Cryptography and Network Security II

Second Course

Lecture .2: Web security, IP security, Firewalls

Overview of the course

- I. CRYPTOGRAPHY
  - Secret-key cryptography
    - Classical encryption techniques
    - DES, AES, RC5, RC4
  - Public-key cryptography
    - RSA
  - Key management
- II. AUTHENTICATION
  - MAC
  - Hashes and message digests
  - Digital signatures
  - Kerberos
- III. NETWORK SECURITY
  - Email security
  - Web security (SSL, secure electronic transactions)
  - IP security
  - Firewalls
  - Wireless security
- IV. OTHER ISSUES
  - Viruses
  - Digital cash
  - Secret sharing schemes
  - Zero-knowledge techniques

Web security

Threats

Secure naming

Authenticated connections

Threats on the Web

- Reports that the webpages of a company or government institution have been replaced with some new pages
  - Yahoo, US army, CIA, NASA, New York Times, etc.
  - Minor harm done – webpages repaired quickly
- Denial of service attacks
  - Servers flooded with traffic (i.e., requests to answer to “ping” from illegal addresses – not only does the server need to answer the ping request but also to deal with its own returned ping answers)
  - Unable to continue normal business – big costs
- Swedish cracker 1999 – broke into Hotmail and created a mirror website that allowed anyone to type in username and password and then read all of that person’s emails
- Thefts of credit card numbers – used both for financial gains and as a war tool
- California student (23-year old) e-mailed a press release to a news agency falsely claiming that Emulex Corporation was going to post a large quarterly loss and CEO was resigning – company’s stock went down quickly causing stockholders to lose over 2 billion $ - the attacker gained some 250.000 $ in stock speculation

Table 17.1 A Comparison of Threats on the Web [RUBI97]

	Threats	Consequences	Countermeasures
Integrity	•Modification of user data •Trojan horse browser •Modification of memory •Modification of message traffic in transit	•Loss of information •Compromise of machine •Vulnerabilty to all other threats	Cryptographic checksums
Confidentiality	•Eavesdropping on the Net •Theft of info from server •Theft of data from client •Info about network configuration •Info about which client talks to server	•Loss of information •Loss of privacy	Encryption, web proxies
Denial of Service	•Killing of user threads •Flooding machine with bogus requests •Filling up disk or memory •Isolating machine by DNS attacks	•Disruptive •Annoying •Prevent user from getting work done	Difficult to prevent
Authentication	•Impersonation of legitimate users •Data forgery	•Misrepresentation of user •Belief that false information is valid	Cryptographic techniques

Secure naming

- A basic request: Alice needs to visit Bob’s website
  - She types in the URL into her browser and a few seconds later a webpage appears – is it really Bob’s page and how difficult is it really to fake that webpage?
  - Trudy may intercept all communications and she could return to Alice a fake page – e.g., Trudy could slash the prices on that webpage to convince Alice to use her credit card to buy something from “Bob”
    - This attack requires Trudy to tap Alice’s or Bob’s line
  - Much simpler ways are available: Trudy may break the DNS system to replace Bob’s IP address with an address of her choice – how difficult is this really?

DNS spoofing

- Assuming Trudy broke the DNS system, the attack could not be any simpler

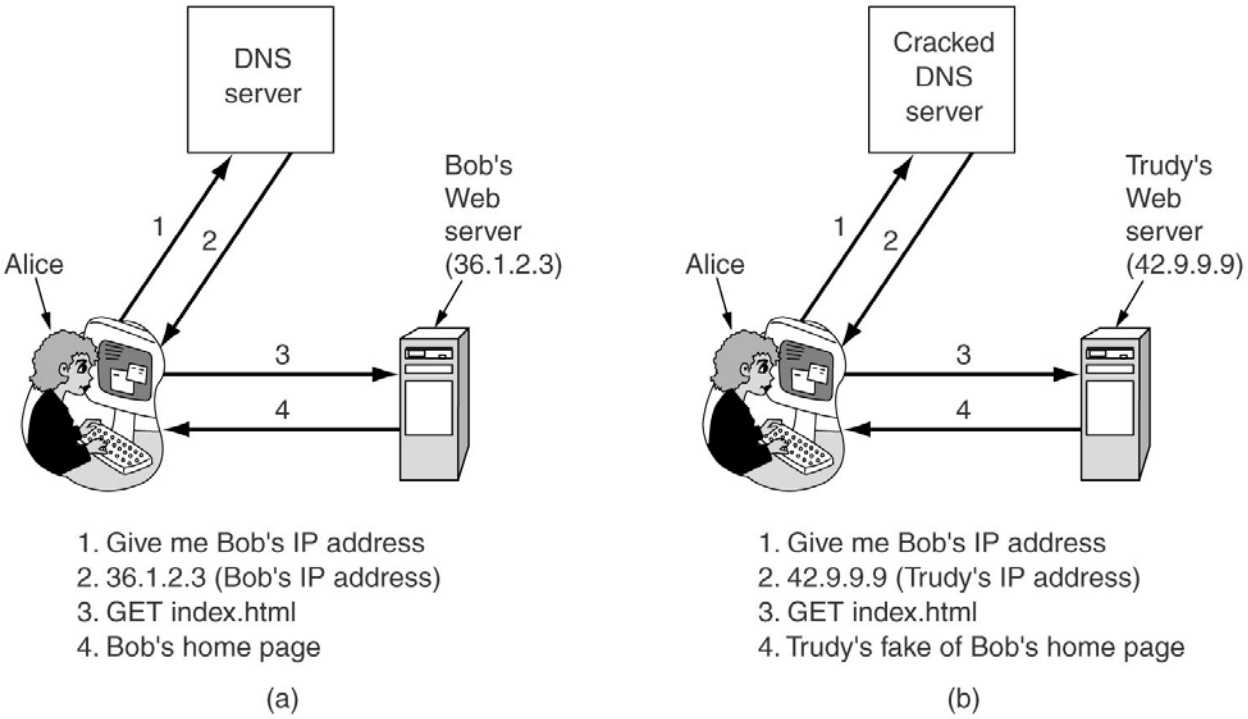

1. Give me Bob's IP address
2. 36.1.2.3 (Bob's IP address)
3. GET index.html
4. Bob's home page

(a)

1. Give me Bob's IP address
2. 42.9.9.9 (Trudy's IP address)
3. GET index.html
4. Trudy's fake of Bob's home page

(b)

Cracking DNS

- How can Trudy fool the DNS?
  - Trudy will trick the DNS to send a query to look up Bob's address
  - In reply, Trudy will inject a false IP address into the DNS server's cache
- Details of the attack (assuming the DNS does not have an entry for Bob's website – if it has one, Trudy will wait until it times out or use other tricks)
  - Trudy sends a request to Alice's Internet Service Provider (ISP) to look-up bob.com – DNS will query the top server for .com domain; each query has a certain sequence number attached to it
  - Trudy registers a domain trudy-the-intruder.com with IP address say 42.9.9.9 – her DNS server is dns. trudy-the-intruder.com with the same address
  - Trudy makes Alice's DNS aware of her domain – ask Alice's ISP for foobar.trudy-the-intruder.com – DNS will look-up dns.trudy-the-intruder.com, get it from the .com domain server and save it in the cache
  - Trudy will query Alice's ISP for www.trudy-the-intruder.com - DNS will send the query with the current query number
  - Immediately, Trudy ask Alice's ISP for www.bob.com
  - Trudy will immediately send herself a forged reply back to DNS with "bob.com is 42.9.9.9"; the reply will carry the sequence number with one higher (she may also send one with a sequence number two higher, etc – one of the forged replies is bound to be accepted as legitimate)
  - All the other replies (including the legitimate one) will be rejected by DNS since no query is outstanding anymore
- Whenever Alice looks-up bob.com, DNS will send her the IP address 42.9.9.9 – Trudy's
  - Successful attack with no line-tapping – Trudy can remain in her own living room!
  - Other ways exist also!

Cracking DNS

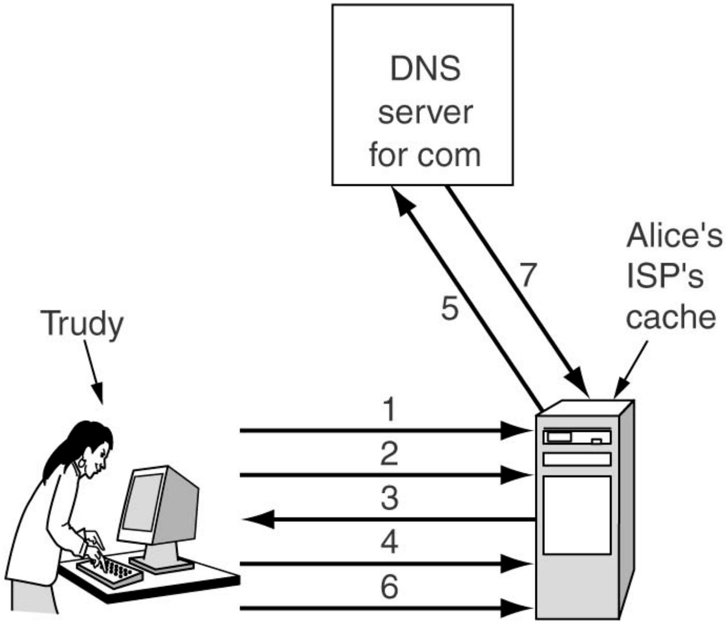

1. Look up foobar.trudy-the-intruder.com (to force it into the ISP's cache)
2. Look up www.trudy-the-intruder.com (to get the ISP's next sequence number)
3. Request for www.trudy-the-intruder.com (Carrying the ISP's next sequence number, n)
4. Quick like a bunny, look up bob.com (to force the ISP to query the com server in step 5)
5. Legitimate query for bob.com with seq = n+1
6. Trudy's forged answer: Bob is 42.9.9.9, seq = n+1
7. Real answer (rejected, too late)

Secure DNS

- This specific attack can be foiled by having DNS servers use random IDs in their queries rather than just counting
- 1994, IETF sets up a project on **DNSsec (DNS security) – RFC 2535**
- Concept of DNSsec
  - Every DNS zone has a public/private key pair
  - All information sent by a DNS server is signed with the originating zone's private key so that the receiver can verify its authenticity
- Services offered by DNSsec
  - Proof of where the data originated
  - Public-key distribution
  - Transaction and request authentication
- DNSsec has not been fully deployed yet – insecure servers exist and even secure servers have to be able to communicate with them
  - .org top-level domain signed with DNSsec in June 2010, .com, .net, .edu throughout 2010 and 2011
  - country top-level domains starting from May 2010
  - November 2011: more than 25% of top-level domains signed with DNSsec

Self-certifying names

- Other solutions than DNSsec exist for securing names
- Assume that each web server has a public/private key pair
- Central idea: each URL contains a SHA1-hash of the server’s name and public key as part of the URL – this would be a self-certifying URL
  - The hash will be represented by sequences of 32 digits and lower-case letter (do not include ‘l, ‘o’, ‘1’, ‘0’) – each character stands for 5 bits of the hash using some agreed-upon representation

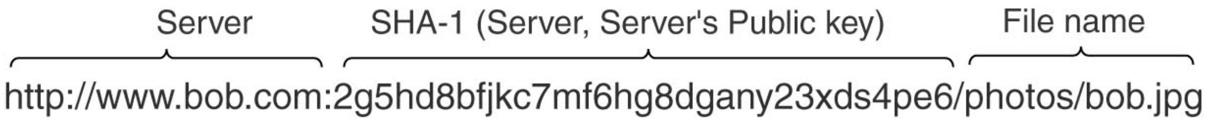

Server

SHA-1 (Server, Server's Public key)

File name

http://www.bob.com:2g5hd8bfjkc7mf6hg8dgany23xds4pe6/photos/bob.jpg

Self-certifying names

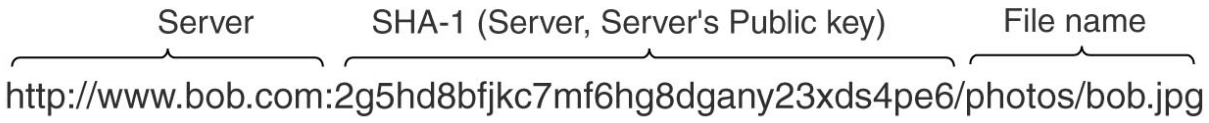

http://www.bob.com:2g5hd8bfjkc7mf6hg8dgany23xds4pe6/photos/bob.jpg

- Alice types in the above URL to get the jpeg she wants – her browser will ask for Bob's public-key and authenticate it based on the hash above
- To check that Bob has the corresponding private key, Alice constructs a message with an AES session key, a random number, and a timestamp, encrypted with Bob's public key
  - Bob replies encrypting the random number with the AES key
  - Alice knows it is Bob she is communicating with
- Alternatively, URLs could present X.509 certificates signed by a CA
- Also, one could do the above protocol with a self-certifying search engine and then have the engine provide various self-certifying URLs that the user can simply click on
- If Trudy performs the same attack of DNS spoofing, she can never read any secure traffic (because she has no access to Bob's private key), but she might achieve a denial-of-service attack

Secure connections

- Secure naming is a good start – next step is secure connections
- How does one ensure having a secure connection to an e-shop?
- 1995, Netscape Communications Corp introduced a security package called SSL (Secure Sockets Layer) for this aim – this is widely used nowadays
  - There is a TLS (transport layer security) working group within IETF to develop a common standard
  - First published version of TLS is essentially an SSLv3.1 – very often one says SSL/TLS
  - Discuss here SSLv3
- SSL builds a secure connection between two sockets (all of these discussed already)
  - Parameter negotiation between client and server
  - Mutual authentication of client and server
  - Secret communication
  - Data integrity protection
- When HTTP is used on top of SSL it is called HTTPS (secure HTTP)

SSL/TLS

- SSL (version 3) consists of two sub-protocols, for setting a secure connection, and for using it
- Connection establishment is shown to the right
  - Alice is sending her preferences regarding compression and cryptographic algorithms
  - Premaster key is a 384-bit key – combined with the two random numbers will generate a much larger internal key used for the encryption (e.g., with RC4)
  - Alice will know who Bob is but Bob will not – Bob could ask Alice to authenticate herself, but it is unlikely for all users to have public/private keys; a login/password could be required

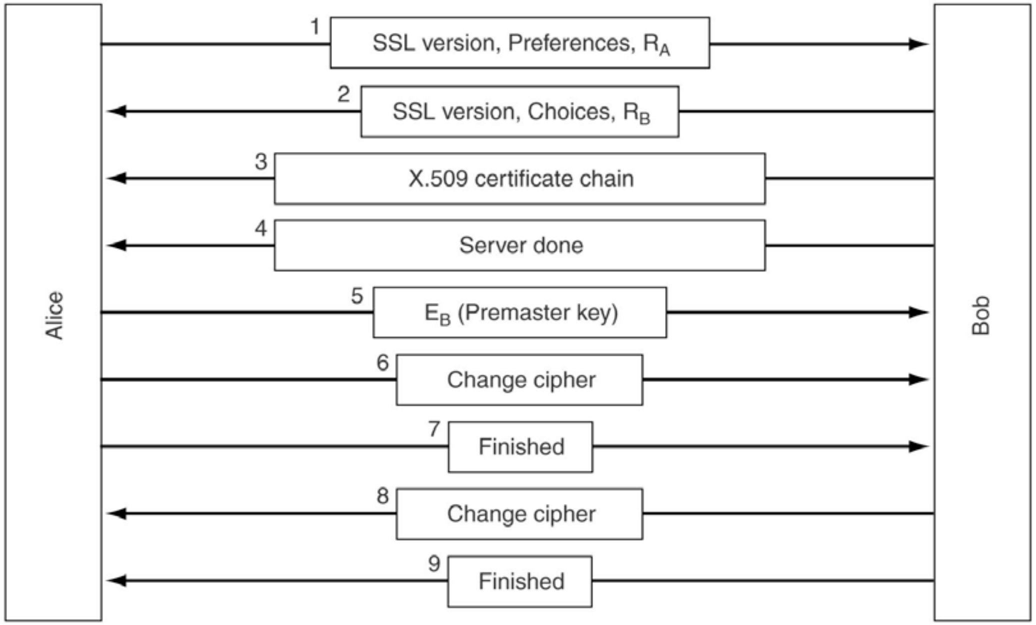

SSL/TLS

- Secure connections are used as shown on the right
  - Messages broken into parts of at most 16kb
- SSL supports multiple cryptographic algorithms
  - Strongest one is AES with 256-bit key for block ciphers; RC4-128 for stream ciphers; SHA-1 for data integrity
  - For normal e-commerce one usually uses RC4 with 128-bit key and MD5
  - The export version of SSL uses 128-bit key but 88 bits of those are made public – result is a 40-bit key encryption, more of a joke than a secure protocol
- 1996, SSL was turned over to IETF for standardization – result is TLS (Transport Layer Security) – also known as SSL 3.1 (1999)
  - Small changes with respect to SSL 3
  - Still, the two are incompatible!
  - Still unclear if TLS will replace SSL in practice

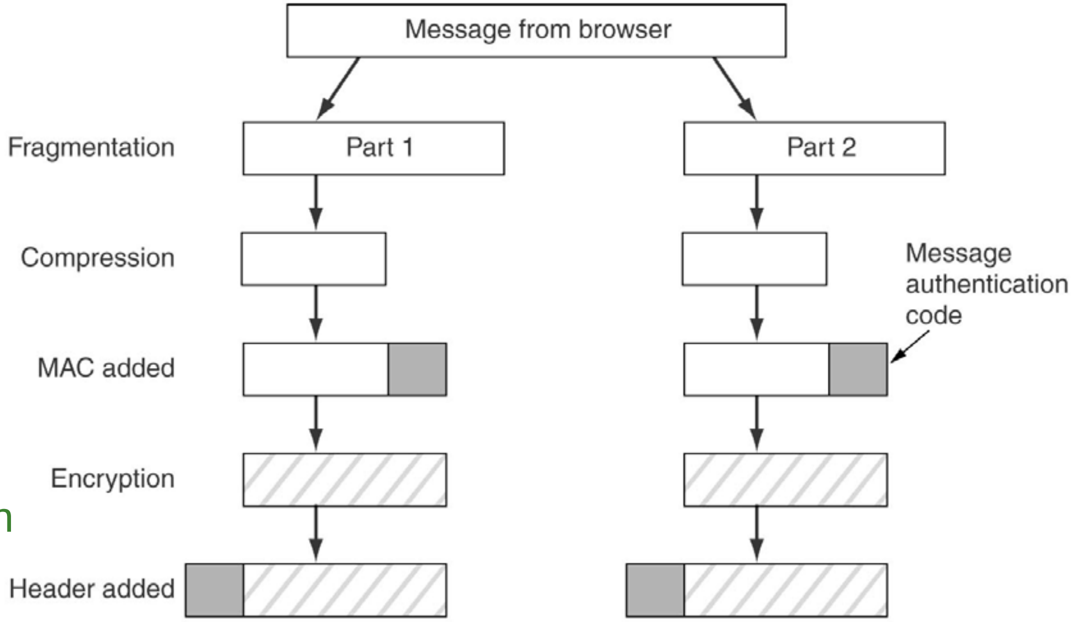

IP security

Security for security-ignorant applications

- Applications providing security have been developed for email (S/MIME, PGP), client/server (Kerberos), Web access (SSL), etc.
- *Question*: Can security be provided also for the security-ignorant applications?
- Also, it is desirable to implement security in such a way that users are not required to pass special training and are not required to do security tasks on a daily basis
- *Solution*: Implement security at IP level – in this way, all communications are secured and/or authenticated without changing the existing applications and protecting the security-unaware users

IPsec

- Issued in 1998, the design for a new network layer was called IPsec (IP security) – RFC 2401, 2402, 2406, and others
  - IPsec can encrypt and/or authenticate all traffic at IP level – remote logon, client/server, email, file transfer, Web access, etc.
  - Major services of IP sec: secrecy, data integrity, and protection from replay attacks – all of these are based on symmetric-key cryptography
  - Flexible design: one cryptographic protocol can be easily replaced with another, a null algorithm is also specified for those users who really do not want (cannot afford) security
- Applications of IPsec
  - Secure branch office connectivity over the Internet: a company can build a secure virtual private network over the Internet or over a public WAN – this enables a business to rely heavily on Internet rather than on private networks
  - Secure remote access over the Internet: an end user having a system equipped with IPsec protocols can connect to the company’s servers using the local Internet service provider – reduces costs for traveling employees
  - Enhances e-commerce security – use of IPsec beneath the usual security protocols should enhance the security and the public trust in this industry

IPsec scenario

- The company maintains various LANs with nonsecure IP traffic inside the LAN
- For traffic offsite, secure IP traffic is used
- The protocols are used in the networking devices, e.g., in the firewalls – these devices will typically encrypt/decrypt all traffic – these operations are transparent to the other applications

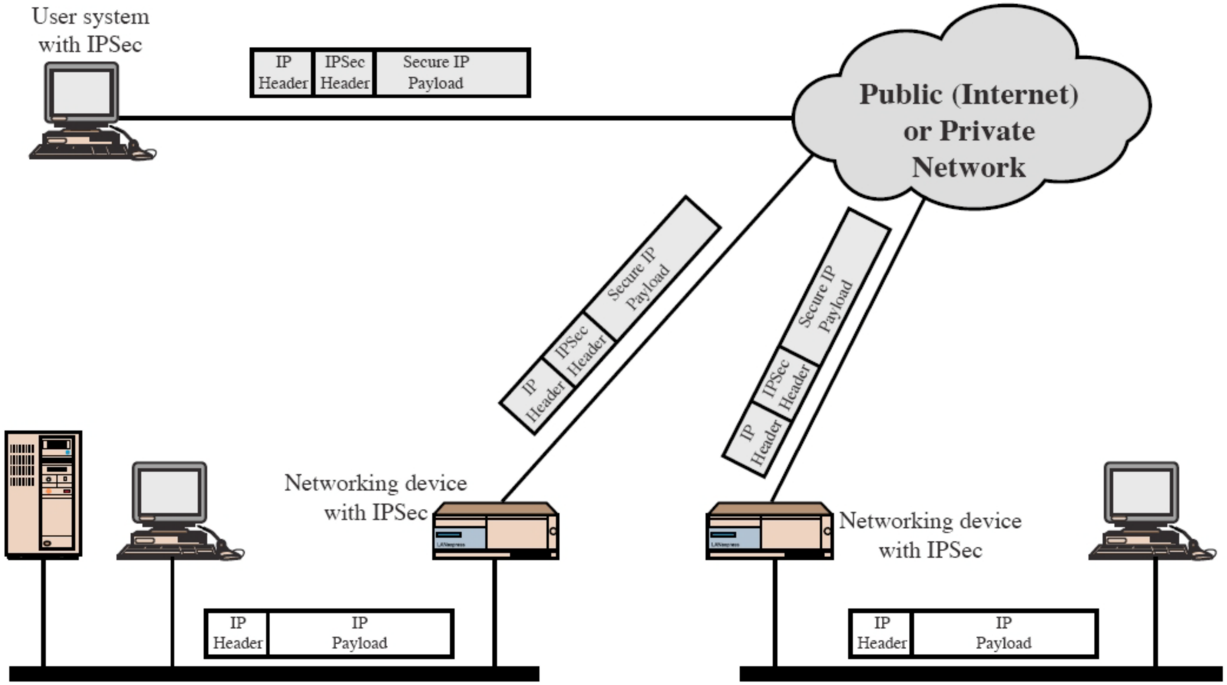

Figure 16.1 An IP Security Scenario

Briefly about the Internet Protocol (IP)

- IPv4 has been the keystone of the TC/IP protocol for many years
  - Header shown to the right – length 20 bytes
  - In particular, IPv4 only allows for addresses on 32 bits
  - Feb 2011: last batch of address blocks assigned
- 1996 – specifications for a next-generation IP: IPv6
  - Header shown bellow – length 40 bytes
  - Addresses on 128 bits
  - Longer header but fewer fields – faster processing in the routers
  - Migration to IPv6 expected to take many years
  - Oct 2011: 3% of domain names, 12% of the networks on the internet have IPv6 support
- IP sec optional in IPv4, mandatory in IPv6

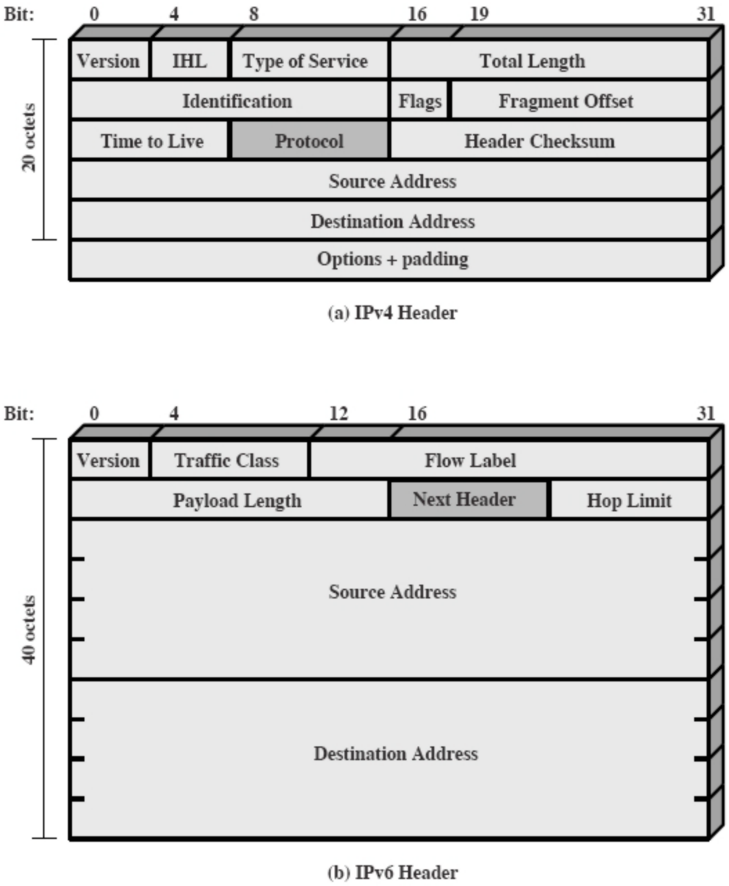

(a) IPv4 Header

(b) IPv6 Header

Figure 16.14 IP Headers

IP sec

- IP sec has two parts
  - The first describes two new headers that can be added to packets to carry security identifiers, integrity control data, and other information
  - The second part, ISAKMP (Internet Security Association and Key management Protocol) deals with establishing keys
- Surprisingly, even though IP sec is in the IP layer (below the transport layer), it is connection-oriented: to have security, a key has to be established and used for some period of time – a kind of connection
  - Connection in context of IP sec is called SA (Security Association) – a simple connection between two end points having a security identifier associated to it
  - SA is unidirectional – if secure traffic is needed in both directions, two SAs will be setup

IP sec – two new headers and modes of operation

- The two new headers specified in IP sec are AH (Authentication header) and ESP (Encapsulating Security Payload)
  - AH (the older one) only allows authentication, but no secrecy
  - ESP allows secrecy and optionally, also authentication
- IP sec can be used (both AH packets and ESP packets) in two modes
  - Transport mode: the IP sec header is inserted just after the IP header – this contains the security information, such as SA identifier, encryption, authentication
    - Typically used in end-to-end communication
    - IP header not protected
  - Tunnel mode: the entire IP packet, header and all, is encapsulated in the body of a new IP packet with a completely new IP header
    - Typically used in firewall-to-firewall communication
    - Provides protection for the whole IP packet
    - No routers along the way will be able (and will not need) to check the content of the packet

IP sec – Authentication Header

- AH – the first new header in IP sec
- Provides integrity checking and anti-replay security, but no encryption
  - *Next header*: the type of header immediately following this header (most often it is 6 for TCP)
  - *Payload length*: the number of 32-bit words in AH minus 2
  - *SPI*: identifies the security association (SA) to be used
  - *Sequence number*: used to number all packets sent on an SA – all packets get a unique number, even retransmissions, to detect replays – if all 2^{32} are exhausted, a new SA must be setup
  - *Authentication data*: variable-length field containing the payload's digital signature – the algorithm to be used and the shared key are negotiated when SA is established. Usually HMAC is used here.

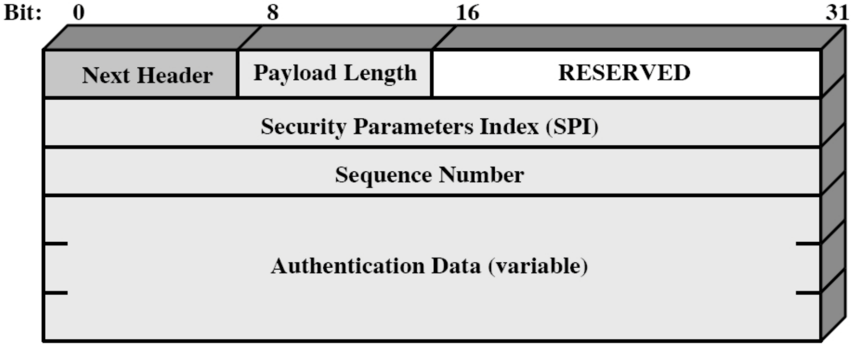

Bit: 0 8 16 31

Next Header | Payload Length | RESERVED

Security Parameters Index (SPI)

Sequence Number

Authentication Data (variable)

Figure 16.3 IPSec Authentication Header

IP sec – scope of AH authentication

- Shown to the right the AH authentication for IPv4 and IPv6 in transport mode and tunnel mode
  - Authentication covers only those fields that do not move in routing – those that can change are set to 0 when computing the HMAC

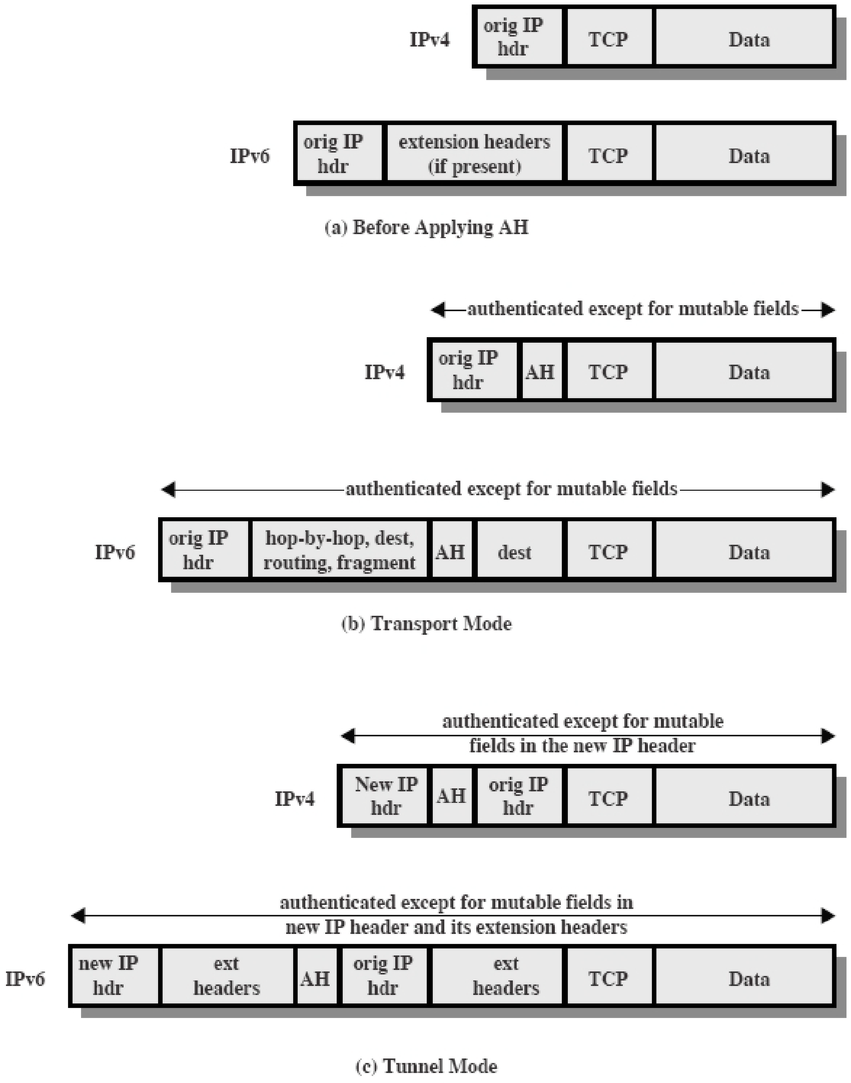

IPv4
orig IP hdr
TCP
Data
IPv6
orig IP hdr
extension headers (if present)
TCP
Data
authenticated except for mutable fields
IPv4
orig IP hdr
AH
TCP
Data
authenticated except for mutable fields
IPv6
orig IP hdr
hop-by-hop, dest, routing, fragment
AH
dest
TCP
Data
authenticated except for mutable fields in the new IP header
IPv4
New IP hdr
AH
orig IP hdr
TCP
Data
authenticated except for mutable fields in new IP header and its extension headers
IPv6
new IP hdr
ext headers
AH
orig IP hdr
ext headers
TCP
Data

(a) Before Applying AH

(b) Transport Mode

(c) Tunnel Mode

Figure 16.6 Scope of AH Authentication

IP sec – Encapsulating Security Payload

- The second new header in IP sec is ESP (Encapsulating Security Payload) – provides security and optionally, also authenticity
  - *Payload data*: the protected part (encrypted and possibly signed)
  - *Padding*: required for the encryption, also needed to protect the real length of the packet
  - *Pad length*: number of pad bytes
  - *Next header*: type of data contained in the payload data (e.g., TCP)
  - *Authentication data*: contains the digital signature of all packet minus the authentication data field

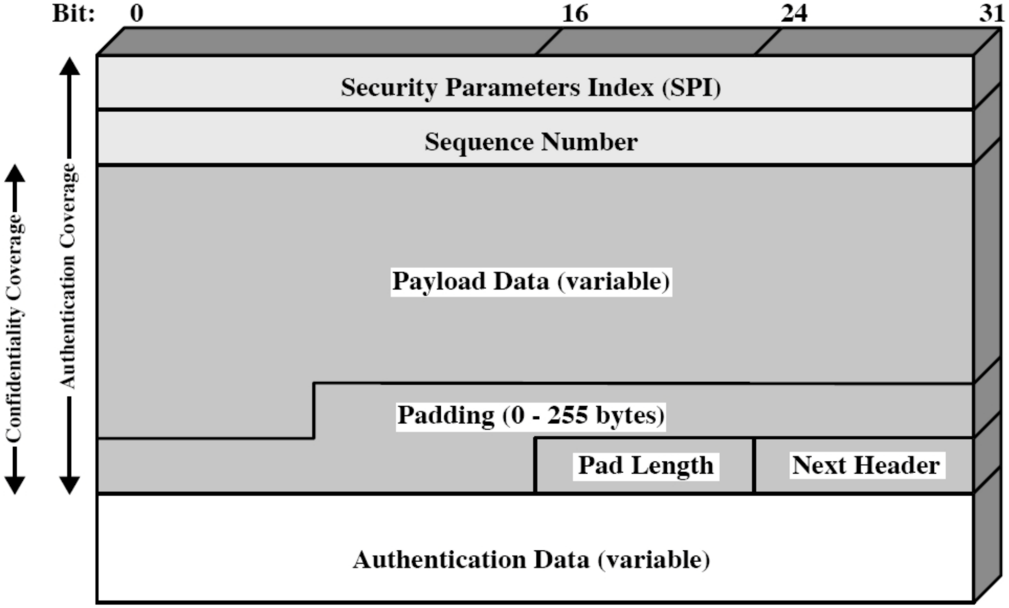

Figure 16.7 IPSec ESP Format

IP sec – the EPS header

- Encryption uses AES in counter mode (see lecture 4)
- Transport and tunnel mode ESP shown in the right
- The signature is placed at the end – better for hardware implementations – compute as the bits arrive and compare values in the end

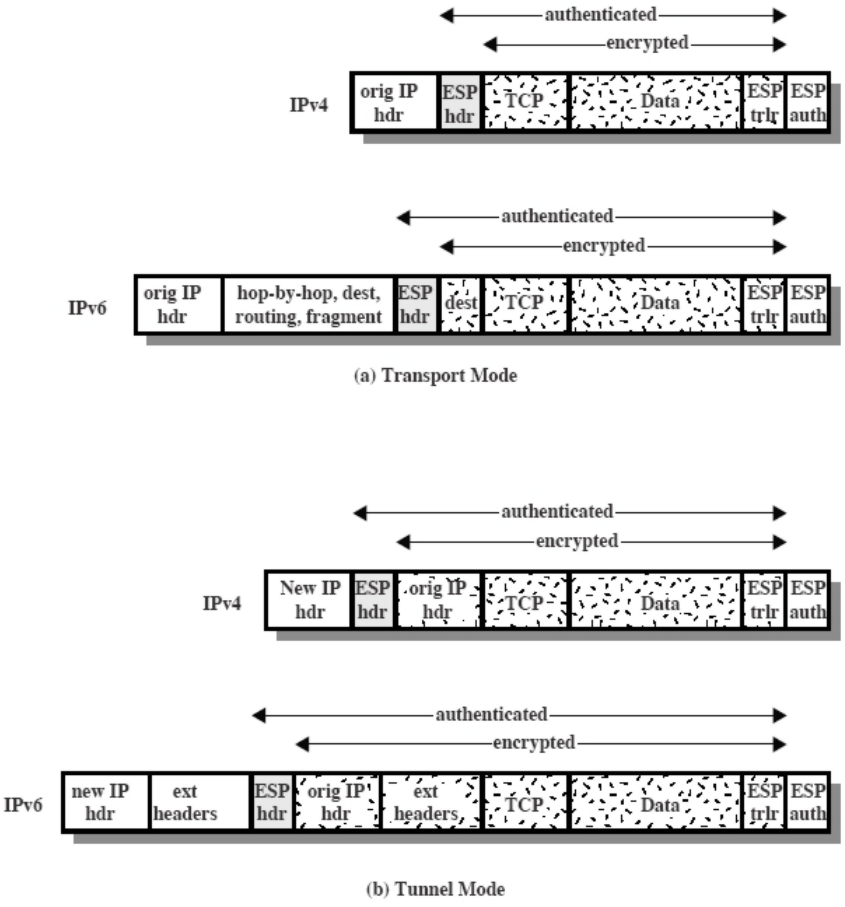

Figure 16.9 Scope of ESP Encryption and Authentication

Combining Security Associations

- Case 1: security provided for end-users implementing IP sec
- Case 2: security provided only between gateways
- Case 3: adds to Case 2 end-to-end security
- Case 4: support for a remote host user using the Internet to reach the company's firewall and then access resources behind the firewall

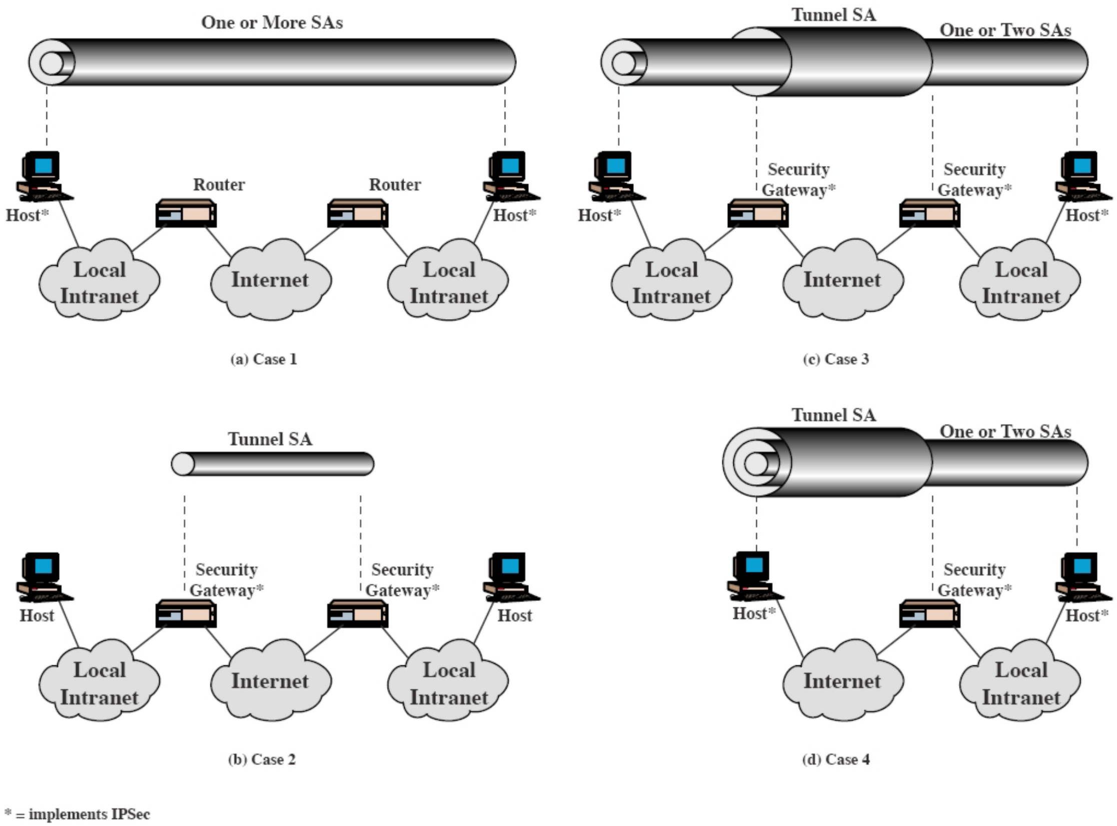

Figure 16.10 Basic Combinations of Security Associations

Firewalls

Firewalls

- Firewalls are a modern adaptation of the old medieval security standby: dig a deep moat around your castle – anybody entering or leaving the castle must pass over a single drawbridge where they can be inspected
- Same trick with networks: many LANs can be interconnected arbitrarily, but all traffic to or from the company is forced through an electronic drawbridge – firewall, see example bellow

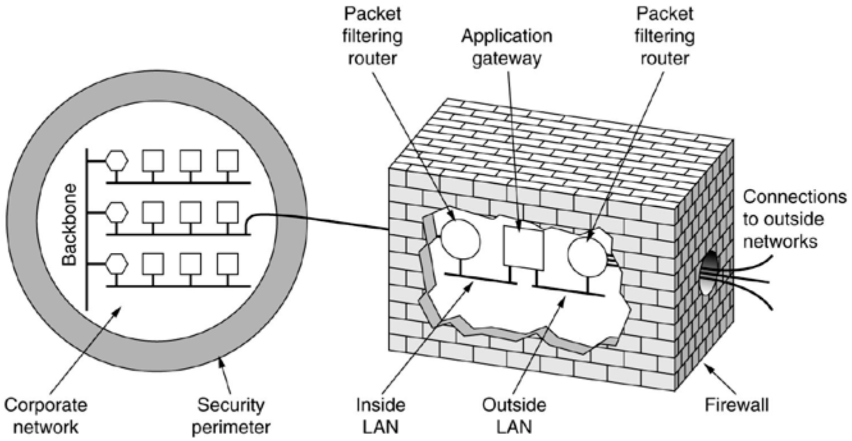

Firewall characteristics

- All traffic from inside to outside and vice-versa must pass through the firewall – this is achieved by physically blocking all access to the local network except via the firewall
- Only authorized traffic, as defined by the local security policy, is allowed to pass
- Firewall itself is immune to penetration – use a trusted system with a secure operating system

Types of firewalls

- Packet-filtering router
- Application-level gateway
- Circuit-level gateway

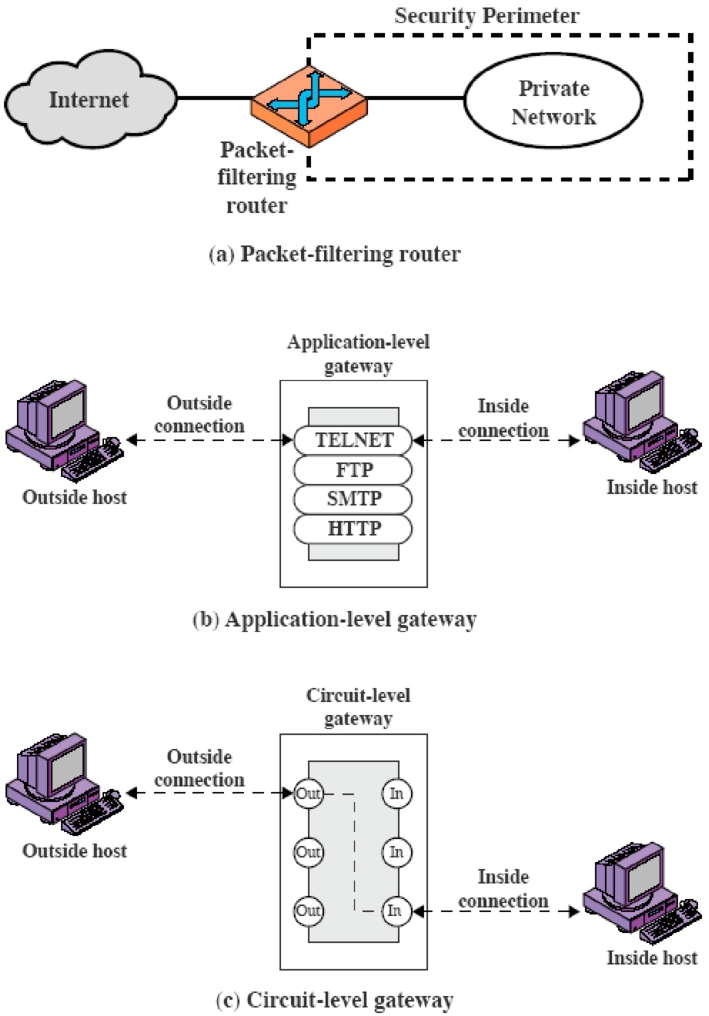

Figure 20.1 Firewall Types

Types of firewalls

- **Packet-filtering router**
  - Applies a set of rules to each incoming or outgoing IP packet and then forwards or discards the packet

Examples

A. Inbound mail is allowed (port 25 is for SMTP incoming) but only to a gateway host, mail from a particular external host (SPIGOT) is blocked

B. Default policy – this is always the last rule in all rule sets

C. Any inside host can send mail to the outside – an outside machine could attack by having other applications linked to port 25

D. Solves the problem above – it allows incoming packets on port 25 that include ACK flag in the TCP segment

E. To deal with FTP connections

Table 20.1 Packet-Filtering Examples

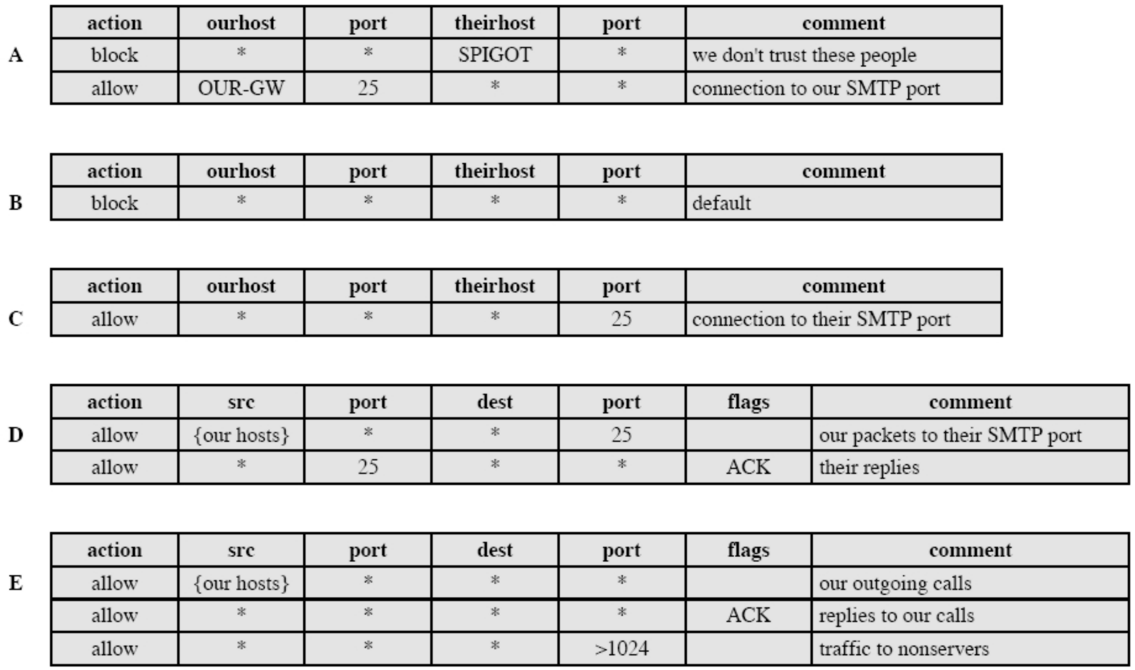

	action	ourhost	port	theirhost	port	comment
A	block	*	*	SPIGOT	*	we don't trust these people
	allow	OUR-GW	25	*	*	connection to our SMTP port

	action	ourhost	port	theirhost	port	comment
B	block	*	*	*	*	default

	action	ourhost	port	theirhost	port	comment
C	allow	*	*	*	25	connection to their SMTP port

	action	src	port	dest	port	flags	comment
D	allow	{our hosts}	*	*	25		our packets to their SMTP port
	allow	*	25	*	*	ACK	their replies

	action	src	port	dest	port	flags	comment
E	allow	{our hosts}	*	*	*		our outgoing calls
	allow	*	*	*	*	ACK	replies to our calls
	allow	*	*	*	>1024		traffic to nonservers

Types of firewalls

- **Application-level gateway** – also called a proxy server
  - Acts as a relay of application-level traffic
  - User contacts the gateway using a TCP/IP application (say *ftp*) – the gateway allows the connection only if it supports that specific application
  - More secure than packet filters – rather than checking many combinations on the TCP and IP level, the application-level gateway supports a number of applications and nothing else
  - Disadvantage – extra processing on each connection

Types of firewalls

- **Circuit-level gateway** – either a stand-alone system or a function performed by an application-level gateway for certain applications
  - Sets up two TCP/IP connections: one between itself and the outside host and one between itself and the inside host
  - No direct connections are allowed
  - Checks which connections are allowed

Bastion hosts

- These are systems identified by the system administrator as critical points for the network’s security – typically hosts for an application-level or circuit-level gateway
  - Bastion host hardware executes a secure version of its operating system, making it a ***trusted system***
  - Only the essential services are installed on the bastion host
  - Each proxy maintains detailed audit information by logging all traffic, each connection, duration of each connection
  - Each proxy module is a very small software package specifically designed for network security – being relatively simple, it is easier to check such modules for security flaws
    - Typical UNIX mail application has about 20.000 lines of code, mail proxy may have less than 100

Firewall configuration – examples

a) For both incoming and outgoing traffic, the firewall only allows IP packets destined for the bastion host, possible exception for a web-server

- bastion host performs authentication and proxy functions
- If the router falls, then traffic flows directly between Internet and the private network

b) Physically prevents the security breach above

c) Most secure of the three systems – in the middle there is an isolated sub-network (the bastion host and perhaps a web server)

- Outside router only advertises the existence of the screened subnet – internal network invisible from the Internet
- Inside router only advertises the existence of the screened subnet – internal network cannot construct direct connections to the Internet

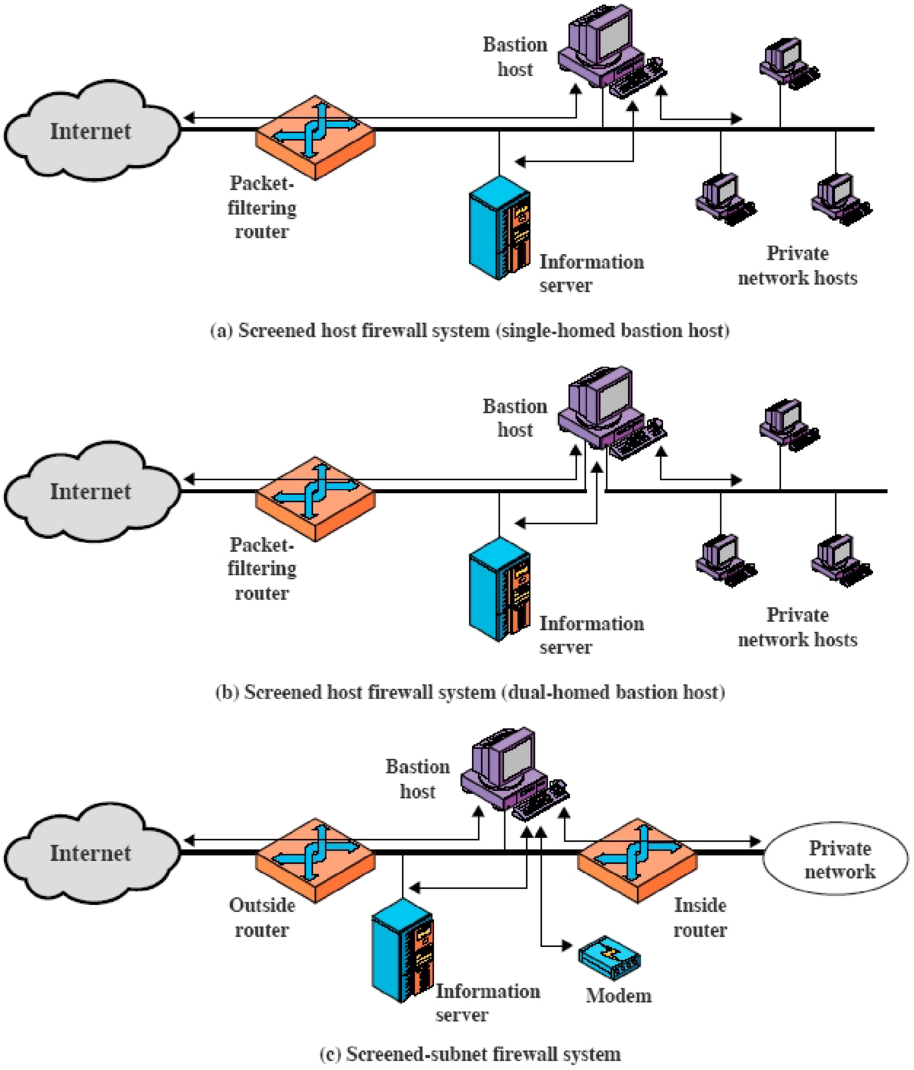

(a) Screened host firewall system (single-homed bastion host)

(b) Screened host firewall system (dual-homed bastion host)

(c) Screened-subnet firewall system

Figure 20.2 Firewall Configurations