1)
12	0, 1	Parameter problem	Sent when the IP header is invalid (Code 0) or when an IP header option is missing (Code 1)

Destination Unreachable

Code	Description	Reason for Sending
0	Network Unreachable	No routing table entry is available for the destination network.
1	Host Unreachable	Destination host should be directly reachable, but does not respond to ARP Requests.
2	Protocol Unreachable	The protocol in the protocol field of the IP header is not supported at the destination.
3	Port Unreachable	The transport protocol at the destination host cannot pass the datagram to an application.
4	Fragmentation Needed and DF Bit Set	IP datagram must be fragmented, but the DF bit in the IP header is set.

Example:
ICMP Port Unreachable

- RFC 792: If, in the destination host, the IP module cannot deliver the datagram because the indicated protocol module or process port is not active, the destination host may send a destination unreachable message to the source host.

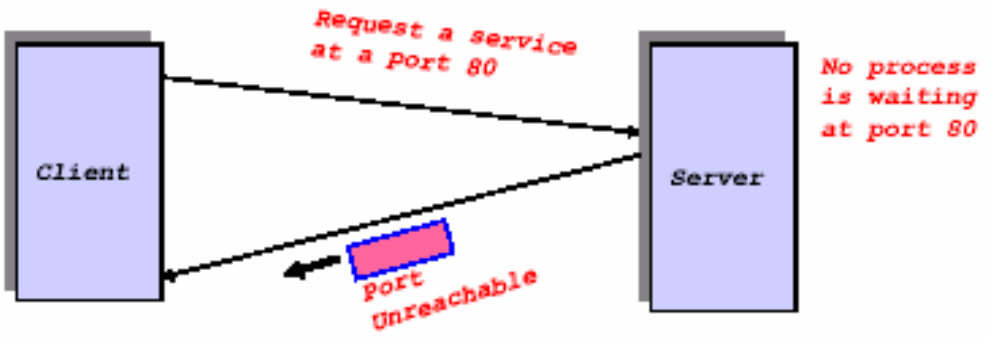

ARP

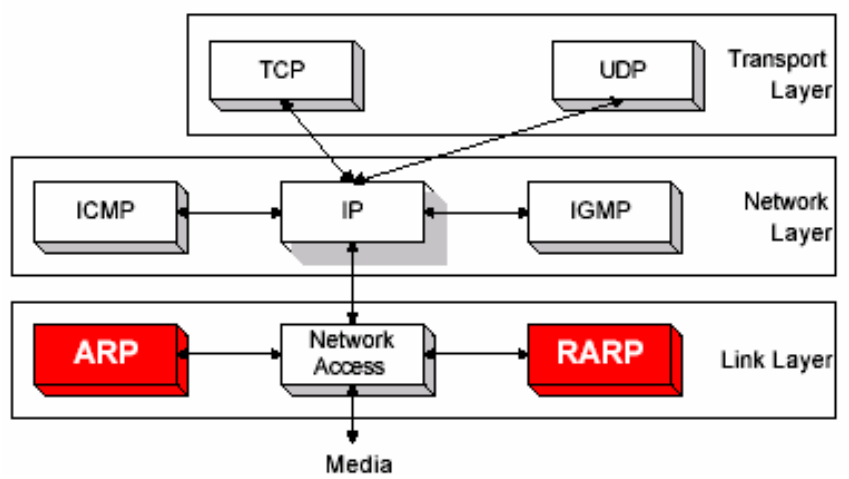

- The Internet is based on IP addresses
- Data link protocols (Ethernet, FDDI, ATM) may have different (MAC) addresses
- The ARP and RARP protocols perform the translation between IP addresses and MAC layer addresses

Address Resolution Protocol

- ARP performs a lookup service that finds a MAC address for a given IP address.
- A system that needs a MAC address for a given IP address broadcasts a query which contains the IP address to all systems on the network.
- If a system receives the query and the IP address in the message matches its own IP address, it sends its MAC address to the sender of the query.
- IP and MAC addresses are the usual but not the only formats available to ARP

Operation of ARP

- Each host maintains a table, the *ARP cache*, temporarily stores the results from previous address resolutions.
- ARP Request is broadcast to all systems on the network.
- In Ethernet, a frame is broadcast when the destination MAC address is set to broadcast address ff:ff:ff:ff:ff:ff.
- A broadcast frame is received and processed by all hosts.
- If a system receives the ARP request and the IP address in the message matches its own IP address, it issues an ARP Reply message to the sender of the query.

Gratuitous ARP

- Every host that sees an ARP Request verifies its ARP cache checking for the sender IP of the ARP Request.
- If such an entry exists, it updates the MAC address with the address in the ARP Request.
- Since ARP Requests are broadcast message, these updates are made by all systems each time an ARP Request is transmitted on the network.
- This feature is exploited in a concept that is called *gratuitous ARP*.

ARP Vulnerabilities

- ARP may be used to redirect traffic intended for a certain IP address to another on the NW
- Broadcast ARP replies with invalid MAC addresses insert incorrect entries into ARP caches.

ARP Packet Formats

Hardware type (2 bytes)		Protocol type (2 bytes)
Hardware address length (1 byte)	Protocol address length (1 byte)	Operation code (2 bytes)
Source hardware address*		
Source protocol address*		
Target hardware address*		
Target protocol address*		

* Depends on length of Datalink and Network layer addresses

Note: ARP Packet Formats - I

- Ethernet carries ARP, with type set to 0x8060
- ARP message, in IP and Ethernet scenario is 28 bytes (48 bit MAC + 32 bit IP)
- Hardware type is datalink protocol
  - Ethernet = 0x0001, 802 = 0x0006
- Protocol type field is the network layer used
  - IP = 0x8000
- Operation code is 0x0001 for ARP requests and 0x0002 for ARP replies
- Hardware address length and Protocol address length specify length of addresses (MAC-6, IP-4).

Note: ARP Packet Formats - II

- The next four fields contain the hardware address and the network address of the sender and the intended receiver of the ARP packet.
- The former is referred to as the source and the latter is referred to as the target.

An ARP, the Router, and My PC

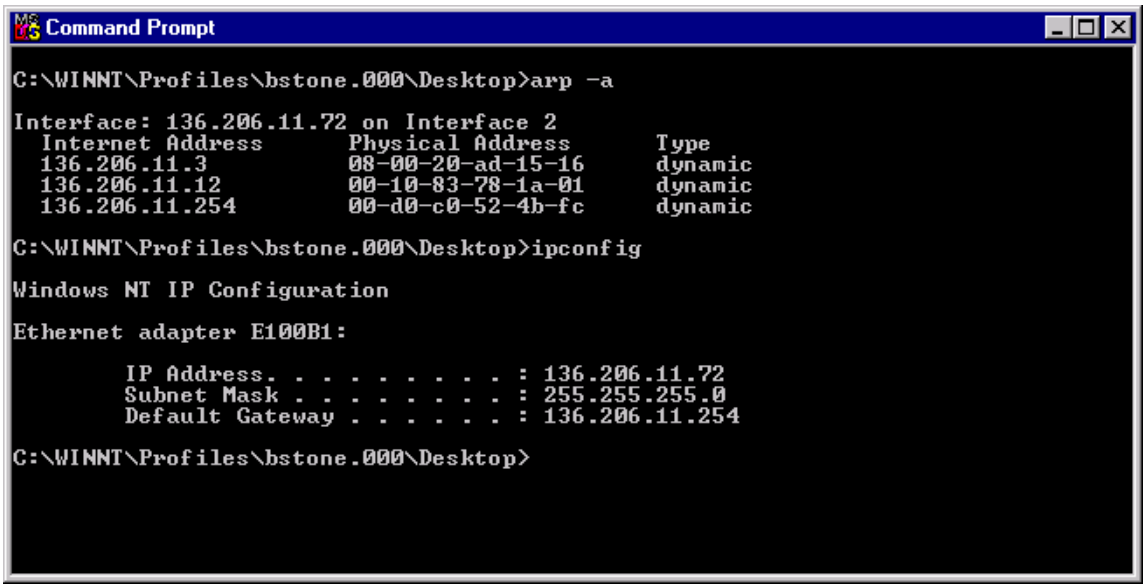

Arp Cache

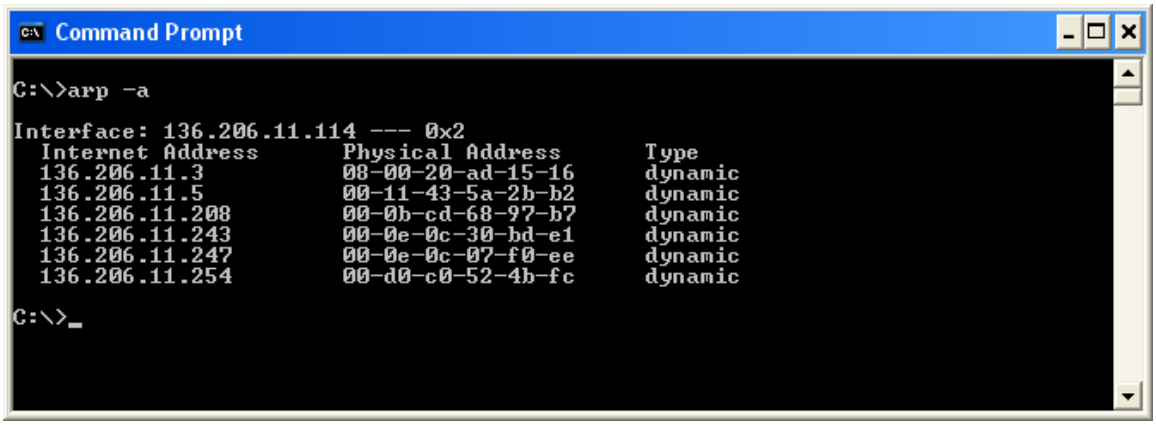

Sample NW and test

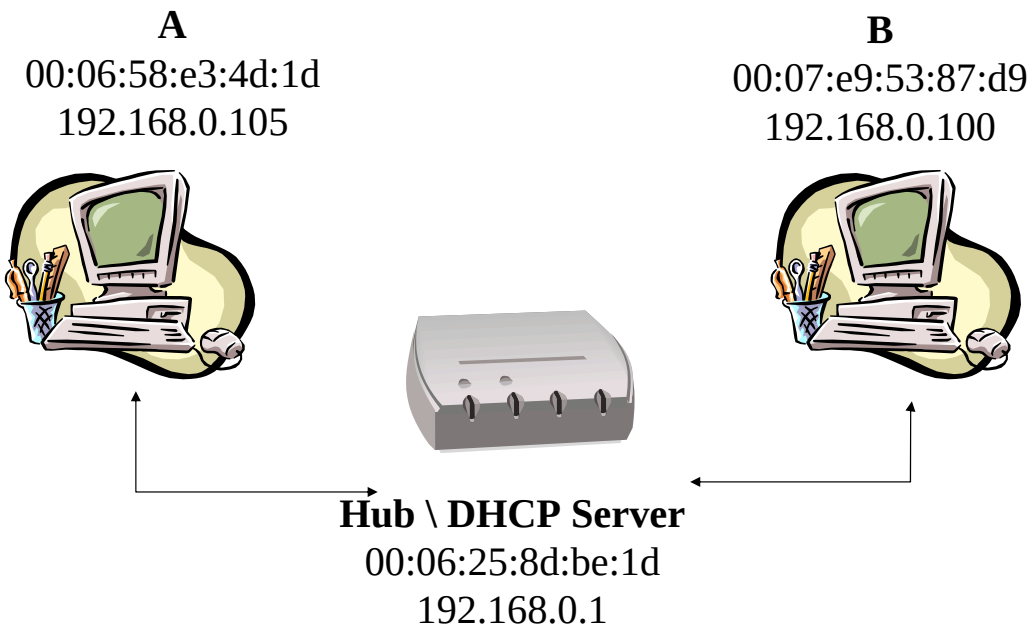

`arp` Tool

- Issue `arp` command
- Tells you how to use it
- Issue `arp –a` it dumps its cache

```
>arp –a
Interface 192.168.0.105 --- 0x10004
  Internet Address      Physical Address      Type
  192.168.0.1           00-06-25-8d-be-1d     dynamic
  192.168.0.100         00:07:e9:53:87:d9     dynamic
```


Make Data

- Issue `ping –n 1 192.168.0.100` from A
- Machine A sends a request message to B
- Check out **`arp.cap`** for results
- Note the source and dstination addreses used in this trace.

Make More data

- Now delete the arp cache with
  - `arp –d 192.168.0.100`
- Do second `ping –n 1 192.168.0.100`
- A issues arp request in packet 3 (broadcast addr)
- B replies in packet 4, replenishing arp cache of A and allowing it to issue a ping request in packet 5
- Finally issue third
  - `ping –n 1 192.168.0.100`
- Results are in packet 7 and 8

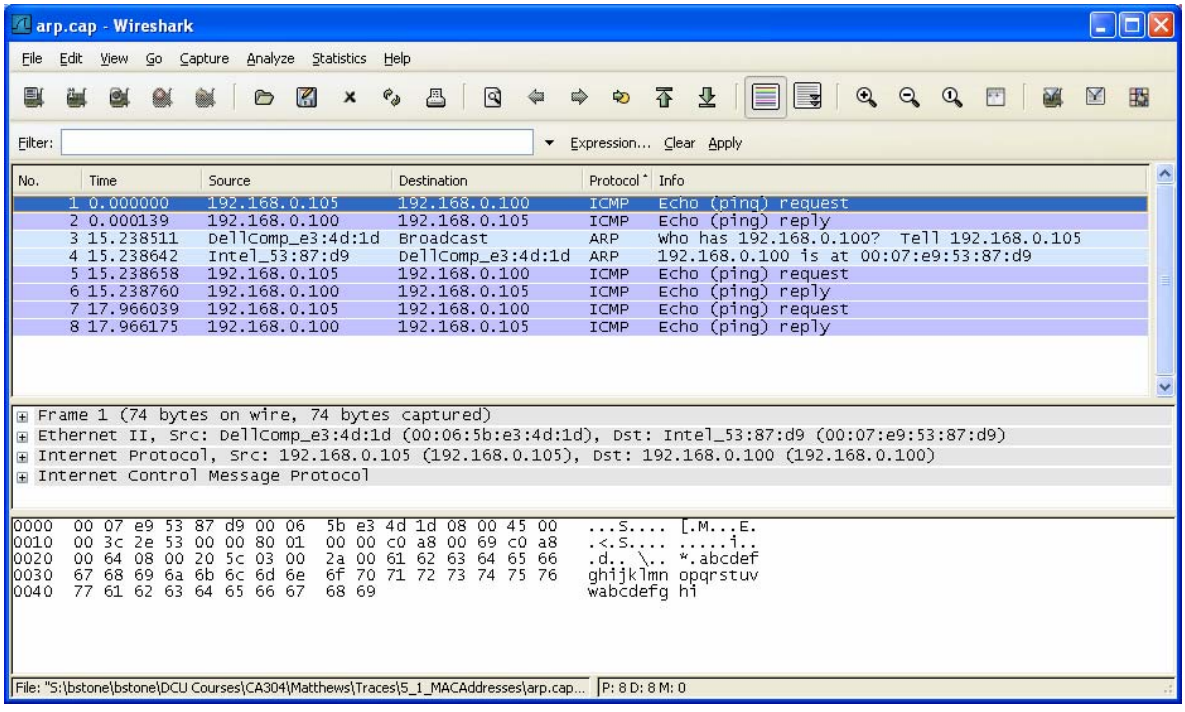

ARP Exercise

- Replicate the previous experiments on machines in the lab (ground floor is equipped with Wireshark).
- Use `arp` and `ipconfig` to find out the IP and MAC addresses of the machines, clear the caches etc as done in the experiments. You will need 2 machines to do this.
- Something new:
  - Check out the CRC calculations in the frames and account for any discrepencies.
- Use appropriate filters in Wireshark to limit captured traffic to that of interest for the experiment.
- Save your traces in Wireshark to a file.

Proxy ARP

- Proxy ARP is a configuration option for IP routers, where an IP router responds to ARP
- Request that arrive from one of its connected networks for a host that is on another of its connected networks.
- Without Proxy ARP enabled, an ARP Request for a host on a different network is unsuccessful, since routers do not forward ARP packets to another network.

RARP

- Given an Ethernet address, what is the IP?
- RFC 903 – RARP sovves this problem – Broadcasts MAC gets back IP from RARP server.
- Broadcast address stays within 1 domain (router)
- Needs to get further or else have 1 RARP server in each MAC broadcast domain.
- Solution – use BOOTP

BOOTP

- RFCs 951, 1048, 1084
- Use UDP messages, broadcasts forwarded over routers!
- Also provides
  - info on IP of file server with disk image
  - IP address of default router
  - Subnet mask
- Problem: Manual config of IP – MAC, gives rise to errors.

DHCP

- Allows manual IP assignment & auto assignment
- Replaces RARP & BOOTP
- Uses RARP server, not necessarilly on same LAN
- DHCP relay agent exists on each LAN

DHCP On a LAN

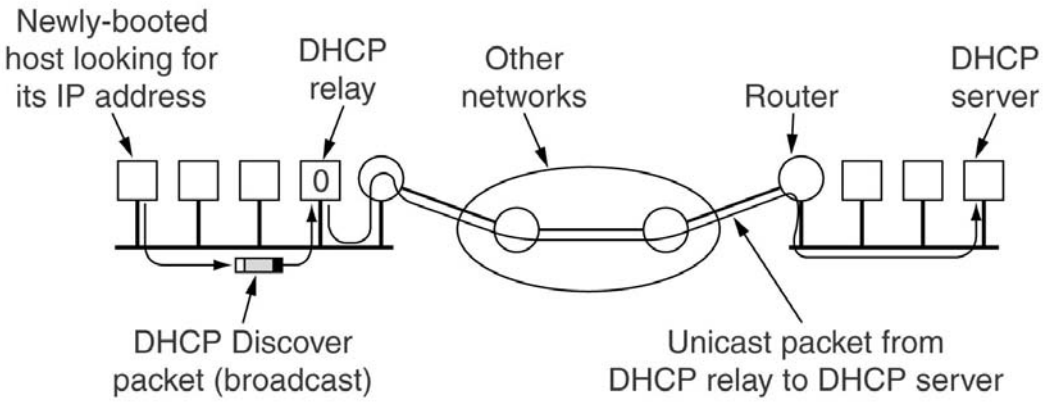

DHCP Operation

- Broadcast DHCP DISCOVER packet, relay agent unicasts to server if not on same LAN.
- Relay agent needs only IP address of DHCP server, possibly on remote LAN.
- Question: How long should IP addres be allocated?
  - Answer: Leasing and renewals.
  - If host fails to attempt to renew lease just before it expires, IP address is withdrawn when lease expires.

DHCP Experiment

- First we opened a CMP window and scrubbed all IP address from machine and reestablished them while running *Wireshark*
  - **`ipconfig /release`**
  - **`ipconfig /renew`**
  - **`ipconfig /renew`**
  - **`ipconfig /release`**
- Results were stored in **dhcp_isolated.cap**

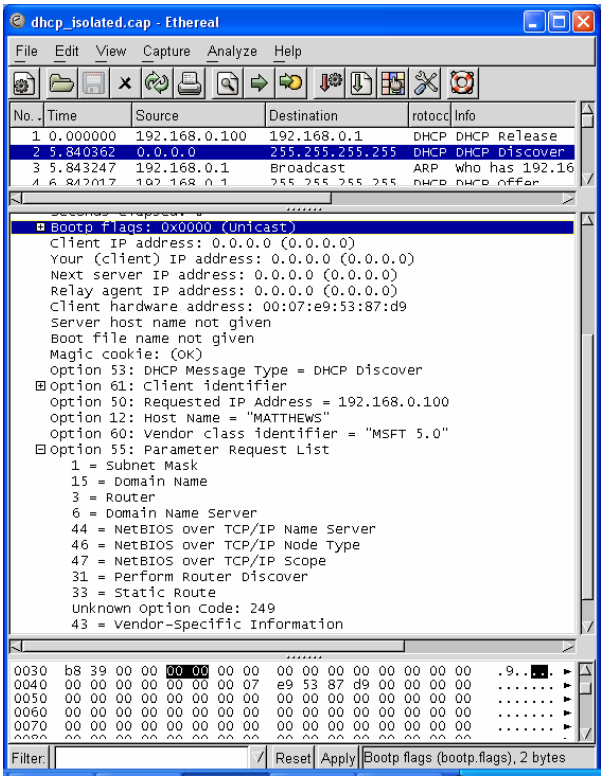

The Packets

- Packet 1 releases current IP address
  - [192.168.0.100]
- Packet 2 is the *DHCP DISCOVERY* message has no IP address so uses [0.0.0.0]. Destination IP is [255.255.255.255] broadcast
- Client includes info about itself
  - Hostname (Matthews)
  - MAC address
  - This may lead server to recognise a known client and enforce policy (same or different IP as before, force registration of MAC etc)

DHCP Discovery (cont.)

- Client lists info required
  - Addres of local router
  - Subnet mask
  - Domain name
- Server responds with *DHCP OFFER* msg
  - Broadcast so that IPless station will read it
  - Contains IP address, local router, subnet mask, domain name & local domain name server

DHCP Discovery (cont.)

- Client in packet 5 indicated acceptance of address by echoing *DHCP REQUEST* with same information.
- There are ARP messages in packets 3, 7 – 12.
- In packet 3 DHCP server asks if anyone has 192.168.0.100.
- Client does likewise three times! WHY?

DHCP Leases

- IP addresses are leased only, not forever.
- Packets 14 & 15 show process of lease renewal.
- They happened because of our second **`ipconfig /renew`** command
- DHCP ACK includes duration of lease renewal (one day).
- If a lease expires, DHCP server is free to reallocate that IP address.
- The final **`ipconfig /release`** allows DHCP server to reallocate out IP address, thus recycling it.