What I am going to talk about

- NAT
- IPSec and NAT
- Proxy, Reverse proxy
- P2P
- Firewall

Problems

- A concern that has spanned decades to the 1980s is the exhaustion of available IP addresses (mobile phones). This was the driving factor for move from classfull networks to CIDR addressing.
- Private network - network that uses RFC 1918 IP address space

Private network IP blocks

- 10.0.0.0 – 10.255.255.255 single class A, 10.0.0.0/8, 16M
- 172.16.0.0 – 172.31.255.255, 16 contiguous class Bs, 172.16.0.0/12, 1M
- 192.168.0.0 – 192.168.255.255, 256 contiguous class Cs, 192.168.0.0/16, 64K
- 169.254.0.0 – 169.254.255.255, single class B, 169.254.0.0/16, 64K

RFCs

- RFC 3022 – Traditional IP Network Address Translator (Traditional NAT)
- RFC 4008 – Standards Track – Definitions of Managed Objects for Network Address Translators (NAT)
- IPSec UDP encapsulation RFC 3948
- More here http://www.wikipedia.org/

NAT

- Most systems using NAT do so in order to enable multiple hosts on a private network to access the Internet using a single public IP address (see gateway). According to specifications, routers should not act in this way, but many network administrators find NAT a convenient technique and use it widely.

How often used

NATs are widely deployed throughout the business world - near 100% of all business networks. About 50% of domestic network edge devices have NAT functionality in them (e.g. ADSL modems). More correctly these devices should be call NAPTs – Network Address and Port Translators.

What exactly going on in NAT

In a typical configuration, a local network uses one of the designated "private" IP address subnets, and a router on that network has a private address (such as 192.168.1.55) in that address space. The router is also connected to the Internet with a one "public" address ("overloaded" NAT) or multiple "public" addresses assigned by an ISP. As traffic passes from the local network to the Internet, the source address in each packet is translated on the fly from the private addresses to the public address(es). The router tracks basic data about each active connection (particularly the destination address and port). When a reply returns to the router, it uses the connection tracking data it stored during the outbound phase to determine where on the internal network to forward the reply; the TCP or UDP client port numbers are used to demultiplex the packets in the case of overloaded NAT, or IP address and port number when multiple public addresses are available, on packet return. To a system on the Internet, the router itself appears to be the source/destination for this traffic.

TCP/IP header and NAT

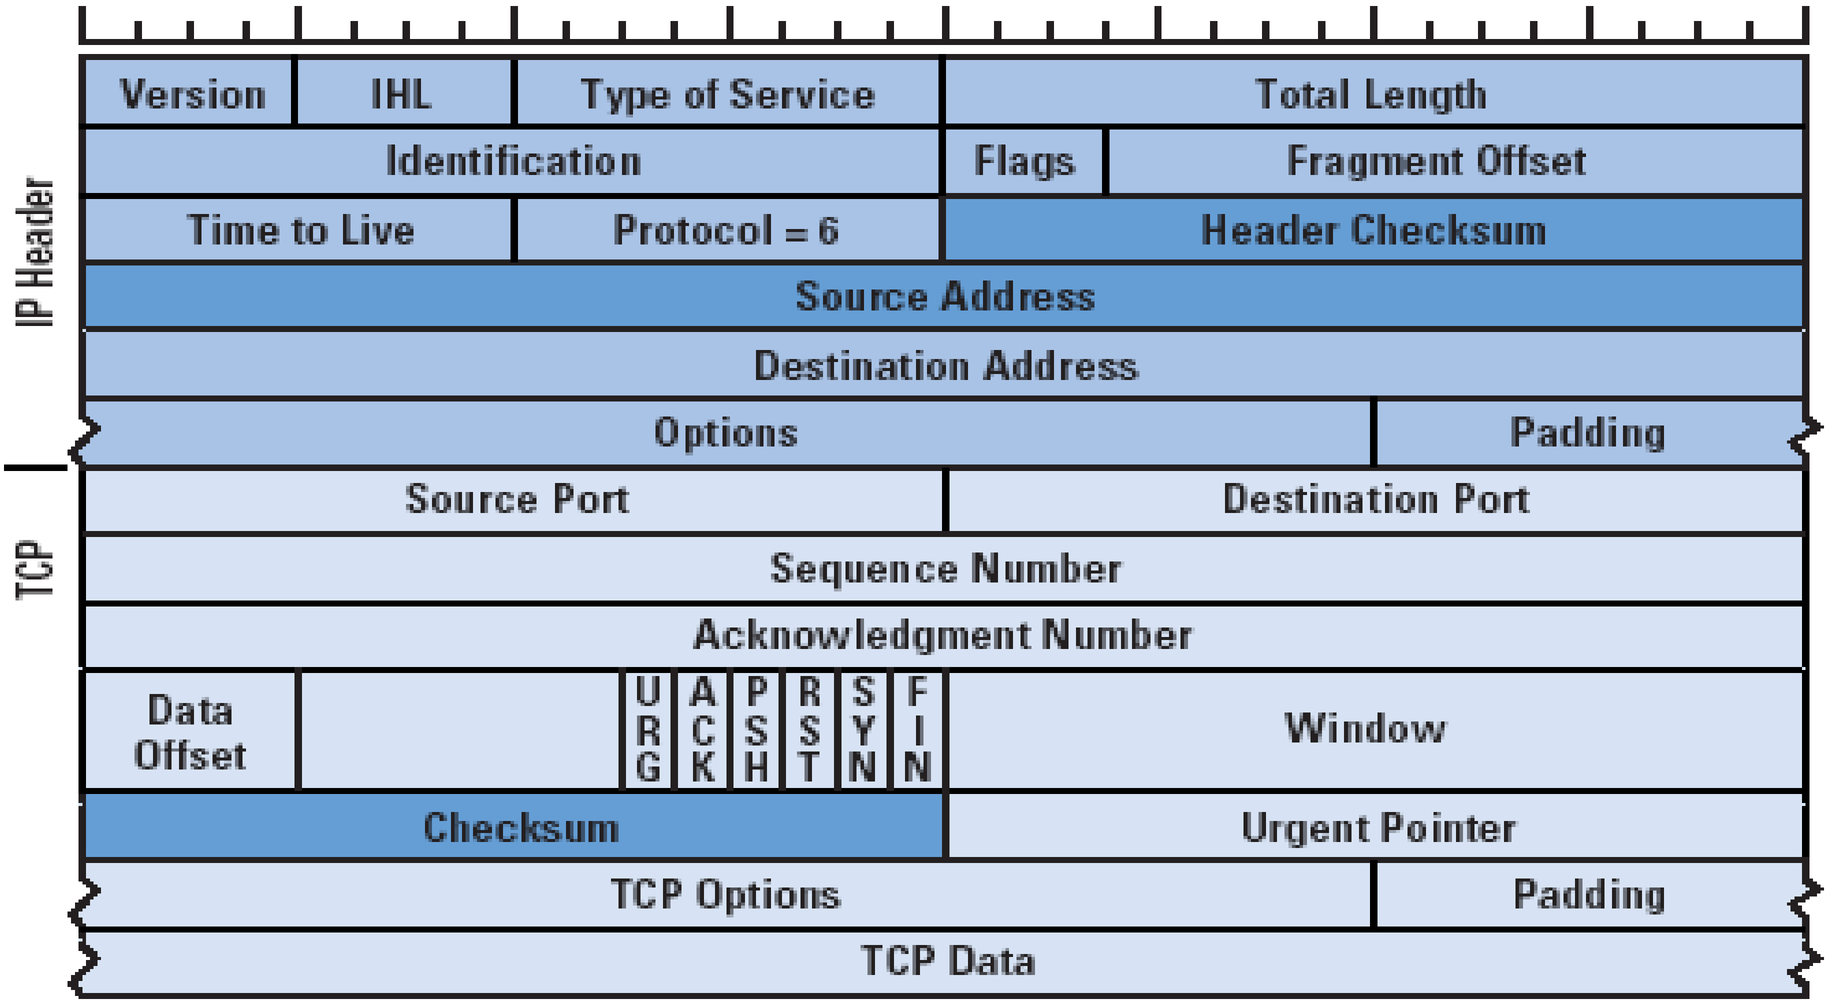

IP Header

Version | IHL | Type of Service | Total Length

Identification | Flags | Fragment Offset

Time to Live | Protocol = 6 | Header Checksum

Source Address

Destination Address

Options | Padding

TCP

Source Port | Destination Port

Sequence Number

Acknowledgment Number

Data Offset | | URG | ACK | PSH | RST | SYN | FIN | Window

Checksum | Urgent Pointer

TCP Options | Padding

TCP Data

Different NATs

- Static NAT - maps an unregistered IP address to a registered IP address on a one-to-one basis
- Dynamic NAT - maps an unregistered IP address to a registered IP address from a group (one ?) of registered IP addresses.
- Overloading - maps multiple unregistered IP addresses to a single registered IP address by using different ports (PAT)
- Overlapping - IP addresses used in the LAN are registered IP addresses in another network

PAT, NAT, NATP, etc.

PAT is similar to NAT, however, any IP address change involves the PAT device’s outside IP address rather than a pool of addresses as in NAT.

More of names: Hidden NAT (Check Point), SNAT/MASQUERADE (Linux iptables), Internet Connection Sharing (Microsoft), port level multiplexed NAT

"One-to-one NAT" or "basic NAT" or "static NAT" involves only address translation, not port mapping

Additional complexities

- TCP/UDP
- TCP handshake – SYN, ACK, FIN
- How to keep “holes” for UDP (gaming, P2P)
- TCP session ID
- Timeout or not timeout
- Stateless protocols - UDP
- DNS, DHCP
- FTP and SIP send network layer address information inside application payloads.
- TCP keep alive
- SCTP

TCP handshake

- SYN (32 bits initial sequence number)
- SYN-ACK
- ACK (SYN-ACK-ACK)
- Data exchange
- Termination 1 (most common): FIN, FIN&ACK, ACK

NAT and ICMP

Internet Control Message Protocol (ICMP) message contains part of the original IP packet in the body of the message, so for the NAT to behave as transparently as possible, the IP address of the IP header contained in the data part of the ICMP packet should be modified according to the NAT binding state, as well as the IP header Checksum field of this inner packet header.

```
INVITE sip:12125551212@211.123.66.222 SIP/2.0
Via: SIP/2.0/UDP 211.123.66.223:5060;branch=a71b6d57-507c77f2
Via: SIP/2.0/UDP 10.0.0.1:5060;received=202.123.211.25;rport=12345
From: <sip:2125551000@211.123.66.223>;tag=108bcd14
To: sip: 12125551212@211.123.66.222
Contact: sip: 2125551000@10.0.0.1
Call-ID: 4c88fd1e-62bb-4abf-b620-a75659435b76@10.3.19.6
CSeq: 703141 INVITE
Content-Length: 138
Content-Type: application/sdp
User-Agent: HearMe SoftPHONE

v=0
o=kayote 0 0 IN IP4 10.0.0.1
s=This is about the meeting
c=IN IP4 10.0.0.1
t=0 0
m=audio 8000 RTP/AVP 4
a=ptime:90
a=x-ssrc:00aea3c0
```

NAT and IP fragmentation

The TCP or UDP header is resident only in the initial IP fragment, and subsequent IP packet fragments do not contain a copy of the transport layer packet header.

- Perform the NAT translation only when the original IP packet has been reassembled (the packet is too large ?)
- Rely on a stored packet fragment translation state (out-of-order packets ?)

Pros and cons

End-to-end connectivity has been a core principle of the Internet

Prevents malicious activity initiated by outside hosts from reaching those local hosts

Practical solution to the impending exhaustion of IPv4 address space

NAT Traversal

- STUN
- ICE
- UPnP
- Bonjour
- SSL/TLS
- TISPAN

UPnP

Pushed by Microsoft. Using this technology , a client queries the NAT via UPnP asking what mapping it should use if it wants to receive on port x. The NAT responds with the address: port pair that someone on the public Internet would need to address were they to reach the client on this port x. Contrary to prevalent security policies, it is the UPnP client (and not the firewall) that controls the opening of pinholes to the outside world.

STUN

Simple Traversal of Udp through NATs is a protocol for setting up of NAT Probe and determining which kind of NAT the client is behind. When STUN server receives a packet, it returns a message from the same port to the source of the received packet containing the address:port pair that it sees as the source of that packet.

ICE

ICE empowers the endpoints to determine the types of NAT’s that exist between them and come up with a list of IP addresses through which the endpoints can communicate.

Full cone NAT

With full cone NAT, also known as one-to-one NAT, all requests from the same internal IP address and port are mapped to the same external IP address and port. An external host can send a packet to the internal host, by sending a packet to the mapped external address.

Full cone NAT

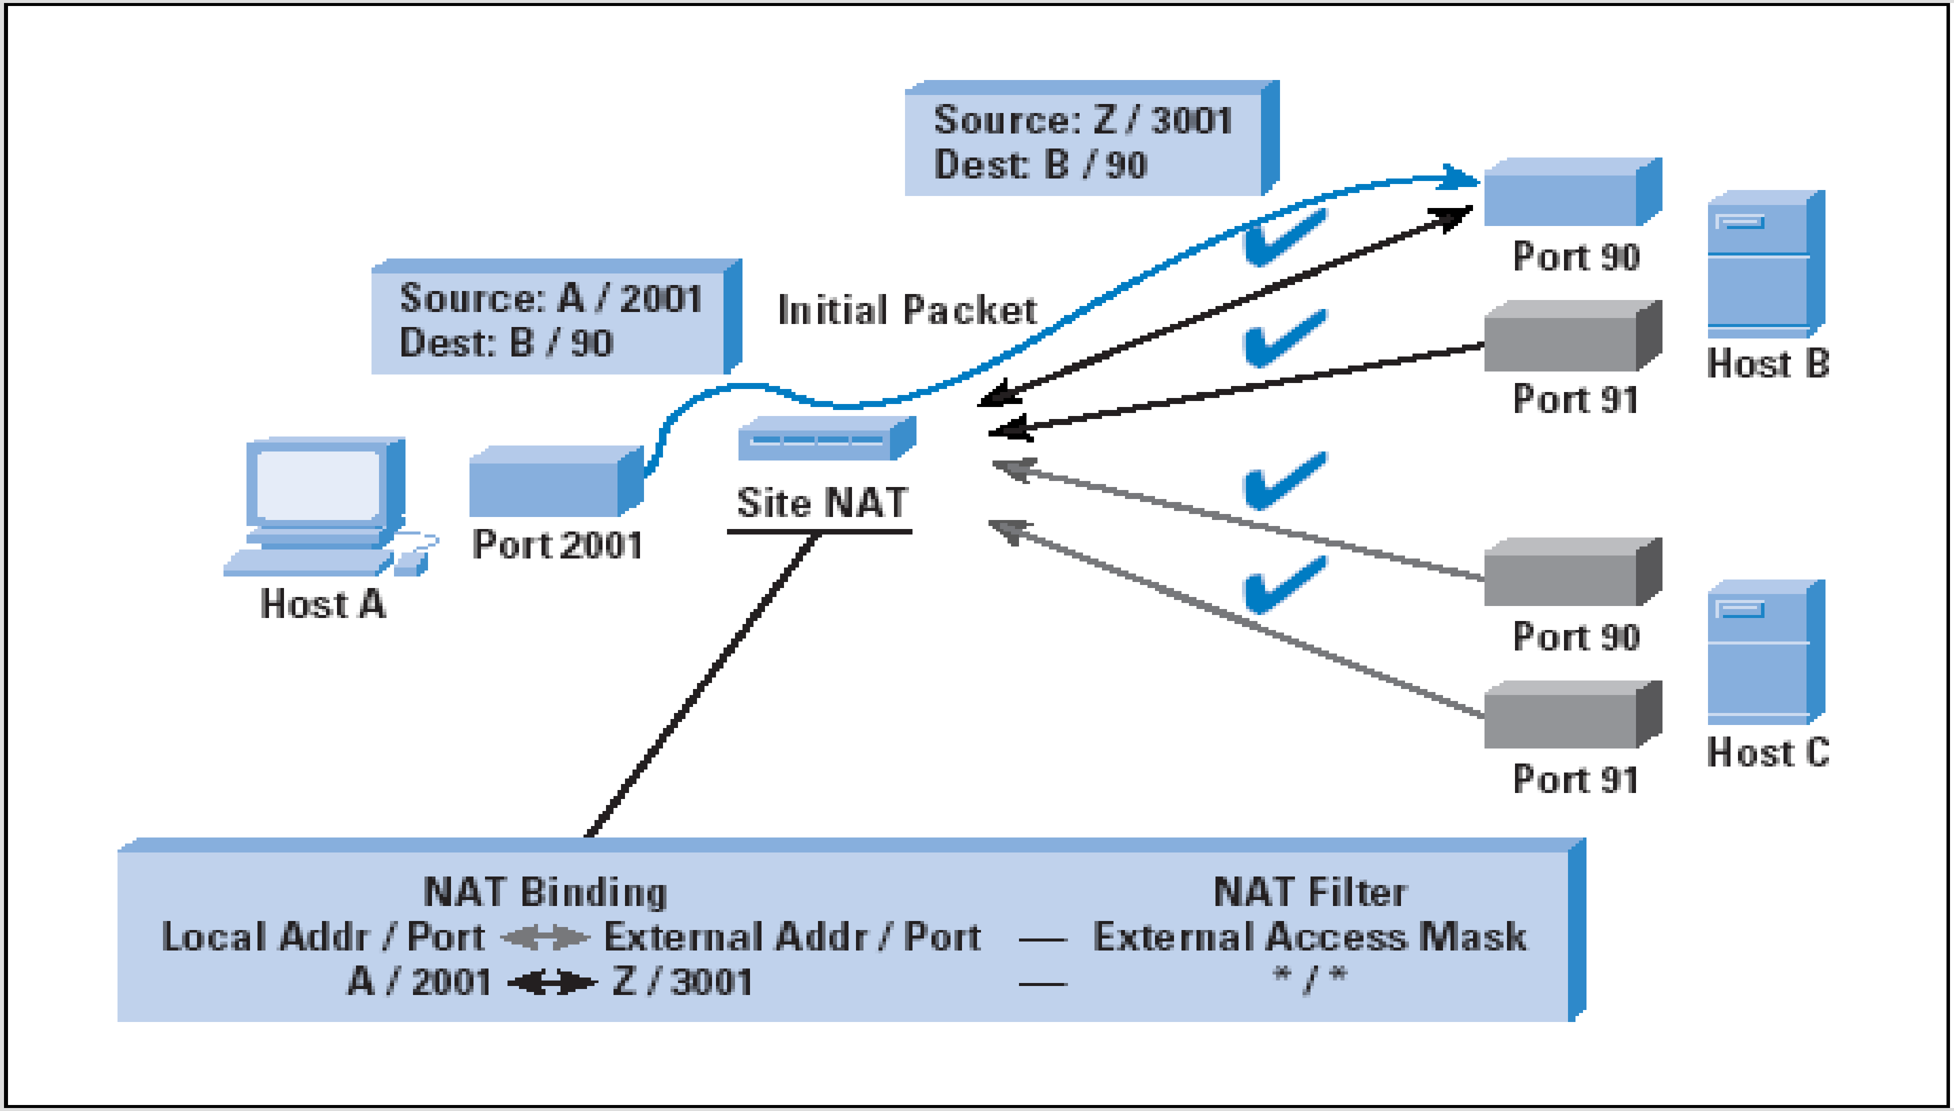

Restricted cone

With restricted cone NAT, all requests from the same internal IP address and port are mapped to the same external IP address and port. Unlike a full cone NAT, an external host can send a packet to the internal host only if the internal host had previously sent a packet to it.

Restricted cone

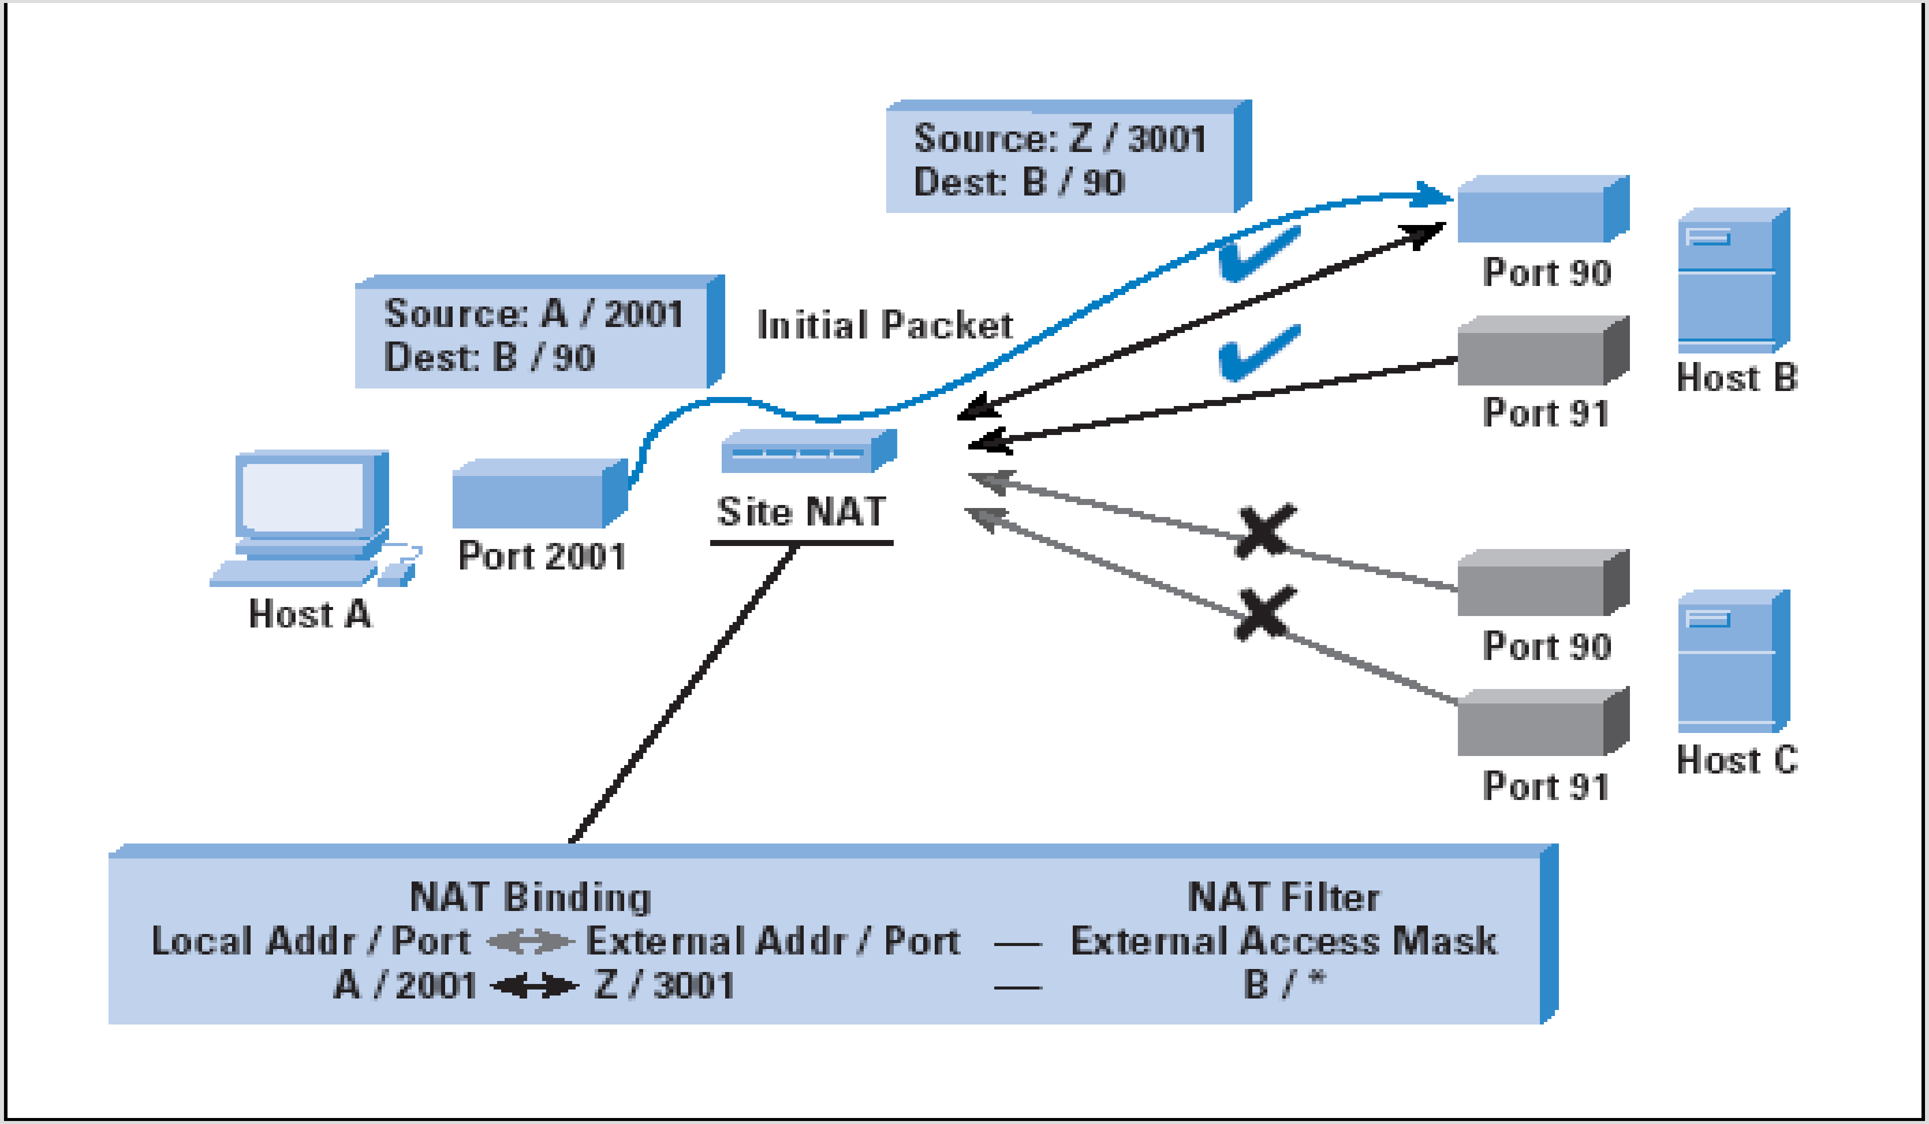

Port restricted cone

Port restricted cone NAT or symmetric NAT is like a restricted cone NAT, but the restriction includes port numbers. Specifically, an external host can send a packet to a particular port on the internal host only if the internal host had previously sent a packet from that port to the external host.

Port restricted cone

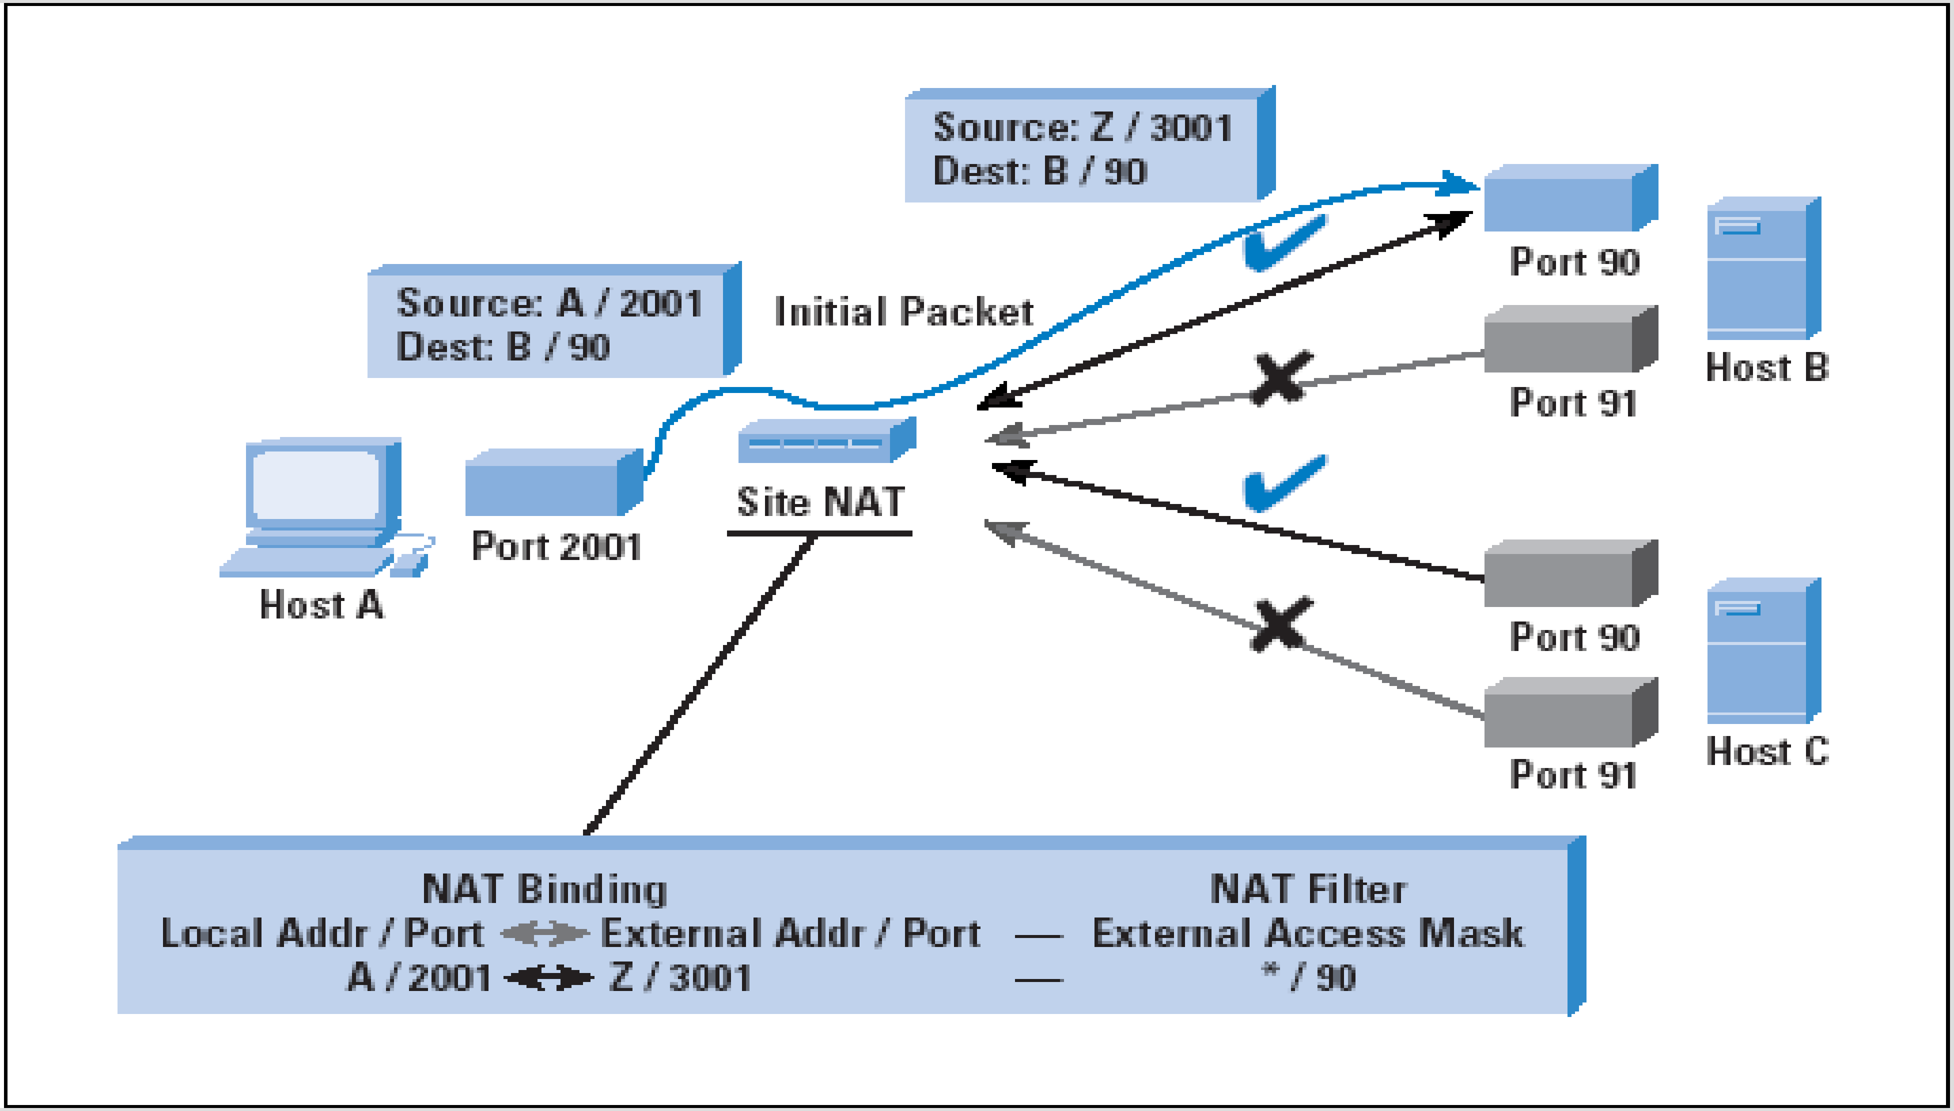

Source: Z / 3001
Dest: B / 90

Source: A / 2001
Dest: B / 90

Initial Packet

Host A

Port 2001

Site NAT

Port 90

Port 91

Host B

Port 90

Port 91

Host C

NAT Binding				NAT Filter
Local Addr / Port	↔	External Addr / Port	—	External Access Mask
A / 2001	↔	Z / 3001	—	* / 90

Symmetric NAT

With symmetric NAT all requests from the same internal IP address and port to a specific destination IP address and port are mapped to a unique external source IP address and port. If the same internal host sends a packet with the same source address and port to a different destination, a different mapping is used. Only an external host that receives a packet can send a UDP packet back to the internal host.

Symmetric NAT

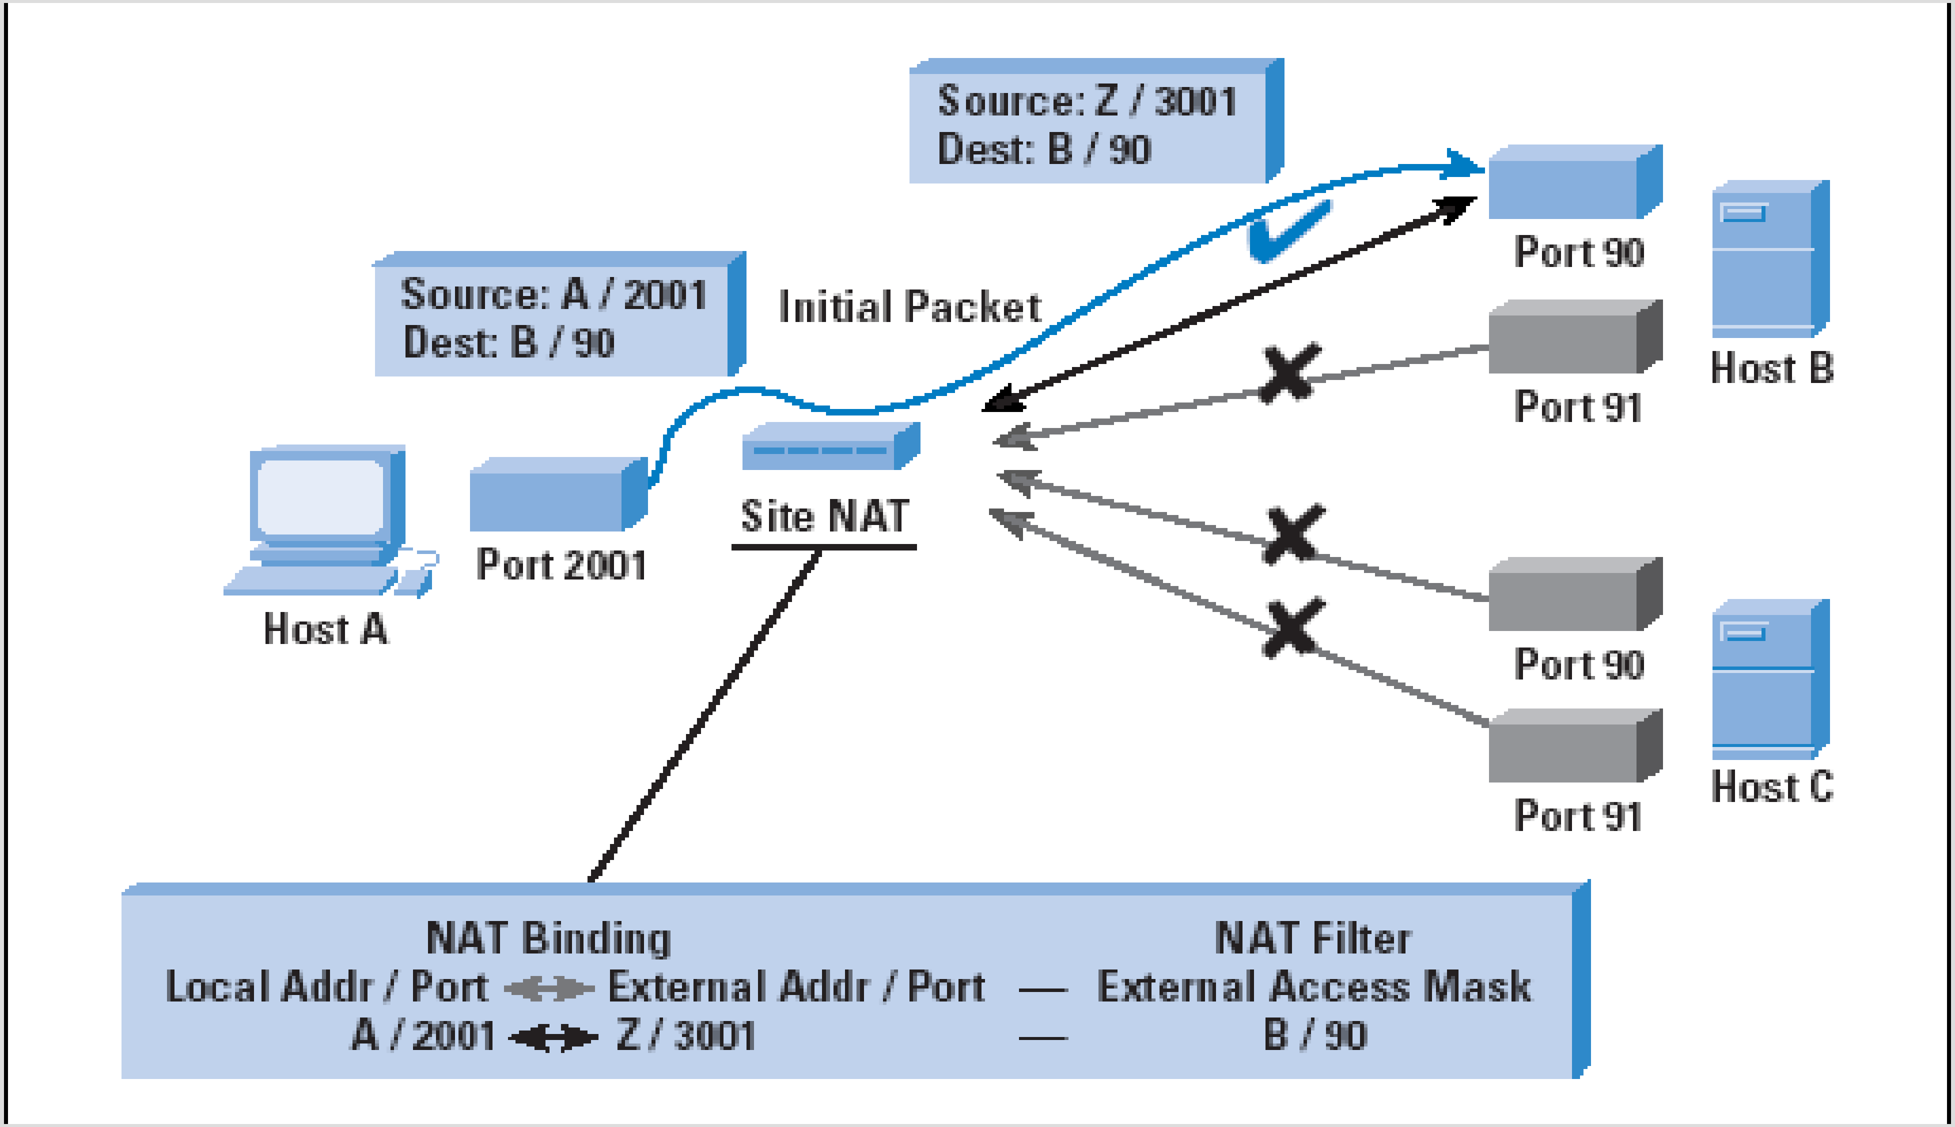

Source: Z / 3001
Dest: B / 90

Source: A / 2001
Dest: B / 90

Initial Packet

Host A

Port 2001

Site NAT

Port 90

Port 91

Host B

Port 90

Port 91

Host C

NAT Binding
Local Addr / Port ⟷ External Addr / Port
A / 2001 ⟷ Z / 3001

NAT Filter
— External Access Mask
— B / 90

Hairpins and Determinism

- Hairpin operation - local host directs a packet to the public address and port of an already mapped local host, or even to its own mapped address and port. The NAT bindings, when created, are available to either side of the NAT
- Nondeterministic NATs change their binding behavior when there is a binding conflict

Port preservation

Many NAT implementations follow a port preservation design. For most communications, they will use the same values as internal and external port numbers.

Important for P2P and VoIP applications

More schemes

- Port overloading - undertake port preservation at all times, so that when a different local host establishes a binding using a port that is already being preserved, the new binding will usurp the existing binding
- Port multiplexing - use external entity to demux the port. if two local systems use the same source port, the NAT preserves the source port in the two bindings

Binding timer refresh

- Bidirectional - packets from either the local hosts or an external host the NAT binding expiration time to be reset
- Outbound (inbound) - resets the expiration timer only when packets pass from the local host to the external host (keep alive is required ?)
- Transport Protocol state - timer is refreshed by any session packet in either direction (bidirectional), with the exception of packets with the TCP RST or FIN flags set (DoS ?)

Port forwarding

Port forwarding (sometimes referred to as tunneling) is the act of forwarding a network port from one network node to another. This technique can allow an external user to reach a port on a private IP address (inside a LAN) from the outside via a NAT-enabled router.

Applications

- Load Balancing: Destination NAT can redirect connections pointed at some server to randomly selected servers.
- Transparent proxying: NAT can redirect HTTP connections targeted at the Internet to a special HTTP proxy which can cache content and filter requests (used by ISPs)
- Multi-homing (Border Gateway Protocols IBGP, EBGP)

NAT and IPSec

We are going to talk about:

- Problem
- AH/ESP
- Transport/Tunneling
- TISPAN/TLS

Authentication header (AH)

In transport mode IPsec AH protects both the payload and the IP header fields and inserts a new header between the original IP header and the payload. In tunnel mode the whole IP packet is encapsulated within an AH and new IP header.

Encapsulating Security Layer (ESP)

Transport mode ESP encapsulates just the payload (compare to tunnel mode - encapsulates the whole IP packet). There is no check that the encrypted portion matches the non-encrypted portion, and so NATs can be readily traversed.

AH and EPS

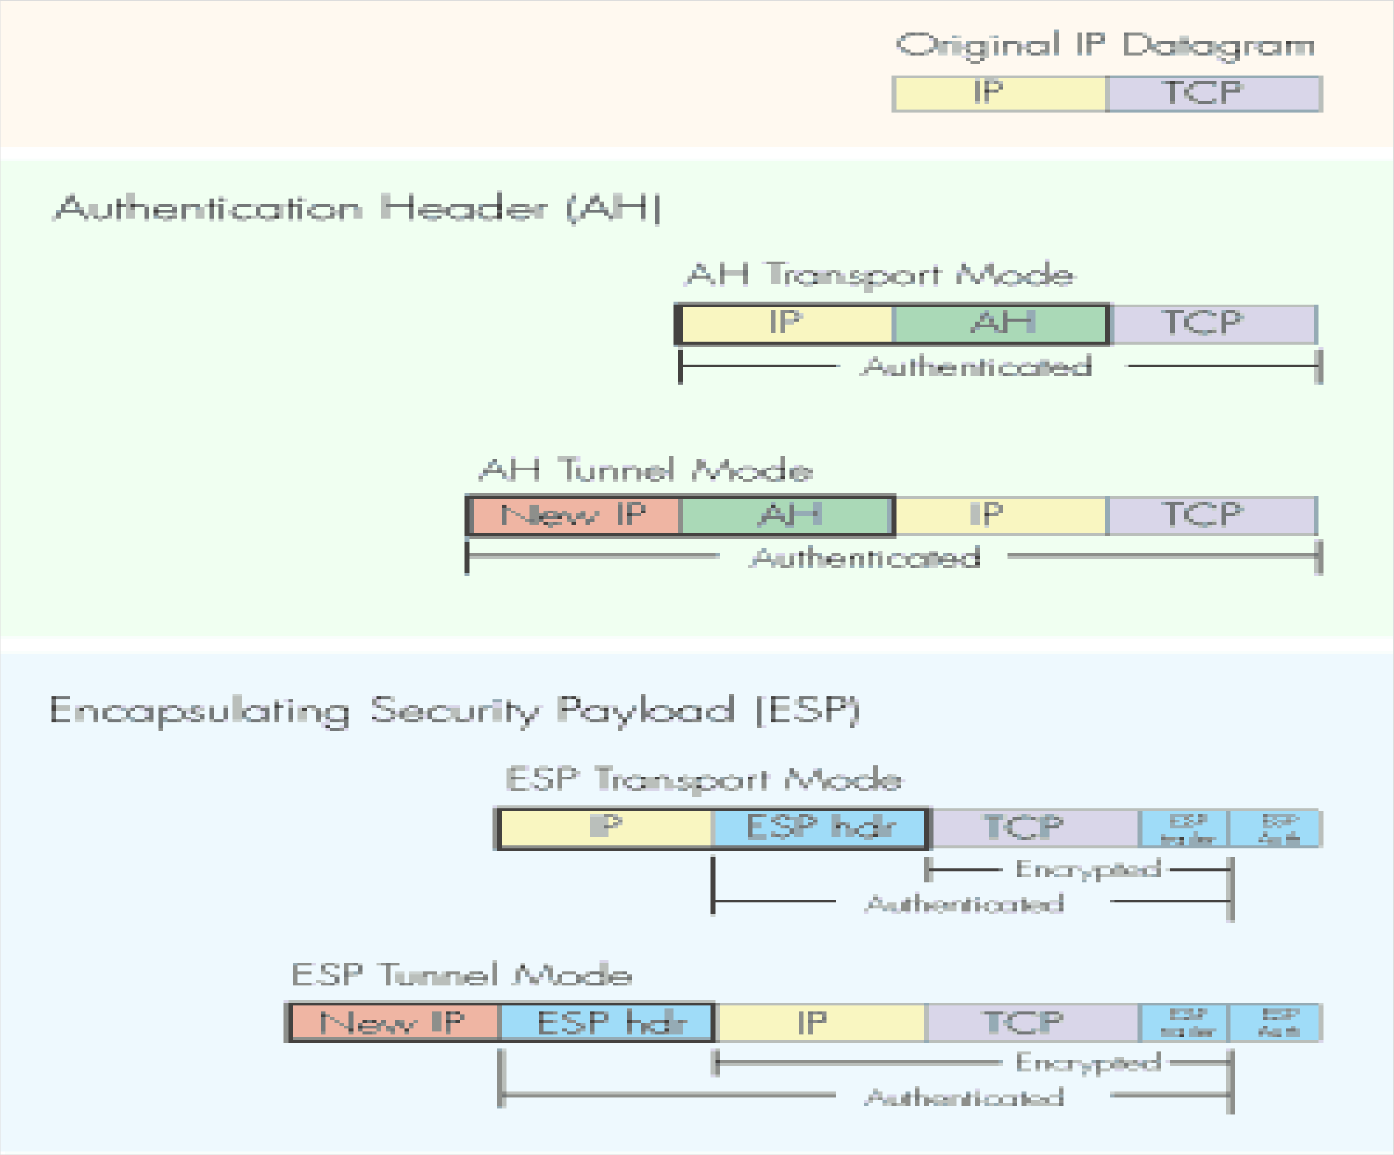
Original IP Datagram
IP
TCP
Authentication Header (AH)
AH Transport Mode
IP
AH
TCP
Authenticated
AH Tunnel Mode
New IP
AH
IP
TCP
Authenticated
Encapsulating Security Payload (ESP)
ESP Transport Mode
IP
ESP hdr
TCP
ESP trailer
ESP Auth
Encrypted
Authenticated
ESP Tunnel Mode
New IP
ESP hdr
IP
TCP
ESP trailer
ESP Auth
Encrypted
Authenticated

IPSec and NAT

IPsec assumes that the end-to-end connection does not have to traverse devices which alter the authenticated packets.

When a NAPT device encounters an IPsec ESP packet it no longer has access to the transport layer ports and will usually revert to NAT-only operation and translates the private IP address to the public IP address. This enables the single device to communicate through the NAT to the far end.

“VPN Compatibility”

In practice most NAPT products with a "VPN compatibility" mode use the binding created by the last outbound packet as the destination for inbound packets. Thus, one user will receive all the signalling.

Tunneling

A tunneling protocol is a network protocol which encapsulates one protocol or session inside another. Protocol A is encapsulated within protocol B, such that A treats B as though it were a data link layer. Tunneling may be used to transport a network protocol through a network which would not otherwise support it. Tunneling may also be used to provide various types of VPN functionality such as private addressing.

Transport or tunnel

In general transport mode is used to secure end-to-end communications between two devices, whilst tunnel mode is used to connect two networks

A gateway (router, firewall, etc.) is only required to support Tunnel mode. Supporting Transport mode is useful when creating an endpoint to the gateway itself, as in the case of network management functions.

TISPAN

The Telecoms & Internet converged Services & Protocols for Advanced Networks (TISPAN) focus is to define the European view of the Next Generation Networking (NGN).

TISPAN has agreed to use UDP encapsulated IPsec according to RFC 3948.

UDP encapsulation

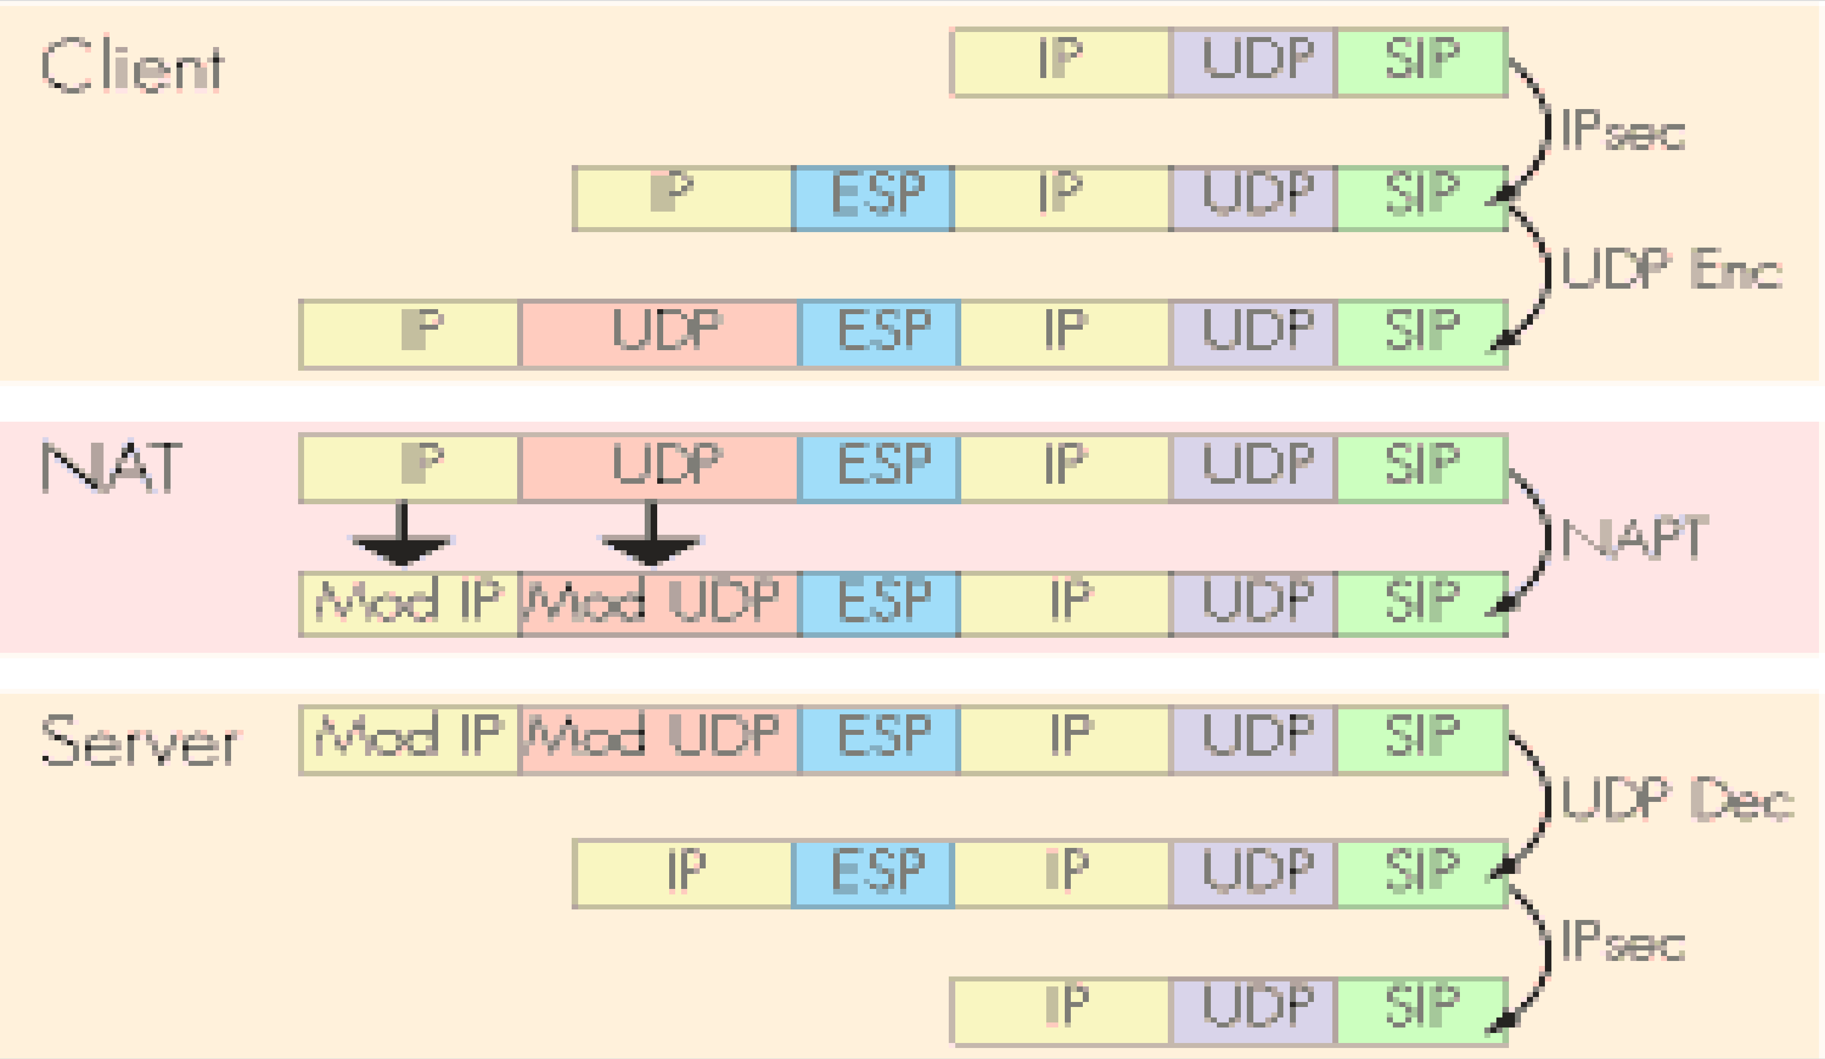

TLS tunnel

The TLS protocol exchanges records; each record can be optionally compressed, encrypted and packed with a message authentication code (MAC). Each record has a content_type field that specifies which upper level protocol is being used.

When the connection starts, the record level encapsulates another protocol, the handshake protocol, which has content type 22.

TLS security features

- Numbering all the records and using the sequence number in the MACs
- The message that ends the handshake ("Finished") sends a hash of all the exchanged data seen by both parties.
- The pseudorandom function splits the input data in half and processes each one with a different hashing algorithm (MD5 and SHA), then XORs them together if one of these algorithms is found to be vulnerable.

Proxy

A proxy server is a computer that offers a service allowing clients to make indirect network connections to other network services. A client connects to the proxy server, requests a connection, file, or other resource on a different server. The proxy provides the resource either by connecting to the specified server or by serving it from a cache. Proxy may alter the client's request or the server's response for various purposes.

Reversed proxy

A proxy server that is installed in the neighborhood of one or more web servers. All traffic coming from the Internet and with a destination of one of the web servers is going through the proxy server.

Why proxy ?

- Security
- Encryption / SSL acceleration
- Load balancing
- Serve/cache static content
- Compression
- Spoon feeding

P2P

A peer-to-peer (or P2P) computer network is a network that relies on the computing power and bandwidth of the participants in the network rather than concentrating it in a relatively low number of servers. P2P networks are typically used for connecting nodes via largely ad hoc connections.

P2P networks

- UDP – Gnutella, BearShare
- TCP – eDonkey, BitTorrent, Freenet, DirectConnect, Napster, Usenet, Kad
- From www.cachelogic.com : the biggest draw on network bandwidth is Peer-to-Peer file sharing applications consuming up to 75% of total off-network traffic. This creates a dilemma for Service Providers

BitTorrent

- Up to 60% of all Internet traffic
- Typically 40-80 coexisting TCP connections and up to 400-500 are required for 4Mbits/s symmetrical connection
- Popular content – TV shows and new releases
- Typical number of potential peers (size of swarm) – hundreds to thousands
- Typical connection time - 24/7
- Protocol is able to utilize any connection

ALG

An Application Layer Gateway (ALG) software module running on a NAT firewall device updates any payload data made invalid by address translation. ALGs obviously need to understand the higher-layer protocol that they need to fix, and so each protocol with this problem requires a separate ALG.

Internet security and Botnets

Botnet is a jargon term for a collection of software robots, or bots, which run autonomously or a collection of compromised machines running programs, usually referred to as worms, Trojan horses, or backdoors, under a common command and control infrastructure.

Botnets

- Oct 2005. The Dutch police found a 1.5 million node botnet
- Nov 2005. A Los Angeles man was arrested and charged with creating a 400,000-PC-strong botnet
- Spring 2006. A program was discovered at a foreign coast guard agency that searched for documents that had shipping schedules and forwarded them to an e-mail address in China
- Georgia Institute of Technology: botnet programs are present on about 11 percent of the more than 650 million computers attached to the Internet

How much is your computer worth ... to someone else ?

- Account theft - $0.29/mon
- Installation of adware - $0.67/mon
- DDoS for hire - $0.01/mon
- Spam relay proxies - $0.20/mon
- Internet dialers (USA) - $45/mon
- Click fraud - ?
- Password cracking - ?

10,000 infected systems at $1.17/mon puts the “bot herder” solidly into the upper middle class income range

Personal information in danger

- Summer 2006. The data collected during a 30-day period came on just one server from 793 infected computers and it generated 54,926 log-in credentials and 281 credit-card numbers. The stolen information affected 1,239 companies, he said, including 35 stock brokerages, 86 bank accounts, 174 e-commerce accounts and 245 e-mail accounts.

Viruses, worms, trojans, ...

- A computer worm is a self-replicating computer program. It uses a network to send copies of itself to other nodes and it may do so without any user intervention. It does not need to attach itself to an existing program (use 3rd party firewall).
- Virus - A computer virus is a computer program written to alter the way a computer operates, without the permission or knowledge of the user, by hiding in other program files. A true virus must do these two things: replicate and execute itself (use VM)
- Trojan horse is a program that contains or installs a malicious program - payload or 'trojan'. (change OS)

How do you know that you've been rooted ?

Check if there is a "start" icon in your left lower corner of the screen. if so - yes, chances are you have caught a virus.

Firewall

A logical barrier designed to prevent unauthorized or unwanted communications between sections of a computer network. Other names: Border protection device or BPD, packet filter.

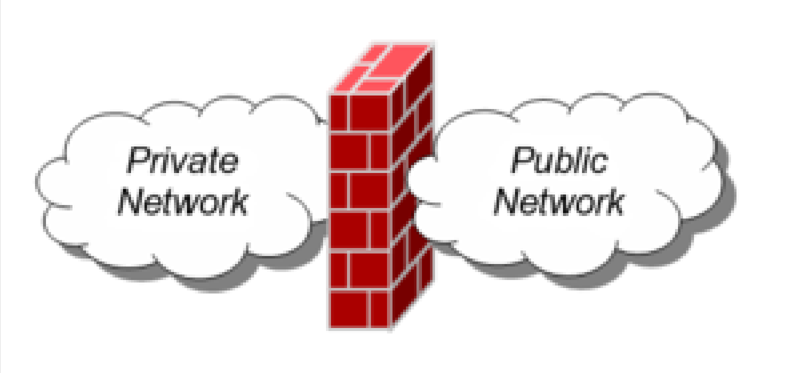

Firewall goal

The ultimate goal is to provide controlled interfaces between zones of differing trust levels through the enforcement of a security policy and connectivity model based on the least privilege principle and separation of duties.

“The primary function of a firewall is to reduce visibility, not add security”. Ideally, you should be able to turn your firewall off and still not be any more vulnerable (what OS you use ?)

History of firewalls

- Packet filters
- Circuit level firewalls
- Application layer (proxy)
- Deep packet inspection (Allot, P-Cube)
- Intrusion-prevention systems

Firewall families

- Stateful packet inspection firewalls (SPI)
- Stateless firewalls

Attack and defense

- DDoS and “Port knocking”
- Brute force SSH “cracking” and two-factor authentication

Second face of Firewalls

Internal (inside LAN) firewalls

- Typical application – University Campuses
- MAC filtering
- DHCP relay

Tunneling to bypass firewalls

Tunneling can also be used to bypass a system firewall. In this case, firewall-blocked data is encapsulated inside a commonly allowed protocol such as HTTP. One example of this type of use is HTTP-Tunnel.