Tik-110.350 Computer Networks (3 cr)
Spring 2000

Address Resolution Protocol (ARP), Reverse ARP, Internet Protocol (IP)

Professor Arto Karila
Helsinki University of Technology
E-mail: Arto.Karila@hut.fi

Contents

1st lecture: ARP and RARP

- **IP addresses (already familiar from Tik-110.300)**
- **Mapping IP addresses to physical addresses**
- **Address Resolution Protocol (ARP)**
- **Reverse ARP (RARP)**

Break (everybody needs one)

2nd lecture: Internet Protocol

- **IPv4**
- **IPv6 ("IP Next Generation", "IPng", mentioned here, lectured in detail on April 4)**

IP addresses

- **Internet is a logical network consisting of a number of physical networks interconnected via routers**
- ***Universal communication service* = any-to-any communication**
- **Host identifier can be:**
  - **Name (answer to: ”what?”)**
  - **Address (answer to: ”where?”)**
  - **Path - (answer to: ”how to get there?”)**
- **Internet is based on compact, standardized, binary addresses - namely *IP addresses***
- **Every Internet host has a unique 32-bit IP address that is used in communications with this host**
- **DNS maps the domain names onto IP addresses**
- **Routers route a packet based on its destination IP address**
- **Routing protocols help routers find a path to the destination**

IP addresses

- IP address is written as four decimal numbers, each between 0 and 255, separated with points
- For example the binary address “11000001 11010010 00001001 00111111” is written “193.210.9.63” and it stands for node “63” of the class C network “193.210.9”
- “0” as the number of a node or a network stands for “this” and ”-1” (number with all bits 1) stands for ”all” - for this reason, they are not used as node or network numbers
- There can be at most 126 class A networks and each of them can have at most 16,777,214 nodes
- There can be at most 16,382 class B networks and each of them can have at most 65,534 nodes
- There can be at most 2,097,150 class C networks and each of them can have at most 254 nodes

Structure of IP addresses

- IP addresses are 32-bit integers and they are divided into three main classes: A, B and C
- Class D was added for multicasting and class E reserved for future use

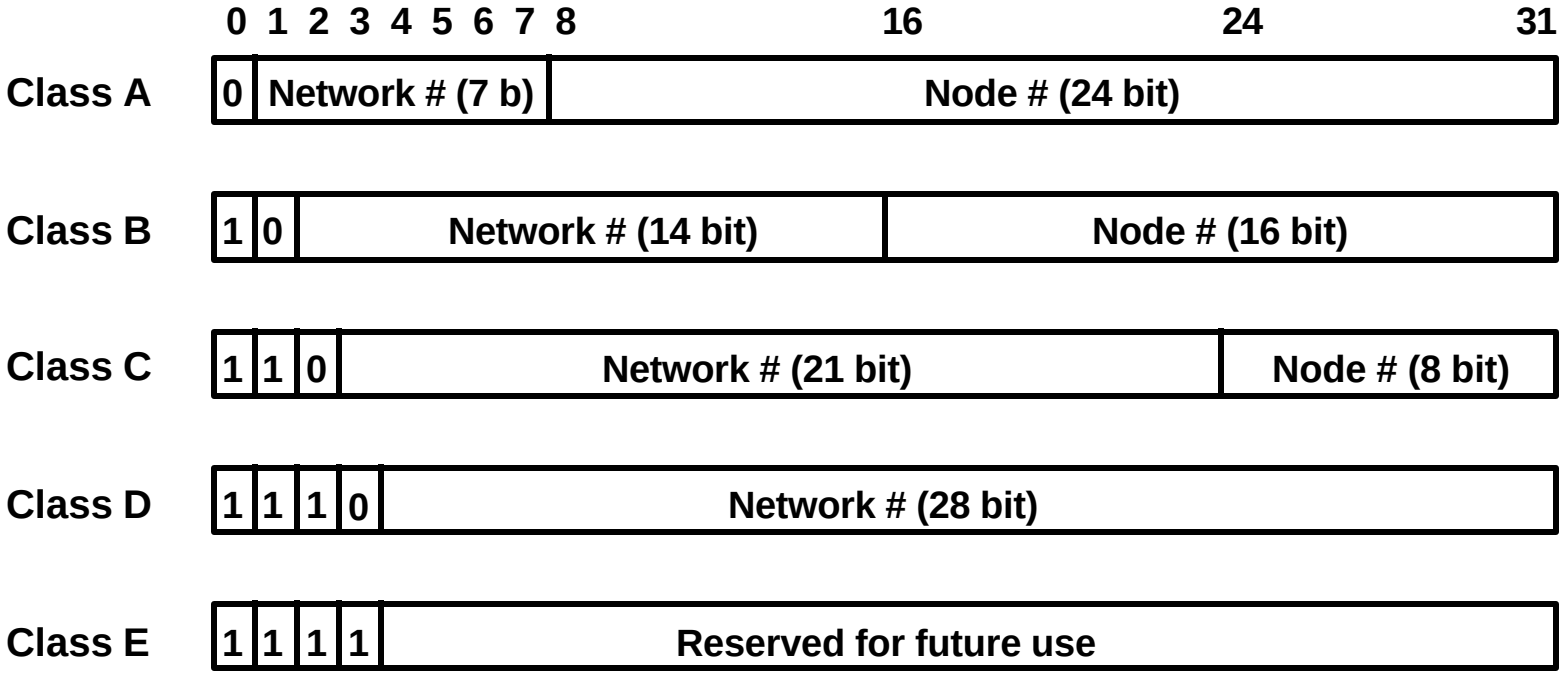

IP addresses

- An IP address conceptually consists of a pair: (net-id, host-id)
- It is easy to separate the net and host parts from an address
- An IP address does not determine a host but a network interface of a host
- For example a router has an IP address for each of the networks it is attached to
- An IP address can also indicate an entire network
- An address with all host-id bits "0" indicates a network
- An address with all host-id bits "1" is a broadcast address
- Directed broadcast = all hosts in a given network
- Limited broadcast (local network broadcast) = all hosts in this network
- We normally try to limit broadcasts to as small a scope as possible
- Multicasting is a non-trivial function to implement

Some IP addresses

- Special addresses
  - Limited broadcast: "255.255.255.255"
  - Directed broadcast:
    "A.255.255.255", "B.B.255.255", "C.C.C.255"
  - Loopback: "127.X.Y.Z"
  - All systems on this subnet: "224.0.0.1"
  - All routers on this subnet: "224.0.0.2"
- Some examples of IP addresses:
  - "0.0.0.0" = some host (source only)
  - "255.255.255.255" = any host (destination only)
  - "129.34.0.3" = host #3 in n/wk 129.34
  - "129.34.0.0" = some host in n/wk 129.34 (source only)
  - "129.34.255.255" = any host in n/wk 129.34 (destination)
  - "255.255.0.3" = host #3 in this n/wk
  - "127.0.0.1" = this host (local loopback)

Problems with IP addresses

- IP address maps into a network interface => when a host moves into another network its IP address changes (this is a problem e.g. in portable PCs)
- When a network grows, changing the network number and IP addresses becomes very difficult
- If a host has several IP addresses, the path to and reachability of the host depend on the address used
- The biggest single problem with IP addresses is its inadequate 32-bit address space (the ”ROADS” problem, Running Out of Address Space)
- Especially the number of class B networks is too small
- Classless Inter-Domain Routing (CIDR) and Network Address Translation (NAT) provide an interim solution to the problem
- IPv6 (IP Next Generation, IPng) solves the problem by using 128-bit addresses but moving into IPv6 is very difficult

Registration of IP addresses

- In an intranet it is (at least in principle) possible to use unregistered IP addresses
- Network connected into the Internet need to be registered
- Network numbers are governed by the Internet Assigned Number Authority (IANA)
- In practice numbers are assigned by the Internet Network Information Center (INTERNIC)
- Today most companies and private people get their IP addresses from an Internet Service Provider (ISP)
- Who ”owns” the addresses?
  - normally the ISP, which poses limitations to competition
- Private Internet Address Space (RFC-1597, including the old ARPANET class A network 10), some registered addresses, and NAT provide a working solution for most companies

Byte order

- The byte order of a 32-bit double word varies by processor:
  - Big Endian (1-2-3-4)
  - Little Endian (4-3-2-1)
  - Others (3-4-1-2, 2-1-4-3)
- A standardized byte order, used by all IP hosts when sending binary numbers, is needed
- In the Internet integers are always sent starting from the most significant byte (Big Endian)
- In most physical networks frames are transmitted starting from the most significant byte and bytes starting from the least significant bit
- From the IP point of view the byte or bit order of a physical network doesn't matter

Mapping IP addresses to physical addr.

- **Two hosts connected to the same physical network can communicate if they know each other's physical addresses**
- **In IP terminology " physical address" usually means, depending on the network, either the network address (e.g. in X.25) or MAC address (e.g. in LANs)**
- **In the physical layer of OSI no addresses are used**
- **Assume that A and B are connected to the same physical network and they have IP addresses I_A and I_B and physical addresses P_A and P_B, respectively**
- **A mapping: I_A => P_A and I_B => P_B is needed**
  - the *Address Resolution Problem*
- **If the physical addresses can be chosen (e.g. in proNET), a direct mapping is possible: $P_A = f(I_A)$**
- **In the general case, a direct mapping is not possible**

Address Resolution Protocol (ARP)

- In LANs fixed 48-bit MAC-level addresses are used
- A 32-bit IP address cannot be directly mapped into a 48-bit MAC address
- The Address Resolution Protocol (ARP) is a mechanism, through which a host can find the physical address of another host residing in the same network based on its IP address
- The source host sends the inquiry to the local network in a broadcast message, which the target host answers to
- The source host caches the physical addresses into a lookup table in order to avoid unnecessary use of the ARP
- At the same time the target host can cache the physical address of the source host
- ARP is a low-level protocol that helps hide the physical network addresses

Implementing ARP

- An ARP implementation consists of two parts:
  - A sending part mapping IP addresses to physical addresses for outbound IP datagrams
  - A receiving part processing incoming ARP messages
- There are practical problems associated with implementing the sending part:
  - The target host may be out of operation
  - Ethernet messages may get lost
  - While waiting for a response, new messages may arrive for the same target host
  - The physical address of the target host may change (e.g. when changing the LAN interface board)

Implementing ARP

- **The receiving part handles two kinds of ARP messages:**
  - **ARP request**
  - **ARP replies**
- **ARP messages are not IP messages but network-dependent low-level messages**
- **ARP can also be used with protocols other than IP**

Determining the IP address at startup

- The physical address is a low-level machine-dependent address
- Every IP host needs at least one 32-bit IP address, that is independent of the physical address, to be able to operate
- Normally the IP address of the host is stored on its hard disk
- How can a diskless workstation determine its IP address?
- The host needs an IP address to be able to boot from a file server
- We want to keep the ROMmed bootstrap loader the same for all machines
- Reverse ARP (RARP) is a mechanism, through which a host can inquire its IP address from a server on the same network
- The host broadcasts its physical address to the local network
- A RARP server returns an IP address for the requesting host

Internet Protocol (IP)

- Comer chapter 7
- IP-protocol was designed to provide for connectionless datagram service in internetworks
- From the user's point of view, internet is a virtual network interconnecting all the hosts connected to it - the internal structure of the internet is unessential
- Conceptually, a TCP/IP-based internet provides three levels of services, presented in figure 7.1 of Comer (the protocols named in parentheses implement these services):
  - Application level services (e.g. SMTP, FTP, HTTP)
  - Reliable transport service (TCP)
  - Connectionless packet service (IP)
- The success of TCP/IP is largely based on the good functioning and adaptability of the architecture (layer structure)

IP (continued)

- The Internet Protocol was designed for use in interconnected systems of packet-switched computer communication networks
- The internet protocol is specifically limited in scope to provide the functions necessary to deliver a package of bits (an internet datagram) from a source to a destination over an interconnected system of networks
- There are no mechanisms to augment end-to-end data reliability, flow control, sequencing, or other services commonly found in host-to-host protocols
- The internet protocol implements two basic functions:
  - Addressing
  - Fragmentation

IP (continued)

- The IP-based packet-forwarding service provided by an internet can be characterized as follows:
  - The service is *unreliable*, because IP does not use acknowledgements and the delivery of a datagram is not guaranteed
  - The service is *connectionless*, because every packet is routed separately and independently of other packets
  - The service is based on *best-effort delivery*, because the internet software is doing its best to deliver every packet
- IP defines three significant things:
  - The structure of the Internet datagram
  - The routing of datagrams
  - A number of rules implementing the principles of connectionless packet service
- IP is the key technology on which the Internet is based

IPv4 datagram header

- The figure depicting the header of an IPv4 datagram is taken directly from RFC-791

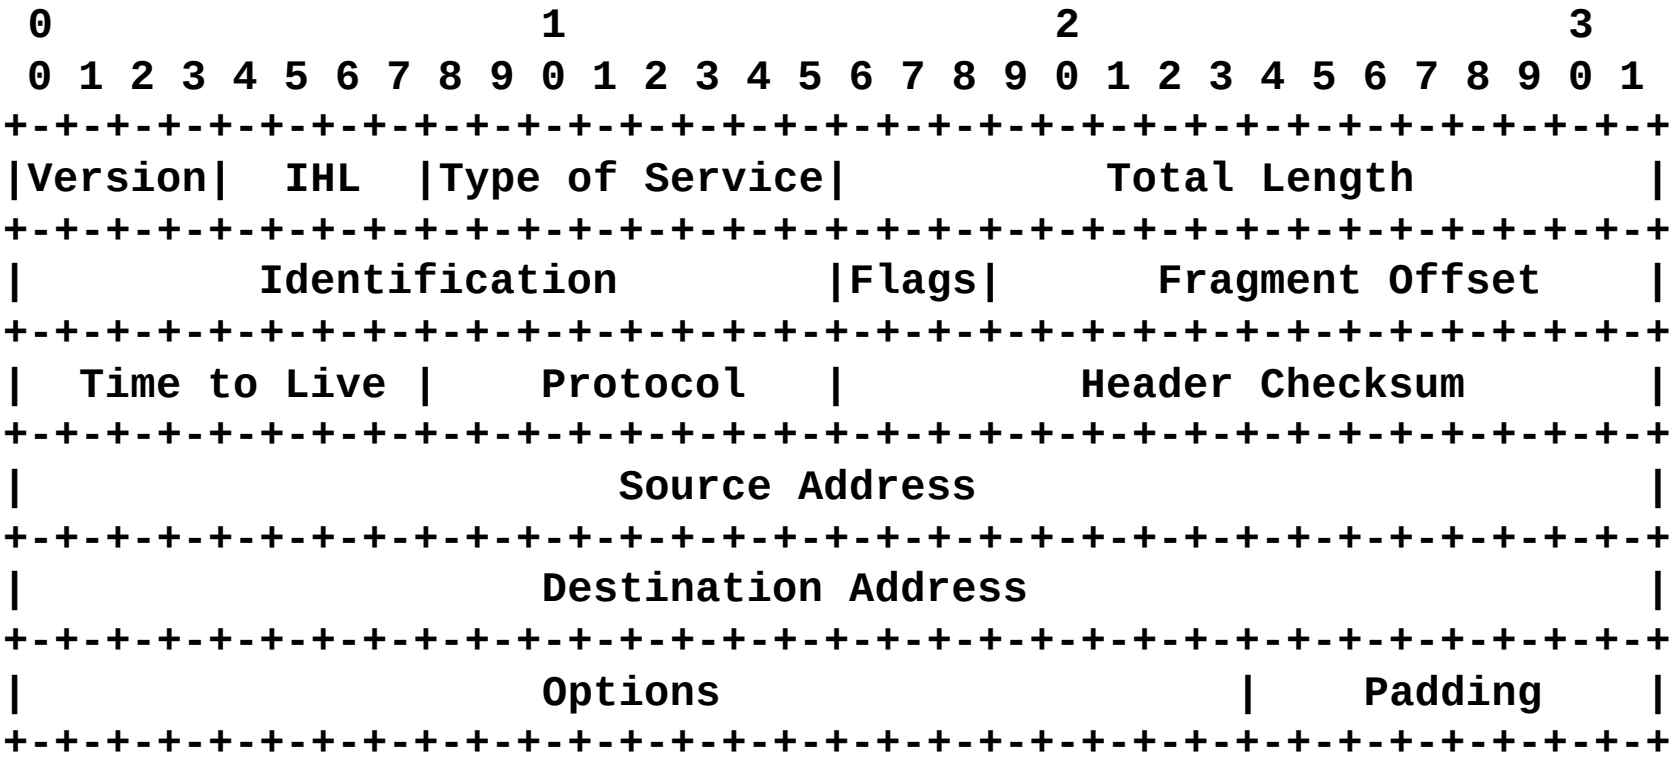

Fields of the IPv4 header

- `Version` (4 bit) specifies the protocol version, which currently is 4; IP software does not handle wrong IP versions
- `IHL` (Internet Header Length, 4 bit) specifies the length of the header in 32-bit double-words; the minimum value is 5
- `Type of Service` (8 bit) specifies the desired quality of service for the datagram; it consists of the `Precedence` part, defining the priority, and three quality of service bits: `D (delay)`, `T (throughput)` and `R (reliability)`
- In practice, the payload carried by IP uses priority 0 and the higher priority 1 is (unfortunately) not used
- The routing algorithm uses quality of service bits as an indication that helps choose the optimal route for each datagram; in practice also these bits are currently not used

Type of Service field (RFC791)

```
Bits 0-2:  Precedence.
Bit    3:  0 = Normal Delay,      1 = Low Delay.
Bits   4:  0 = Normal Throughput, 1 = High Throughput.
Bits   5:  0 = Normal Reliability, 1 = High Reliability.
Bit  6-7:  Reserved for Future Use.
```

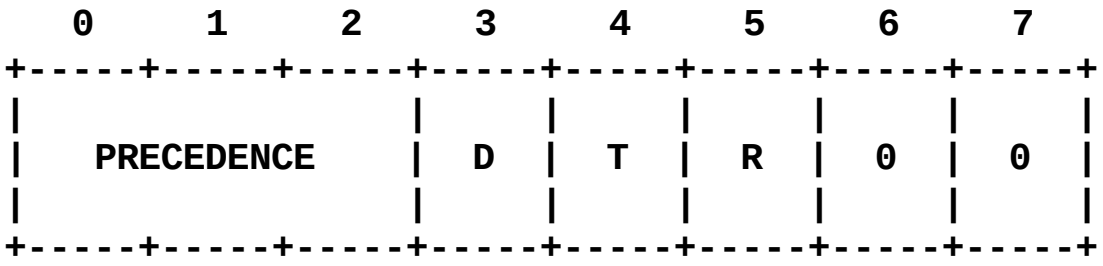

Precedence

- 111 - Network Control
- 110 - Internetwork Control
- 101 - CRITIC/ECP
- 100 - Flash Override
- 011 - Flash
- 010 - Immediate
- 001 - Priority
- 000 - Routine

Fields of the IPv4 header (continued)

- **`Total Length` (16 bit) indicates the length of the entire datagram (header and data) in octets**
- **The maximum length for a datagram is 65535 octets**
- **Every host connected to the network must accept at least 576 octet datagrams (512 octets of data and 64 octets of header); the sending of datagrams longer than this is only recommended, if the sender knows that the recipient can receive them**
- **`Identification` (16 bit) is a running number chosen by the sender and used for combining fragments**
- **`Flags` (3 bit) is associated with fragmenting datagrams:**
  - **Bit 0: reserved, must be zero**
  - **Bit 1: (DF) 0 = May Fragment, 1 = Don't Fragment**
  - **Bit 2: (MF) 0 = Last Fragment, 1 = More Fragments**

Fields of the IPv4 header (continued)

- **`Fragment Offset` (13 bit) specifies the offset for the first data octet of the fragment from the beginning of the datagram in multiples of 8; the offset of the first fragment is 0**
- **`Time to Live` (8 bit) specifies the maximum life-time left for the datagram in the internet in seconds; every router decrements the remaining life-time by at least one**
- **`Protocol` (8 bit) specifies the upper layer protocol (such as TCP or UDP), protocol numbers are defined in RFC-1700, ”Assigned Numbers”, (1994)**
- **`Header Checksum` (16 bit) is the one’s complement of the one’s complement sum of all the 16-bit words of the header; this field is checked and recalculated at each node that handles the header**

Fields of the IPv4 header (continued)

- **`Source Address` (32 bit) contains the IP address of the original sender of the datagram**
- **`Destination Address` (32 bit) contains the IP address of the ultimate recipient of the datagram**
- **`Options` (varying length) may belong or not belong to each datagram; handling of IP options must be implemented in all computers and routers of an internet**
- **There are two types of options:**
  - **Single `Option-type` octets**
  - **Sequences: `Option-type` octet, `Option-length` octet and the actual `Option-data` octets**

Option-type field

- An `Option-type` octet has three fields:
  - `copied flag` (1 bit) indicates if the option has been copied to all the fragments of a datagram
  - `option class` (2 bit)
    - 0 = control
    - 1 = reserved for future use
    - 2 = debugging and measurement
    - 3 = reserved for future use
  - `option number` (5 bit)
- The table on the next foil list the defined options

Defined options (RFC791)

CLASS	NUMBER	LENGTH	DESCRIPTION
0	0	-	End of Option list. This option occupies only 1 octet; it has no length octet.
0	1	-	No Operation. This option occupies only 1 octet; it has no length octet.
0	2	11	Security. Used to carry Security, Compartmentation, User Group (TCC), and Handling Restriction Codes compatible with DOD requirements.
0	3	var.	Loose Source Routing. Used to route the internet datagram based on information supplied by the source.
0	9	var.	Strict Source Routing. Used to route the internet datagram based on information supplied by the source.
0	7	var.	Record Route. Used to trace the route an internet datagram takes.
0	8	4	Stream ID. Used to carry the stream identifier.
2	4	var.	Internet Timestamp.

IPv6

- IP Next Generation (”IPng”) or IPv6 was born after a lengthy process in December 1995 and is defined in RFC-1883
- IPv6 differs from IPv4 mainly in the following respects:
  - Sufficient address space
  - Simpler and more flexible header
  - Better extensibility and options
  - Flow labeling
  - Security features (IPSEC)
- The migration from IPv4 to IPv6 will be the biggest single change in the history of the Internet
- Automatic tools for facilitating the migration are being developed
- The transfer may not happen for a while
- The Network Address Translation (NAT) function implemented in routers and firewalls extends the life-time of IPv4

The structure of IPv6 datagram

- An IPv6 datagram includes a Base Header, zero or more Extension Headers and data (figure 29.1 of Comer)
- With one exception, most routers only need to handle the Base Header
- The exception is the "Hop-by-Hop Options" -header, which needs to be handled by every router
- Routers participating in source routing also have to examine the Routing Header
- The structure of the Base Header is depicted in figure 29.2 of Comer
- Flow labels can be used to define data flows and specify the quality of service per data flow (the concept of data flow was entirely missing in IPv4)
- The base header is of fixed length (40 octets)
- Extension headers are of varying length

IPv6 extension headers

- The extension headers of IPv6 are recommended to be placed in the following order:
  - Hop-by-Hop Options header
  - Destination Options header (intermediate destinations)
  - Routing header
  - Fragment header
  - Authentication header
  - Encapsulating Security Payload header
  - Destination Options header (final destination)
  - upper-layer header (compare with ”protocol” field of IPv4)
- Each header may appear only once with the exception of the "Destination Options" header, which may appear twice
- Only the sender (not a router) may fragment a datagram
- Fragmentation can also be handled through tunneling (figure 29.6 of Comer)

Addresses of IPv6

- The address architecture of IPv6 is defined in RFC-1884 (December 1995)
- IPv6 addresses are divided into three main types:
  - Unicast - a normal address specifying a network interface
  - Anycast - one (usually the ”closest”) of a number of network interfaces (e.g. a server cluster)
  - Multicast - a group of recipients
- The most significant difference from the old IP addresses is that the IPv6 addresses are 128-bit long
- RFC-1884 defines the division of IPv6 address space
- IPv4 will be automatically converted into IPv6 addresses as a necessary condition for changing over to IPv6

Addresses of IPv6

- IPv6 addresses are normally written in hexadecimal notation so that the 16-bit parts are separated with colons, such as:
  FEDC:BA98:7654:3210:FEDC:BA98:7654:3210
  1080:0:0:0:8:800:200C:417A
- Some examples of IPv6 addresses:

1080:0:0:0:8:800:200C:417A	unicast
FF01:0:0:0:0:0:0:43	multicast
0:0:0:0:0:0:0:1	loopback
0:0:0:0:0:0:0:0	unspecified

- The above addresses can also be written in short-hand:

1080::8:800:200C:417A	unicast
FF01::43	multicast
::1	loopback
::	unspecified

IPv4 vs. IPv6

- The transition from IPv4 to IPv6 will be a big one
- CIDR and NAT have removed the immediate need for making the transition
- NAT violates the end-to-end principle of the Internet and there is a NAT-considered-harmful discussion going on
- IPv6 offers several advantages over IPv4, such as:
  - flow labels support real-time voice and video
  - conventional routers can process IPv6 much faster than IPv4 (modern fast routers process both at wire-speed)
  - mobile IP becomes much easier with IPv6
  - IPSEC is a mandatory option of IPv6
- For these reasons, some new development projects (for example those based on mobile IP) are better done using IPv6
- IPv6 will ultimately win but IPv4 will continue to dominate the Internet for some years