Network Address Translation (NAT)

Adapted from
Tannenbaum’s Computer Network Ch.5.6;
computer.howstuffworks.com/nat1.htm;
Comer’s TCP/IP vol.1 Ch.20

Long term and short term solutions to Internet scalability

- Internet will eventually run out of IPv4 addresses, a few years from now
- The long term solution is the migration to IPv6 which has 128 bit addresses
- The IPv4 to IPv6 transition is slowly occurring, but it will take years before the process completes. As a consequence some people felt the need for a quick fix, which came in the form of **NAT (Network Address Translation).**
- For reference see RFC3022 and Dutcher 2001

NAT

Network Address Translation Technology provides transparent IP-level access to the Internet from a host with a private address

Translation

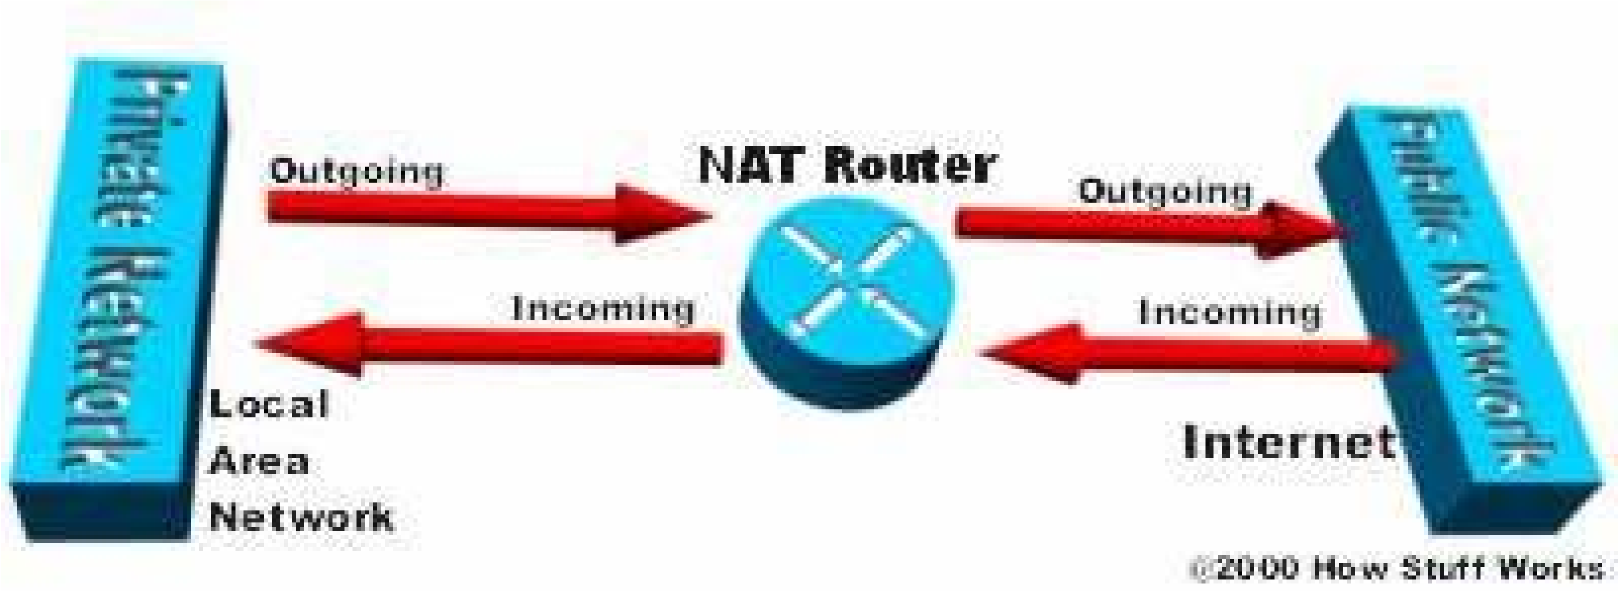

Network Address Translation allows a single device, such as a router, to act as an agent between the Internet (or "public network") and a local (or "private") network so that a single, unique IP address is required to represent an entire group of computers.

Reserved Addresses

- To make this possible three ranges of IP addresses have been declared as private no packets containig these addresses may appear on the Internet
- The three reserved ranges are:
  - **10.0.0.0   to 10.255.255.255/8**
  - **172.16.0.0 to 172.31.255.255/12**
  - **192.168.0.0 to 192.168.255.255/16**
- Note that the first block is nothing but a single class A network number, while the second block is a set of 16 continuous class B network numbers, and the third block is a set of 255 continuous class C network numbers.

NAT basics

- Assign to each company a single public IP address
- Within the company every computer gets a unique private ID which is used for routing the intramural traffic
  (private means that addresses not used in the Internet)
- When the packet exits the company and goes to the Internet Service Provider an address translation takes place and the packet exits with one of the compay IPs
- Typically with the company IP and the reference to the source computer coded in the payload (in the TCP or UDP port field)
- This is called Port-Level Multiplexed NAT; other NAT techniques: Static NAT, Dynamic NAT, Overlapping

Port-level multiplexed NAT

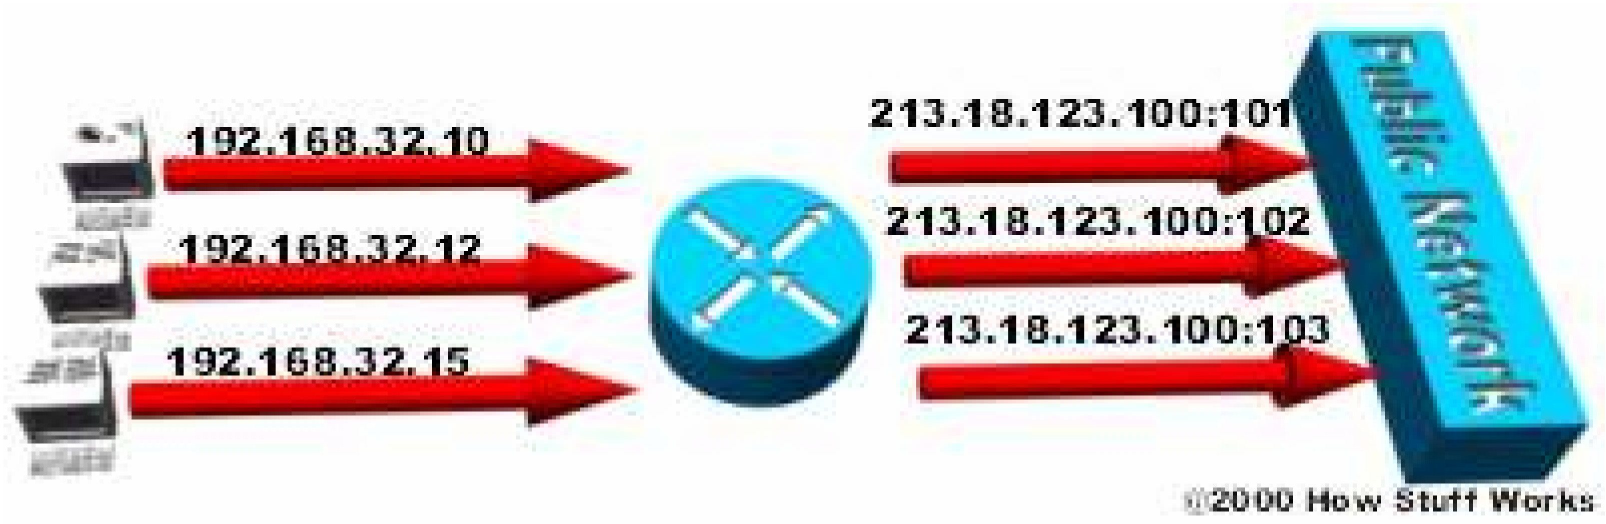

Overloading is a form of dynamic NAT that maps multiple unregistered IP addresses to a single registered IP address by using different ports.

In overloading, each computer on the private network is translated to the same IP address (213.18.123.100), but with a different port number assignment.

Overloading is known also as PAT, Port Address Translation, or Network Address Port Translation (NAPT), or single address NAT or port-level multiplexed NAT.

Address:port

NAT overloading utilizes a feature of the TCP/IP protocol stack, **multiplexing**, that allows a computer to maintain several concurrent connections with a remote computer (or computers) using different TCP or UDP **ports**.

An IP packet has a header that contains the following information:

- **Source Address -** The IP address of the originating computer, e.g. 201.3.83.132
- **Source Port -** The TCP or UDP port number assigned by the originating computer for this packet, e.g. Port 1080
- **Destination Address -** The IP address of the receiving computer, e.g. 145.51.18.223
- **Destination Port -** The TCP or UDP port number that the originating computer is asking the receiving computer to open, e.g. Port 3021

The addresses specify the two machines at each end, while the port numbers ensure that the connection between the two computers has a unique identifier. The combination of these four numbers defines a single TCP/IP connection. Each port number uses 16 bits, which means that there are a possible 65,536 (2^{16}) values. Realistically, since different manufacturers map the ports in slightly different ways, you can expect to have about 4,000 ports available.

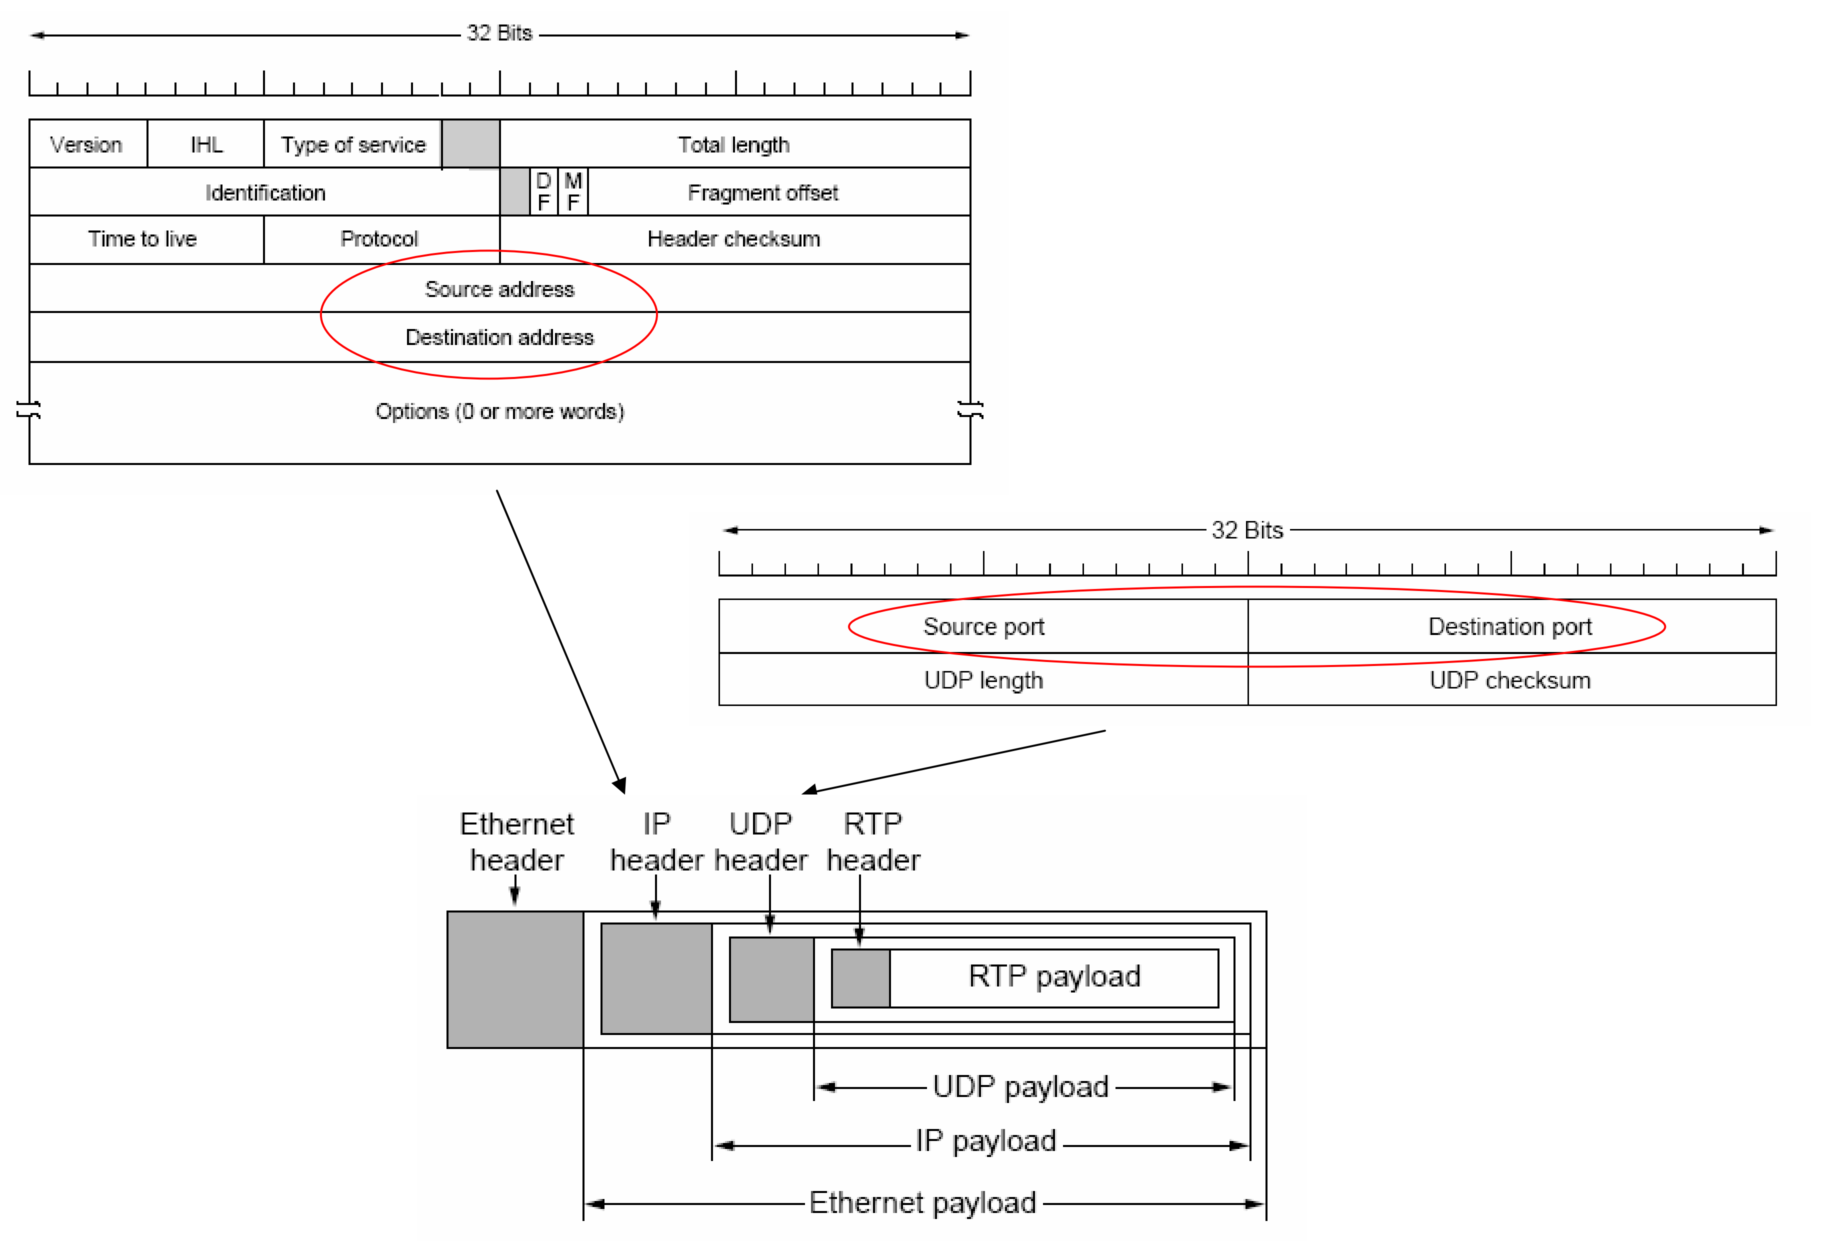
32 Bits
Version
IHL
Type of service
Total length
Identification
D F
M F
Fragment offset
Time to live
Protocol
Header checksum
Source address
Destination address
Options (0 or more words)
32 Bits
Source port
Destination port
UDP length
UDP checksum
Ethernet header
IP header
UDP header
RTP header
RTP payload
UDP payload
IP payload
Ethernet payload

Stub domain

The internal network is usually a **LAN (Local Area Network)**,
commonly referred to as the **stub domain**.
A stub domain is a LAN that uses IP addresses internally.
Most of the network traffic in a stub domain is local,
so it doesn't travel outside the internal network.
A stub domain can include both registered and unregistered
IP addresses.

Of course, any computers
that use unregistered IP addresses must use NAT
to communicate with the rest of the world.

How it works

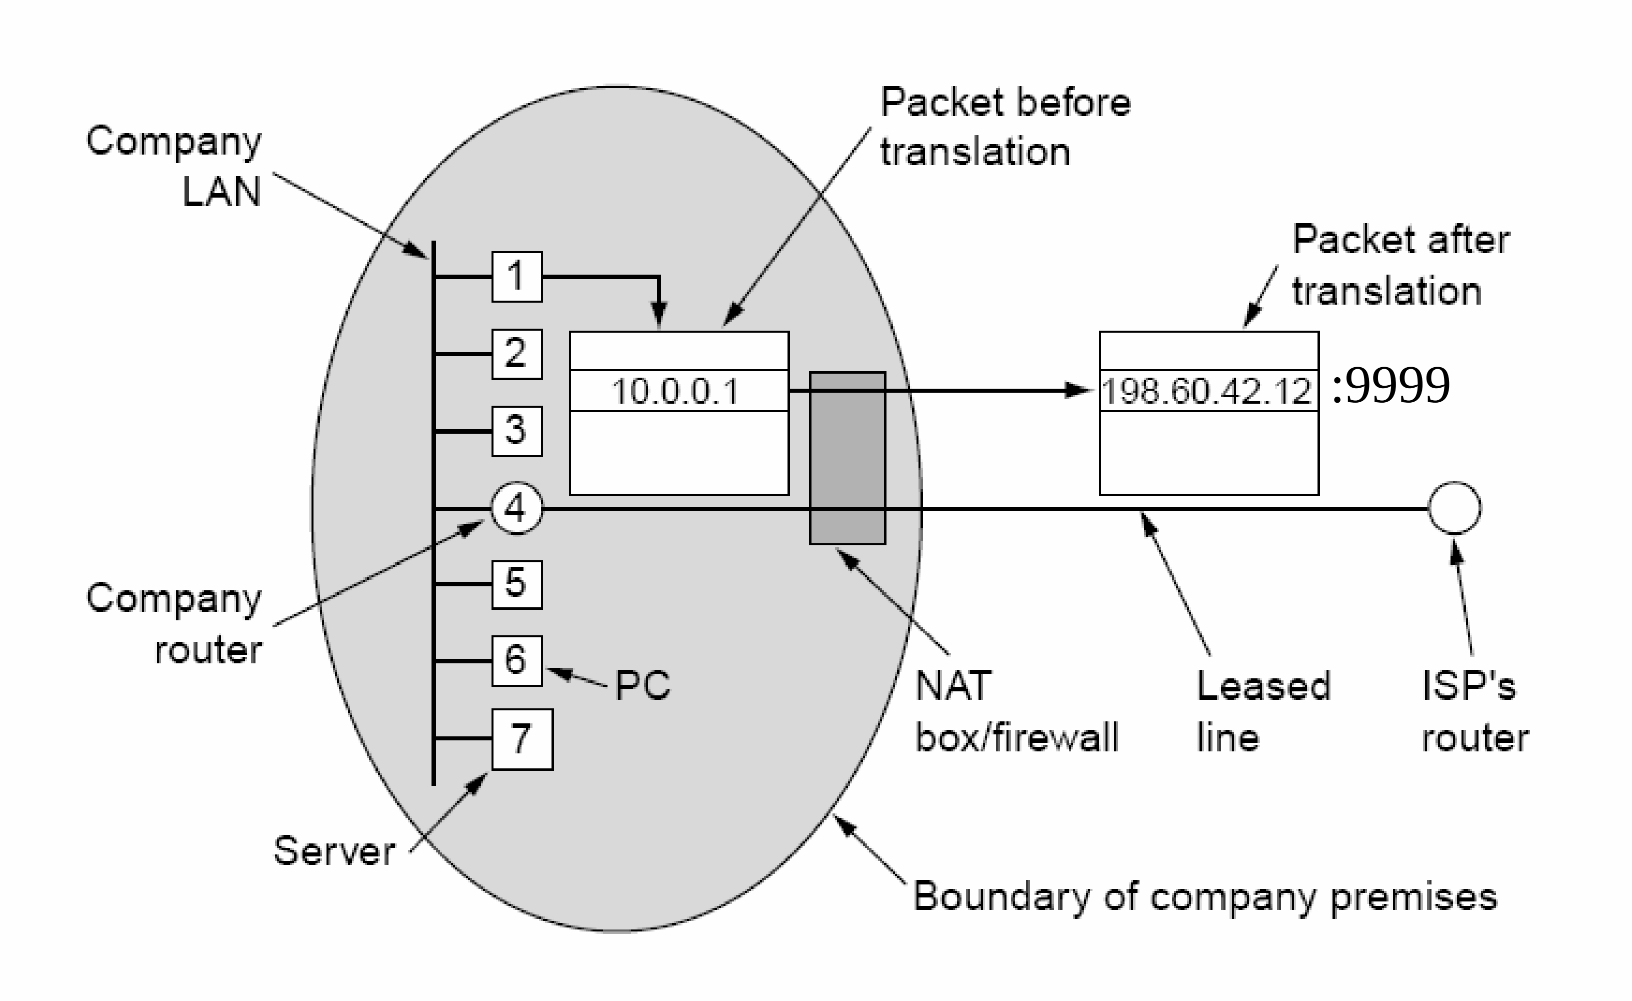
Company LAN
Packet before translation
Packet after translation
1
2
3
10.0.0.1
198.60.42.12
:9999
4
5
6
7
Company router
PC
Server
NAT box/firewall
Leased line
ISP's router
Boundary of company premises

Outgoing packets

Whenever an outgoing packet enters the NAT box:

- The 10.x.y.z source address is replaced by the company true IP address
- The source port field is replaced by a number corresponding to an index of the NAT box's 65536-entry translation table
- This table entry contains the original IP address contains the true original source address and source port

Incoming packets

When a packet arrives at the NAT box from the ISP:

- the source port is extracted and used as an index into the NAT box’s mapping table.
- From the entry located, the internal IP address and original TCP Source Port are extracted and inserted into the packet
- The packet is then passed to the company router for normal delivery using the 10.x.y.z address

NAT Table Example

```
Private    Private   NAT     Protocol
Address    Port      Port    Used
--------------------------------------
10.0.0.5   21023     14003   tcp
10.0.0.1     386     14010   tcp
```

NAT creates a firewall

- Implementing dynamic NAT automatically creates a **firewall** between your internal network and outside networks, or between your internal network and the Internet: NAT only allows connections that originate inside the stub domain.
- Essentially, this means that a computer on an external network cannot connect to your computer unless your computer has initiated the contact. You can browse the Internet and connect to a site, and even download a file; but somebody else cannot latch onto your IP address and use it to connect to a port on your computer.
- In specific circumstances, Static NAT, also called inbound mapping, allows external devices to initiate connections to computers on the stub domain. For instance, if you wish to go from an inside global address to a specific inside local address that is assigned to your Web server, Static NAT would enable the connection.

NAT vs Proxy

NAT is sometimes confused with **proxy servers**,
but there are differences between them:

- NAT is transparent to the source and to destination computers.
  Neither one realizes that it is dealing with a third device.

- But a proxy server is not transparent. The source computer knows
  that it is making a request to the proxy server and must be configured to do so.
  The destination computer thinks that the proxy server **IS** the source computer,
   and deals with it directly.

- Also, while NAT is a layer 3 (network) protocol, proxy servers usually work
  at layer 4 (transport, e.g. TCP proxies) or higher (e.g. http proxies)

.

- Working at a higher layer makes proxy servers
   slower than NAT devices in most cases.

NAT advantages

- Use of a single IP for many hosts (up to 64k)
- Creates a hidden network
- Improves security
- Simplicity and reliability
- Can be used to realize load balancing
- You can move your Web server or FTP server to another host computer without having to worry about broken links by simply changing the inbound mapping at the router
- NAT + RFC 1918 has slowed down IPv4 address exhaustion: 70% of Fortune1000 companies use NAT

Objections to NAT(0)

- The opposers of NAT object also that this fix to the address exaustion problem is is lowering the pressure over passing to the real solution: IPv6
- Since the TCP source port field is 16 bits, at most 65,536 machines can be mapped onto an IP address (actually 4096 ports are reserved)

Objections to NAT(1)

- NAT violates the number one rule of protocol layering:
  - layer k may not make any assumptions about what layer k+1 has put into the payload field
    - for instance if TCP is upgraded to TCP-2 with different header layout, say with a 32-bit port field, NAT will fail.
- NAT destroys the independence between layers and the interchangeability of protocols:
  - processes on the Internet are not required to use TCP or UDP, the use of a protocol different from TCP and UDP with NAT is a problem.

Objections to NAT(2)

Some applications insert IP addresses in the body of the text,
the receiver then extracts these addresses and uses them.
Since NAT knows nothing about these addresses, it cannot replace them.

FTP works this way and can fail in presence of NAT,
unless special precautions are taken.

Break some security protocols and some Qos Function

Objections to NAT(3)

- NAT changes the Internet from a connectionless network to a kind of connection-oriented network: ideed NAT must maintain information of each connection passing through it
- If the NAT box crashes and the mapping table is lost, all its TCP connections are destroyed (in absence of NAT, router crashes have no effect on TCP: the sending process just times out and within a few seconds retransmits the unacknowledged packets)

Objections to NAT(4)

- NAT violates the architectural model of IP, which states that every IP address identifies a single machine worldwide.
  - thowsends of machines use the address 10.0.0.1
  - impossible to merge a privately-addressed network with another privately addressed network, complex and costly network renumbering operation.
  - The Peer-to-Peer model requires end nodes to have unique, globally-routable addresses. (example : File-sharing and Voice-over-IP).
    NAT difficulty in using these applications. Organisation choosing private-addressing now limits the scope of future deployments of peer-to-peer applications or any other end-to-end services.

The internet was not meant to be this way

- mid 70’s: The Pouzin/Cerf concept implied that the network addresses were by definition unique
- The IPv4 internet was supposed to be free of any artifacts impeding end to end reachability.
- In 1994 with an address shortage looming, IETF defined private, non unique network addresses : the beginning of the end of p2p

- In all fairness, this is what likely allowed 20 years of spectacular internet growth without the often predicted collapse or implosion
- IPv4 was introduced 21 years ago (January 1st 1983) and has served us well but is now bursting at the seams. It is our duty to usher a new growth phase in the history of the internet.

Limitations of a one way internet

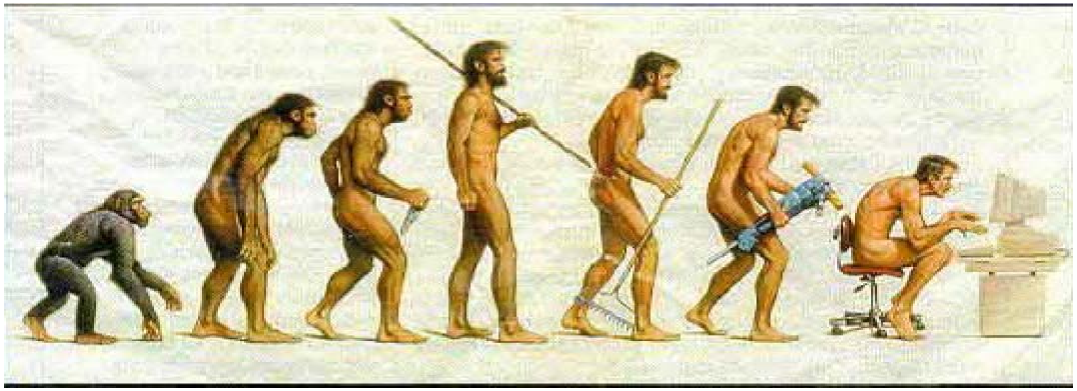

- The internet today :
  - 500 million people and devices who cannot reach each other directly.
    A couple of million major servers sitting in the network cloud mediate everything and are the biggest security risk.

- Removing the shackles
  - Telephony growth was stunted by lack of scalability. Growth was only possible with direct dial, in other words peer 2 peer.
    - The first dial tone was in Germany in 1908.
    - It took until the 70's to go from local DD to DDD to generalized IDDD
  - Now Internet growth is stunted by the lack of scalability.

Is the internet in an impasse?

- IPv4 addresses are effectively being rationed and will likely run out by 2008-2010

- The shortage is hidden by the proliferation of NAT's which allow re-use of addresses.
  - Worse than having an extension number behind a PBX
  - Like having a manually patched phone call nearly a century ago.

- Telephony in 1920 needed permanent phone numbers and peer 2 peer communications

- Internet in 2004 needs permanent addresses and peer 2 peer communications

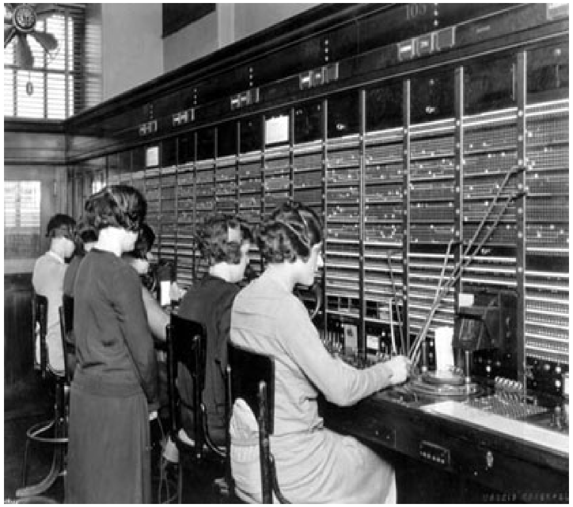

But why are IPv6 and P2P so crucial?

- Peer to peer implies permanent addresses. How else can I be reached anywhere, anytime with any type of medium, by people and communicating devices?
- The issue : The internet is running out of permanent addresses. New revenue opportunities and successful recovery from the recession risk derailment.
- The remedy:
  - IPv6 which has a much larger address field than the currently used IPv4 protocol (128 bits vs 32) and is standardized by IETF
  - The other crucial advantages of IPv6 are mandatory support of end to end security, plug and play mode to add devices to the network and scaleable support of mobile devices.