



# Reti

(già "Reti di Calcolatori")

**Livello Rete**  
**ARP – ICMP - DHCP**

Renato Lo Cigno

<http://disi.unitn.it/locigno/teaching-duties/computer-networks>

- *Credits*
  - *Part of the material is based on slides provided by the following authors*
    - *Jim Kurose, Keith Ross, “Computer Networking: A Top Down Approach,” 4th edition, Addison-Wesley, July 2007*
    - *Douglas Comer, “Computer Networks and Internets,” 5th edition, Prentice Hall*
    - *Behrouz A. Forouzan, Sophia Chung Fegan, “TCP/IP Protocol Suite,” McGraw-Hill, January 2005*
- *La traduzione, se presente, è in generale opera (e responsabilità) del docente*



- Spazio di indirizzamento
- Indirizzi IP e loro uso
- Consegna dei pacchetti
- **Configurazione dei PC e delle reti**
- Instradamento e Routing



# ARP: ADDRESS RESOLUTION PROTOCOL

Protocollo di supporto a IP per mappare gli indirizzi IP sulle interfacce fisiche, ovvero sugli indirizzi MAC (Ethernet)




- A crucial step of the forwarding process requires a translation:
  - forwarding uses IP addresses
  - a frame transmitted must contain the MAC address of the next hop
  - IP must translate the next-hop IP address to a MAC address
- The principle is:
  - IP addresses are abstractions
    - provided by protocol software
  - The Data-Link does not know how to locate a computer from its IP address
    - the next-hop address must be translated to an equivalent MAC address



- Translation from a computer's IP address to an equivalent hardware address is known as address resolution
  - And an IP address is said to be resolved to the correct MAC address
- Address resolution is local to a network
  - simple for Point-to-Point connections
  - need a protocol in the **general case** of shared access medium
- A server-based solution introduces delays and a weak point
- Local communications are cheap and often the medium is broadcast
- A “broadcast and select” solution is the one chosen by IETF



- One computer can resolve the address of another computer only if both computers attach to the same physical network
  - Direct delivery
  - A computer never resolves the address of a computer on a remote network
  - Address resolution is always restricted to a single network





- How can a host know if the address to resolve is local?
  - if it is local, the dest. IP address should have the same NetID (prefix) of the source IP address
- What happens if the address is not local?
  - **Indirect delivery**
  - Give the packet to a machine router that is on the way to the destination → next topic
  - Must in any case translate the IP of the Router into its MAC address





- Suppose B needs to resolve the IP address of C
- B broadcasts a request that says:
  - “I’m looking for the MAC address of a computer that has IP address C”
- The broadcast only travels across one network
- An ARP request message reaches all computers on a network
- When C receives a copy of the request it sends a directed reply back to B that says:
  - “I’m the computer with IP address C, and my MAC address is M”



# ARP Message Format

0	8	16	24	31
<b>HARDWARE ADDRESS TYPE</b>		<b>PROTOCOL ADDRESS TYPE</b>		
<b>HADDR LEN</b>	<b>PADDR LEN</b>	<b>OPERATION</b>		
<b>SENDER HADDR (first 4 octets)</b>				
<b>SENDER HADDR (last 2 octets)</b>		<b>SENDER PADDR (first 2 octets)</b>		
<b>SENDER PADDR (last 2 octets)</b>		<b>TARGET HADDR (first 2 octets)</b>		
<b>TARGET HADDR (last 4 octets)</b>				
<b>TARGET PADDR (all 4 octets)</b>				




- **HARDWARE ADDRESS TYPE**
  - 16-bit field that specifies the type of hardware address
  - the value is 1 for Ethernet
- **PROTOCOL ADDRESS TYPE**
  - 16-bit field that specifies the type of protocol address
  - the value is 0x0800 for IPv4
- **HADDR LEN**
  - 8-bit integer that specifies the size of a hardware address in bytes
- **PADDR LEN**
  - 8-bit integer that specifies the size of a protocol address in bytes



- OPERATION
  - 16-bit field that specifies whether the message
    - “request” (1) or “response” (2)
- SENDER HADDR
  - HADDR LEN bytes for the sender's hardware address
- SENDER PADDR
  - PADDR LEN bytes for the sender's protocol address
- TARGET HADDR
  - HADDR LEN bytes for the target's hardware address
- TARGET PADDR
  - PADDR LEN bytes for the target's protocol address



- An ARP message contains fields for two address bindings
  - one binding to the sender
  - other to the intended recipient, ARP calls it **target**
- When a request is sent
  - the sender does not know the target's hardware address (that is the information being requested)
    - field TARGET HADDR in an ARP request is filled with “0”
- In a response
  - the target binding refers to the initial computer that sent the request



- When it travels across a physical network an ARP message is encapsulated in a hardware frame
  - e.g., Ethernet
- An ARP message is treated as data being transported
  - the network does not parse the ARP message or interpret fields



- The **type** field in the frame header specifies that the frame contains an ARP message
- A sender must assign the appropriate value to the type field before transmitting the frame
- A receiver must examine the type field in each incoming frame
- Ethernet uses type field **0x806** to denote an ARP message
- The same value is used for both ARP requests/ responses
  - Frame type does not distinguish between types of ARP messages
  - A receiver must examine the OPERATION field in the message to determine whether an incoming message is a request or a response



- Sending an ARP request for each datagram is inefficient
  - Three frames traverse the network for each datagram
    - an ARP request, ARP response, and the data datagram itself
- Most communications involve a sequence of packets
  - a sender is likely to repeat the exchange many times
- To reduce network traffic
  - ARP software extracts and saves the information from a response
    - so it can be used for subsequent packets
  - The software does not keep the information indefinitely
    - Instead, ARP maintains a small table of bindings in memory





- ARP manages the table as a cache
  - an entry is replaced when a response arrives
  - the oldest entry is removed whenever the table runs out of space or after an entry has not been updated for a long period of time
  - ARP starts by searching the cache when it needs to bind an address
- ARP entries expire after  $\sim 30$ s to avoid sending packets to the wrong destination if the mapping IP-MAC changes



- If the binding is present in the cache
  - ARP uses the binding without transmitting a request
- If the binding is not present in the cache
  - ARP broadcasts a request
  - waits for a response
  - updates the cache
  - send the packet
- The cache is updated when an ARP message arrives
  - either a request or a response
  - since traffic is normally two-way updating the cache on requests reduces overhead

# ICMP: INTERNET CONTROL MESSAGE PROTOCOL

Messaggi di controllo, segnalazione, errore al livello IP



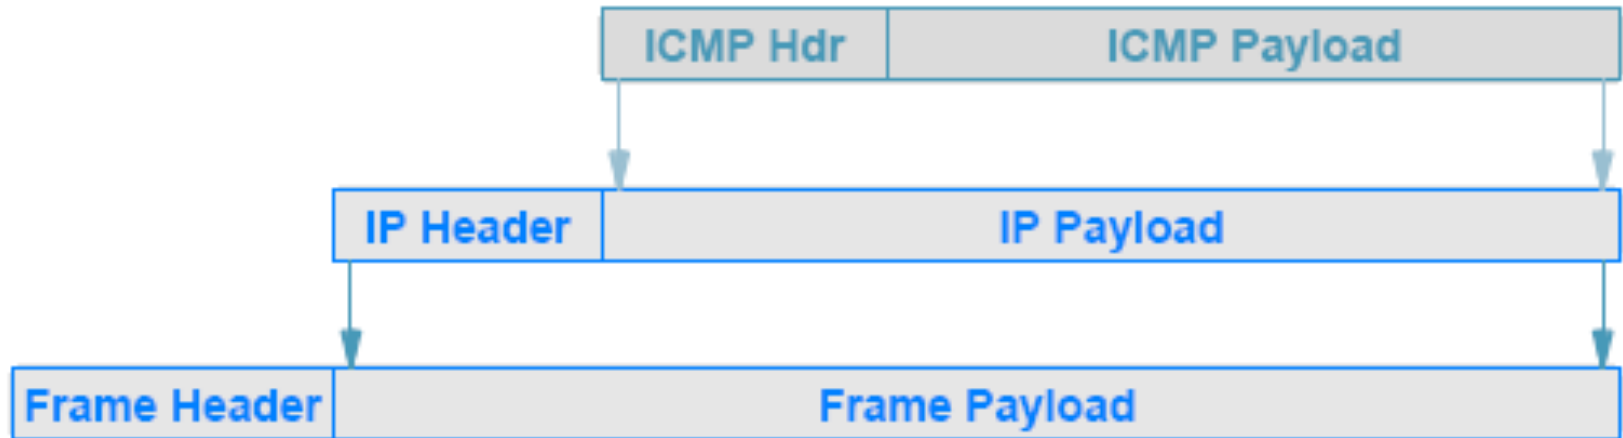
- IP includes a companion protocol, ICMP
  - It is used to report errors back to the original source
- IP and ICMP are co-dependent
  - IP depends on ICMP to report errors
  - and ICMP uses IP to carry error messages
- ICMP can be seen as a signaling protocol for network management and maintenance
- Many ICMP messages have been defined



Number	Type	Purpose
0	Echo Reply	Used by the ping program
3	Dest. Unreachable	Datagram could not be delivered
5	Redirect	Host must change a route
8	Echo	Used by the ping program
11	Time Exceeded	TTL expired or fragments timed out
12	Parameter Problem	IP header is incorrect
30	Traceroute	Used by the traceroute program



- ICMP contains two message types:
  - messages used to **report errors**
    - e.g., **Time Exceeded** and **Destination Unreachable**
  - messages used to **obtain information**
    - e.g., **Echo Request** and **Echo Reply**
- Echo Request/Reply are used by the ping application to test connectivity
  - When a host receives an echo request message
    - ICMP software on a host or router sends an echo reply that carries the same data as the request



- ICMP uses IP to transport messages:
  - when a router has an ICMP message to send
    - creates an IP datagram and encapsulates the ICMP message in it
  - the ICMP message is the payload area of the IP datagram
  - the datagram is forwarded as usual



- ICMP messages do not have special priority
  - They are forwarded like any other datagram, with one minor exception
- If an ICMP error message causes an error
  - no error message is sent
- The reason should be clear:
  - the designers wanted to avoid the Internet becoming congested carrying error messages about error messages






- Comando “ping”
  - Echo Request + Echo Replay
- Comando traceroute
  - Il mittente invia normali pacchetti IP con TTL settato a 1, 2, 3, ...
  - Con TTL = 1, il primo router decrementa TTL che arriva a 0, quindi il pacchetto viene scartato e il router manda (dovrebbe mandare) un messaggio ICML Time Exceeded
  - Con TTL= 2 il primo router decrementa e inoltra, il secondo ...
  - E così via
- Esempi “live”
  - Con ping misuro RTT, con Traceroute capisco che strada fa il mio pacchetto

# DHCP: DYNAMIC HOST CONFIGURATION PROTOCOL

Come bootstrappare una rete senza dover configurare i singoli host



- Once a host or router has been powered on, OS is started and the network software is initialized
- How does the network software in a host or router begin operation?
- For a router, the configuration manager must specify initial values for items such as
  - the IP address for each network interface
  - the protocol software to run
  - and initial values for a forwarding table
  - the configuration is saved, and a router loads the values during startup
- Host configuration usually uses a two-step process, known as **bootstrapping**
  - DHCP is used to take care of most configuration needs

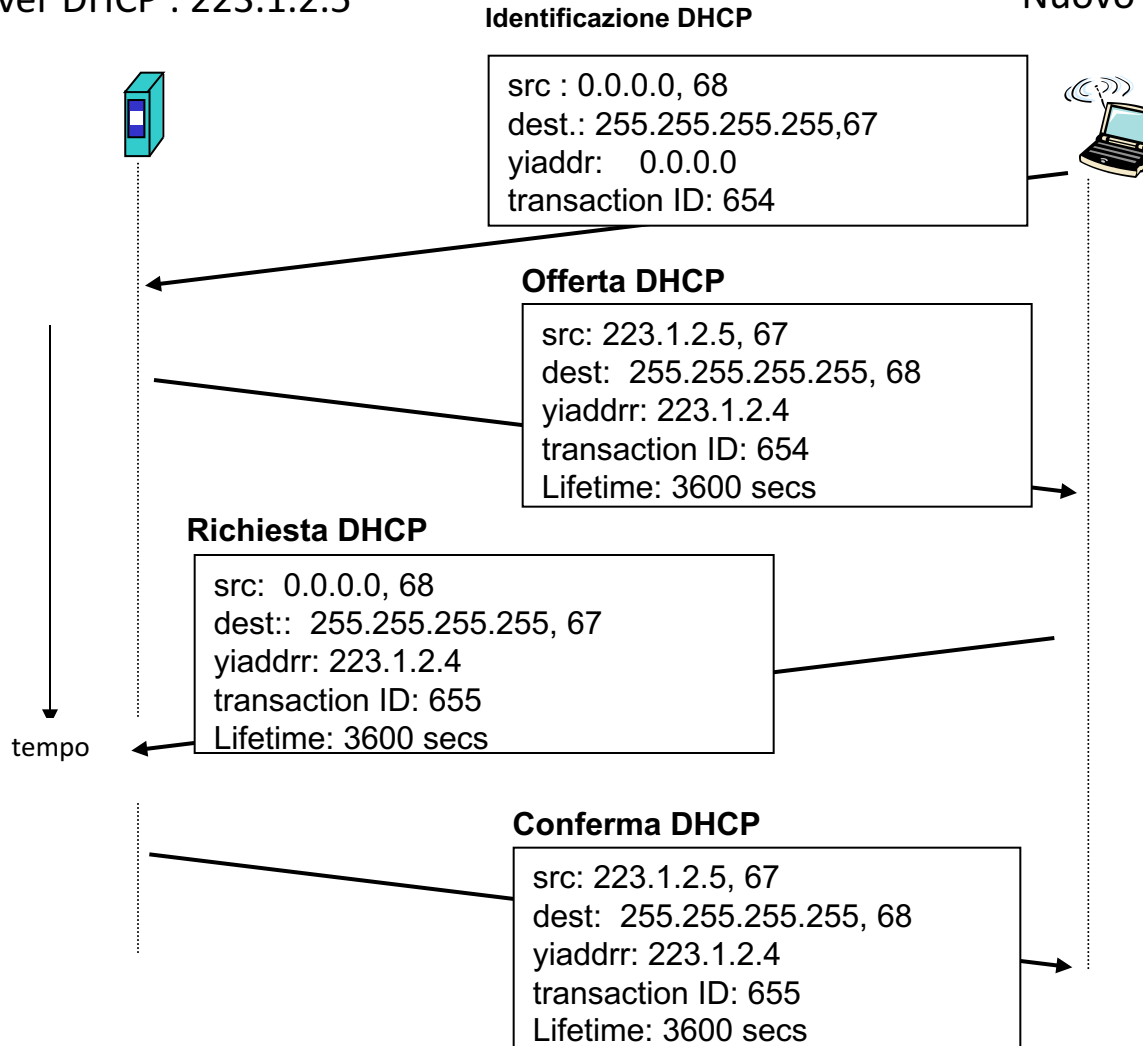




- When a computer boots
  - the DHCP client broadcasts a DHCP Request
  - the server(s) send a DHCP Reply
    - a server reply is called **offer**
    - the server is offering an address to the client
- We can configure a DHCP server to supply two types of addresses:
  - permanently assigned addresses
  - a pool of dynamic addresses to be allocated on demand
- Typically, a permanent address is assigned to a server, and a dynamic address is assigned to an arbitrary host
- Addresses assigned on demand are not given out for an arbitrary length of time

server DHCP : 223.1.2.5

Nuovo host





- DHCP issues a lease on the address for a finite period
  - The use of leases allows a DHCP server to reclaim addresses
- When the lease expires
  - the server places the address to the pool of available addresses
- When a lease expires, a host can choose to relinquish the address or renegotiate with DHCP to extend the lease
  - Negotiation occurs concurrent with other activity
- Normally, DHCP approves each lease extension
  - A computer continues to operate without any interruption
  - However, a server may be configured to deny lease extension for administrative or technical reasons
  - DHCP grants absolute control of leasing to a server
  - If a server denies an extension request
    - the host must stop using the address

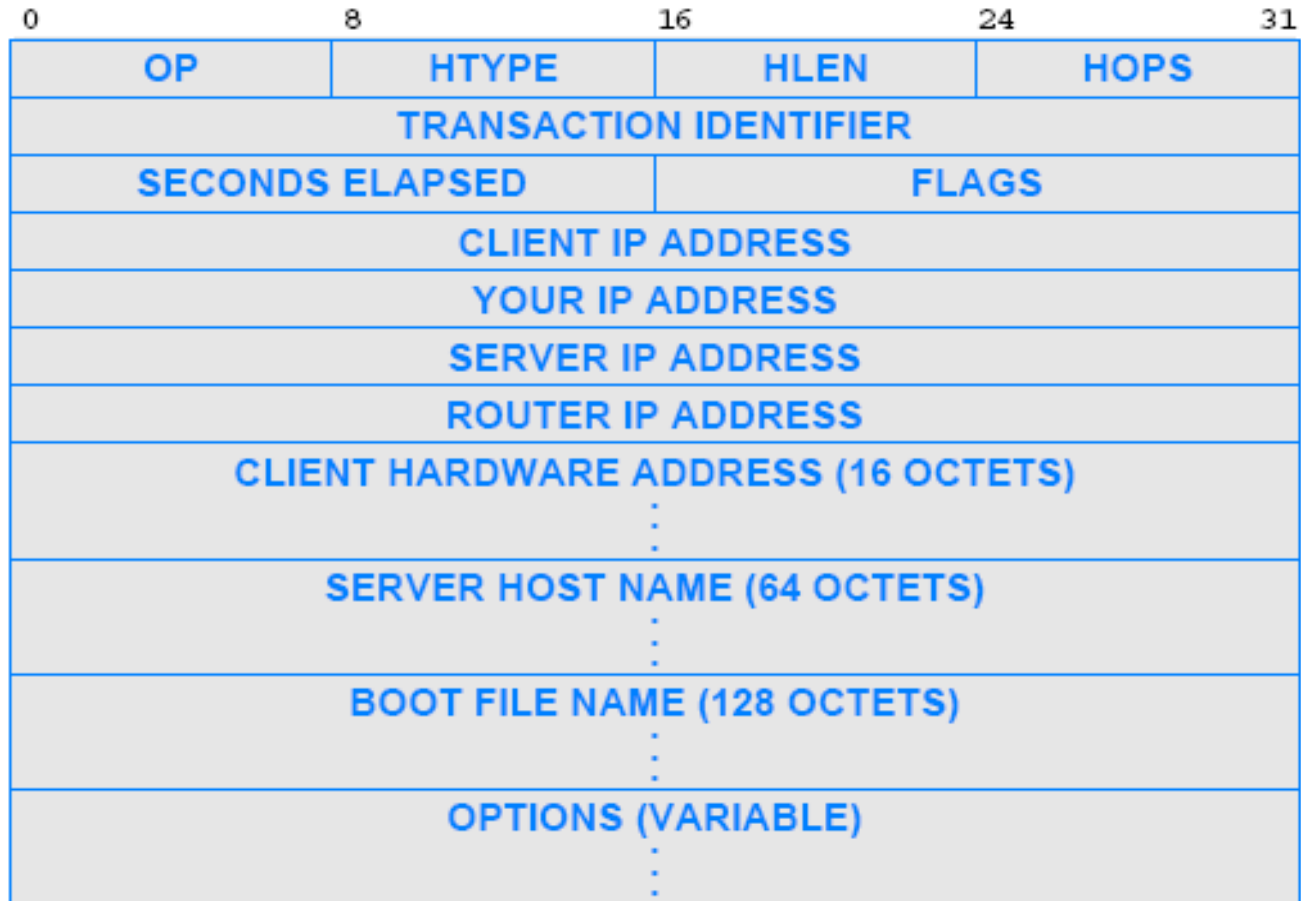


- Recovery from loss or duplication
  - DHCP is designed to insure that missing or duplicate packets do not result in misconfiguration
  - If no response is received
    - a host retransmits its request
  - If a duplicate response arrives
    - a host ignores the extra copy
- Caching of a server address
  - once a host finds a DHCP server
    - the host caches the server's address
- Avoidance of synchronized flooding
  - DHCP takes steps to prevent synchronized requests





# DHCP Message Format





- OP specifies whether the message is a Request or a Response
- HTYPE and HLEN fields specify the network hardware type and the length of a hardware address
- FLAGS specifies whether it can receive broadcast or directed replies
- HOPS specifies how many servers forwarded the request
- TRANSACTION IDENTIFIER provides a value that a client can use to determine if an incoming response matches its request
- SECONDS ELAPSED specifies how many seconds have elapsed since the host began to boot
- Except for OPTIONS (OP), each field in a DHCP message has a fixed size



- Later fields in the message are used in a response to carry information back to the host that sent a request
  - if a host does not know its IP address, the server uses field YOUR IP ADDRESS to supply the value
  - server uses fields SERVER IP ADDRESS and SERVER HOST NAME to give the host information about the location of a server
  - ROUTER IP ADDRESS contains the IP address of a default router
  - Options may include (and normally do) the local DNS server
- DHCP allows a computer to negotiate to find a boot image
  - To do so, the host fills in field BOOT FILE NAME with a request
  - The DHCP server does not send an image