

Protocol vs Hardware Addresses

- Upper-level IP addresses are virtual and are maintained by software. They hide hardware details.
- LAN hardware does not understand IP addresses.
- Network hardware needs to know the hardware address of a frame in order to transmit it across a physical network.
- Consequently, a protocol address must be translated into a hardware address for delivery.

Address Resolution

- address resolution - finding the hardware address for a given protocol address.
- Address resolution is local to a particular physical network.
- A physical network component can resolve an address only for other components on the same physical network.
- Example.

Address Resolution Techniques

- binding - an association between a protocol address and a hardware address.
- Three techniques for resolving protocol addresses:
 - Table lookup.

Bindings are stored in memory with addresses as keys.

Software looks up protocol addresses to find hardware addresses.

Address Resolution Techniques

- Closed-form computation.

The protocol address assigned to a computer is based on the computer's hardware address.

A hardware address is computed from the protocol address.

- Message exchange.

Computers exchange messages across a network to resolve an address.

A computer needing an address resolution sends a request. The destination responds with the hardware address.

Table Lookup

- Make a simple table containing IP addresses and corresponding hardware addresses.
- Search the IP addresses for a given IP address and extract the corresponding hardware address.
- On a particular physical network, all IP addresses have the same prefix. Can save space by dropping the prefix.
- For small networks, a sequential or serial search is adequate. $O(n)$.
- For large networks, a more efficient search technique such as hashing should be used.

Closed-Form Computation

- If there is flexibility in choosing IP and hardware addresses, the two addresses can be made to correspond.
- The hardware address can then be computed as a function of the IP address.
- Example: Bacchus has IP address 138.234.44.50. Suppose we can assign its hardware address to be 50. Then
hardware_address = IP_address & 0xff

Message Exchange

- The first two methods do a local computation on a single computer.
- Message exchange uses the entire physical network to resolve an IP address.
- The computer needing to resolve an IP address sends a message across the physical network and receives a reply.

Message Exchange

- Two designs:
 - Server-based: A computer sends a request to a server which sends a reply.
 - Distributed: A computer broadcasts a request to all computers on the physical network and the matching computer sends a reply.

Advantages:

- Server-based: centralized and easier to configure, manage, and control.
- Distributed: needs no dedicated computers and no administration.

ARP

- Address Resolution Protocol (ARP) is a protocol in the TCP/IP suite.
- ARP specifies how message exchange should be carried out. It has two parts:
 - ARP specifies the form of a request.
 - ARP specifies the form of a reply.
- ARP is analogous to DNS.
- A major difference between DNS and ARP is that DNS resolves host names for hosts anywhere on the Internet, whereas ARP resolves IP addresses only for hosts on the same LAN.

ARP Message Exchange

- A computer needing to resolve an IP address places an ARP request message in a hardware frame and broadcasts it on the LAN.
- Each computer on the LAN receives the request and examines the IP address.
- The computer whose IP address is in the request sends a response message. All other computers discard the request.
- The responding computer places an ARP response message in a hardware frame and sends it only to the requesting computer.

ARP Message Exchange

- The requesting computer then can extract the hardware address from the reply and send an IP packet to the destination.
- Example: UNIX has an **arp** utility which uses the ARP protocol.

ARP Message Format

- ARP can be used with arbitrary protocols and hardware types, but it is almost always used to bind IP addresses (32 bits) and Ethernet hardware addresses (48 bits).
- HARDWARE ADDRESS TYPE = 1 for Ethernet.
- PROTOCOL ADDRESS TYPE = 0x0800 for IP.
- OPERATION = 1 for a request, 2 for a response.
- A message contains bindings for both the sender and the target.

Sending an ARP Message

- A sender forms an ARP message.
- The ARP message is carried as data in a hardware frame—called encapsulation.
- Question: How does a computer know whether an incoming frame contains an ARP message?
- Answer: The frame header contains a frame type specifying an ARP message.
- Ethernet uses the type value 0x0806.

Caching ARP Responses

- Using ARP for each IP packet adds two packets of overhead for each packet sent.
- To deal with this, ARP software caches ARP responses in a table for subsequent use.
 - An ARP cache is not permanently stored. For example, the cache is cleared on system startup.
 - Older entries are discarded periodically.

Processing an ARP Message

- So here's what happens when an ARP message arrives.
- The receiver extracts the sender's hardware address and updates its local ARP cache table.
- The receiver checks the OPERATION field to see if the message is a request or a response.
- If response:
 - Add sender's address to the local cache.
 - Send any pending IP packets.

Processing an ARP Message

- If request:
 - Compare the TARGET PADDR field with the local IP address.
 - If equal, the receiver is the target. Form a response.

Reverse the sender's and target's bindings.

Insert the local hardware address in the SENDER HADDR field.

Change the OPERATION field to 2.

- Send response to sender.
- Add sender's address to the local cache.

Processing an ARP Message

- Remarks:

- Computers have a limited storage for ARP cache.
- A computer adds a binding only if it's the target of a request.
- All computers on the network receive an ARP broadcast, and so everyone could add the sender's binding.

Inefficient.

A computer not the target updates the binding of the sender only if the binding is already in its cache.

Address Resolution & Layering

- Address resolution (ARP) takes place in the network interface layer of the TCP/IP 5-layer model.
- Software in higher layers uses protocol addressing (IP).
- ARP hides the ugly details of hardware addressing and allows software in higher layers to use protocol addressing.