

INF3510 Information Security

Lecture 11: Network Perimeter Security



Audun Jøsang


University of Oslo
Spring 2016

Outline

- Firewalls
 - Routers
 - Proxies
 - Architectures
- Intrusion Detection Systems
 - Host-based
 - Network based
 - Dealing with false alarms
- Wireless LAN Access Control
 - Evolution & history
 - WPA2: Robust Security Network architecture (RNS)

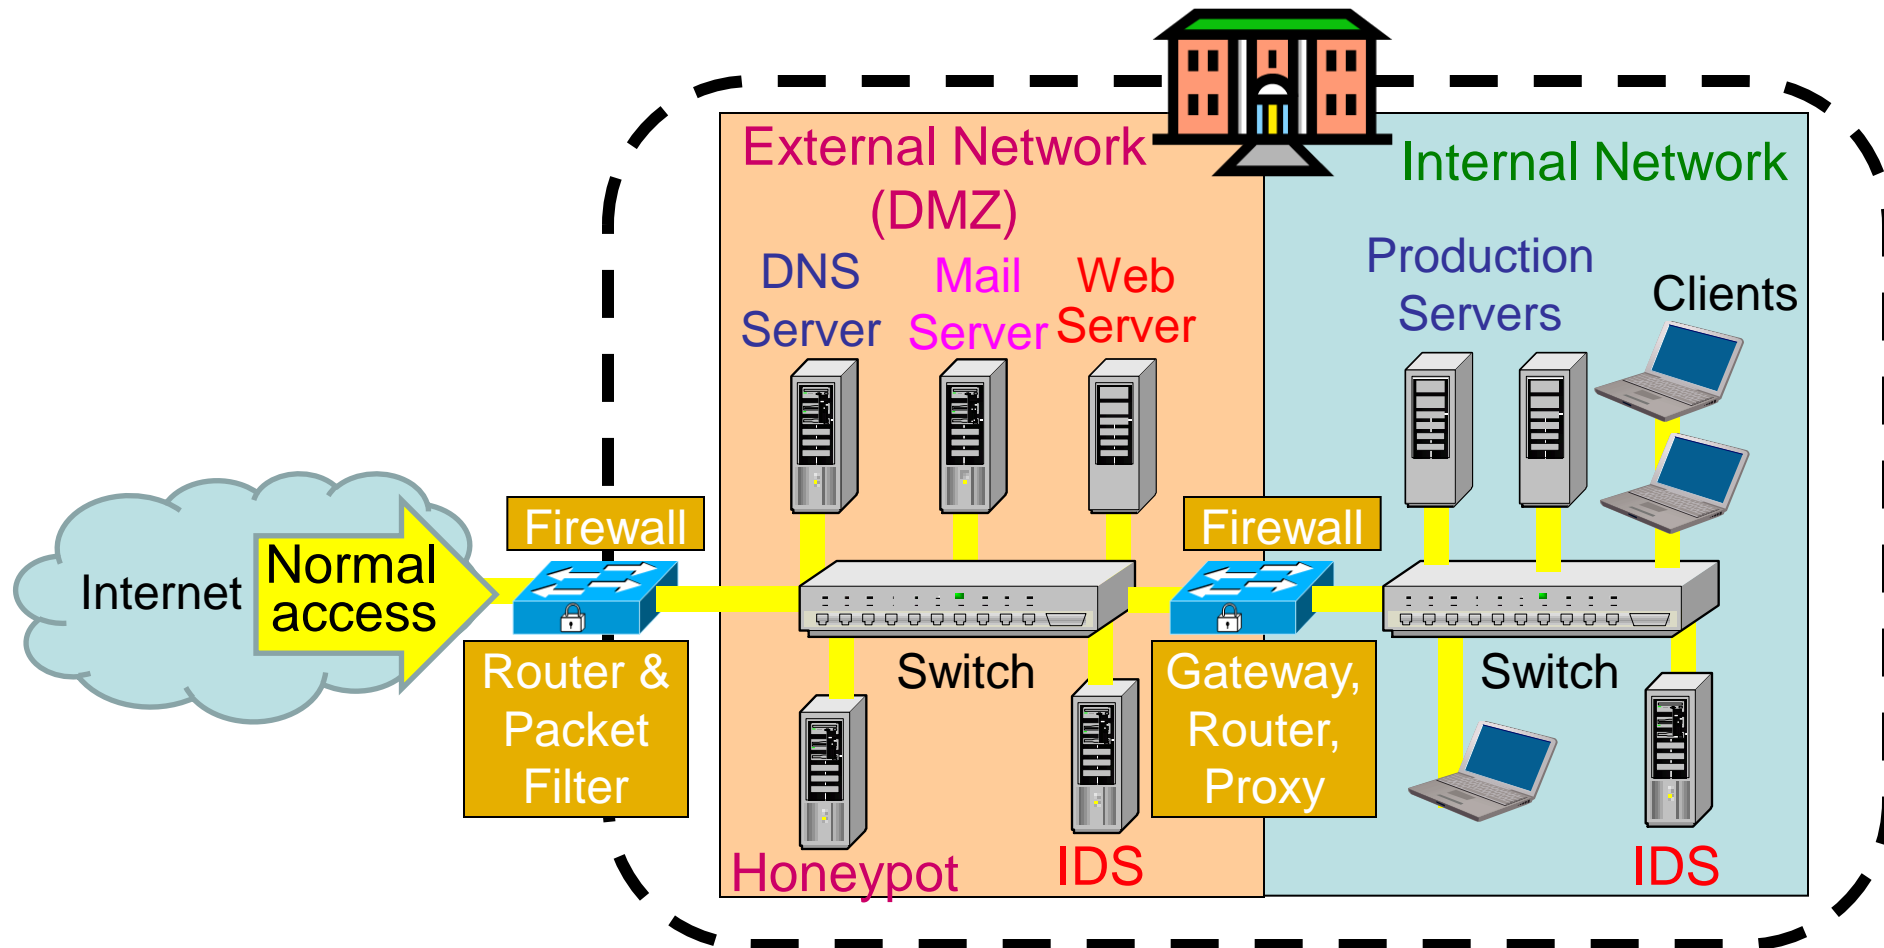
Perimeter security analogy

Medieval Castle Defences



Defending local networks


Network Perimeter Security



Firewalls

Network perimeter security method: Firewalls

- A firewall is a check point that protects the internal networks against attack from outside networks
- The check point decide which traffic can pass in & out based on rules



Firewalls: Overview 1


- If the risk of having a connection to the Internet is unacceptable, the most effective way of treating the risk is to avoid the risk altogether and disconnect completely.
- If disconnection from the Internet is not practical, then firewalls may provide an effective level of protection that can reduce the risk to an acceptable level.
- Firewalls are often the first line of defence against external attacks, but should not be the only defence.
- A firewall's purpose is to prevent unauthorized access to or from a private network.

Firewalls: Overview 2


- All traffic entering or leaving must pass through firewall
- The network owner must define criteria for what is (un)authorized
- The effectiveness of firewalls depends on specifying authorized traffic in terms of rules
 - The rules defines what to let pass through;
 - The rules defines what to block.
- Firewalls must be effectively administered, updated with the latest patches and monitored.
- Firewalls can be implemented in both hardware and software, or a combination of both.

Types of Firewall Technology (vehicle analogy)


- Router Packet Filters




- Stateful Packet Filters



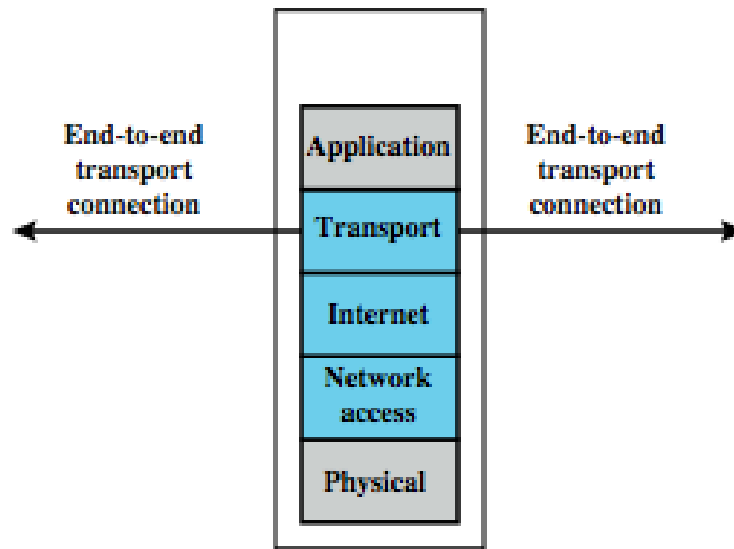
- Application Layer Proxy



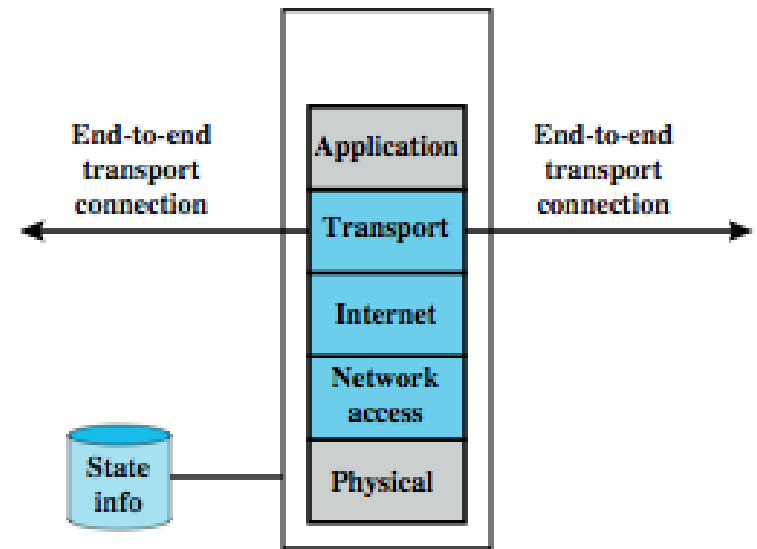
- Next Generation Firewall



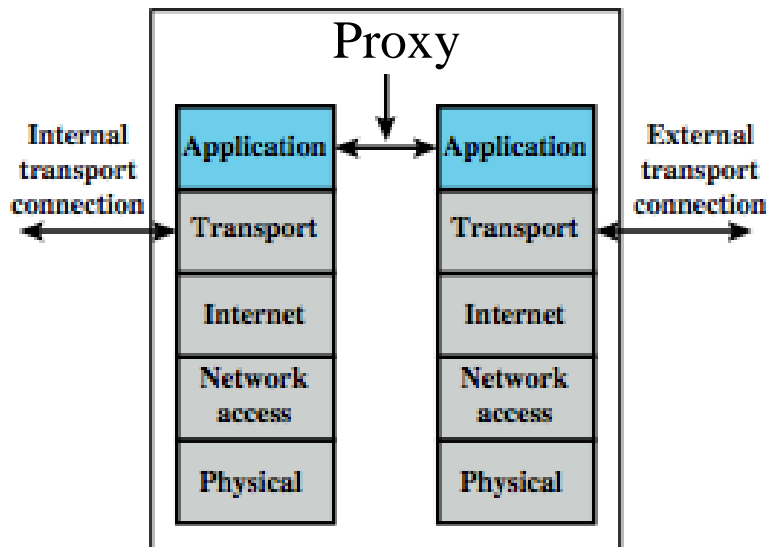
Types of firewalls



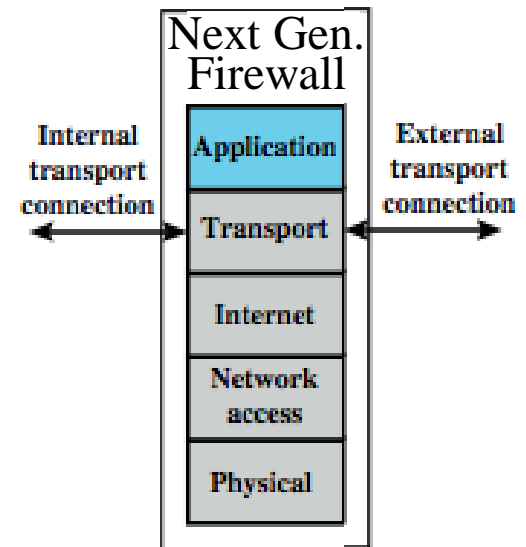
Simple Packet Filter



Stateful Packet Filter



Application Proxy Firewall



Next Generation Firewall

Router-based Packet Filter

- A packet filter is a network router that can accept/reject packets based on headers
- 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
- Unaware of session states at internal or external hosts
- High speed, but primitive filter

Host-based Packet Filters

- A host can also perform packet filtering, in addition to performing other host tasks such as web serving
 - in this case the packet filter is designed to protect the host itself, not other hosts on the network
- Common packet filter software includes:
 - IPChains for Linux (superseded)
 - TCP Wrappers for various Unix
 - IP Filter for Sun Solaris

Stateful Packet Filters

- Stateful packet filters track current state of a connection
 - More ‘intelligent’ than simple packet filters.
- Stateful packet filters keep track of sessions
 - Recognise if a particular packet is part of an established connection by ‘remembering’ recent traffic history.
 - Will add a temporary rule to allow the reply traffic back through the firewall.
 - When session is finished, the temporary rule is deleted.
- This makes the definition of filtering rules easier to accomplish and therefore potentially more secure.
 - High speed, can use relatively advanced filter rules
- Requires memory
 - So can be subject to DOS (Denial of Service) attacks

Packet Filter Strengths and Weaknesses

- **Strengths:**
 - Low overhead and high throughput
 - Supports almost any application
- **Weaknesses:**
 - Unable to interpret application layer data/commands
 - may allow insecure operations to occur
 - Allows direct connection between hosts inside & outside firewall
 - Non-stateful packet filters only: primitive and more difficult to write complex rules

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 and Windows servers ship with Internet Connection Firewall (ICF).

IPv4 Network Address Translation (NAT)

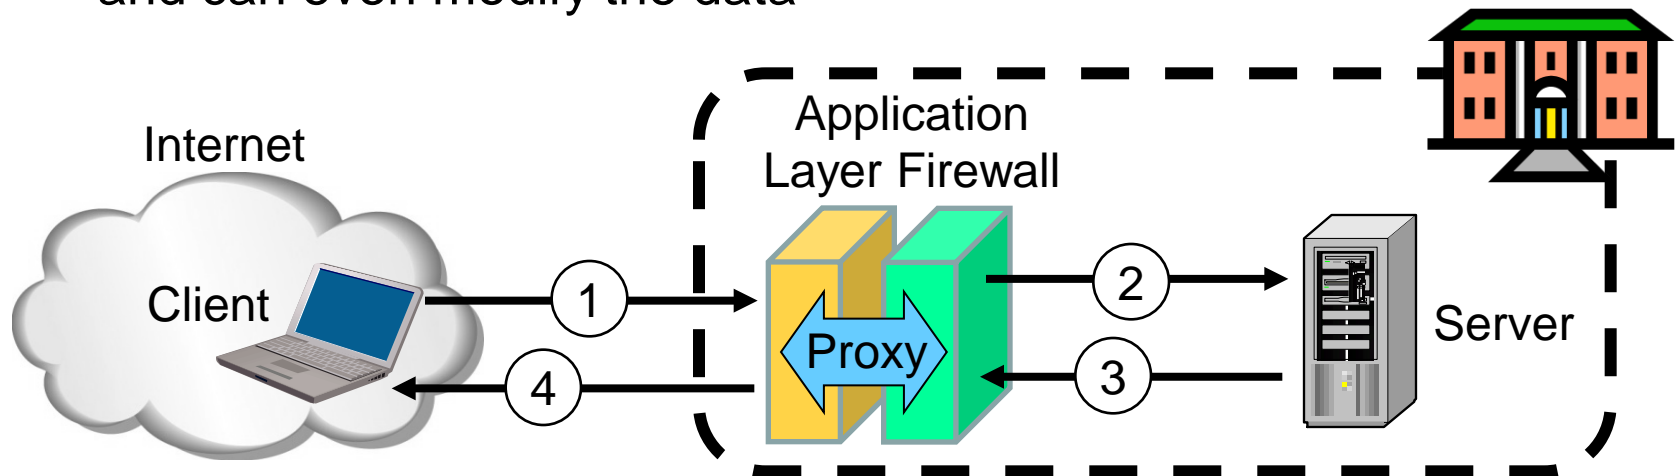
- NAT used to increase IPv4 address space
- Translates public IP addr. ↔ private IP addr. and ports
- Each local network can reuse private IP address ranges
 - Artificially increases the number of usable IP addresses
- Possibilities:
 - Static mapping
 - permanent mapping of public to private address (no gain)
 - 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

IPv4 NAT: + & -

- Advantages
 - Helps enforce control over outbound connections
 - Helps restrict incoming traffic
 - Helps conceal internal network configuration
 - Makes port scanning more difficult
- Can't be used with:
 - protocols that require a separate back-channel
 - protocols that encrypt TCP headers such as IPSec
 - embedded TCP address info
 - (Not recommended with) IPv6

Application Layer Proxy

1. External client sends a request to the server, which is intercepted by the outwards-facing firewall proxy
 2. Inwards-facing proxy sends request to server on behalf of client.
 3. Server sends reply back to inwards-facing firewall proxy.
 4. Outwards facing proxy sends reply to the client.
- Client and server both think they communicate directly with each other, not knowing that they actually talk with a proxy.
 - The proxy can inspect the application data at any level of detail, and can even modify the data



Next Generation Firewalls (NGFW)

- Inspects payload in end-to-end application connection
- Can support specific application protocols
 - e.g. http, telnet, ftp, smtp etc.
 - each protocol supported by a specific proxy HW/SW module
- Can be configured to filter specific user applications
 - E.g. Facebook, Youtube, LinkedIn
 - Can filter detailed elements in each specific user application
- Very high processing load in firewall
 - High volume needs high performance hardware, or else will be slow



High performance NGFWs



High range model: *PA-7050*

Up to 120 Gbps throughput

Prices starting from: US\$ 200,000



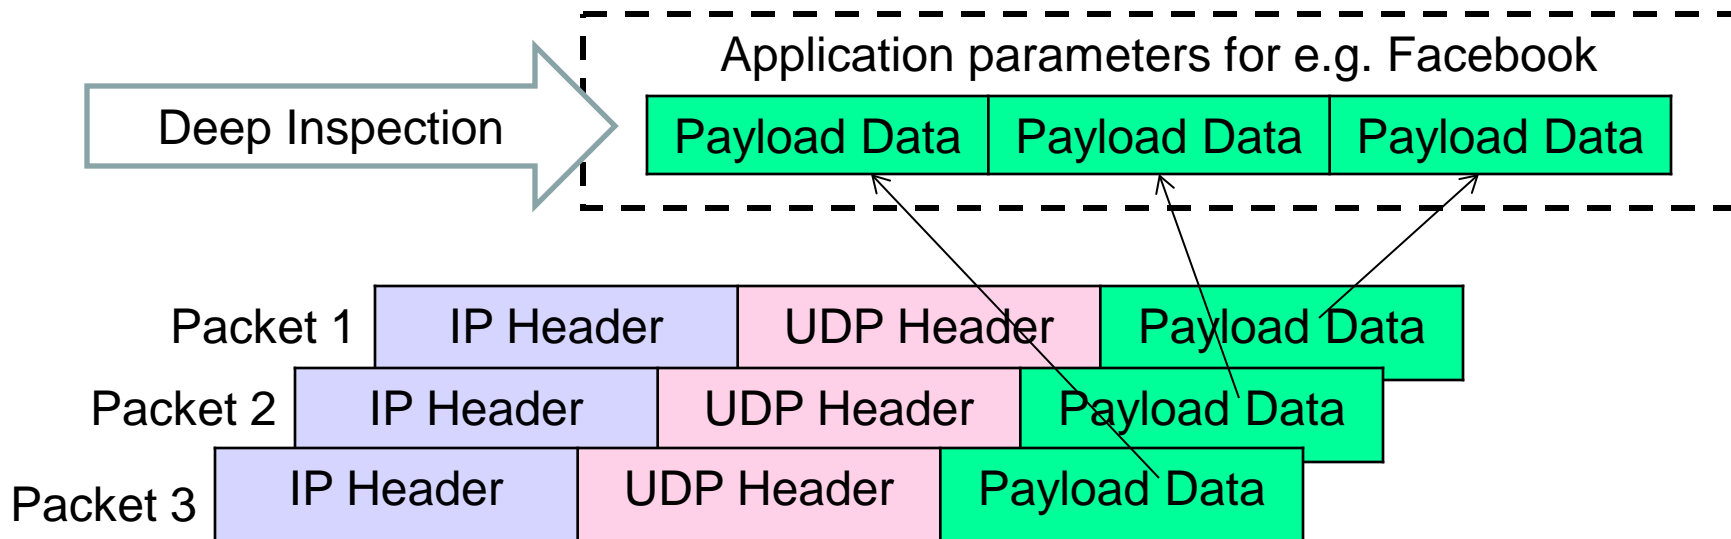
High range model: *61000 Security system*

Up to 400 Gbps throughput

Prices starting from: US\$ 200,000

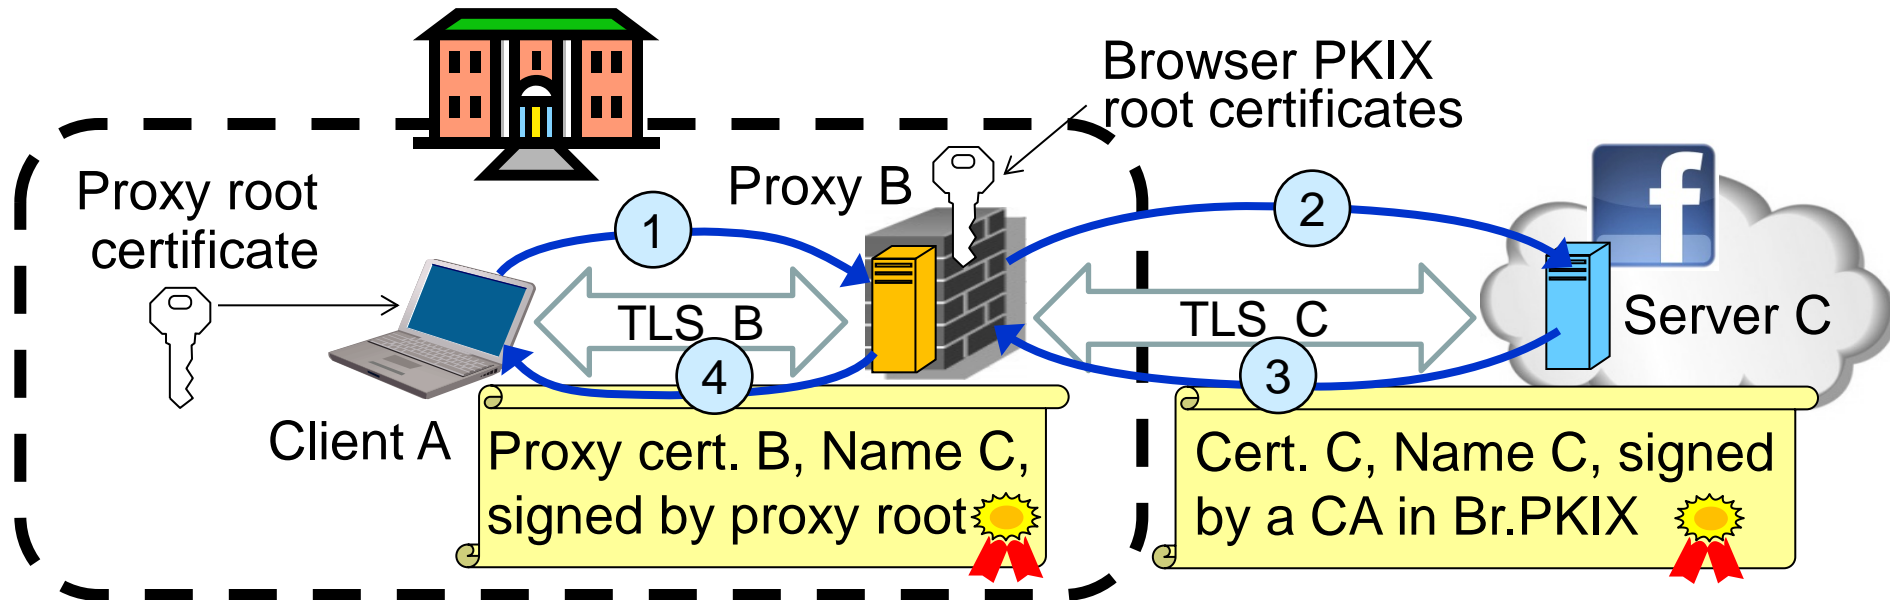
Inline Deep Packet Inspection

- 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



TLS/SSL content inspection in firewalls

- TLS designed for end-to-end encryption, normally impossible to inspect
- In order to inspect TLS, proxy must pretend to be external TLS server
- Proxy creates proxy server certificate with the name of external server (e.g. facebook.com), signed by proxy root private key
- Assumes that proxy root certificate is installed on all internal hosts
- The proxy server certificate is automatically validated by internal client, so user may believe that he/she has TLS connection to the external server

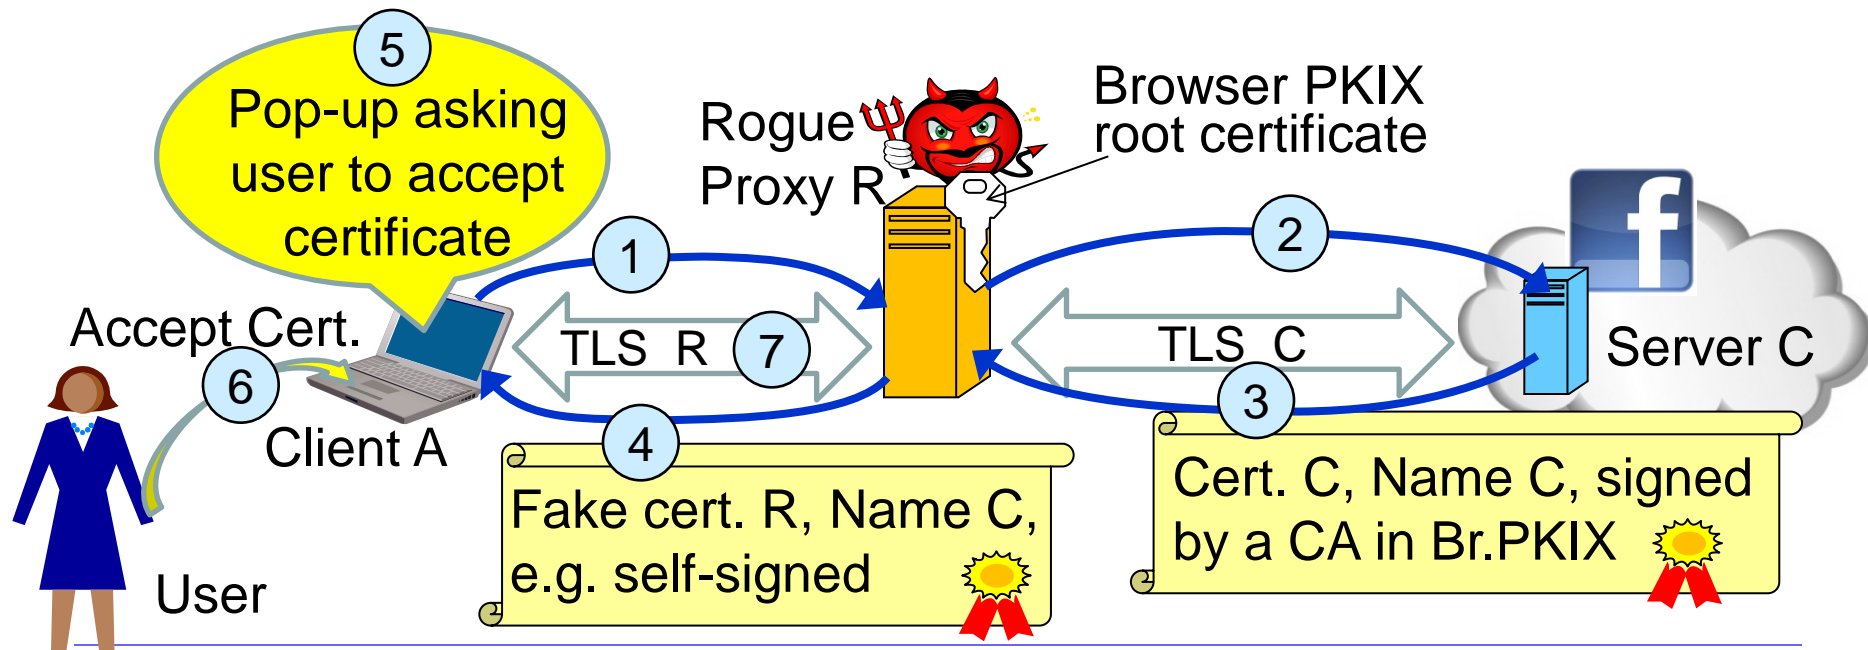


Application Proxy Firewalls + & -

- **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 adapting to new applications
 - Much slower than packet filters
 - Much more expensive than packet filters

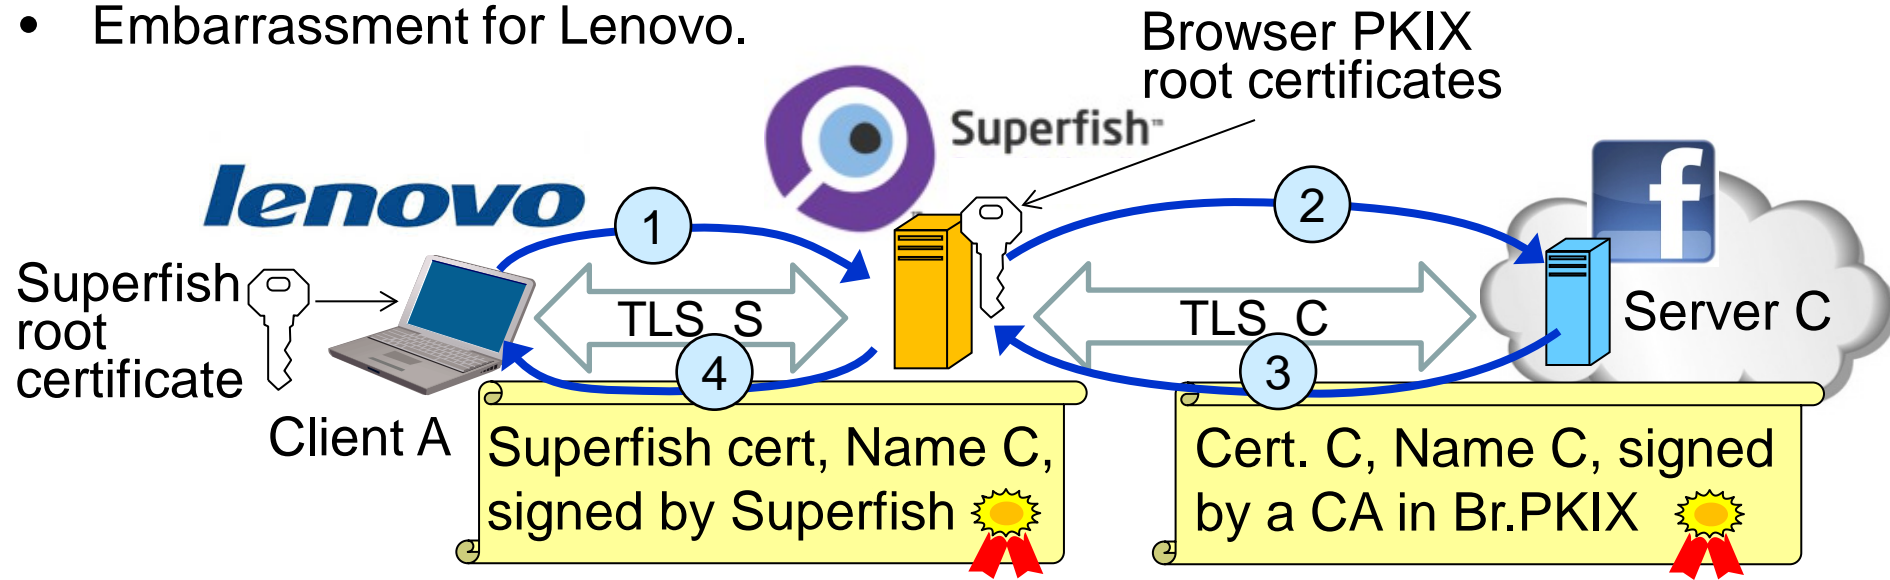
TLS inspection attack with rogue proxy server

- Depending on network, attackers may be able to install rogue proxy
- SSL inspect does **not** assume pre-installed client proxy root certificate
- Proxy creates fake server certificate with the name of external server (e.g. facebook.com), that e.g. can be self-signed
- Fake server certificate is not validated, so browser asks user to accept it
- Fake certificate has (name = domain name), so browser sets up TLS, and user believes that he/she has TLS connection to the external server




Lenovo and the Superfish scam

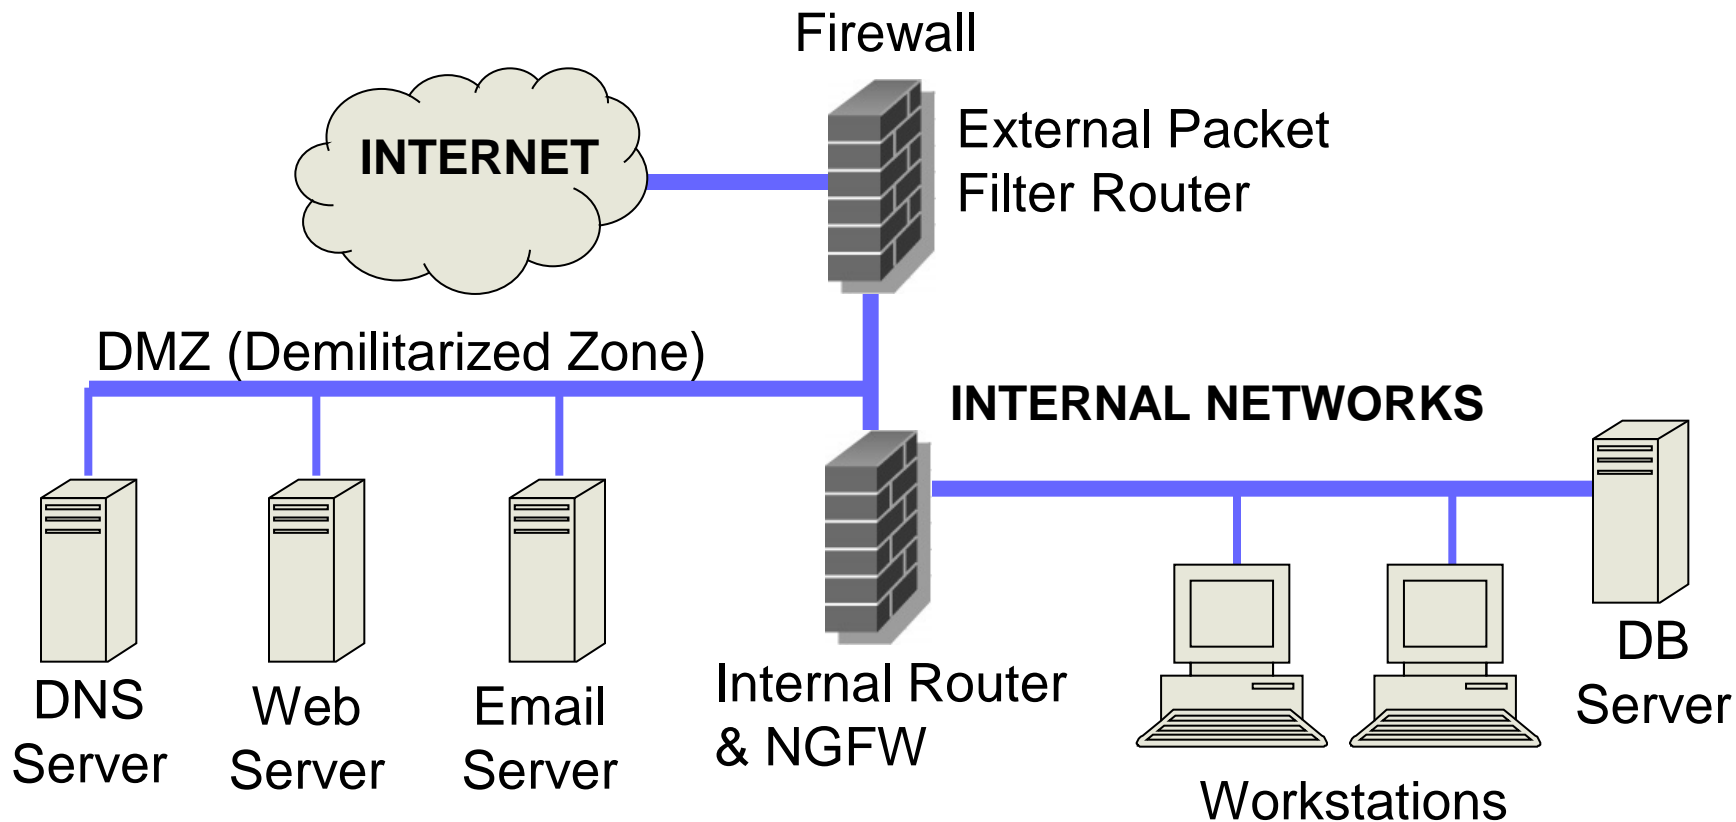
- Superfish root certificate and diversion on some Lenovo models during 2014
- All https connections diverted to Superfish server to inject advertisements.
- Superfish created fake server certificates with names of web servers (e.g. facebook.com), signed by Superfish root private key.
- Fake server certificates were automatically validated, so users got the impression that he/she had https connection to the web server.
- Scam discovered in 2015, Superfish cert. deleted and diversion removed.
- Embarrassment for Lenovo.



Firewalls: Simple Firewall Architecture



Firewalls: DMZ Firewall Architecture



Intrusion Detection Systems

Intrusion Detection and Prevention

- **Intrusion**

- Actions aimed at compromising the security of a target network (confidentiality, integrity, availability of resources)

- **Intrusion detection**

- The identification of possible intrusion through intrusion signatures and network activity analysis
- IDS: Intrusion Detection Systems

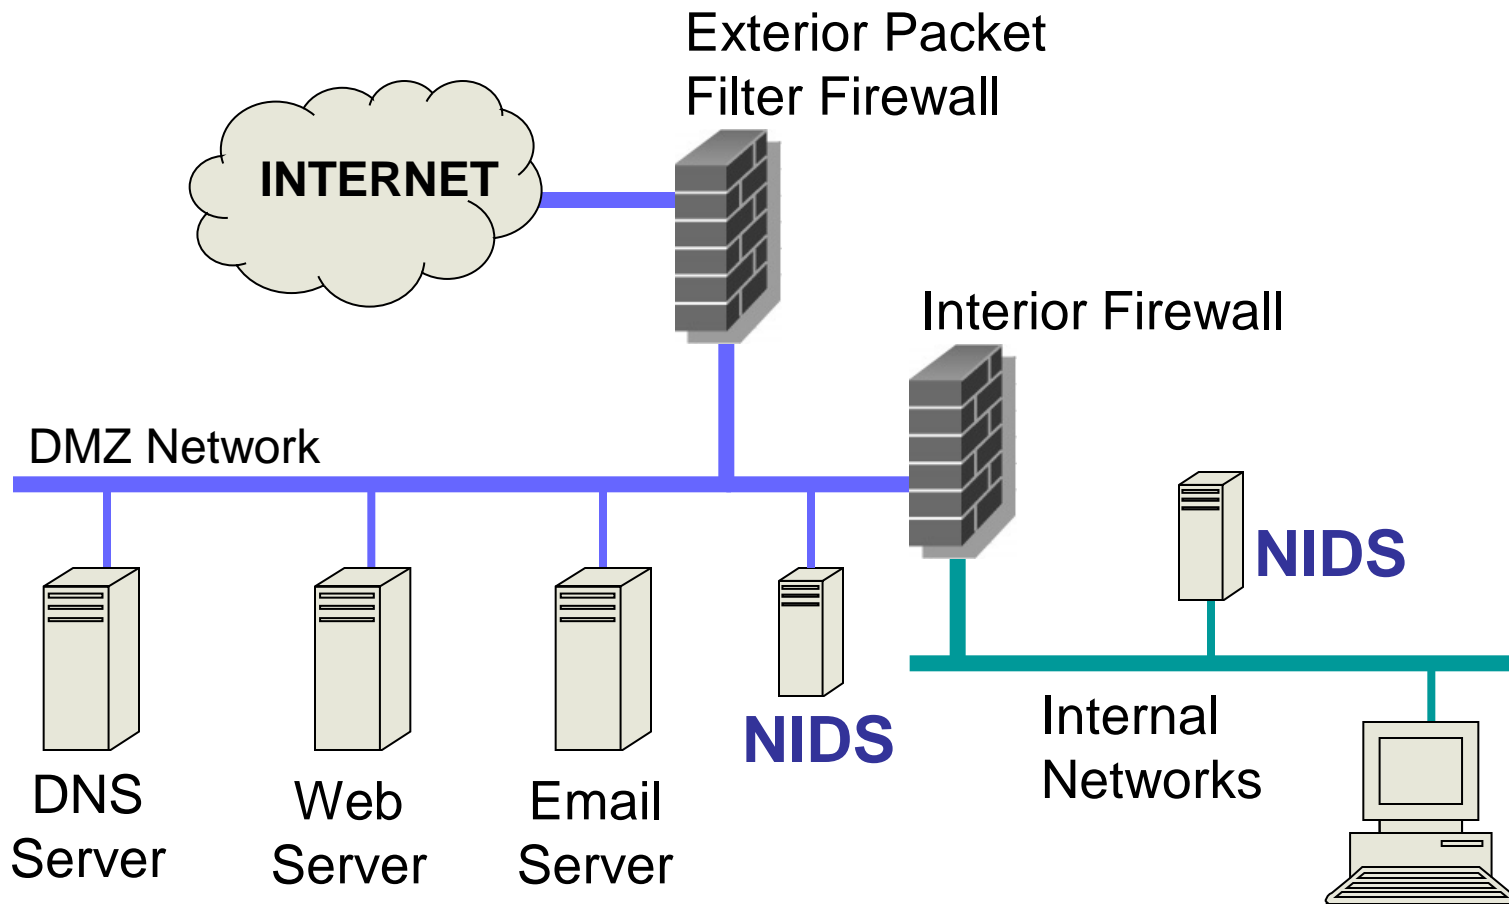
- **Intrusion prevention**

- The process of both detecting intrusion activities and managing automatic responsive actions throughout the network
- IPS: Intrusion Prevention Systems
- IDPS: Intrusion Detection and Prevention Systems

Intrusion Detection Systems:

- IDS are automated systems that detect suspicious activity
- 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 OS files and traffic sent to the host
- Network based IDS (NIDS) detect intrusions on one or more network segments, to protect multiple hosts
 - monitor network/s looking for suspicious traffic
- What can be detected:
 - Attempted and successful misuse, both external and internal agents
 - Malware: Trojan programs, viruses and worms
 - DOS (Denial Of Service) attacks

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
 - Relatively few false positives
- **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
 - Many false positives

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 Interpreter** uses a script defining the network's security policy to decide what to do in response



Port Scanning

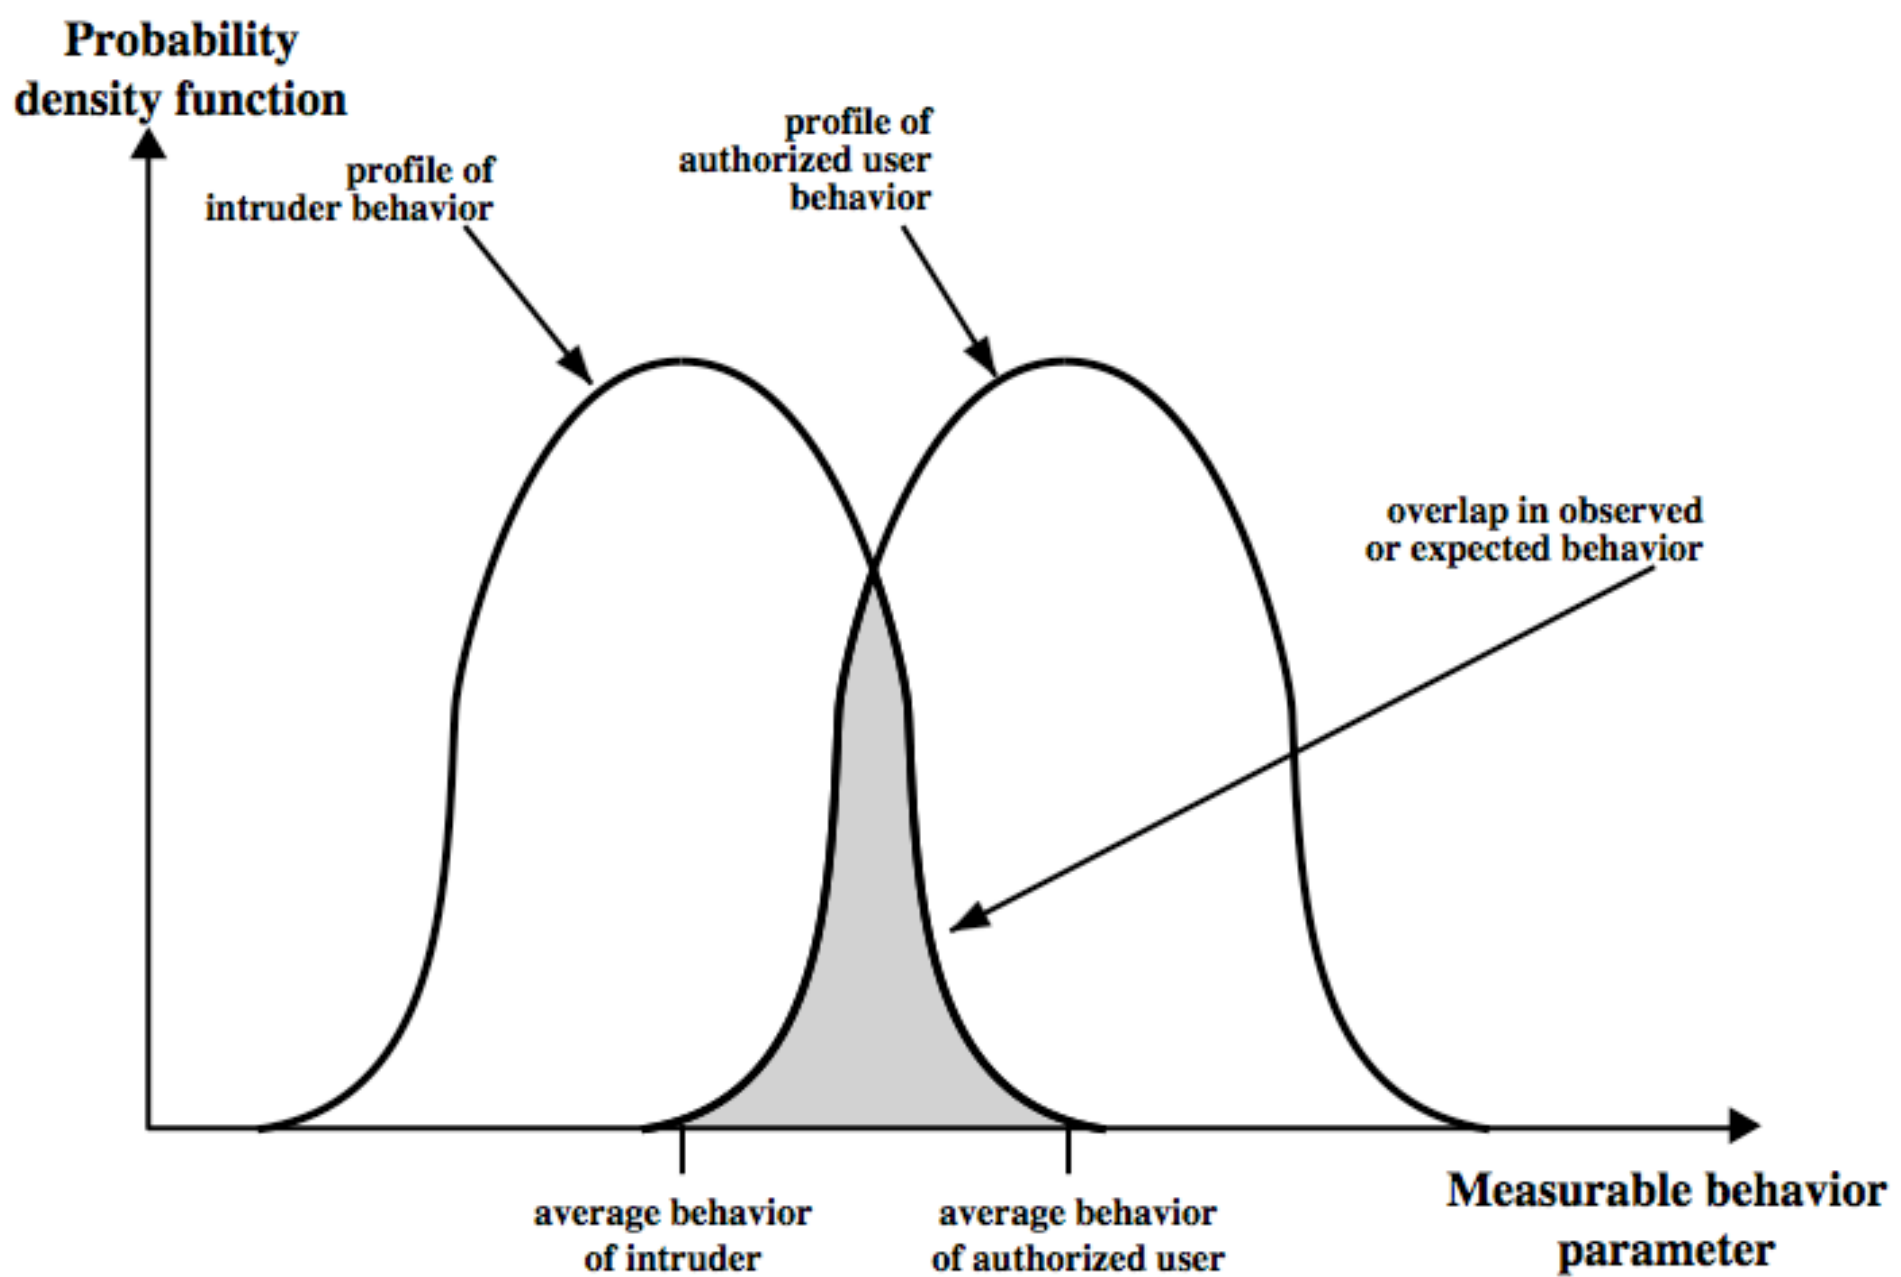
- Many vulnerabilities are OS-specific
 - Bugs in specific implementations, default configuration
- **Port scan** is often a prelude to an attack
 - Attacker tries many ports on many IP addresses
 - For example, looking for an old version of some daemon with an unpatched buffer overflow
 - If characteristic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

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 attack
 - Big problem in statistical anomaly detection
- Both types of IDS suffer from both error types
- Both false positives and false negatives are problematic
 - Attacks are fairly rare events
 - IDS often suffer from “base-rate fallacy”



Base Rate Fallacy

- Consider statements: r : “attack occurs”, s : “signature detected”
 $p(r|s)$: probability of attack, given that signature is detected
 $p(s|r)$: probability of detecting signature, given that attack occurs
 $p(s|\neg r)$: 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 $p(r|s)$
- Base rate fallacy is to assume $p(r|s) \approx 1$ without considering $a(r)$
‘ $p(r|s) \approx 1$ ’ is a good approximation when $a(r) \approx 1$ or $p(s|\neg r) \approx 0$
‘ $p(r|s) \approx 1$ ’ is a bad approximation when $a(r) < 1$ and $p(s|\neg r) > 0$
- Correct $p(r|s)$ requires $a(r)$:
$$p(r | s) = \frac{a(r)p(s | r)}{a(r)p(s | r) + (1 - a(r))p(s | \neg r)}$$

Remarks on Intrusion Detection

- Most alarms are false positives
 - Requires automated screening and filtering of alarms
- Most true positives are trivial incidents
 - can be ignored,
 - the attacks will never be able to penetrate any system
- Serious incidents need human attention
 - Can be dealt with locally
 - May require external expertise
- Potential for improvement through more intelligent IDS
 - Less false positives
 - Better detection of advanced attacks (APT)

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
- HoneyNet is an international security club, see next slide

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
- False positives are problematic, because automated prevention measures can block services

WLAN Security



IEEE 802.11 Standards for WLAN

- IEEE 802.11 formed in 1990's
 - charter to develop a protocol & transmission specifications for wireless LANs (WLANs)
- Since then the demand for WLANs, at different frequencies and data rates, has exploded
- New ever-expanding list of standards issued
 - from 10Mbps to 1Gbps transmission rate



802.11 WiFi Security

- Only authorized terminals (or users) may get access through Wireless LAN
- Should be impossible to set up rogue AP
- Interception of traffic by radios within range should be impossible

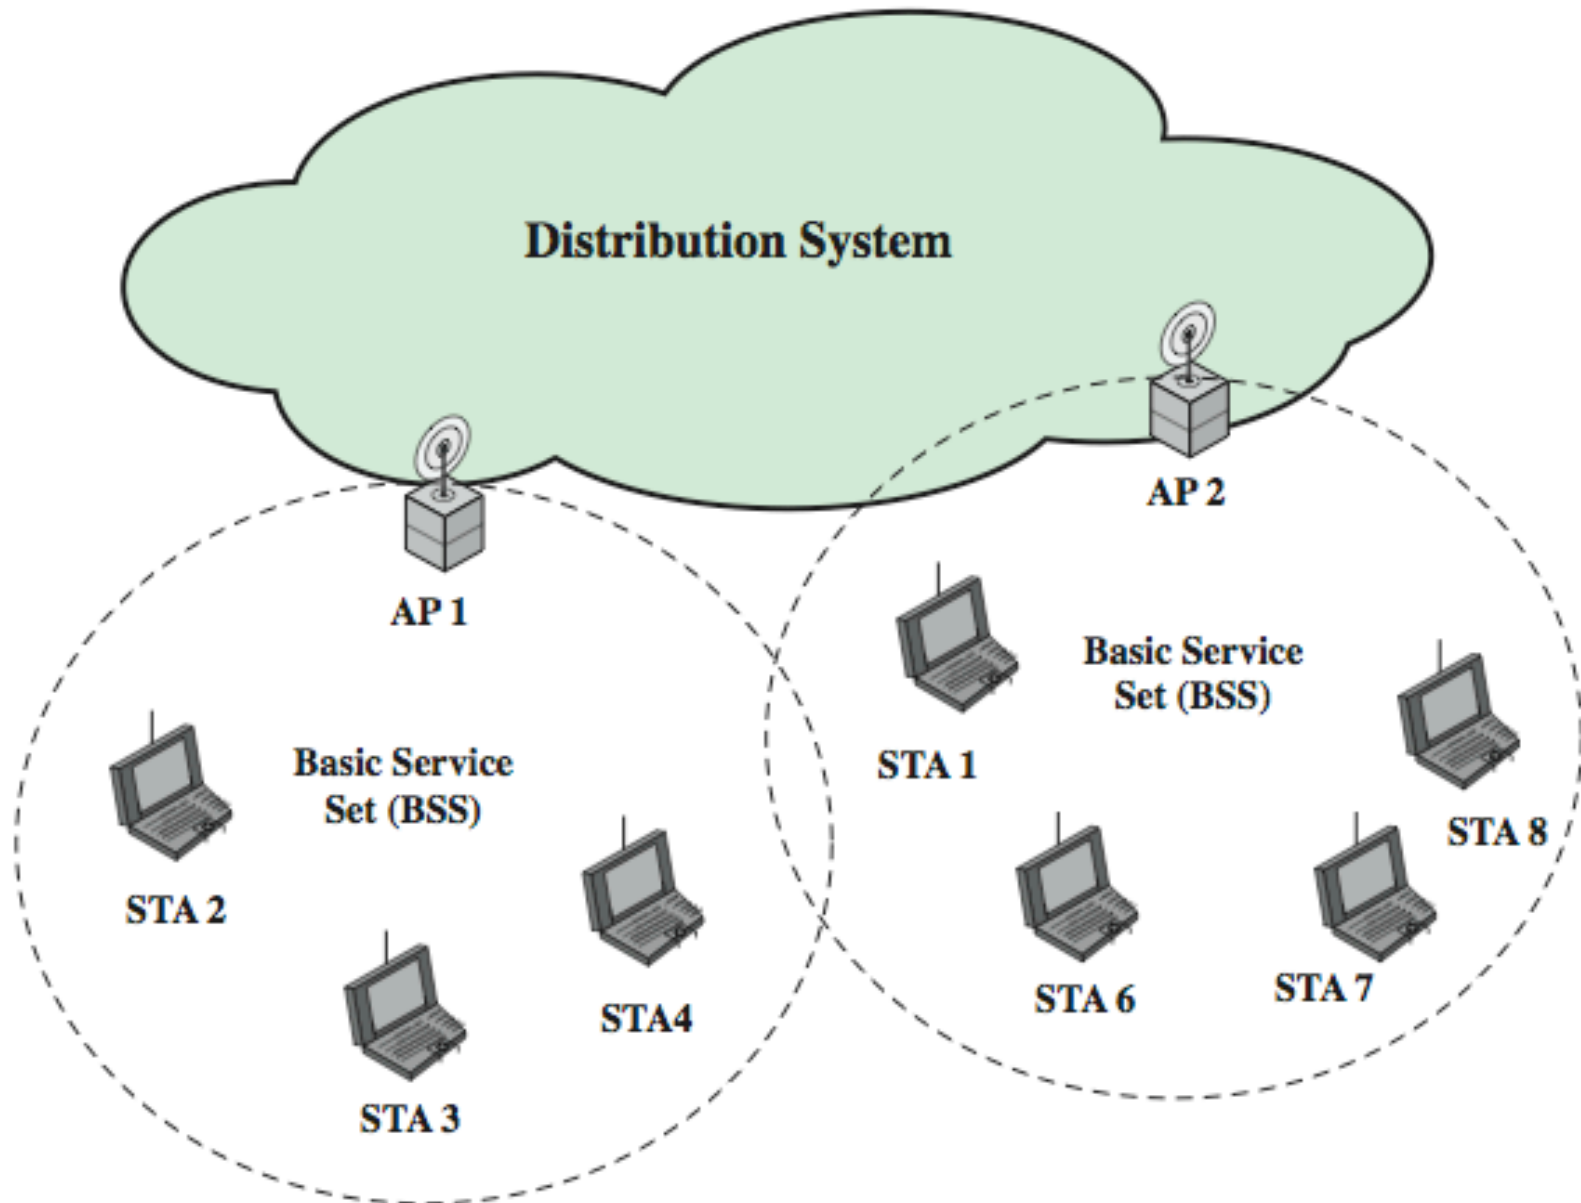
	WEP (1999) 801.11b	WPA (2003) 802.11i (subset)	WPA2 (2004) (aka. RSN) 802.11i (full set)
Auth. & key gen.	WEP	EAP	EAP
Encryption	RC4	RC4+TKIP	CCMP AES CTR (or TKIP)

- WEP: Wired Equivalent Privacy (broken)
- WPA: WiFi Protected Access
- EAP: Extensible Authentication Protocol
- RC4: Rivest Cipher 4 (a stream cipher)
- TKIP: Temporal-Key Integrity Protocol
- CCMP: Counter Mode with CBC Message Authentication Protocol
- RSN: Robust Security Network

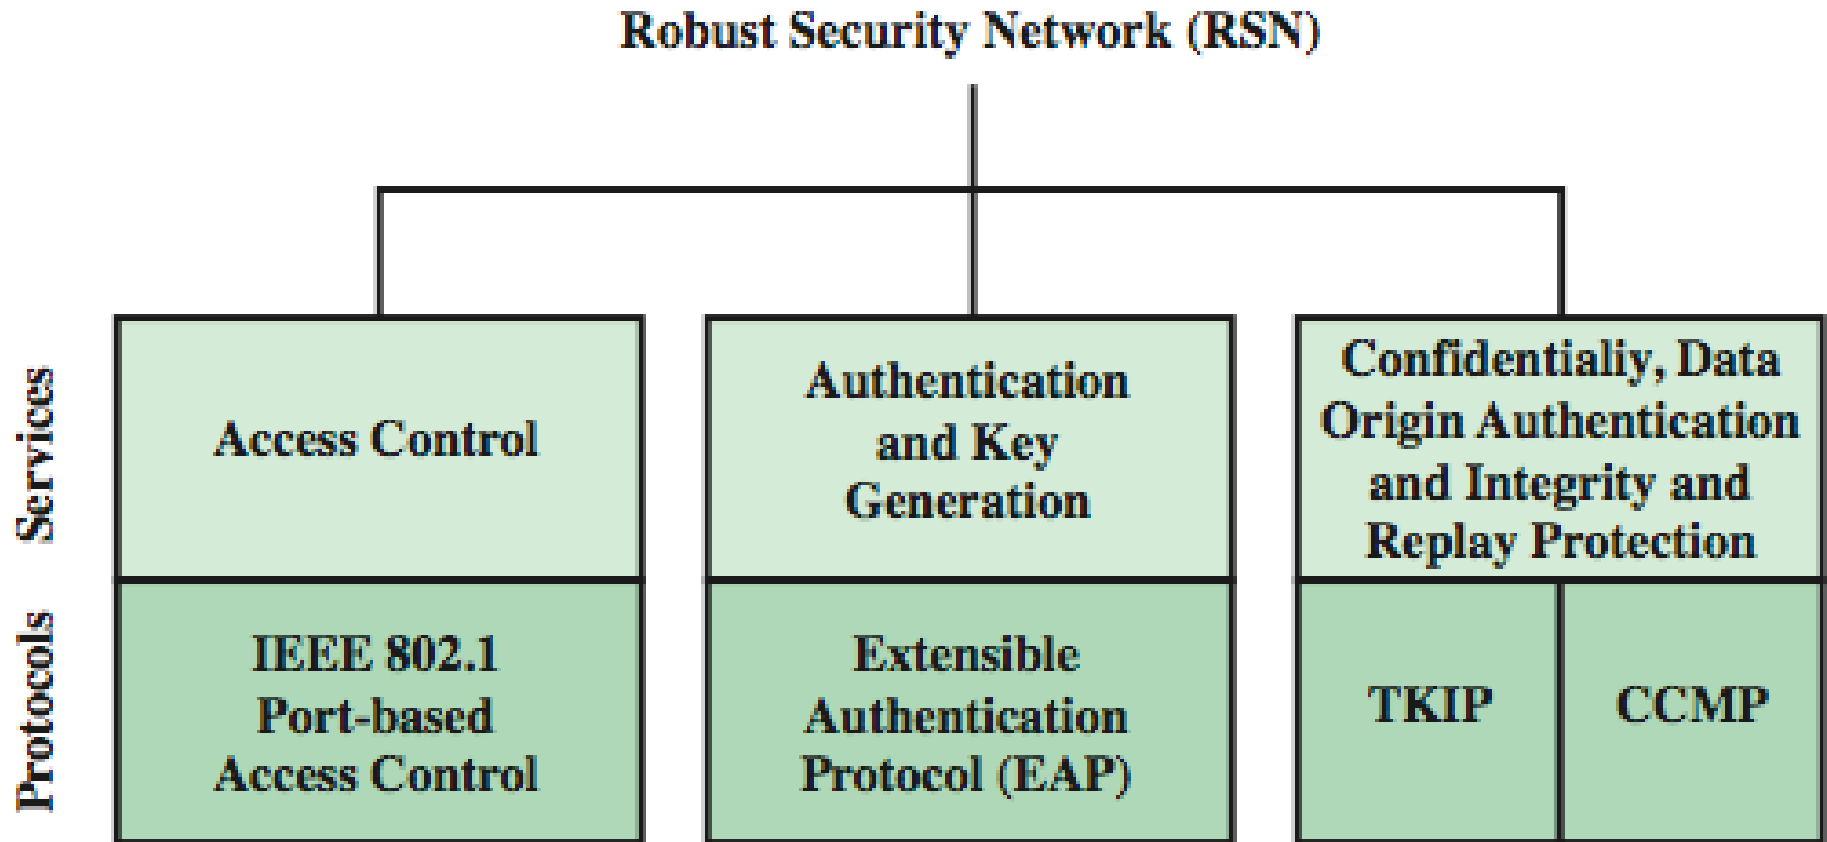
IEEE 802 Terminology

- Station (STA)
 - Wireless terminal that communicates with 802.11 functionality
- Access Point (AP)
 - Receives radio signals and controls access to network
- Basic Service Set (BSS)
 - Set of stations and one AP
- Extended Service Set (ESS)
 - Set of multiple BSSs
- Distribution System (DS)
 - Contains an Authentication Server (AS)
 - Integrates multiple BSSs into one ESS

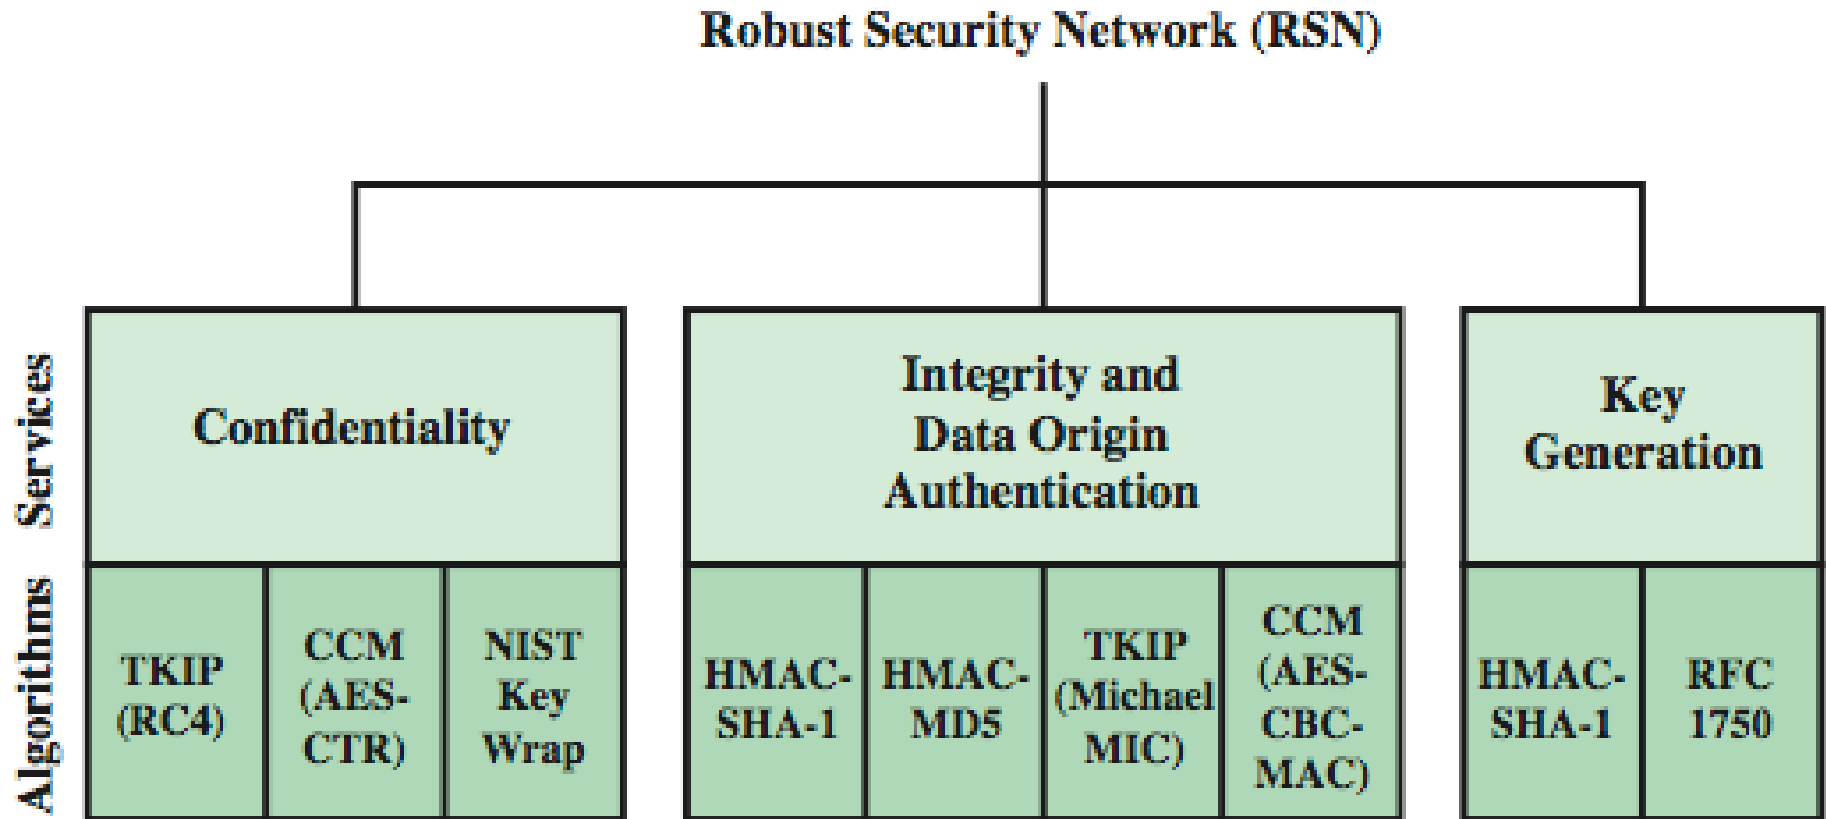
Network Components & Architecture



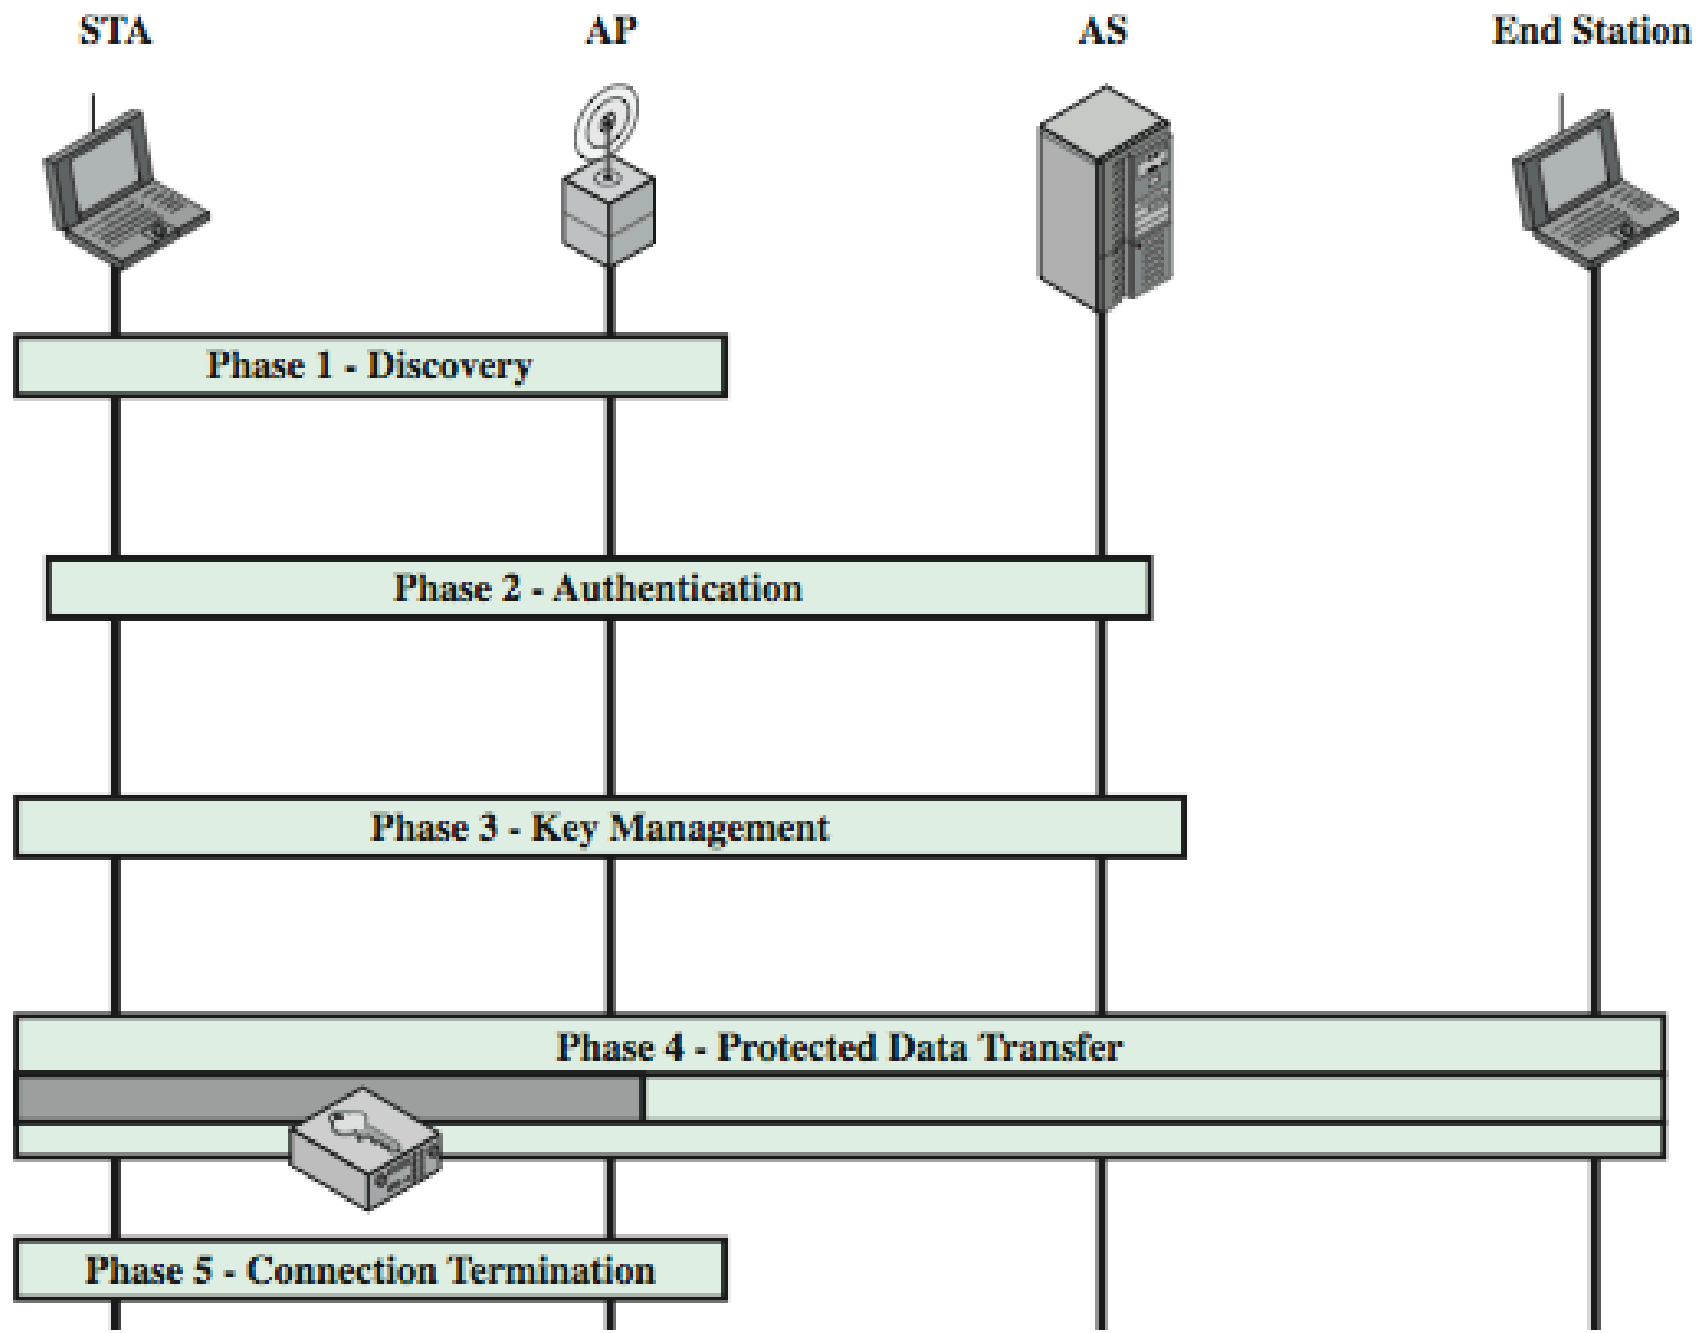
802.11i RSN Services and Protocols



802.11i RSN Cryptographic Algorithms

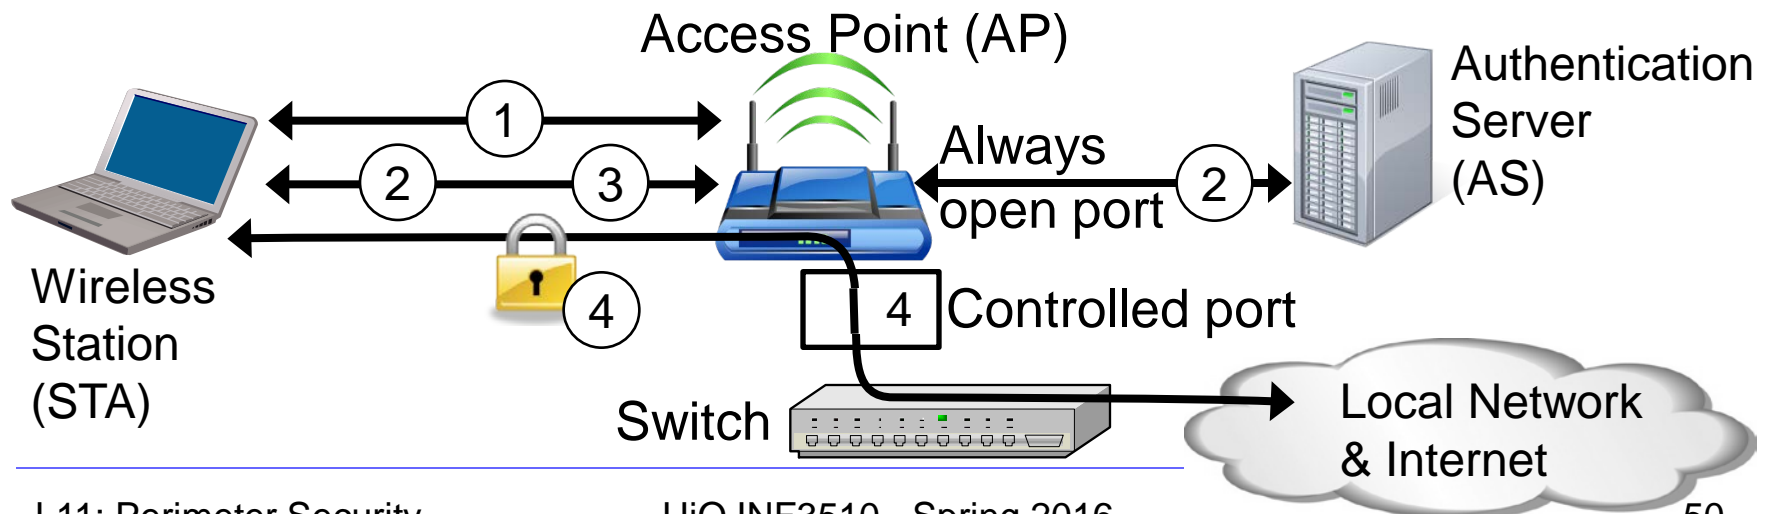


802.11i Phases of Operation



802.11i WiFi Access Control

1. Mutual identity request between STA and AP
 2. Mutual authentication between STA and AS.
 3. Derive pairwise master key (PMK) between STA and AP.
 4. Encrypt radio link and open port (connect) to network access
- Controlled port from AP to network
 - is closed (disconnected) before authentication
 - is open (connected) after successful authentication



When you don't control the WLAN

- Often you want to connect to a wireless LAN over which you have no control, e.g. in café
- Options:
 - If you can, connect securely (WPA2, 802.11i, etc.)
 - Beware of SSL-stripping
 - If unsecured, connect to online resources securely:
 - Use a VPN (Virtual Private Network)
 - IPSEC connection to home gateway
 - TLS/SSL connections to secure web server (with HSTS)
 - Be careful not to expose passwords
 - Watch for direct attacks on untrusted networks

End of Lecture

This lecture presented:

- Firewall techniques
- Intrusion detection techniques
- WLAN Access