UNIK4250 Security in Distributed Systems
University of Oslo
Spring 2012

Part 11

Firewalls and IDS

Audun Jøsang

Outline

- Firewalls
  - Routers
  - Proxies
  - Architectures
- Intrusion Detection Systems
  - Host-based
  - Network based
  - Dealing with false alarms

Perimeter Security
Medieval Castle Defences

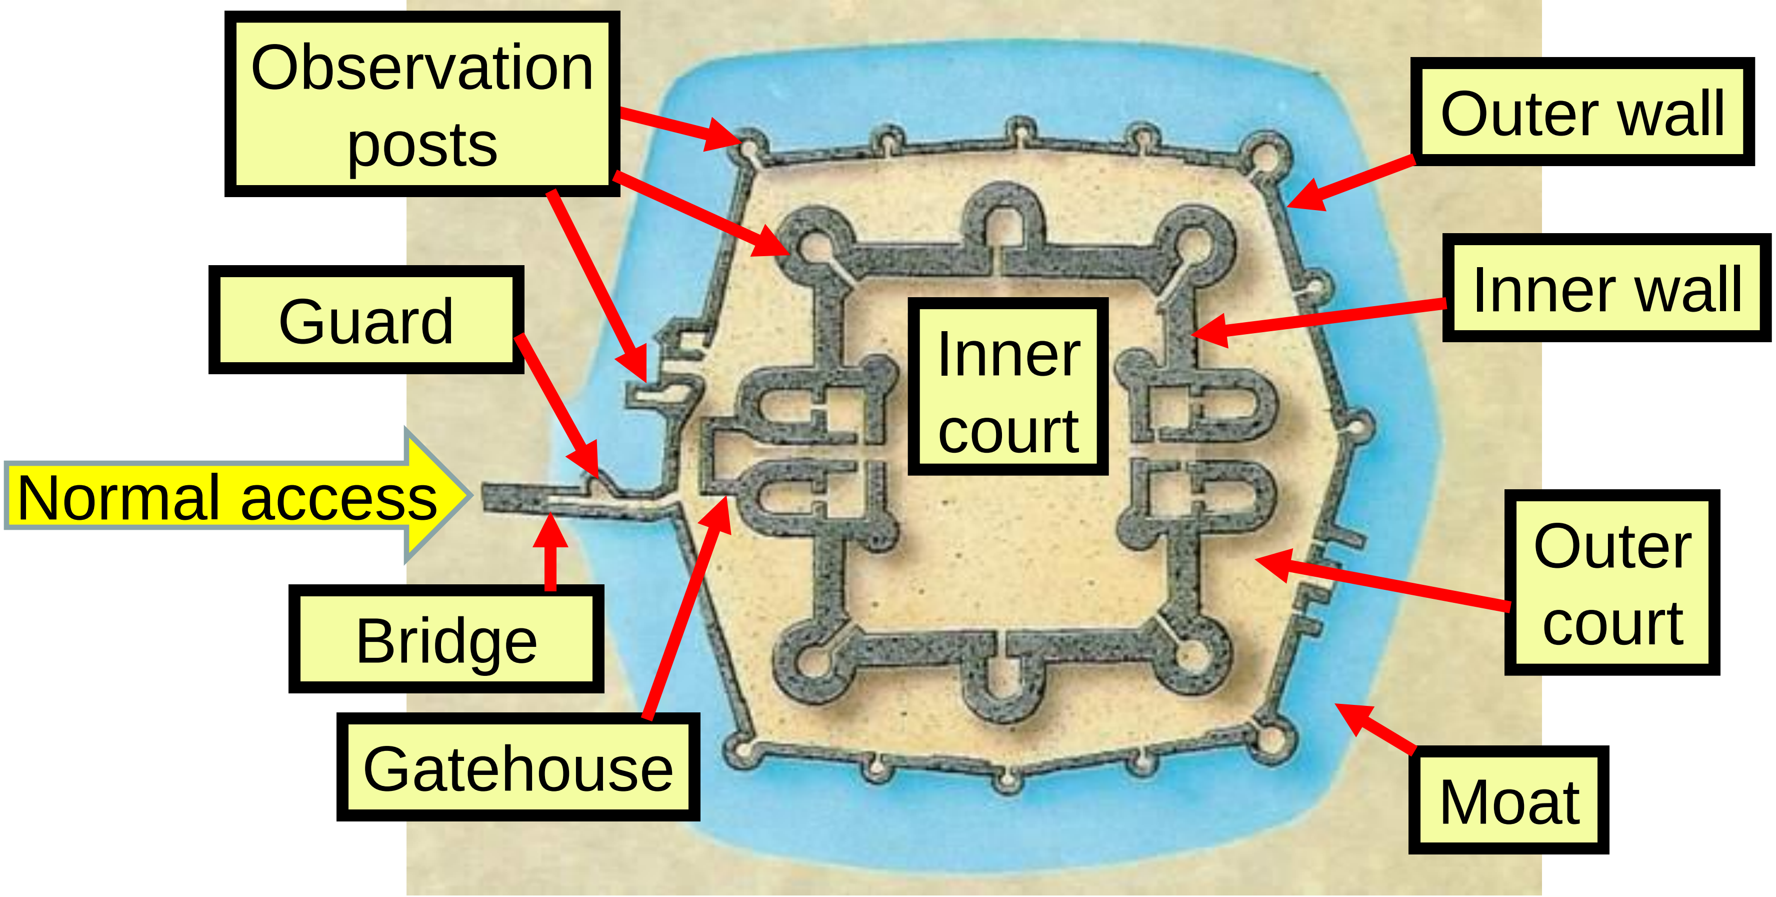

Perimeter Security
Organisation network defences

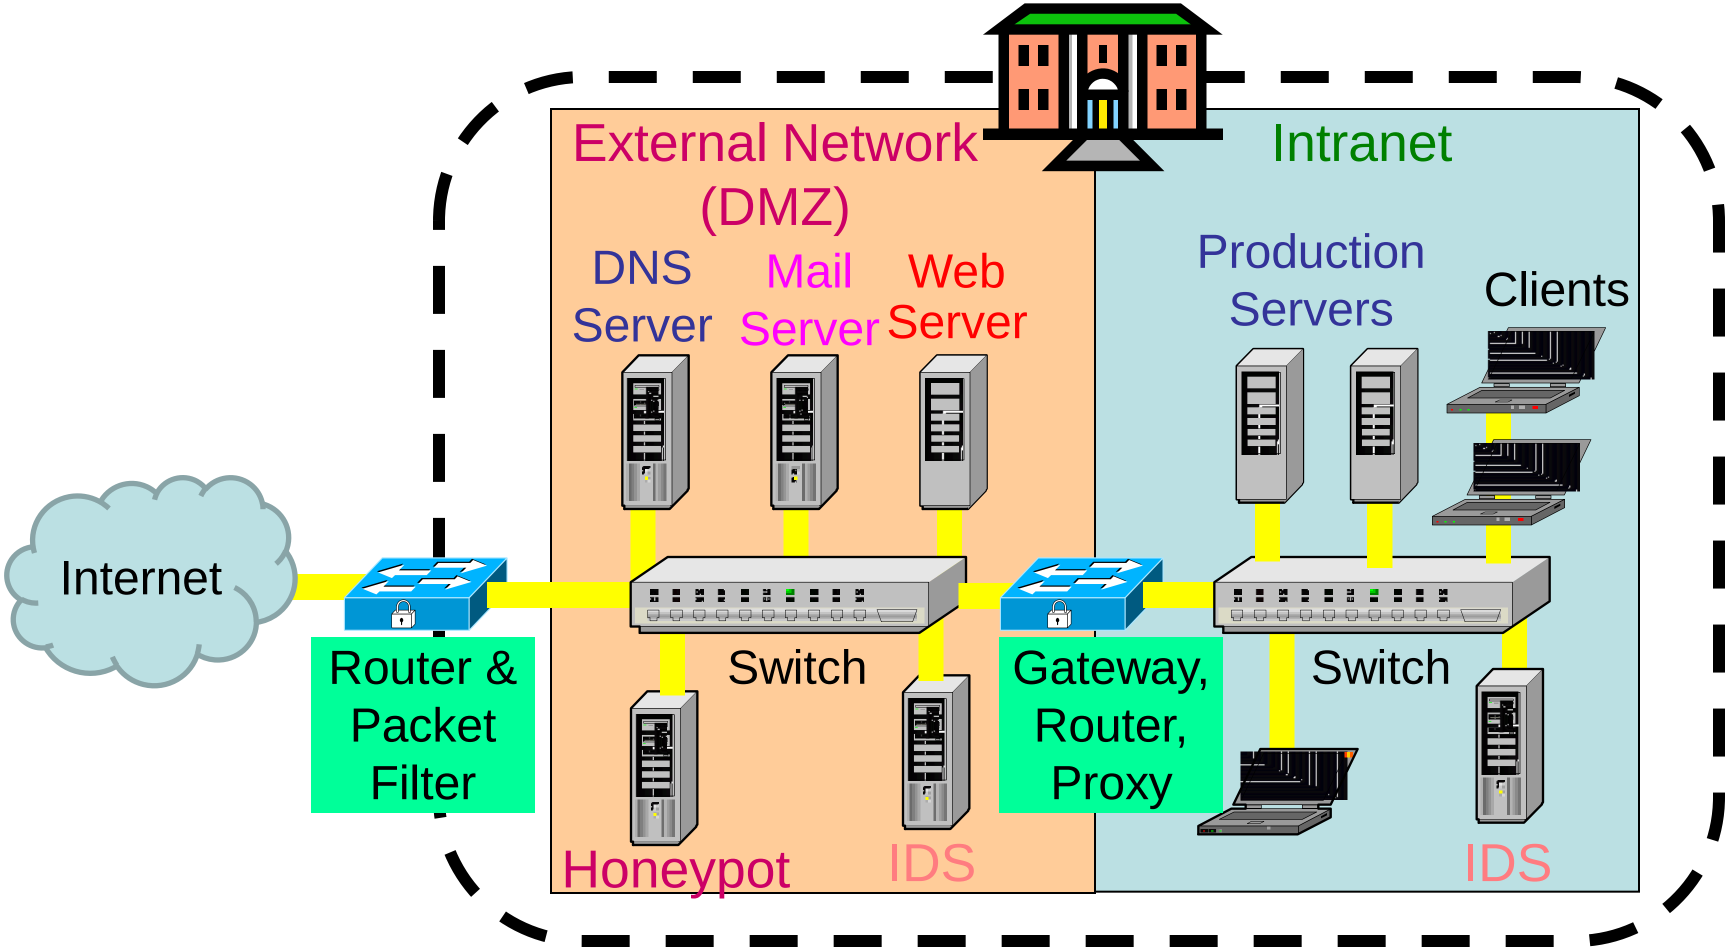

Firewalls

Network Perimeter Security withFirewalls

- Firewall: a misnomer
- A firewall is a check point that protects the internal networks against attack from outside network
- The check point function applies rules to decide which traffic can pass in and out

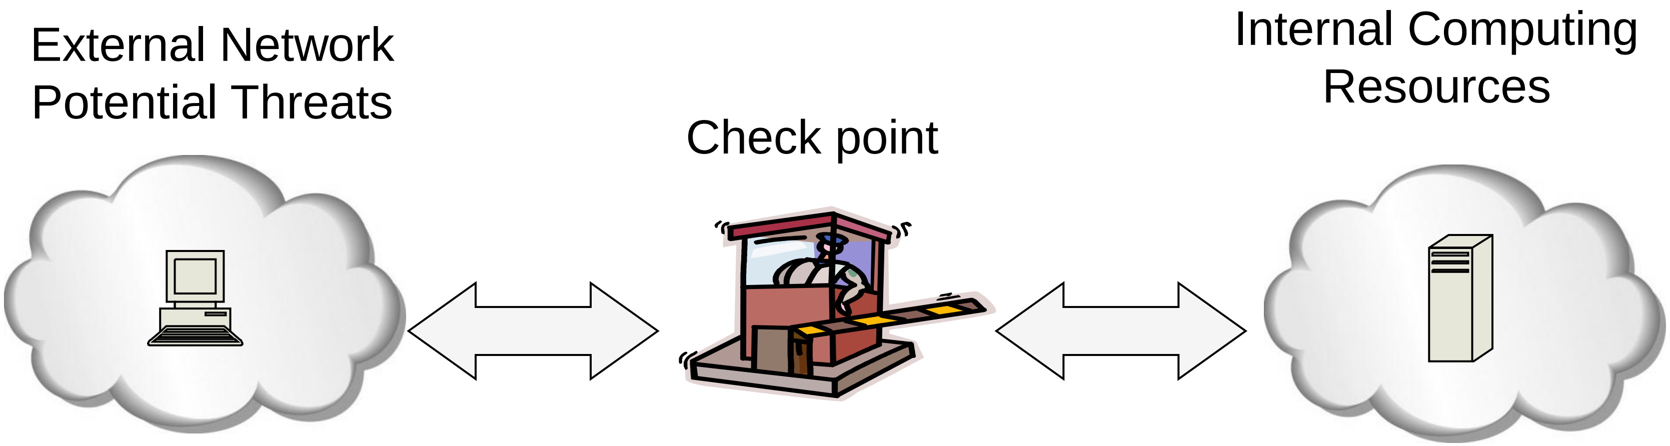

Firewalls: Overview

- If the level of risk associated with maintaining a connection between an organisation’s internal network and the Internet (or some other network(s)) is unacceptable, the most effective way of treating the risk is to avoid the risk altogether and disconnect completely.
- If this is not possible, then firewalls may provide an effective control suitable for reducing the level of risk to an acceptable level.
- Firewalls are often the first line of defence against external attacks, but should not be the only defence

Firewalls: Overview

- A firewall prevents unauthorized access to or from a private network.
- SysAdmin must define criteria for what is (un)authorized
- All traffic entering or leaving must pass through the firewall, which examines each packet and blocks those that do not meet the specified security criteria.
- Firewalls can be implemented in both hardware and software, or a combination of both.
- Firewalls are frequently used to prevent unauthorized Internet users from accessing private networks connected to the Internet, especially intranets.

Firewalls: Overview

- Firewalls must be effectively administered, updated with the latest patches and monitored.
- Firewalls are only effective if they are set up to implement a well formulated security policy
  - i.e. the most important aspect of a firewall is a clear conception of what it is meant to protect.
- Under the security policy, acceptable traffic must be clearly specified so that only specific traffic passes through the firewall; everything else is considered unacceptable and must be stopped by the firewall.

Firewall Limitations

- Can not protect from attacks bypassing it
  - eg sneaker net, utility modems, trusted organisations, trusted services (eg SSL/SSH)
- Can not protect against internal threats
  - eg disgruntled or colluding employees
- Can not protect against access via WLAN
  - if improperly secured against external use
- Can not protect against malware imported via laptop, PDA, storage infected outside

Types of Firewall Technology

- Simple Packet Filters
- Stateful Packet Filters
- Circuit Level Gateways
- Application Gateways

Router Packet Filter

- A network router function that accepts/rejects packets based on headers is referred to as a packet filter.
- Packet filters examine each packet’s headers and make decisions based on attributes such as:
  - Source or Destination IP Addresses
  - Source or Destination Port Numbers
  - Protocol (UDP, TCP or ICMP)
  - ICMP message type
  - And which interface the packet arrived on

Router Packet Filters

- A packet filter examines each packet that attempts to pass through the filter.
  - Done for both directions (entering and leaving a network/host)
- Each packet is examined independently of other packets that may be part of the same connection
  - Unaware of session states at internal or external hosts

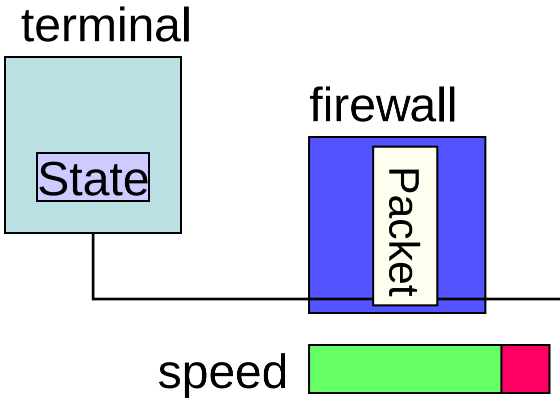

Router Packet Filtering:
Packet header is inspected
Single packet attacks caught
Very little overhead in firewall
High volume filter

Packet Filter example rules

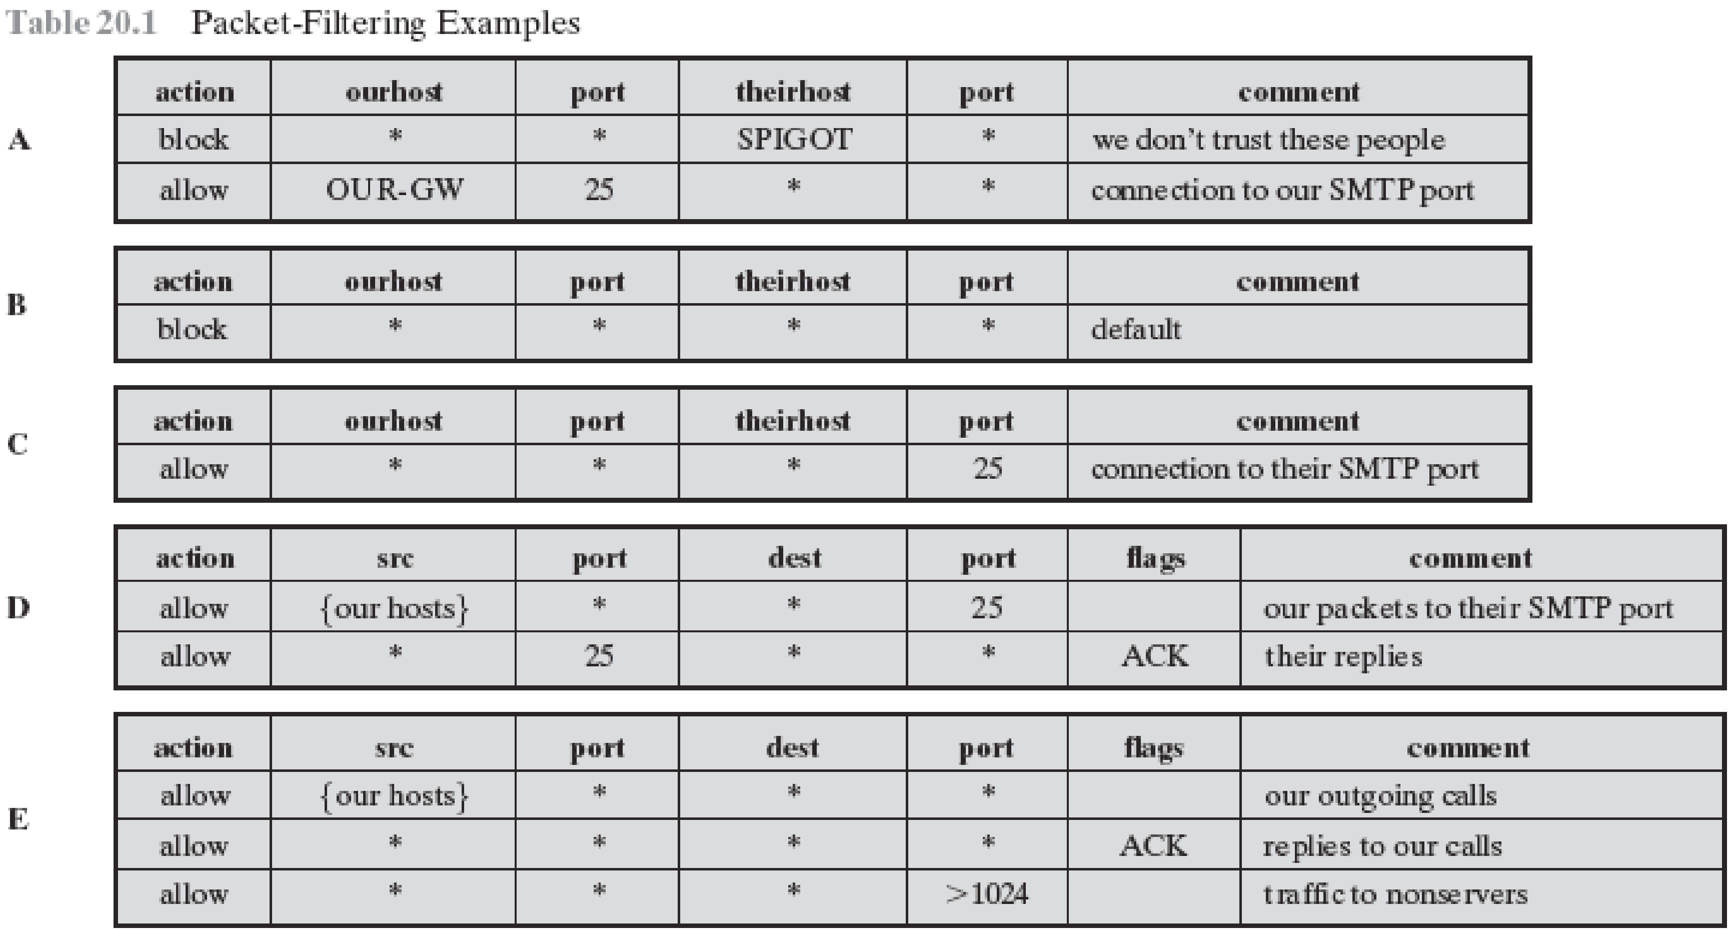

Table 20.1 Packet-Filtering Examples

A

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

B

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

C

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

D

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

E

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

Attacks on Packet Filters

- IP address spoofing
  - fake source address to be trusted
  - add filters on router to block
- Source routing attacks
  - attacker sets a route other than default
  - block source routed packets
- Tiny fragment attacks
  - split header info over several tiny packets
  - either discard or reassemble before check

Stateful Packet Filters

- Stateful packet filters take account of the current state of a connection
- Stateful packet filters are more ‘intelligent’ than simple packet filters.
- Stateful packet filters are able to recognise if a particular packet is part of an established connection by ‘remembering’ recent traffic history.
- This makes the definition of filtering rules easier to accomplish and therefore potentially more secure.

Stateful packet filter

- A statefull packet filter keeps track of sessions
- Requires memory
- Can be subject to DOS (Denial of Service) attacks

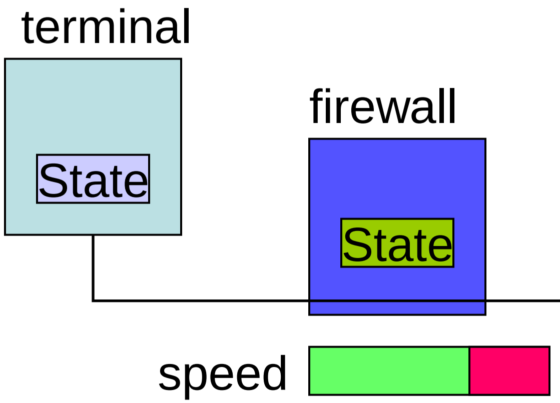

Stateful Inspection
Most multi-packet attacks caught
More fields in packet header inspected
Little overhead in firewall: quick

Stateful Packet Filters

- Sometimes called dynamic packet filters due to their ability to add rules ‘on the fly’. For example:
  - Can recognise an outgoing connection request from an internal client being sent to an external server,
  - And will add a temporary rule to allow the reply traffic back through the firewall.
  - When session is finished, the temporary rule is deleted.
- Common software packages include:
  - IPTables for Linux
  - Checkpoint Firewall-1
  - Cisco PIX (integrated hardware & software)
  - Microsoft Internet Security and Acceleration Server

Packet Filter Strengths and Weaknesses

- Strengths:
  - Low overhead and high throughput
  - Supports almost any application
- Weaknesses:
  - Do not usually interpret application layer data/commands
    - may allow insecure operations to occur
  - Allows direct connection between hosts inside & outside firewall
  - Non-stateful packet filters only: less secure and more difficult to write complex rules

Host-based Packet Filters

- A host can perform packet filtering in addition to other duties, such as web serving
  - in this case the packet filter is designed to protect the host itself, not other hosts
- Advantages:
  - can tailor filtering rules to host environment
  - protection is provided independent of topology
  - provides an additional layer of protection
- Common packet filter software includes:
  - TCP Wrappers for various Unix
  - IP Filter for Sun Solaris

Personal Firewalls

- A personal firewall is a program that is designed to protect the computer on which it is installed.
- Personal firewalls are frequently used by home users to protect themselves from the Internet.
- Personal firewalls are usually a stateful packet filter.
- Some products include anti-virus software as well (usually at extra cost).
  - Vendors such as ZoneAlarm, and Sygate provide a free version of their product for personal use.
  - Windows clients (XP, W7) and Windows servers ship with Internet Connection Firewall (ICF).

Personal Firewalls

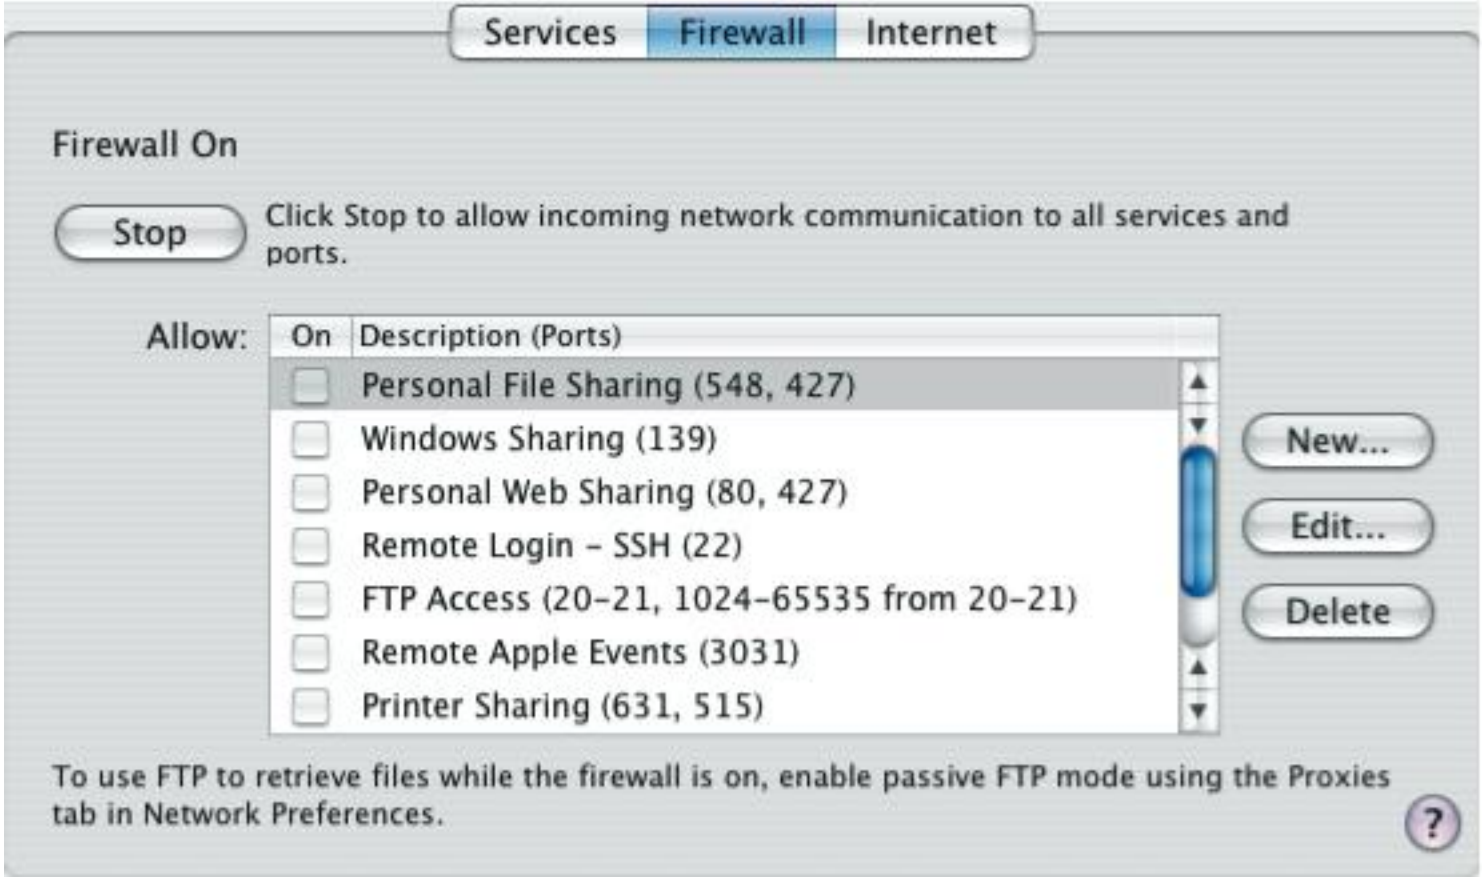

Firewall On

Click Stop to allow incoming network communication to all services and ports.

Allow:

On	Description (Ports)
□	Personal File Sharing (548, 427)
□	Windows Sharing (139)
□	Personal Web Sharing (80, 427)
□	Remote Login – SSH (22)
□	FTP Access (20-21, 1024-65535 from 20-21)
□	Remote Apple Events (3031)
□	Printer Sharing (631, 515)

To use FTP to retrieve files while the firewall is on, enable passive FTP mode using the Proxies tab in Network Preferences.

Circuit Level Gateways

- A circuit level relays two TCP connections.
- It acts as a relay of TCP/UDP layer data rather than application data, and usually no analysis of the application layer data is performed.
- Connections are validated before allowing data to be exchanged.
- Able to identify a particular packet as being part of a particular connection
  - By default stateful inspection
  - Limited additional security checking once packet has been identified

Circuit Level Gateway

- High performance possible due to limited security checking
- Similar strengths and weaknesses to stateful packet filters except
  - Can examine application layer data to a certain extent, but not up to application level gateway standards
  - E.g: Some control/blocking of insecure FTP commands

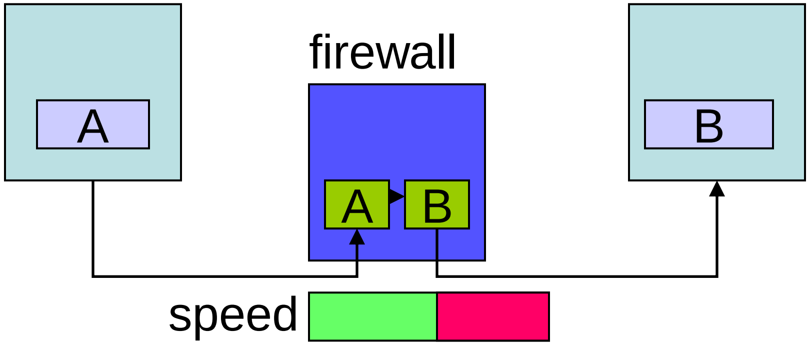

Circuit-Level Firewall:

Packet session terminated and recreated via a Proxy Server

All multi-packet attacks caught

Packet header completely inspected

High overhead in firewall: slow

Circuit Level Gateway Firewall

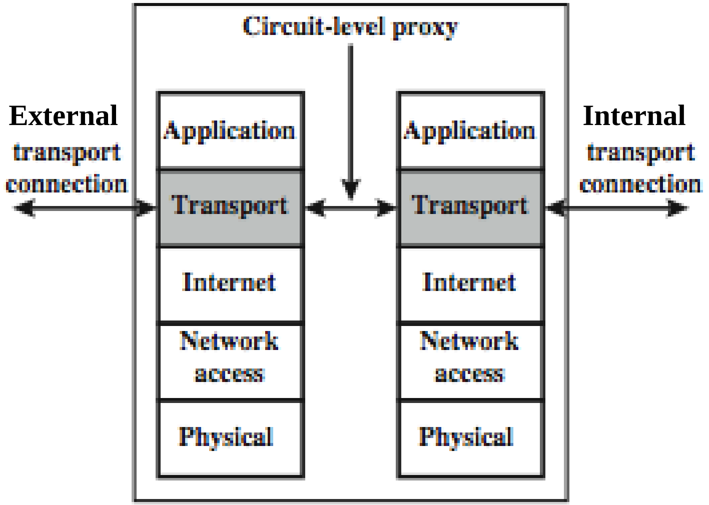

(e) Circuit-level proxy firewall

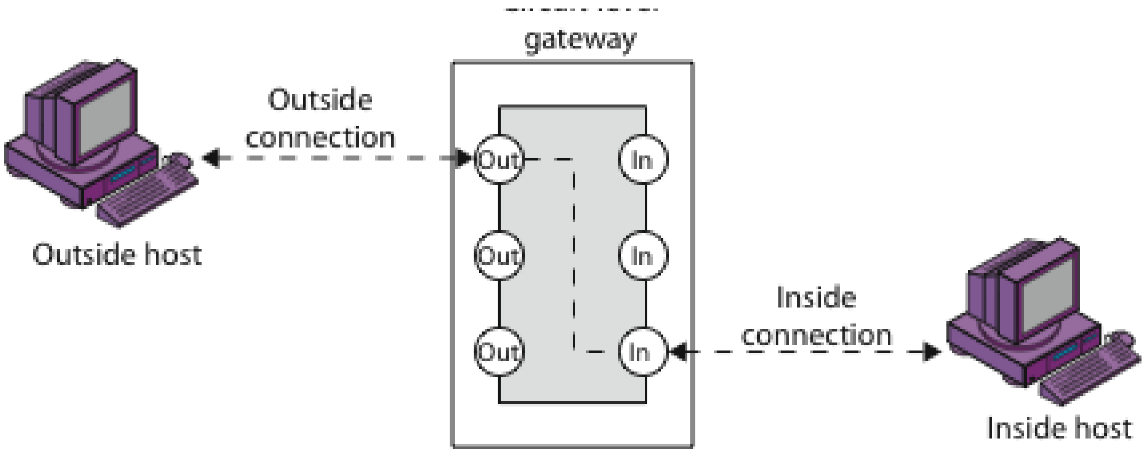

(c) Circuit-level gateway

Bastion Host

- Highly secure host system
- Runs circuit / application level gateways
  - or provides externally accessible services
- Potentially exposed to "hostile" elements
- Hence is secured to withstand this
  - hardened O/S, essential services, extra auth
  - proxies small, secure, independent, non-privileged
- May support 2 or more net connections
- May be trusted to enforce policy of trusted separation between these net connections

Application Level Gateway

- Acts as a relay of application level traffic
- Also known as an application proxy because the firewall needs to act on behalf of the client
- Usually configured to support only specific applications or specific features of an application
  - each application supported by a specific gateway in the firewall

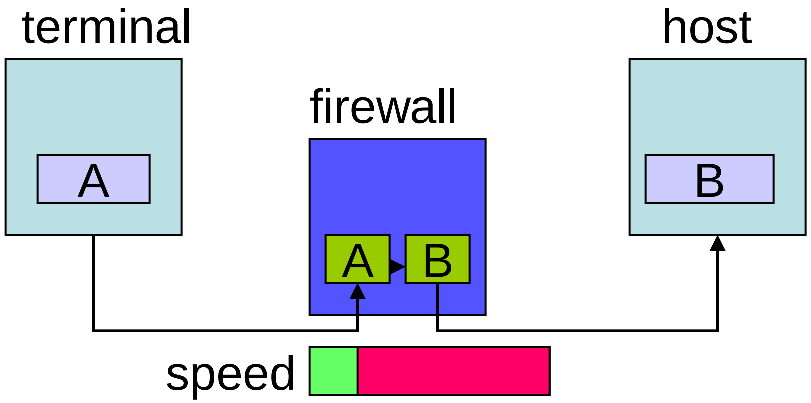

Application-Level Firewall
Packet session terminated and recreated via a Proxy Server
Packet header completely inspected
Application data can be inspected
Highest overhead: slow & low volume

Application Level Gateway (or Proxy)

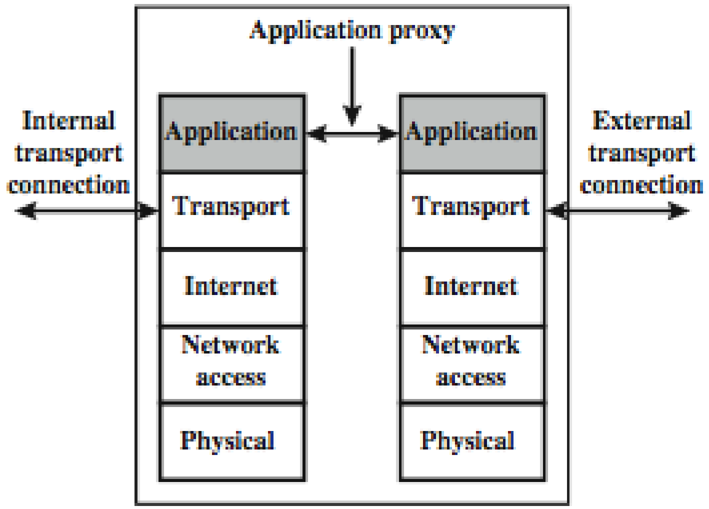

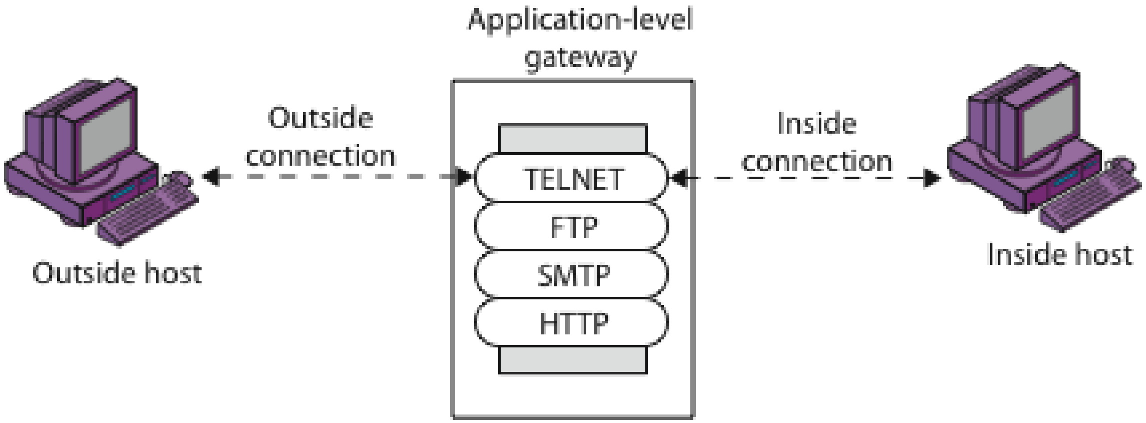

(b) Application-level gateway

Application Level Gateway scenario

- Client sends a request to the server, which is intercepted by the firewall (application gateway)
- Firewall sends the request to the server on behalf of client.
- Server sends reply back to the firewall.
- Firewall sends reply to the client.
- Both client and server think they are communicating with each other, not knowing the firewall exists. It is **transparent**.

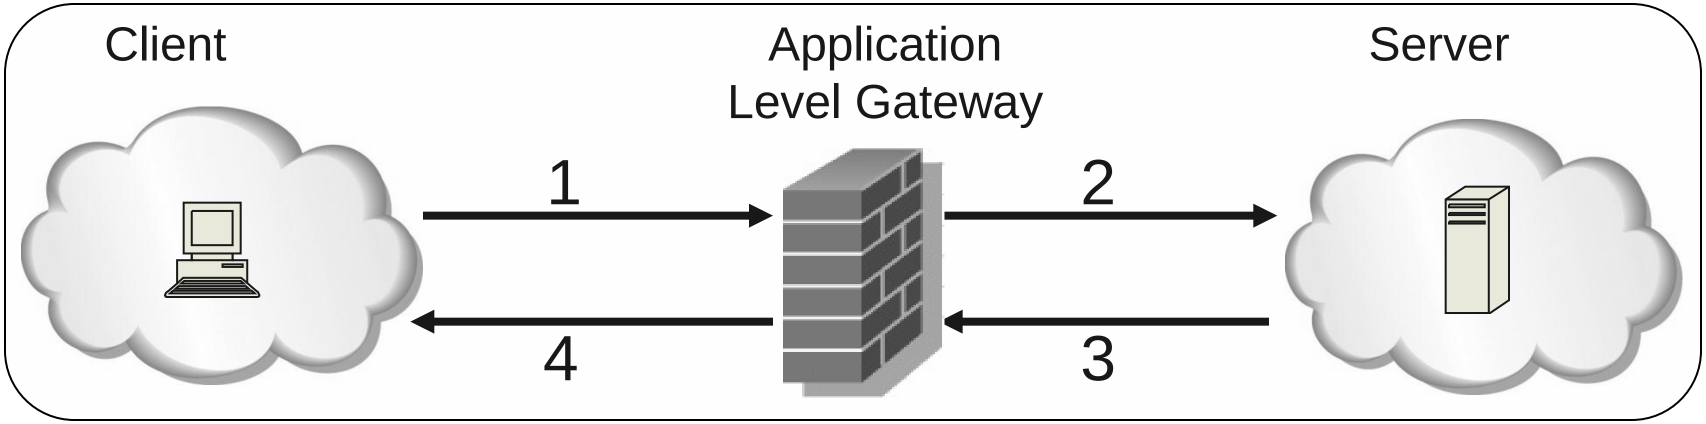

TLS proxy and TLS traffic inspection

- TLS designed for end-to-end encryption, so firewall can not inspect
- In order to inspect TLS, terminate TLS connections at Gateway
- SysAdmin must create Internal PKI Root and issue internal server certificates with the name of external servers (e.g. facebook.com)
- Internal users/hosts will receive server certificate from Gateway, and believe that the certificate comes from the external server
- Causes cleartext gap at Gateway, but it is transparent to users

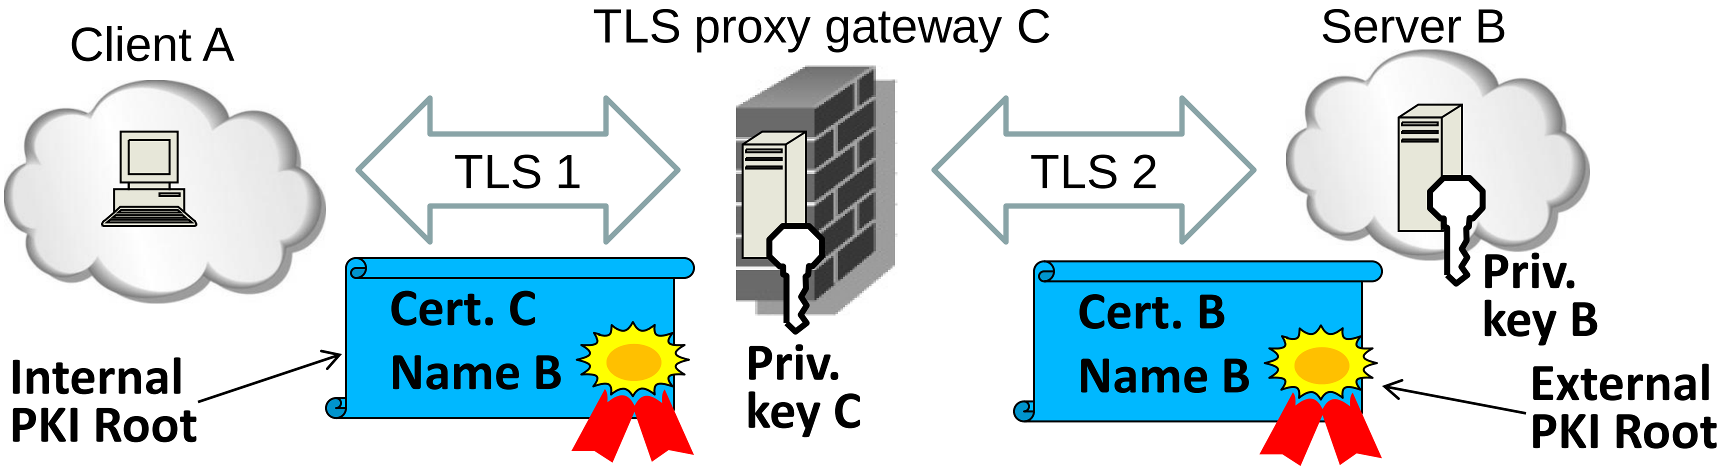

Client A

TLS proxy gateway C

Server B

TLS 1

TLS 2

Cert. C
Name B

Priv.
key C

Cert. B
Name B

Priv.
key B

Internal
PKI Root

External
PKI Root

Deep Inspection Application Gateways

- Deep Packet Inspection looks at application content instead of individual or multiple packets.
- Deep inspection keeps track of application content across multiple packets.
- Potentially unlimited level of detail in traffic filtering

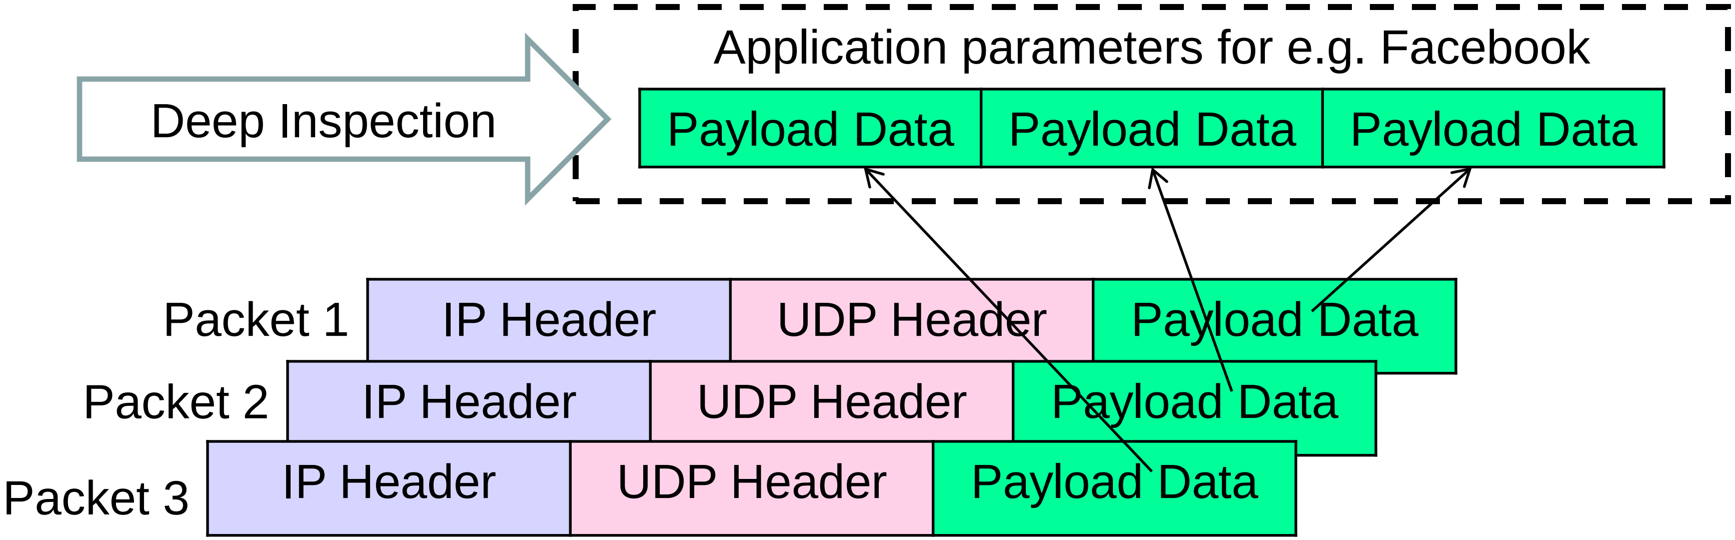

High performance firewalls

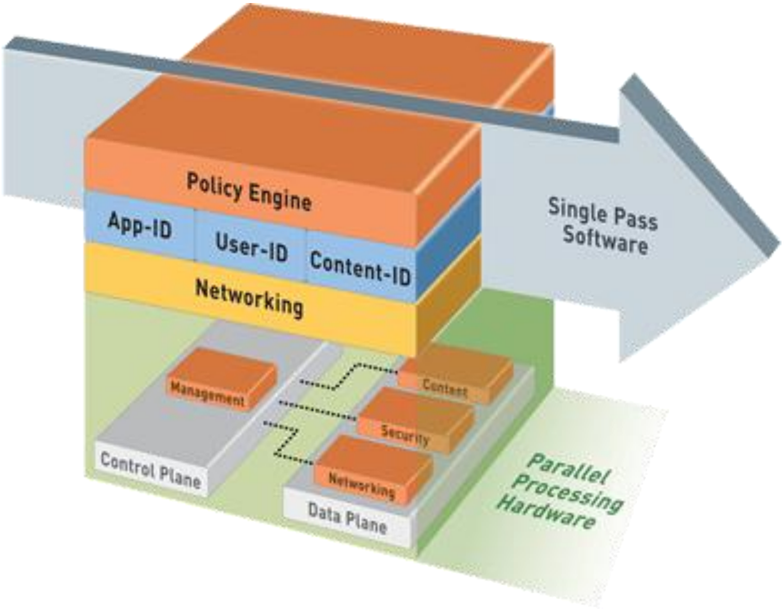

- Modern firewalls can be aware of, and filter according to:
  - Various online applications
  - Users Ids
  - Specific content

Application Gateway

- Strengths:
  - Easy logging and audit of all incoming traffic
  - Provides potential for best security through control of application layer data/commands
- Weaknesses:
  - May require some time for vendor to write new gateways for new applications
  - Requires one more additional connection (including processing resources) for each new connection
  - Slower than packet filters

IPv4 addresses

- IPv4 addresses of 32 bits → 2^{32} = 4,294,967,296 unique addresses
- Represented as 4 decimal bytes separated by dots
  - e.g. www.uio.no = 129.240.8.200
- IANA (Internet Assigned Numbers Authority) assigns address blocks to the 5 different RIR (Regional Internet Registry) in the world
  - Last five remaining 24-bit-blocks (= /8 blocks) assigned in February 2011
  - APNIC RIR (Asia-Pacific Netw. Inf. Centre) ran out of addresses in April 2011
  - The 5 RIRs are:
    - AfriNIC (African Network Information Centre)
    - ARIN (American Registry for Internet Numbers)
    - APNIC (Asia-Pacific Network Information Centre)
    - LACNIC) (Latin America and Caribbean Network Information Centre)
    - RIPE (Réseaux IP Européens Network Coordination Centre)

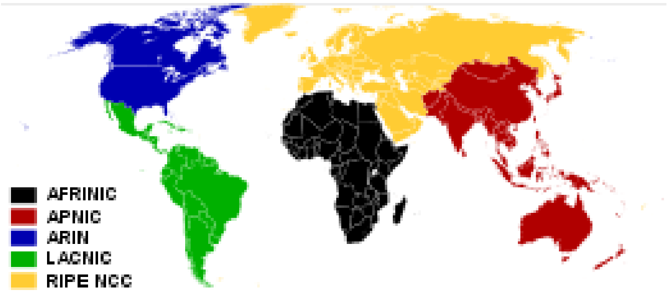

- Private non-routable address ranges:
  - 10.0.0.0 - 10.255.255.255 → 2^{24} = 16,777,216 addresses
  - 172.16.0.0 - 172.31.255.255 → 2^{20} = 1,048,576 addresses
  - 192.168.0.0 - 192.168.255.255 → 2^{16} = 65,536 addresses

Network Address Translation (NAT)

- Translates public ↔ private addresses and ports
- Possibilities:
  - Static mapping
    - permanent mapping of public to private address
  - Dynamic mapping
    - mapping of public to private address when needed
    - unmapped when no longer needed
  - PAT (Port Address Translation)
    - multiple internal addresses mapped to same public address but with different port numbers

Network Address Translation (NAT)

- Advantages
  - Helps enforce control over outbound connections
  - Helps restrict incoming traffic
  - Helps conceal internal network configuration
  - Prevents port scanning
- Can’t be used with:
  - protocols that require a separate back-channel
  - protocols that encrypt TCP headers
  - embed TCP address info
  - IPv6

Firewall Configurations
Screened host

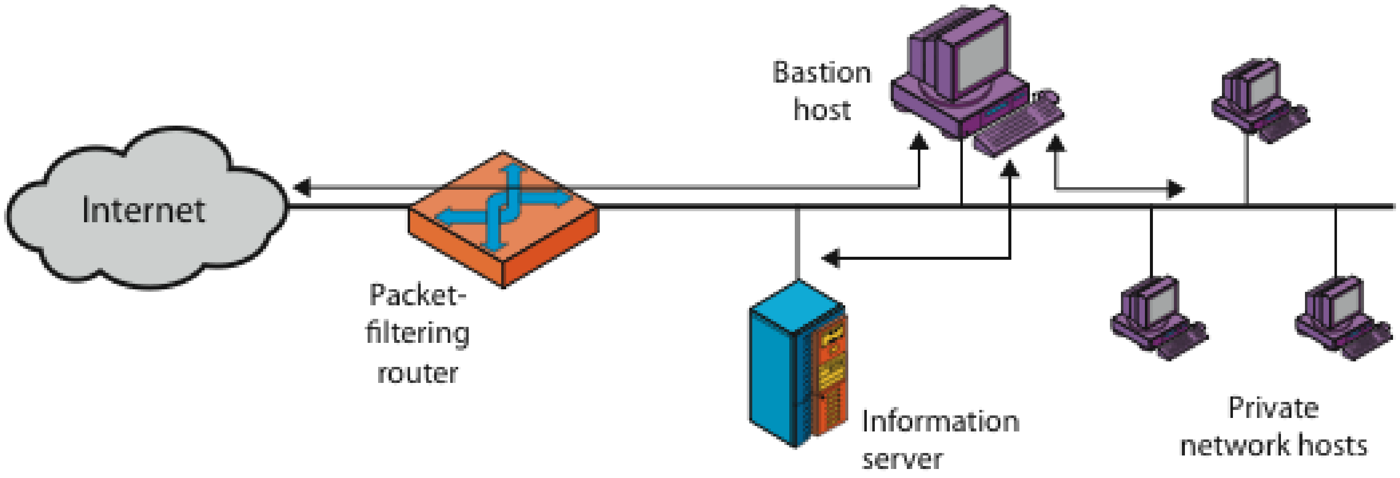

(a) Screened host firewall system (single-homed bastion host)

Firewall Configurations
Screened dual homed bastion host

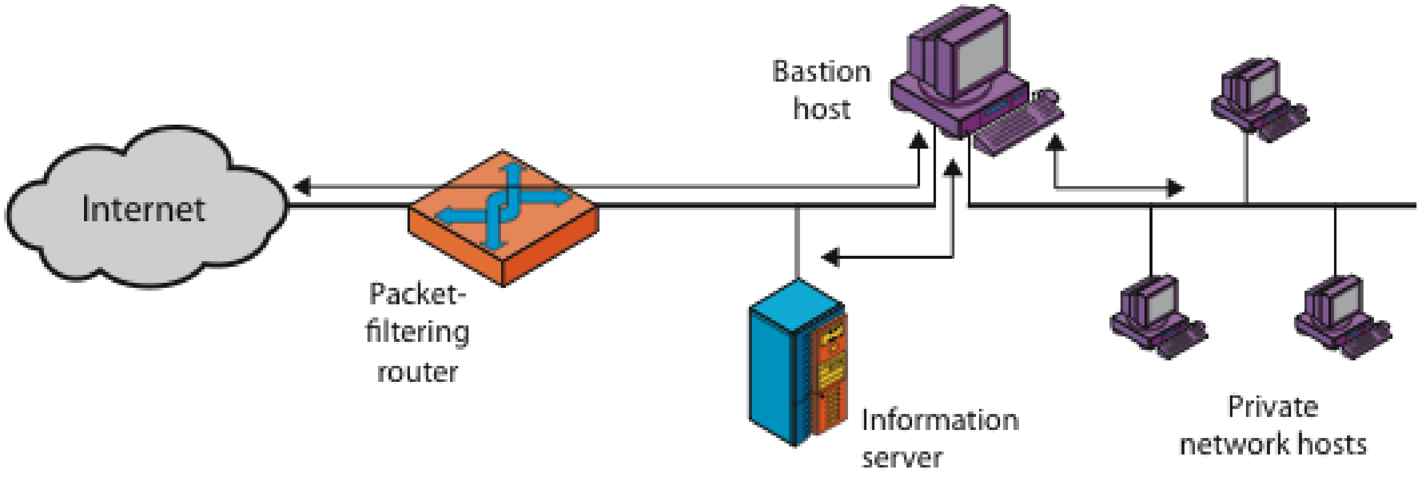

(b) Screened host firewall system (dual-homed bastion host)

Firewall Configurations Screened subnet (DMZ)

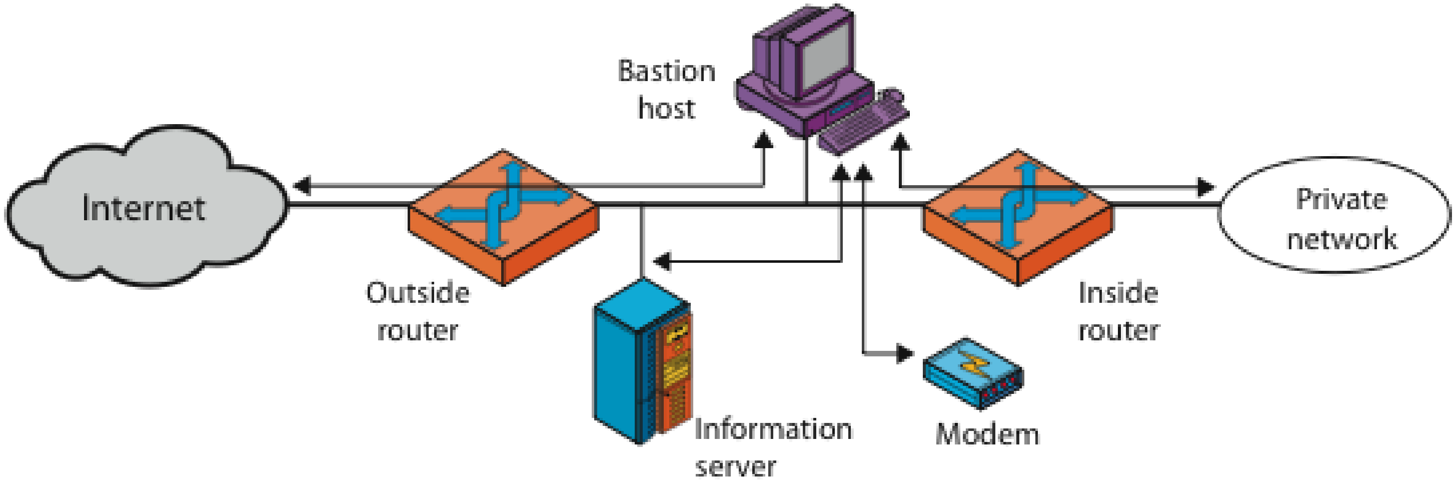

(c) Screened-subnet firewall system

DMZ Networks

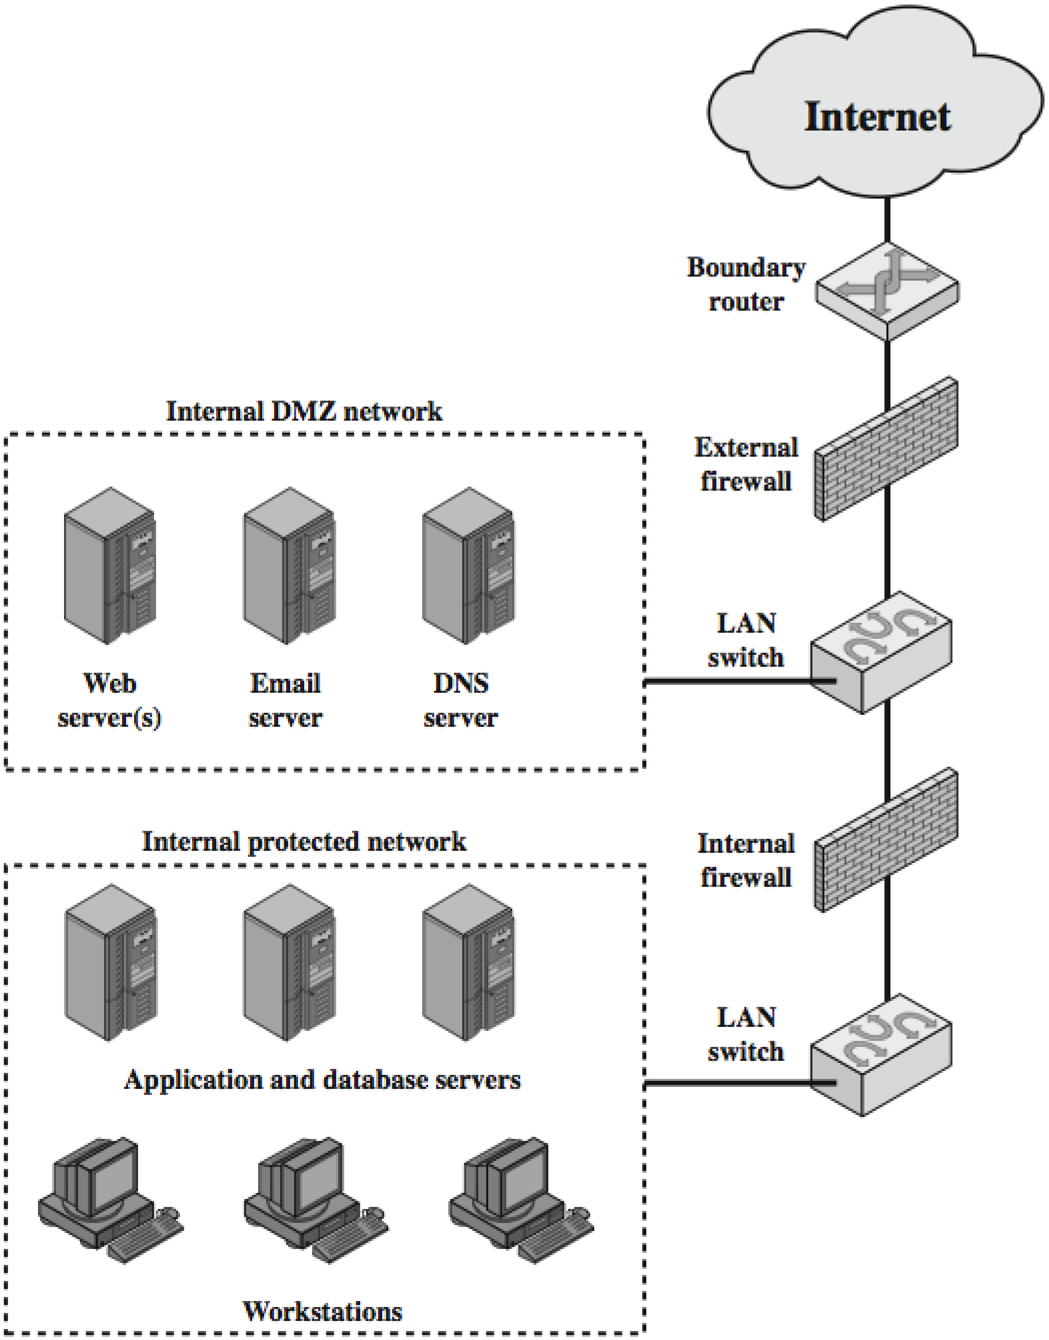

Virtual Private Networks

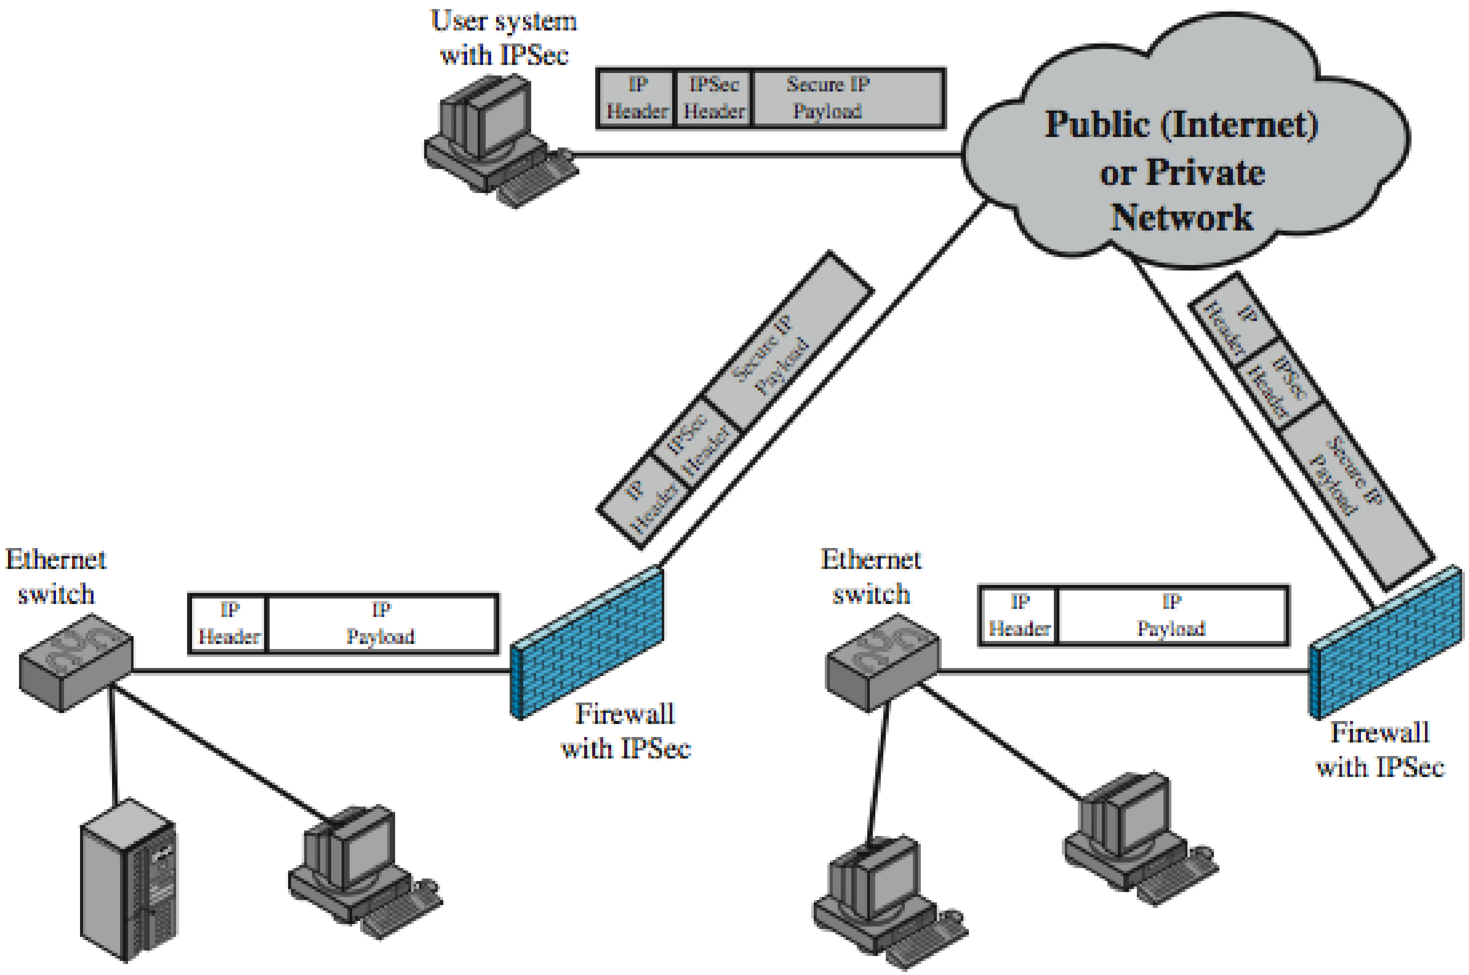

Distributed Firewalls

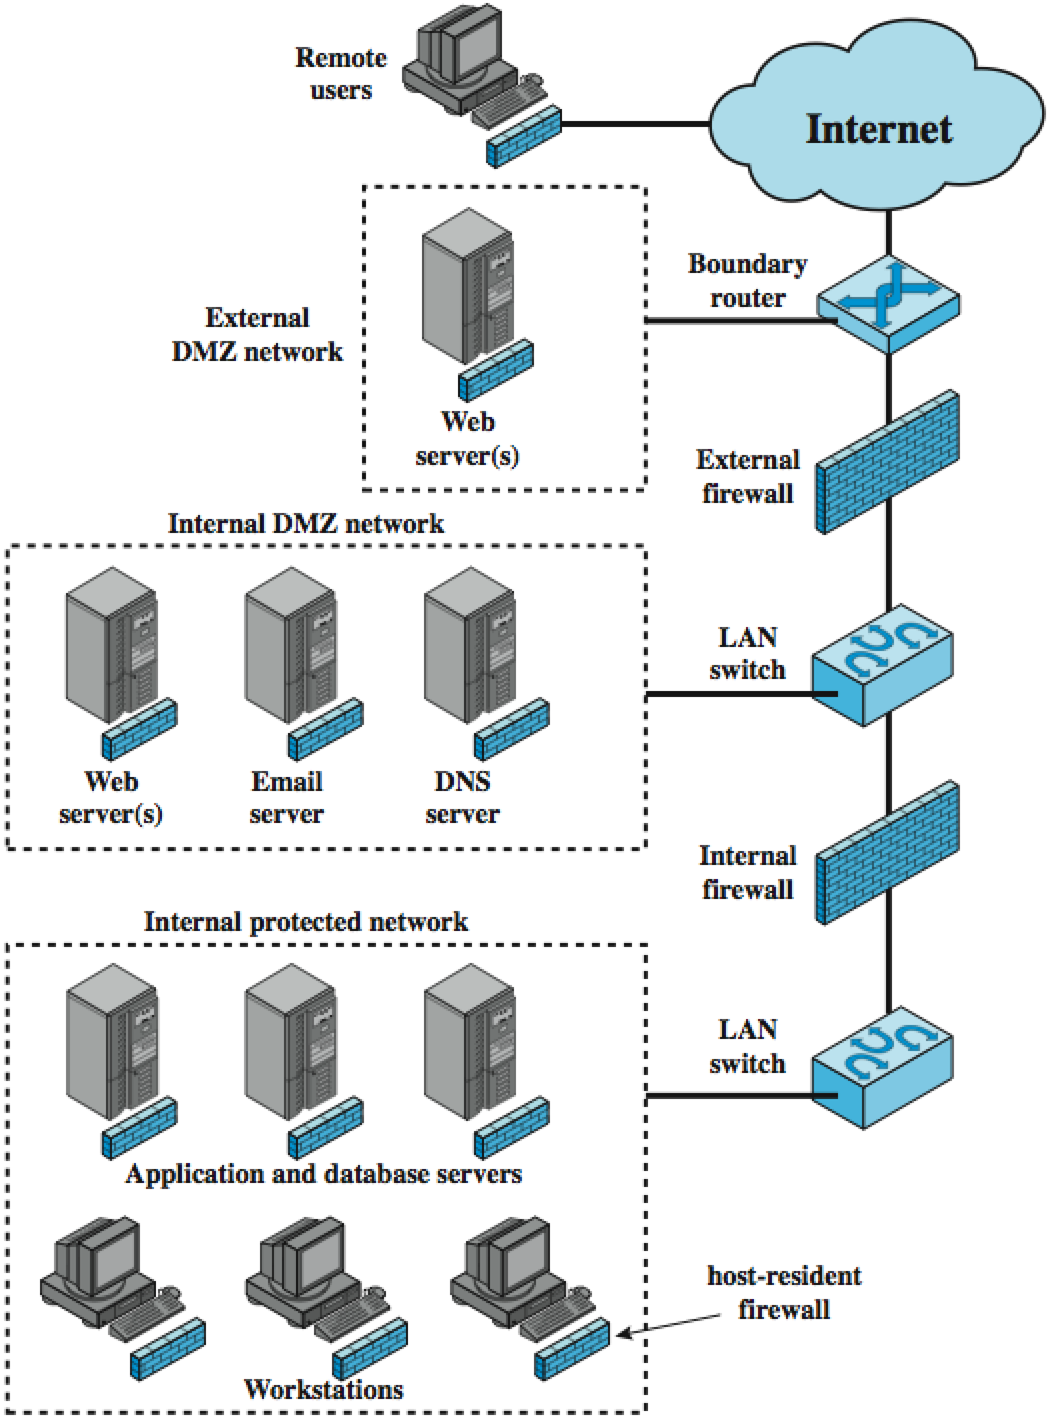

Intrusion Detection Systems

Intrusion Detection Systems: Overview

- Intrusion detection systems (IDS) are automated systems (programs) that detect suspicious events
- IDS can be either host-based or network-based.
- A host based IDS is designed to detect intrusions only on the host it is installed on
  - monitor changes to host’s operating system files and traffic sent to the host
- Network based IDS (NIDS) are designed to detect intrusions on one or more network segments, usually deployed to protect a number of hosts
  - monitor network/s looking for suspicious traffic

What Should Be Detected?

- Attempted and successful break-ins
- Attacks by legitimate users
  - For example, illegitimate use of root privileges
  - Unauthorized access to resources and data
- Trojan horse malware
- Viruses and worms
- Denial of service attacks

IDS:
Network IDS Deployment

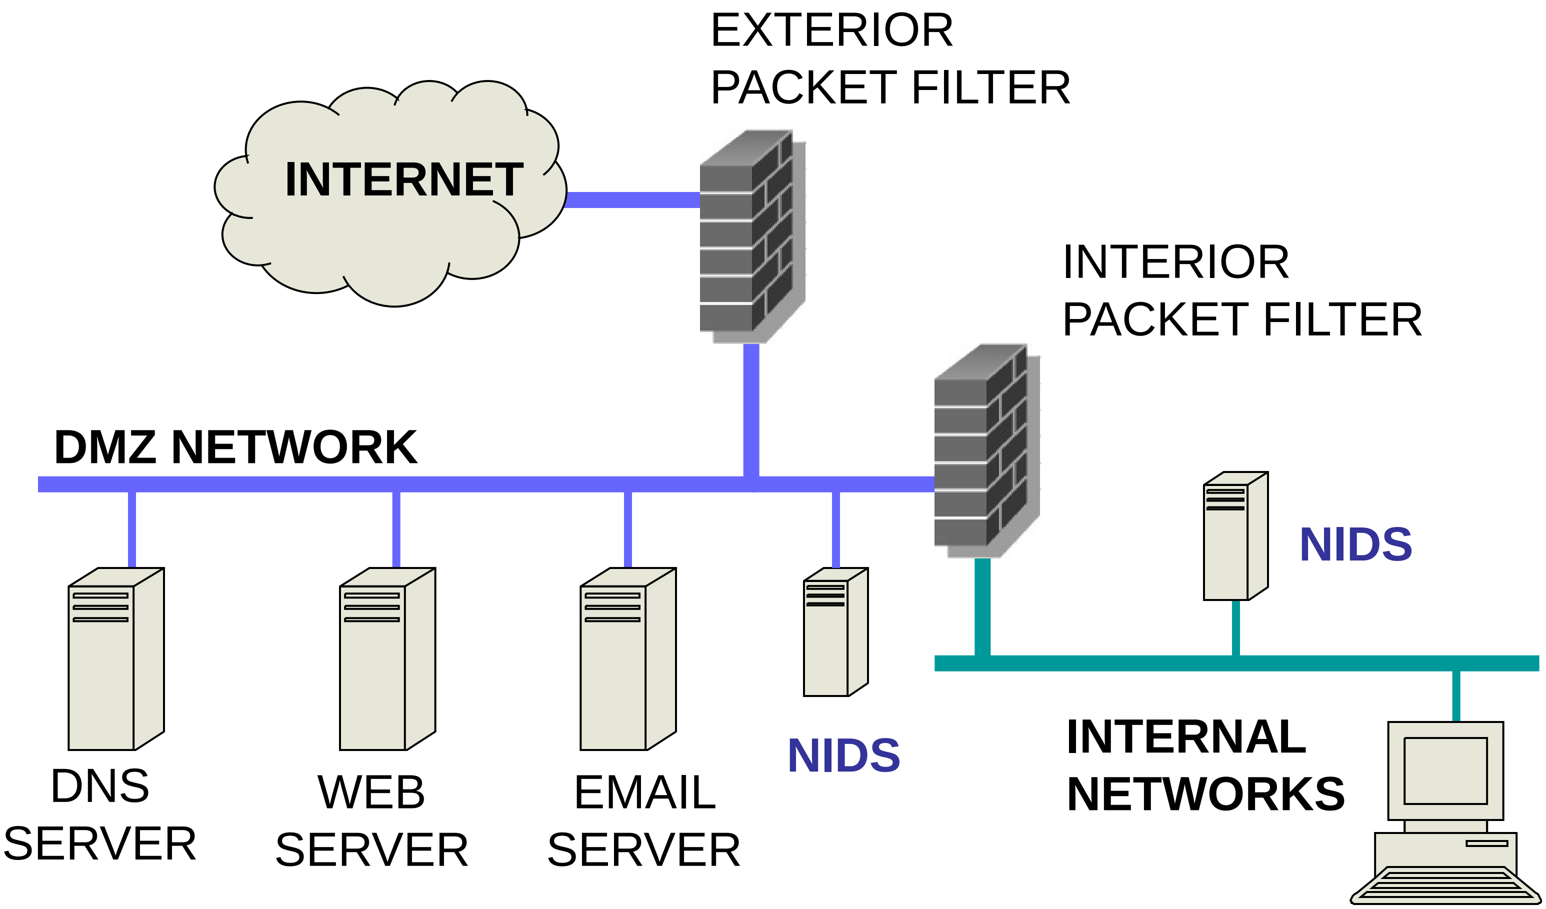

Intrusion Detection Techniques

- Misuse detection
  - Use attack “signatures” (need a model of the attack)
    - Sequences of system calls, patterns of network traffic, etc.
  - Must know in advance what attacker will do (how?)
  - Can only detect known attacks
- Anomaly detection
  - Using a model of normal system behavior, try to detect deviations and abnormalities
    - E.g., raise an alarm when a statistically rare event(s) occurs
  - Can potentially detect unknown attacks
- Which is harder to do?

Level of Monitoring

- Which types of events to monitor?
  - OS system calls
  - Command line
  - Network data (e.g., from routers and firewalls)
  - Processes
  - Keystrokes
  - File and device accesses

Popular NIDS

- Snort (popular open-source tool)
  - Large rule sets for known vulnerabilities, e.g.
    - **2009-03-31:** A programming error in MySQL Server may allow a remote attacker to cause a Denial of Service (DoS) against a vulnerable machine.
    - **2009-03-27:** Microsoft Windows GDI Buffer Overflow: A programming error in the Microsoft Windows kernel may allow a remote attacker to execute code with system level privileges. This may be exploited when specially crafted EMF files are viewed using Microsoft Internet Explorer.
- Bro (developed by Vern Paxson)
  - Separates data collection and security decisions
    - Event Engine distills the packet stream into high-level events describing what’s happening on the network
    - Policy Script Interpeter uses a script defining the network’s security policy to decide what to do in response

Port Scanning

- Many vulnerabilities are OS-specific
  - Bugs in specific implementations, default configuration
- Port scan is often a prelude to an attack
  - Attacker tries many ports on many IP addresses
    - For example, looking for an old version of some daemon with an unpatched buffer overflow
  - If characteristic behavior detected, mount attack
  - “The Art of Intrusion”: virtually every attack involves port scanning and password cracking

Attacking and Evading NIDS

- Overload NIDS with huge data streams, then attempt the intrusion
  - Bro solution: watchdog timer
    - Check that all packets are processed by Bro within T seconds; if not, terminate Bro, use tcpdump to log all subsequent traffic
- Use encryption to hide packet contents
- Split malicious data into multiple packets
  - NIDS does not have full TCP state and does not always understand every command of receiving application
  - Simple example: send “ROB<DEL><BS><BS>OT”, receiving application may reassemble to “ROOT”

Intrusion Detection Problems

- Lack of training data with real attacks
  - But lots of “normal” network traffic, system call data
- Data drift
  - Statistical methods detect changes in behavior
  - Attacker can attack gradually and incrementally
- Discriminating characteristics hard to specify
  - Many attacks may be within bounds of “normal” range of activities
- False identifications are very costly
  - Sysadm will spend many hours examining evidence

Intrusion Detection Errors

- False negatives: attack is not detected
  - Big problem in signature-based misuse detection
- False positives: harmless behavior is classified as an attack
  - Big problem in statistical anomaly detection
- Both types of IDS suffer from both error types
- Which is a bigger problem?
  - Attacks are fairly rare events
  - IDS often suffer from “base-rate fallacy”

Base Rate Fallacy

- Consider statements: r: “attack occurs”, s: “signature detected”
- Base rate fallacy is to assume $p(r \mid s) = 1$ without considering $a(r)$
- Consider the probabilities:

  $p(s \mid r)$: probability of detecting signature given that attack occurs

  $p(s \mid \neg r)$: probability of detecting signature when no attack occurs

  $a(r)$: base rate of attacks (i.e. average rate of attack per connection)
- Learning produces $p(s|r)$ and $p(s|\neg r)$, but detection requires $p(r|s)$
- Deriving $p(r|s)$ requires $a(r)$:

$$p(r \mid s) = \frac{a(r)\,p(s \mid r)}{a(r)\,p(s \mid r) + (1 - a(r))\,p(s \mid \neg r)}$$

  - $p(r \mid s) = 1$, is a good approximation when $a(r) >> 0$ and $p(s|\neg r) \approx 0$
  - $p(r \mid s) = 1$, is a bad approximation when $a(r) \approx 0$

Strategic Intrusion Assessment

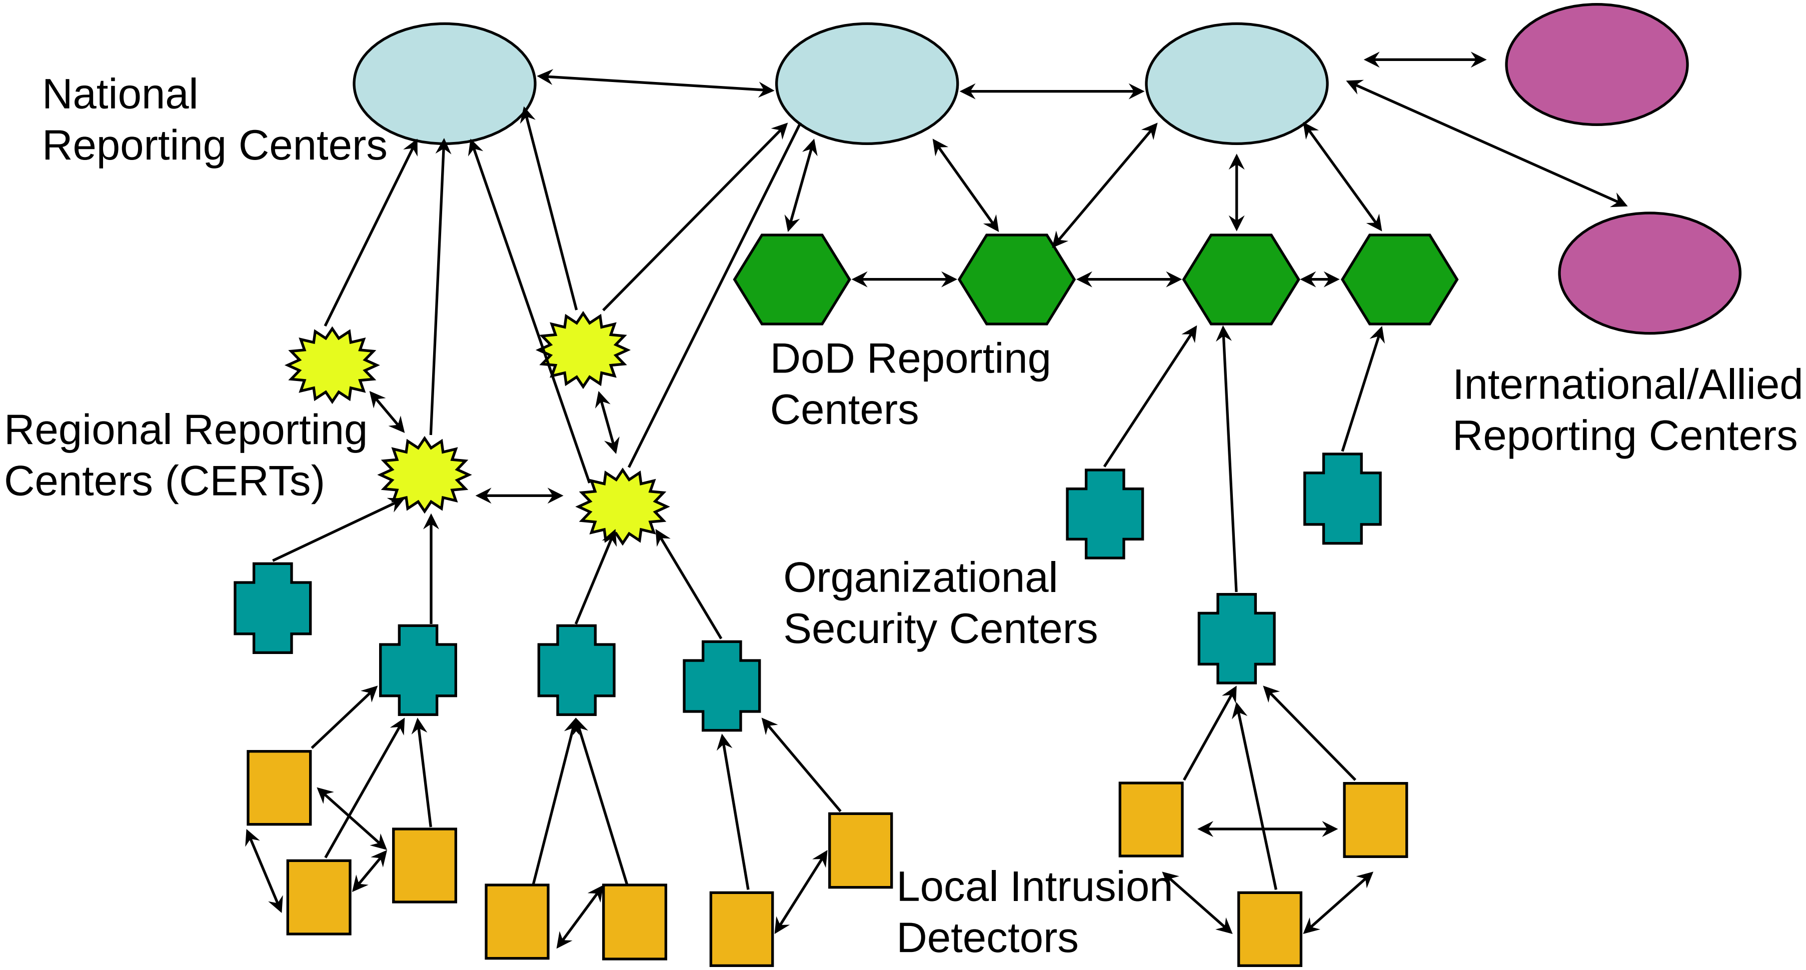

National Reporting Centers

DoD Reporting Centers

International/Allied Reporting Centers

Regional Reporting Centers (CERTs)

Organizational Security Centers

Local Intrusion Detectors

Strategic Intrusion Assessment

- Test over two-week period by Air Force Information Warfare Center
  - Intrusion detectors at 100 Air Force bases alarmed on 2,000,000 sessions
  - Manual review identified 12,000 suspicious events
  - Further manual review => four actual incidents
- Conclusion
  - Most alarms are false positives
  - Most true positives are trivial incidents which can be ignored, i.e. the attacks will never be able to penetrate any system
  - Of the significant incidents, most are isolated attacks to be dealt with locally

Honeypots

- A honeypot:
  - is a computer configured to detect network attacks or malicious behaviour,
  - appears to be part of a network, and seems to contain information or a resource of value to attackers.
- But honeypots are isolated, are never advertised and are continuously monitored
- All connections to honeypots are per definition malicious
- Can be used to extract attack signatures

Intrusion Prevention Systems

- Intrusion Prevention System (IPS) is a relatively new term that can mean different things
- Most commonly, an IPS is a combination of an IDS and a firewall
- A system that detects an attack and can stop it as well
- Can be application specific
  - Deployed on a host to stop attacks on specific applications such as IIS
- Can be an extension of an NIDS
- Can be used to defend systems that cannot be patched

End of Lecture

We have discussed:

- Firewall techniques
- Intrusion detection techniques