IN2120 Information Security
Autumn 2018

Lecture 11:
## Network Perimeter Security

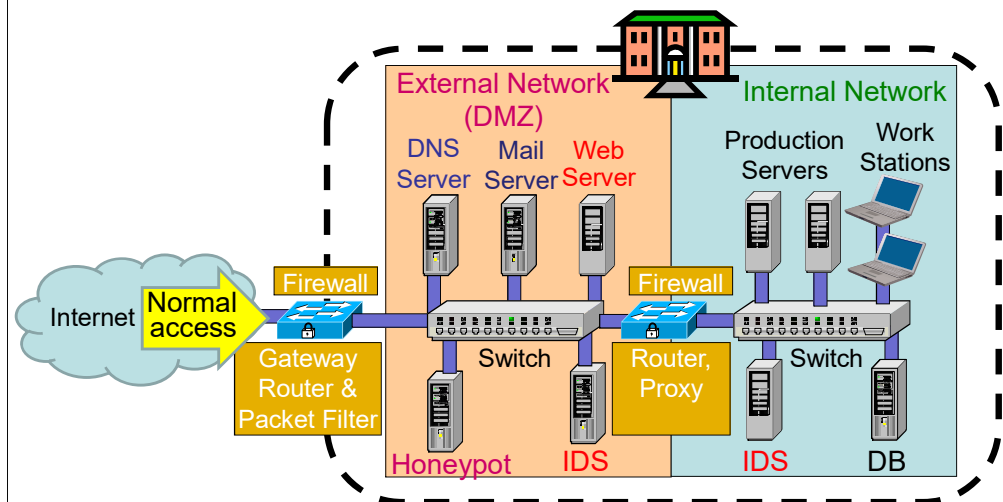*Audun Jøsang*
University of Oslo

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

Observation posts
Outer wall
Inner wall
Guard
Inner court
Normal access
Outer court
Bridge
Gatehouse
Moat

Defending local networks
Network Perimeter Security

External Network (DMZ)
Internal Network
DNS Server
Mail Server
Web Server
Production Servers
Work Stations
Internet
Normal access
Firewall
Firewall
Gateway Router & Packet Filter
Switch
Router, Proxy
Switch
Honeypot
IDS
IDS
DB

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 decides which traffic can pass in & out based on rules

External Network
Potential Threats

Internet

Internal Resources

Firewall
=
Check Point

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)

- Packet Filters

  Inspects packet headers only

- Stateful Packet Filters

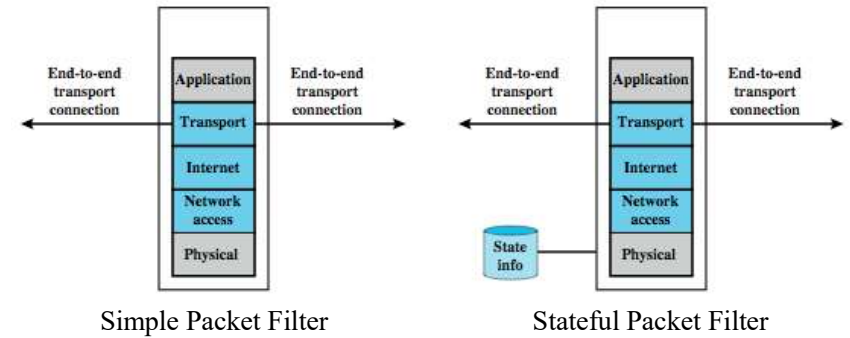 Analyses bi-directional traffic

- Application Layer Proxy

  Splits connection, inspects payload, and analyses traffic
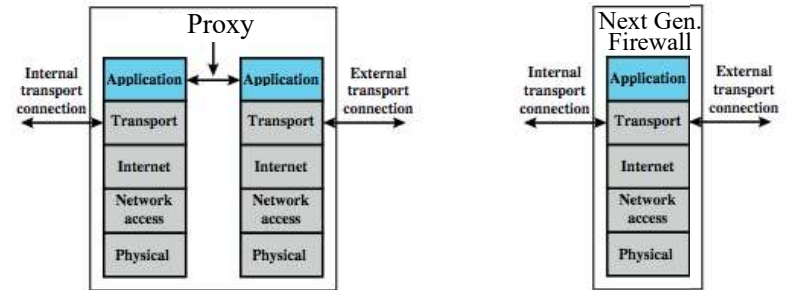
- Next Generation Firewall

  End-to-end connection inspects payload, and analyses traffic

Types of firewalls

Simple Packet Filter

Stateful Packet Filter

Application Proxy Firewall

Next Generation Firewall

(Stateless) 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

(Stateless) Packet Filters

- Widespread packet filter software (Linux):
  - iptables / netfilter
  - nft / nttables
- Examples (iptables)
- `iptables -A FORWARD -s 131.234.142.33 -j ACCEPT`
  - All packets from source IP Address 131.234.142.33 are accepted
- `iptables -A FORWARD -p tcp -d 10.0.0.56 --dport 22 -j ACCEPT`
  - All packets using transport protocol and destination address 10.0.0.56 and destination port 22 are accepted

Problems with Stateless Filtering

- Assume a typical "security policy":
  - Access from internal to external allowed
  - Access from external to internal prohibited
  - Example application: home network
- Naive packet filter configuration:
  - outgoing packet → forward
  - incoming packet → reject

- Most internet applications would not work!

Stateful Filtering
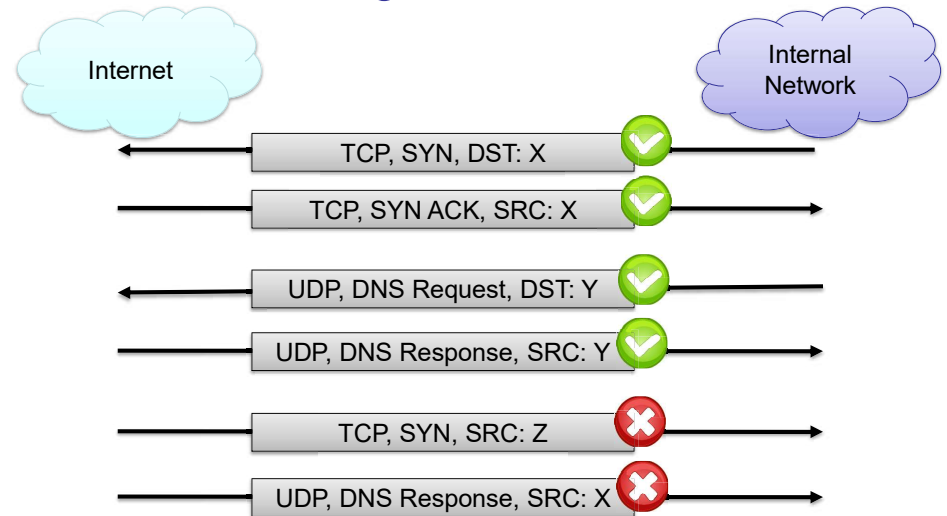
Stateful Packet Filters

- 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

Stateful Packet Filters

- Examples (iptables)
- `iptables -A FORWARD -m state --state NEW -i eth0 -j ACCEPT`
- Accept new connections (i.e. TCP SYN) from network interface eth0 ("from inside")
- `iptables -A FORWARD -m state --state ESTABLISHED,RELATED -j ACCEPT`
- Accept ALL packets which belong to an established TCP connection or are related to an existing UDP communication

(Stateful) Packet Filter: Evaluation

- Strengths:
  - Low overhead and high throughput
  - Supports almost any application
- Weaknesses:
  - Unable to interpret application layer data/commands
    - may allow insecure operations to occur
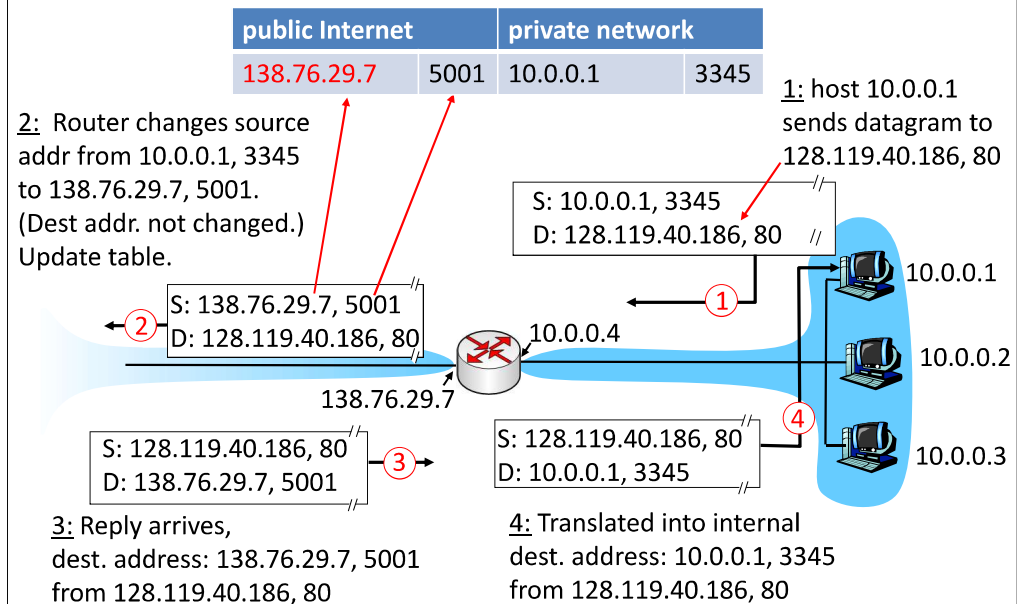 - Allows direct connection between hosts inside & outside firewall

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
- Nowadays for example included in Windows
- Advantage compared to network firewall: rules can take applications into account

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 Network Address Translation (NAT)

public Internet		private network	
138.76.29.7	5001	10.0.0.1	3345

1: host 10.0.0.1 sends datagram to 128.119.40.186, 80

2: Router changes source addr from 10.0.0.1, 3345 to 138.76.29.7, 5001. (Dest addr. not changed.) Update table.

S: 10.0.0.1, 3345
D: 128.119.40.186, 80

S: 138.76.29.7, 5001
D: 128.119.40.186, 80

10.0.0.4

138.76.29.7

S: 128.119.40.186, 80
D: 138.76.29.7, 5001

S: 128.119.40.186, 80
D: 10.0.0.1, 3345

10.0.0.1

10.0.0.2

10.0.0.3

3: Reply arrives, dest. address: 138.76.29.7, 5001 from 128.119.40.186, 80

4: Translated into internal dest. address: 10.0.0.1, 3345 from 128.119.40.186, 80

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 or proxy application connection
- 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
- Can support TLS/SSL encrypted traffic inspection
- Can provide intrusion detection and intrusion prevention
- 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
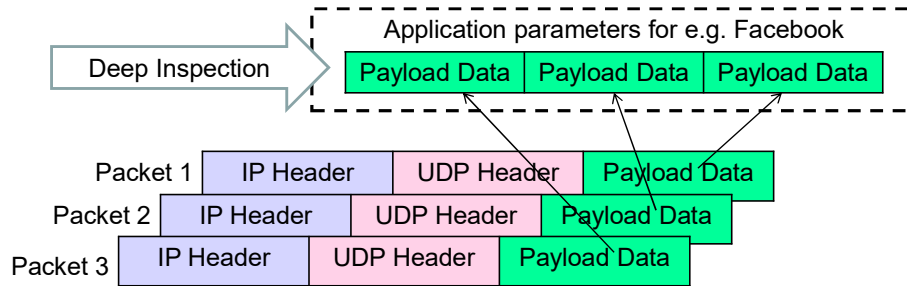
Prices starting from: US$ 200,000

High range model: *61000 Security system*

Up to 400 Gbps throughput
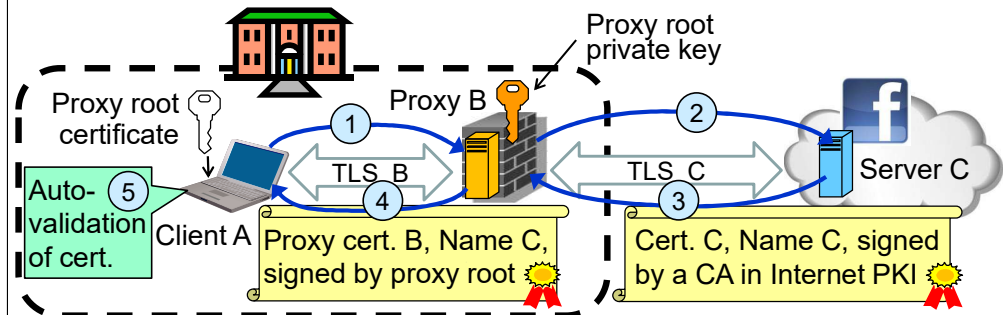
Prices starting from: US$ 200,000

Inline Deep Packet Inspection

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

Application parameters for e.g. Facebook

Deep Inspection → | Payload Data | Payload Data | Payload Data |

Packet 1: | IP Header | UDP Header | Payload Data |
Packet 2: | IP Header | UDP Header | Payload Data |
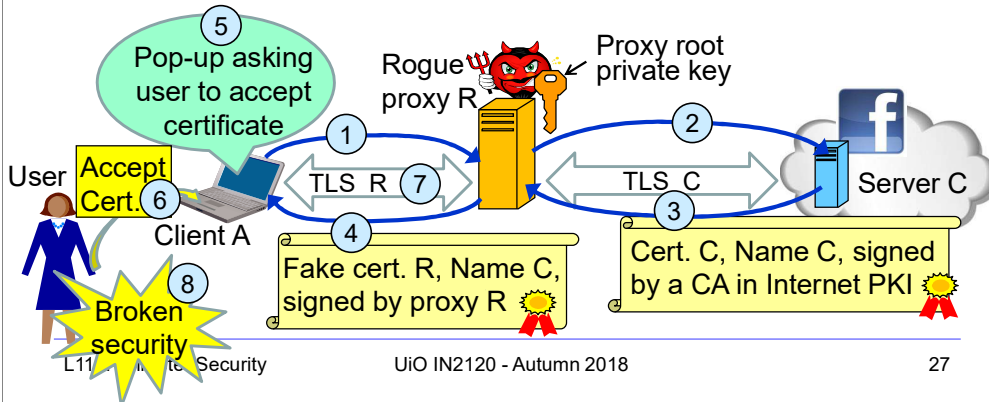Packet 3: | IP Header | UDP Header | Payload Data |

TLS/SSL encrypted traffic 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 local proxy root private key
- Assumes that local proxy root certificate is installed on all local hosts
- The proxy server certificate is automatically validated by local client, so user may believe that he/she has TLS connection to the external server

Proxy root private key

Proxy root certificate

Auto-validation of cert. (5)

Client A

Proxy B (1) (4) TLS_B

Proxy cert. B, Name C, signed by proxy root

(2) (3) TLS_C

Server C

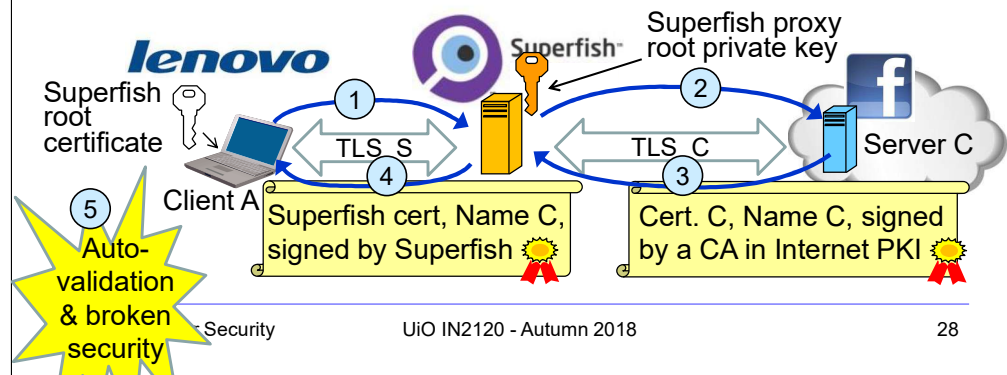Cert. C, Name C, signed by a CA in Internet PKI

TLS inspection attack with rogue proxy server

- Depending on network, attackers may be able to install rogue proxy
- Rogue TLS inspect does **not** assume pre-installed 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 dame), so browser sets up TLS, and user believes that he/she has TLS connection to the external server

(5) Pop-up asking user to accept certificate

Rogue proxy R

Proxy root private key

User

Accept Cert. (6)

Client A

(1) (4) TLS_R (7)

Fake cert. R, Name C, signed by proxy R

(2) (3) TLS_C

Server C

Cert. C, Name C, signed by a CA in Internet PKI

(8) Broken security

---

Lenovo and the Superfish scam

- Superfish root certificate and diversion in shipped 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 believed that he/she had secure end-to-end https connection to the web server.
- Scam discovered in 2015, Superfish cert. deleted and diversion removed.
- Embarrassment for Lenovo. Superfish changed name to JustVisual.

Superfish root certificate

Superfish proxy root private key

Client A

(1) (4) TLS_S

Superfish cert, Name C, signed by Superfish

(2) (3) TLS_C

Server C

Cert. C, Name C, signed by a CA in Internet PKI

(5) Auto-validation & broken security

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

## Firewalls: Simple Firewall Architecture

Firewalls: DMZ Firewall Architecture

DMZ Example

- DMZ = A part of your LAN with other restrictions, e.g. allowing publicly available services (web servers, mail etc.)
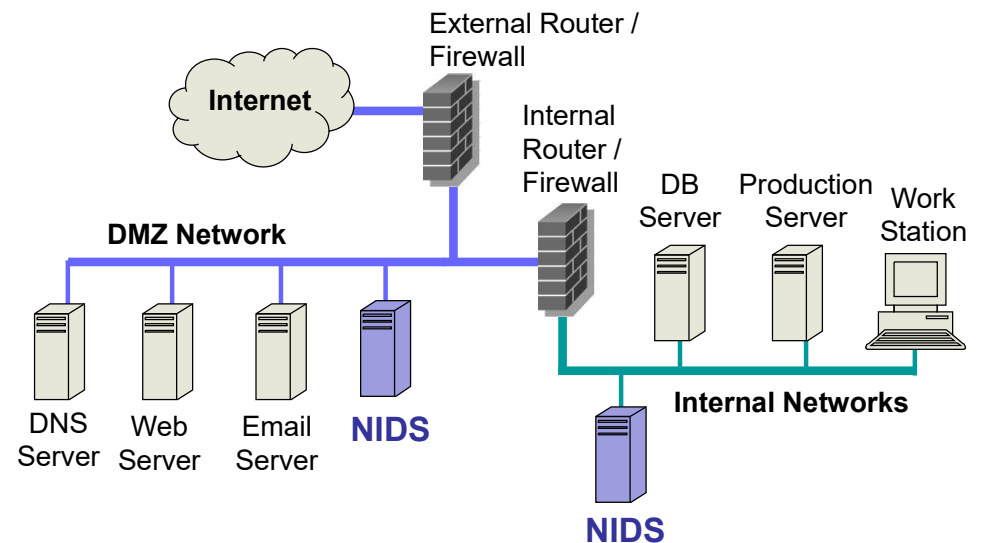
Intrusion Detection Systems

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 events, 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 networks 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 can do, based on known attack patterns
  - 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

Example: Vulnerability + Snort Rule

CVE-2017-0147 Detail

Current Description

The SMBv1 server in Microsoft Windows Vista SP2; Windows Server 2008 SP2 and R2 SP1; Windows 7 SP1; Windows 8.1; Windows Server 2012 Gold and R2; Windows RT 8.1; and Windows 10 Gold, 1511, and 1607; and Windows Server 2016 allows remote attackers to obtain sensitive information from process memory via a crafted packets, aka "Windows SMB Information Disclosure Vulnerability."

Source: MITRE
Description Last Modified: 03/16/2017
+View Analysis Description

```
alert tcp $HOME_NET 445 -> any any ( msg:"OS-WINDOWS Microsoft Windows SMB
possible leak of kernel heap memory"; flow:to_client,established;
content:"Frag",fast_pattern; content:"Free"; content:"|FA FF FF|";
content:"|F8 FF FF|",within 3,distance 5; content:"|F8 FF FF|",within
3,distance 5; metadata:policy balanced-ips alert,policy security-ips
drop,ruleset community; service:netbios-ssn; reference:cve,2017-0147;
reference:url,technet.microsoft.com/en-us/security/bulletin/MS17-010;
classtype:attempted-recon; sid:42339; rev:2; )
```

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 of vulnerability detected, mount attack
  - "The Art of Intrusion": virtually every attack involves port scanning and password cracking

Port Scanning

- Example: network services on a Windows computer

```
Proto. Local Address          Foreign Address       State
TCP    0.0.0.0:80             0.0.0.0:0             LISTEN
TCP    0.0.0.0:135            0.0.0.0:0             LISTEN
TCP    0.0.0.0:445            0.0.0.0:0             LISTEN
TCP    0.0.0.0:554            0.0.0.0:0             LISTEN
TCP    0.0.0.0:623            0.0.0.0:0             LISTEN
TCP    0.0.0.0:2869           0.0.0.0:0             LISTEN
TCP    0.0.0.0:5357           0.0.0.0:0             LISTEN
TCP    0.0.0.0:10243          0.0.0.0:0             LISTEN
TCP    0.0.0.0:16992          0.0.0.0:0             LISTEN
TCP    0.0.0.0:49152          0.0.0.0:0             LISTEN
TCP    0.0.0.0:49153          0.0.0.0:0             LISTEN
TCP    0.0.0.0:49154          0.0.0.0:0             LISTEN
TCP    0.0.0.0:49155          0.0.0.0:0             LISTEN
TCP    0.0.0.0:49157          0.0.0.0:0             LISTEN
TCP    0.0.0.0:56238          0.0.0.0:0             LISTEN
```

Intrusion Detection Problems

- Lack of training data with real attacks
  - But lots of "normal" network traffic, system call data
- Data drift
  - Statistical methods detect (rapid) changes in behavior
  - Attacker can attack gradually and incrementally to avoid detection
- Discriminating characteristics hard to specify
  - Many attacks may be within bounds of "normal" range of activities
- False identifications are very costly
  - Sysadmin will spend many hours examining evidence of attack which might turn out to be a false positive
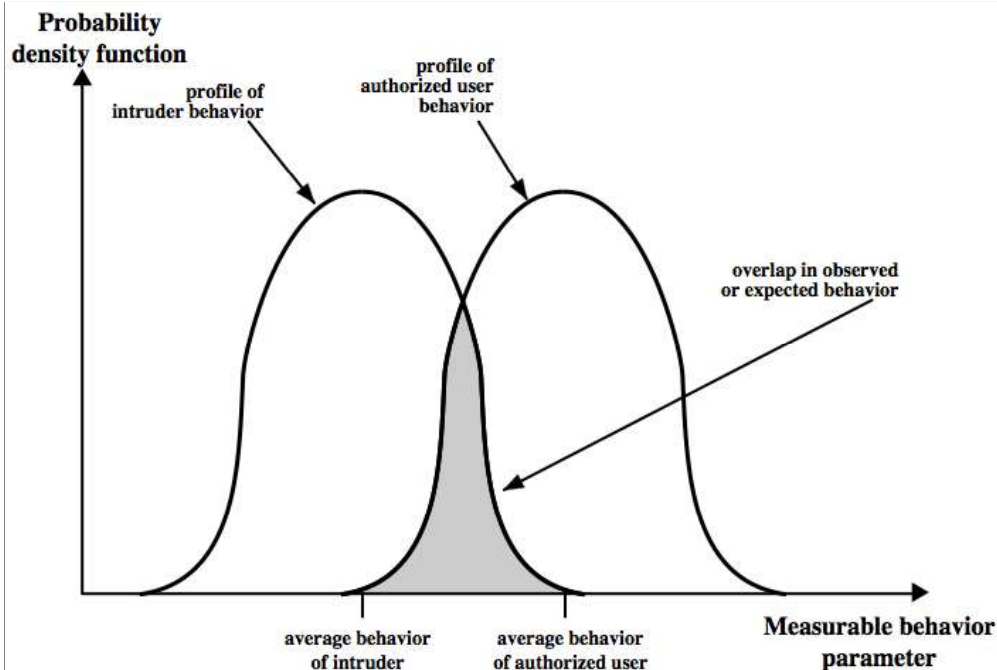
Intrusion Detection Errors

- **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"

Importance of the base-rate of attack

- Consider statements:
  - A: "attack occurs"
  - D: "detection occurs"
- We can measure/estimate:
  - $P(D|A)$: probability of detection, given that attack occurs
  - $P(D|\neg A)$: probability of detection, given that no attack occurs
  - $P(A)$: base-rate probability of attack
  - $P(\neg A) = 1 - P(A)$: base-rate probability of "no attack"
- We want to know the actual probability of attack
  - $P(A|D)$: probability of attack, given that detection triggers
  - Base-rate fallacy (error) is to assume that: $P(A|D) = P(D|A)$
- Bayes' theorem gives the correct answer by including $P(A)$:
  - $P(A|D) = \frac{P(D|A) \cdot P(A)}{P(D)} = \frac{P(D|A) \cdot P(A)}{P(D|A) \cdot P(A) + P(D|\neg A) \cdot P(\neg A)}$

Example of Base Rate Fallacy (error) and correct probability of actual attack

- Example:
  - Assume detection is 99% correct: $P(D|A) = 0.99$, $P(D|\neg A) = 0.01$
  - The base-rate fallacy (error) would be to think that $P(A|D) = 0.99$
  - Assume that the base-rate of attack is: $P(A) = 1 / 10000 = 0.0001$
- $P(A|D) = \frac{0.99 \cdot 0.0001}{0.99 \cdot 0.0001 + 0.01 \cdot 0.9999} = 0.010098$
  - The correct probability of attack given detection is only 0.01

- Conclusion:
  - 1% accuracy
  - 99 false positives per true positive
  - This would be a useless detection system

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)

UTM - Unified Threat Management

- UTM platforms consolidate different security features into a single hardware and/or software system.
- Security features typically combined in a UTM are e.g.:
  - Network firewall
  - Intrusion detection
  - Intrusion prevention (can block traffic classified as malicious)
  - Gateway anti-virus
  - Deep packet inspection
  - Web proxy, content filtering and TLS Inspection
  - Data loss prevention (DLP)
- Can be seen a NGFW which also integrates many features
- UTM systems are also known as TM (Threat Management), ITM (Integrated Threat Management (UTM), and STM (Security Threat Management) systems.

Honeypots

- A honeypot:
  - is a computer configured to detect network attacks or malicious behavior,
  - 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

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

Wireless Station (STA)

WiFi Security

Access Point (AP)

Internet

802.11 WiFi Security

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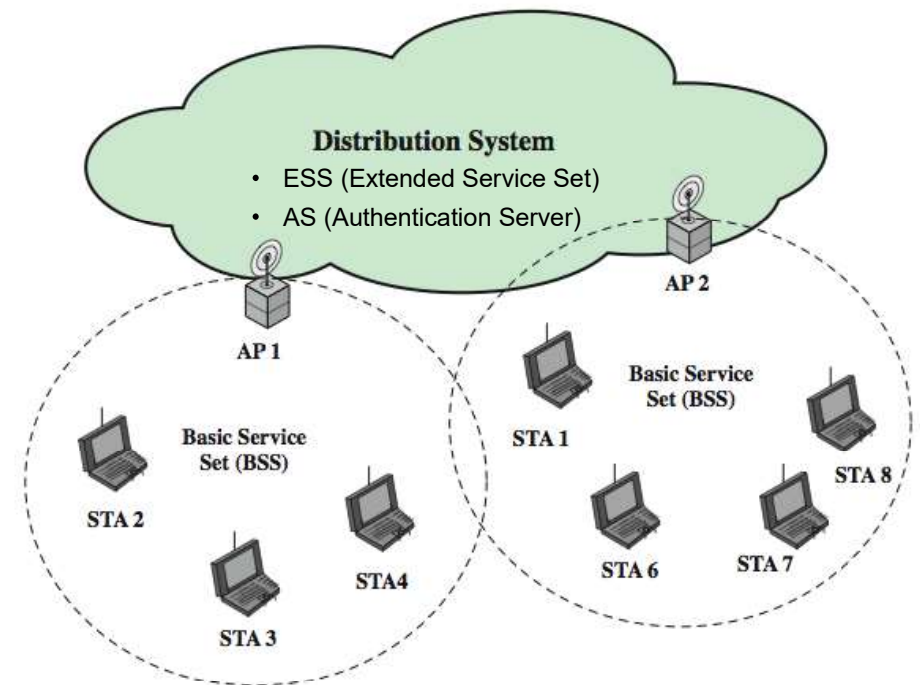- Interception of traffic by radios within range should be impossible

	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
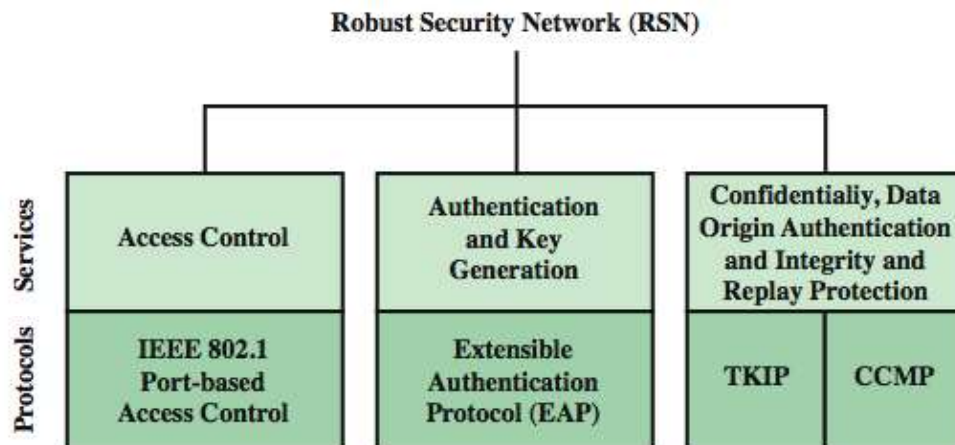- RSN: Robust Security Network

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)
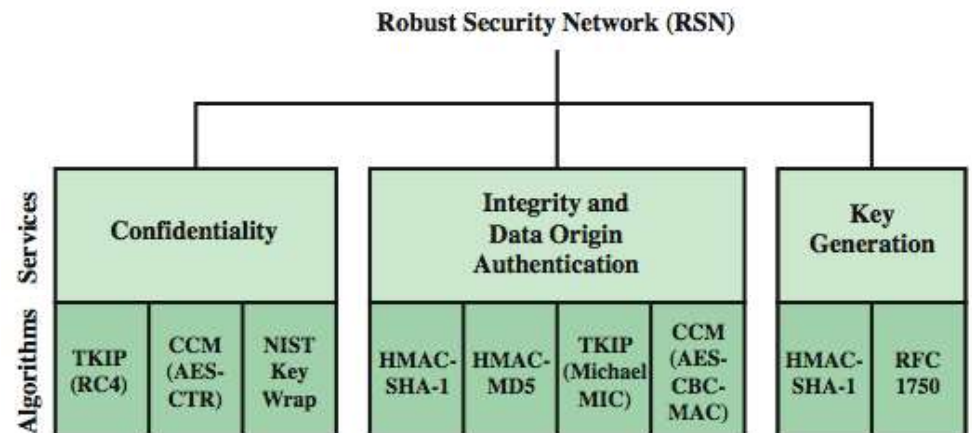 - Integrates multiple BSSs into one ESS

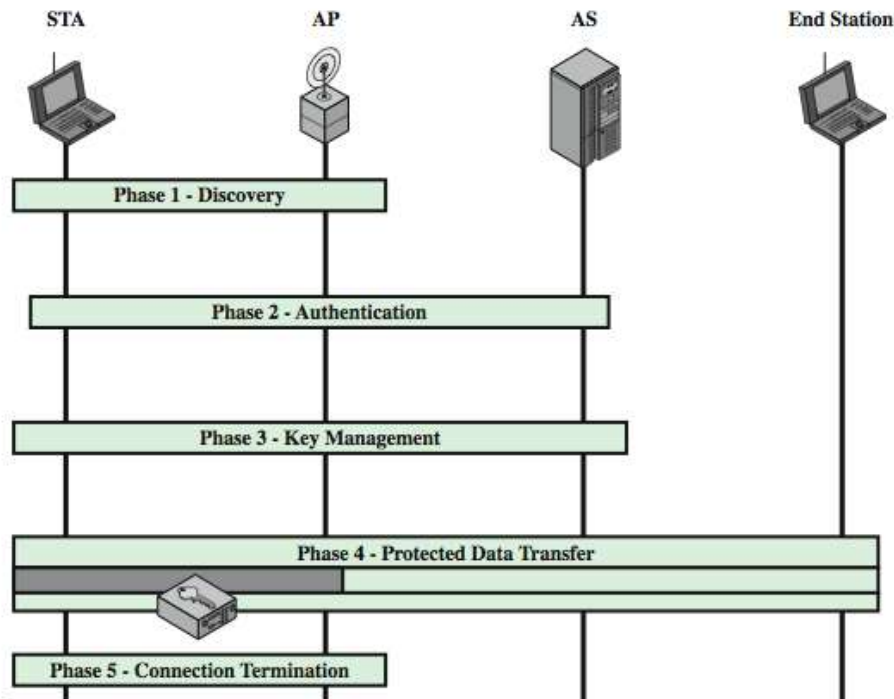# WiFi Network Components & Architecture

Distribution System
- ESS (Extended Service Set)
- AS (Authentication Server)

802.11i RSN Services and Protocols

802.11i RSN Cryptographic Algorithms

802.11i Phases of Operation

STA AP AS End Station

Phase 1 - Discovery

Phase 2 - Authentication

Phase 3 - Key Management

Phase 4 - Protected Data Transfer

Phase 5 - Connection Termination

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
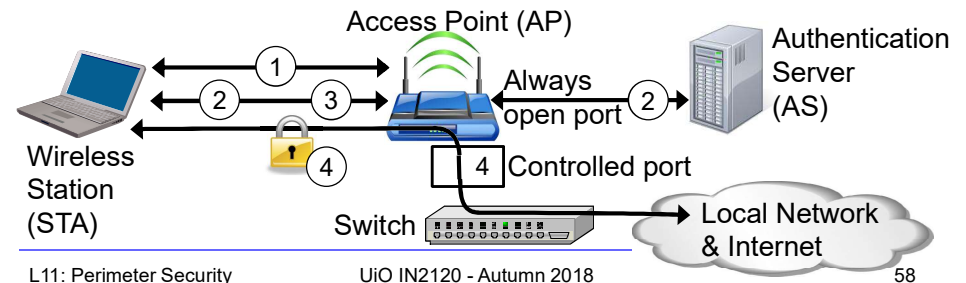
When you don't control the WLAN

- 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