BSc in Computer Engineering
CMP4103
Computer Systems and Network Security

Lecture 4

Physical Security and Access Control

Eng Diarmuid O'Briain, CEng, CISSP

Department of Electrical and Computer Engineering,
College of Engineering, Design, Art and Technology,
Makerere University

Table of Contents

1. Site & Facility Design

1.1 Location

- Emergency Services
- Hazards and threats
- Agency to services

1.2 Threats

- Fire
- Water and flooding
- Storms
- Vandalism
- Sabotage
- Explosions
- Building failure, collapse
- Utility failure and continuity
- Equipment failures
- Access
- Strikes

1.3 Secure Facility Plan

The planning process should begin by involving all stakeholders and posing two fundamental questions:

- What are we securing against?
- What levels of security do we need and are we willing to provide?

Once these questions are answered a list of possible threats should be drawn up.

The plan itself is developed using critical path analysis. With this process you systematically relate the company applications with all the possible threats to it. A Database Server will require, hardware, software, power, temperature control. We must now look critically at the dependencies for this server, what if the electricity goes down, what if the hardware overheats.

1.4 Physical Security Controls

When we consider physical security controls we can group them into the following groups:

- Physical
  - Walls
  - Fences
  - Gates
  - Locks
  - Lighting
  - Guards
  - Guard dogs
- Technical
  - Intrusion detection systems
  - Alarms
  - CCTV
  - Fire detection
  - Fire Suppression
- Administrative
  - Site Management
  - Personnel Access Controls
  - Security Training
  - Procedures in the event of security breaches

1.4.1 Server Rooms

Server rooms should be enclosed, restricted and protected rooms where mission critical equipment should be maintained in a temperature and humidity controlled environment. Halon type oxygen displacement fire detection and extinguishing systems should be available. Human access should be severely restricted to prevent unauthorised access as well as casual human access by employees who have no business there.

1.4.2 Work Areas

In as much as is possible work areas should be designed to prevent shoulder surfing. Shoulder surfing is the act of gathering information by watching someone's monitor and keyboard. The level of access an employee has should determine the work area they have. If they have high levels of access it is important that the proximity of their work area to lower level access employees does not allow for unauthorised access.

2. Physical Access Controls

2.1 Fence

This is usually the first line of defence. The following guidelines should be considered when establishing such a fence.

- 1 metre Deter casual trespassers
- 2 meters Hard to climb easily
- 2.5 meters Delay determined intruders

Another consideration is the planning laws in your locality. These may impact the type or look of the fence you plan.

A grass or gravel clearway along a fence should be considered to deter vehicles from parking near the fence. Bollards are a good method to deter such vehicles.

2.2 Access points

These points can be a weakness in the first layer of defence. By their nature gates provide access through the fence and therefore should be afforded the appropriate management.

2.3 Intrusion detection devices

- Photoelectric beams
- Ultrasonic
- Passive infrared
- Microwave
- Pressure sensitive pads

The use of intrusion detection systems can be mixed. They can either trigger audio or silent alarms or perhaps drown the area in light. One consideration however is the triggering of alarms by non intruders i.e. animals and birds.

2.4 Light

This is an very important consideration in your security plan. If it is dark it makes it easier for the intruder to access undetected. We also need to light areas to allow escape from the building in emergencies.

- Continuous Lighting
  - Fixed lights should be installed 2.5 metres above ground. The light on the ground from the lights should be at least 2 lumens
- Motion sensitive/trip lighting
  - Sensor activated light can be both a good security deterrent and a cost effective alternative to continuous lighting
- Standby lighting
  - Lights that come on in the event of power failure
- Exit lighting
  - Lights to indicate the exit points

2.5 Closed Circuit Television (CCTV)

CCTV is the use of video cameras to transmit a signal to a specific place, on a limited set of monitors. CCTV equipment may be used to observe parts of a process from a central control room; when, for example, the environment is not suitable for humans.

CCTV systems may operate continuously or only as required to monitor a particular event.

A more advanced form of CCTV, utilising Digital Video Recorders (DVR), provides recording for possibly many years, with a variety of quality and performance options and extra features like motion-detection and email alerts.

Points to consider when installing CCTV systems:

- The ability to **detect** an object
- The ability to **recognise** a detected object
- The ability to **identify** object details

2.6 Security Guards

A security guard is a privately and formally employed person who is paid to protect property, assets, and people.

Security officers are uniformed and act to protect property by maintaining a highly overt and visible presence to deter inappropriate access, observing for signs of crime, fire or disorder; then taking action and reporting any incidents to their client and emergency services as appropriate.

Generally the Security Guard will practice the "detect, deter, observe and report" methodology and call on the civil police when a situation is getting beyond their control.

Security officer's primary duty is the prevention and deterrence of crime. Security personnel enforce company rules and can act to protect lives and property. In fact, they frequently have a contractual obligation to provide these actions.

Security personnel may also perform access control at building entrances and vehicle gates, meaning, they ensure that employees and visitors display proper passes or identification before entering the facility. Security officers are often called upon to respond to minor emergencies (lost persons, lockouts, dead vehicle batteries, etc.) and to assist in serious emergencies by guiding emergency responders to the scene of the incident, helping to redirect foot traffic to safe locations, and by documenting what happened on an incident report.

2.7 Access Logs

Company: **Date:**

Name	Company	Name of person visiting	Security Guard	Time in	Time out

Access logs should be maintained either in paper form though more commonly in electronic form to record the comings and goings on non employees.

3. Perimeter Security

3.1 Doors

- Panels and glass should be protected against being kicked in or knocked out
- Install metal lining on exterior wooden doors to resist drilling or sawing
- Secure double doors with heavy duty, multiple-point, long flush bolts. Make sure the frame is as strong as the door
- All exterior doors should be constructed of steel, aluminium alloy, or solid-core hardwood, with minimum 1.5 mm steel on side and rear doors
- Door frames should be securely fixed to the walls
- Glass doors should have burglar-resistant glass installed
- Doors should be secured with a minimum of 3 hinges
- Doors should be clearly lit
- Emergency doors should be clearly marked
- Doors should provide entry and exit in the event of emergencies like power failure
- Doors should have the same fire rating as the walls

3.2 Locks

Exterior swinging doors should have a minimum 25 mm deadbolt lock, 25 mm throw bolt with a hardened insert, and free turning steel or brass tapered-cylinder guard. Steel strike plates should be used on aluminium door frames. All outside hinges should have non-removable hinge pins.

3.2.1 Electronic/Electrical Locks

An electronic or electric lock is a locking device which operates by means of electric current. Electric locks are sometimes stand-alone with an electronic control assembly mounted directly to the lock. More often electric locks are connected to an access control system. The advantages of an electric lock connected to an access control system include:

- Key control, where keys can be added and removed without re-keying the lock cylinder
- Fine access control, where time and place are factors
- Transaction logging, where activity is recorded

3.2.2 Authentication methods

Electronic locks offer a variety of means of authentication some are described below;

Numerical codes, passwords and passphrases
Perhaps the most prevalent form of electronic lock is that using a numerical code for authentication. The correct code must be entered in order for the lock to deactivate. Such locks typically provide a keypad. Combination lengths are usually between 4 and 6 digits long.

Security tokens
Another means of authenticating users is to require them to scan or swipe a security token such as a smart card or similar, or to interact a token with the lock.

Biometrics
As biometrics become more and more prominent as a recognised means of positive identification, their use in security systems increases. Some new electronic locks take advantage of technologies such as fingerprint scanning, retinal scanning and iris scanning, and voiceprint identification to authenticate users.

3.2.3 Padlocks

The most common assaults on padlocks are made with bolt cutters or pry bars. Quality padlocks should have the following features:

- Laminated or solid body case
- Hardened steel shackle with a minimum diameter of 8 mm
- A double locking mechanism providing "heel and toe" locking, and at least 5 pin tumblers in the cylinder

3.3 Turnstiles

A turnstile, also called a baffle gate, is a form of gate which allows one person to pass at a time. It can also be made so as to enforce one-way traffic of people, and in addition, it can restrict passage only to people who insert a security pass, or similar. Thus a turnstile can be used to restrict access to authorised people, for example in the lobby of an office building.

From a security standpoint, they lead patrons to enter single-file, so security personnel have a clear view of each patron. This enables security to efficiently isolate potential trouble or to confiscate any prohibited materials. Thus, turnstiles are a tool which leads to a more safe and secure atmosphere throughout a site.

3.4 Mantrap

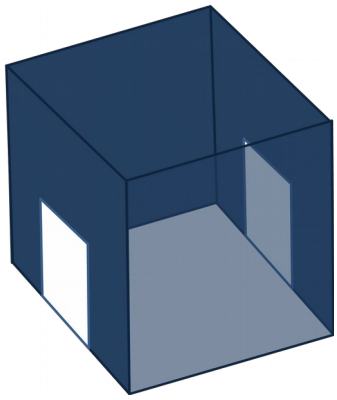

A mantrap refers to a small space having two sets of interlocking doors such that the first set of doors must close before the second set opens. ID may be required for each door, and possibly different measures for each door. For example, a key may open the first door, but a personal identification number entered on a number pad opens the second. Other methods of opening doors include proximity cards or biometric devices such as fingerprint readers or iris recognition scans.

Mantraps may be configured so that when an alarm is activated, all doors lock and trap the suspect between the doors in the "dead-space" or lock just one door to deny access to a secure space.

3.5 Windows

Windows should offer light, ventilation, and visibility, but not easy access. Locks should be designed so they cannot be reached and opened by breaking the glass. First floor windows should be protected with burglar-resistant glass, bars, grilles, grates, or heavy-duty wire screening to provide optimum window security.

3.5.1 Plate Glass

This is the most common type of glass found in windows. It is easy to get and cut for openings and for replacement. One problem is it tends to shatter in shards when broken or subject to an explosion. This presents a safety hazard.

3.5.2 Tempered Glass

This form of glass has been processed by controlled thermal or chemical treatments to increase its strength compared with normal glass. Tempered glass is made by processes which create balanced internal stresses which give the glass strength. It will usually shatter into small fragments instead of sharp shards when broken, making it less likely to cause severe injury and deep lacerations.

3.5.3 Polycarbonate Glass

This is not really glass but thermoplastic polymer moulded to look like glass. Fabrication is done with fine tooth saws. Polycarbonate is the toughest glazing available for windows. It very difficult to cut with a knife, but it is easily scratched, damaging the appearance.

4. Environment and Safety

4.1 Power

The maintenance of a secure and sustained power source is essential for technology businesses, data centres and technical laboratories.

4.1.1 Power problem terms

Fault – This is a momentary loss of power
Blackout – Complete loss of power
Sag – Lowering of the power supply voltage
Brownout – Prolonged period of low voltage
Spike – Momentary increase in voltage
Surge – Prolonged period of high voltage
Noise – A continuous power fluctuation
Transient – A short period of noise
Ground – Electrical earth
Clean – Continuous non fluctuating power
Inrush – Surge of voltage given initially after a device is connected to a power source

4.1.2 Uninterruptible Power Supply (UPS)

A UPS or sometimes called a battery backup, is an electrical device that provides emergency power when the input power source, typically the mains, fails. A UPS differs from an auxiliary or emergency power system or standby generator in that it will provide instantaneous or near-instantaneous protection from input power interruptions by means of one or more attached batteries and associated electronic circuitry. The on-battery runtime of most uninterruptible power sources is relatively short 5 – 15 minutes being typical for smaller units but sufficient to allow time to bring an auxiliary power source on line, or to properly shut down the protected equipment.

4.1.3 Electrical/Electronic Noise

In electronics and communication systems, noise is a random fluctuation or variation of an electromagnetic analogue signal such as a voltage or a current. Electronic noise is a characteristic of all electronic circuits. Depending on the circuit, the noise generated by electronic devices can vary greatly. Noise can be produced by several different effects. Thermal noise and shot noise are inherent to all devices. The other types depend mostly on manufacturing quality and semiconductor defects.

4.2 Water and Fire

4.2.1 Development of a Fire

A fire develops typically in four stages, and fire detectors are designed to detect some characteristic effect of one or more of these stages:

- Incipient stage
  - No visible smoke, no flame and very little heat
  - A significant amount of invisible (but sometimes detectable by smell) combustion particles may be created
  - This stage usually develops slowly
- Smouldering/smoke stage
  - Smoke, but no flame and little heat
- Flame stage
  - Visible flame, more heat, often less or no smoke, particularly with flammable liquids and gas fires
- Heat stage
  - Large amounts of heat, flame, smoke and toxic gases are produced.
  - The transition from the previous stage can be very fast

4.2.2 Fire triangle

The fire triangle is a simple model for understanding the ingredients necessary for most fires.

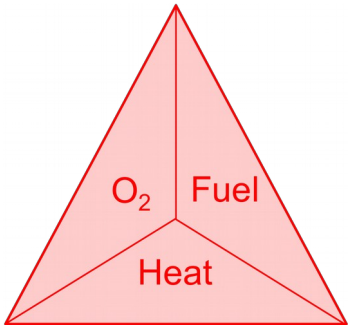

The triangle illustrates the rule that in order to ignite and burn, a fire requires three elements: heat, fuel, and an oxidising agent (usually oxygen). The fire is prevented or extinguished by removing any one of them. A fire naturally occurs when the elements are combined in the right mixture.

Without sufficient heat, a fire cannot begin, and it cannot continue. Without fuel, a fire will stop. Without sufficient oxygen, a fire cannot begin, and it cannot continue.

4.2.3 Fire Classes

European	American	Fuel/Heat source	Agents
Class A	Class A	Ordinary combustibles	Water
Class B	Class B	Flammable liquids	CO_2, Foam
Class C		Flammable gases	Dry chemical, Gas
Class D	Class D	Combustible metals	Dry Powders
Class E	Class C	Electrical equipment	CO_2, Foam, Gas
Class F	Class K	Cooking oil or fat	Wet chemical

Fires are identified according to one or more fire classes. Each class designates the fuel involved in the fire, and thus the most appropriate extinguishing agent. The classifications allow selection of extinguishing agents along lines of effectiveness at putting the type of fire out, as well as avoiding unwanted side effects.

4.2.4 Fire management in data centres and laboratories

With all the electronics in data centres and laboratories fire is a real risk. Typically water is used in fire control but used on electronic equipment will result in further damage to the equipment. For this reason Halon 1301 Gas was used in such environments. Unlike water, Halon 1301 didn't damage equipment however it did damage the ozone layer. The Montreal Protocol of 1987, limited the production of Halon 1301 to roles like aircraft emergency equipment where another alternative did not exist.

Halon 1301 was replaced by a number of extinguishing agents like Argon and Inergen when protecting data centres. They fall into two broad categories:

- Halocarbon gases
  - These work by removing heat from the fire
  - The room must be evacuated before the release of these agents
  - Lower storage space requirement compared to inert gasses
  - Fast fire suppression time (10 sec)
  - Must be very near point of use (max 30M)
  - More expensive than inert gasses
- Inert gases
  - These suppress fires by lowering the oxygen concentration in the room below the level needed to sustain combustion
  - Perform more effectively in rooms that aren't well sealed
  - More gas required than Halocarbon gasses to suppress an area
  - These can be piped long distances (100 – 200M) to a room and still retain their effectiveness

4.2.5 Pre-action Sprinklers

Pre-action sprinkler systems also are an option. The best choice for a particular facility will depend on the system's overall cost, the way in which the system will be used and the space available to house the extinguishing substance. The pipes in pre-action sprinkler systems do not hold water which reduces the risk of leaks that could damage computer or telecommunications equipment.

Instead, a valve within the system is located outside the data centre and keeps water from entering. In order for water to get past the valve, a smoke detector has to let the system know that a fire is occurring; at that point, water moves into the pipes. However, the fire has to grow to a certain temperature before the valve will open and water can discharge into the room. Given that these two events have to occur before water will flow through the pipes that are located within the data centre, the risk of an accidental leak is greatly reduced.

4.3 Water threat

Water damage is a threat in itself. Information systems and paper records can be badly damaged should they get wet particularly if saturated. Data centres and laboratories should be considered for water detection sensors that can trigger an alarm. Such rooms having raised floors to allow time for a water threat to be reacted to are common (though these are also used for conduits to carry room power and network cabling. Water threats are another reason to place such rooms above ground level.

4.4 Heating, Ventilating, and Air Conditioning (HVAC)

HVAC is the technology of indoor environmental comfort plus temperature and humidity control in data centres and laboratories. HVAC is particularly important in the design of medium to large industrial and office buildings such as skyscrapers and in marine environments such as aquariums, where safe and healthy building conditions are regulated with temperature and humidity, as well as "fresh air" from outdoors.

4.4.1 Positive Pressure

By applying greater air pressure in the room or building than is outside which ensures that should there be any leakage it will be out and thus prevent any unwanted air in. Monitoring of air pressure in a controlled room is a method that can be applied to the alarm system. Should the pressure change suddenly it is an indication of the possibility of unauthorised access.

5. Access Control

Access Controls protect confidentiality, integrity and availability of objects [1].

5.1 Preventive access control

Stop unwanted or unauthorised activity from occurring. These include biometrics, fences and locks, data classification, job rotation and separation of duties plus auditing, cryptography and monitoring.

5.2 Deterrent access control

Discourages the violation of security policies, often filling the gap left by preventive controls. They include gates, keyed access and security guards, badges, cameras and intrusion alarms or awareness training, separation of duties and security clearances.

5.3 Detective access control

Discovers unwanted or unauthorised activity but often take effect after an incident has occurred as opposed to before or during its occurrence. Security patrols, Security badges, guard dogs, security cameras, motion detectors and sound alarms or incident investigations, supervisory review, audits, and violation or exception reports.

5.4 Corrective access control

These restore systems to a known-good state following a security-related breach or incident. Access termination, service restarting or system rebooting, the implementation of intrusion detection systems, antivirus programs and malware scanners or business continuity planning, disaster recovery planning and security policies are typical corrective access controls.

5.5 Recovery access control

These are used to repair and restore critically damaged capabilities, functions and resources following a security violation. These are more complex in scale and scope than corrective controls and include backups, rollbacks and restorations, fault-tolerance, redundancy and clustering, or antivirus scanners, database shadowing and data replication.

5.6 Compensation access control

This provides aid to various other existing controls in the enforcement of system-wide security policies. This includes security policies, operational requirements or utilisation criteria or personnel supervision, monitoring and work procedures.

[1] ***Object*** *– Computer/system that may be accessed.*

5.7 Directive access control

This confines and controls the actions of subjects [2] to enforce and encourage strict security policy compliance. Security guards, guard dogs and security cameras, policy requirements, security criteria and posted notifications, or escape routes, employee supervision and awareness training all come under the umbrella of directive access control.

5.8 Administrative access control

This is the defined policies and procedures of the organisation that governs overall access, focusing on personnel and business practices. Workplace policies, procedures and hiring practices, background checks, data classification and security training or work reviews, employee supervision and personnel controls are examples.

5.9 Logical or technical access control

These are hardware and software mechanisms that manage access to and provide protection for shared computer and network resources. Encryption, passwords and smartcards, access control lists, biometrics and constrained interfaces or cryptographic protocols, firewall appliances and network routers.

5.10 Physical access control

These are structural barriers emplaced to prevent direct access to components of a facility, network or system. Security guards, guard dogs and fences, alarm systems, motion detectors and security windows or security lights, security locks and video cameras.

[2] ***Subject*** *– He/she who may try and access the object.*

6. Access Control in a Layered Environment

6.1 Layered / Defence in depth

- This means the use of several forms of access control to provide greater security.

6.2 Identification

- This is the processes and procedures through which a subject is proven to be authentic, accredited with permissions and held accountable for individual actions and activities. This is typically achieved with the use of login names, identity cards.

6.3 Authentication

- This is the process of verifying that a given identity is valid.
    - Type 1 – “Something you know”
        - Password
    - Type 2 – “Something you have”
        - Token
    - Type 3 – “Something you are”
        - Biometric
    - “Something you do”
    - “Somewhere you are”
    - Multi-factor Authentication

6.4 Authorisation

- This is the process of determining the types and extent of activities that are permissible to established users or groups on a protected system.

6.5 Auditing and Accountability

- This is the process of formally examining and reviewing activities, applications and processes initiated by subjects on a system. This systematic monitoring of activity can be used to detect malicious activity and investigate non-compliance issues.

7. Identification and Authentication techniques

7.1 Identification

- Subject must provide an identity to a system to start the Authentication, Authorisation and Accountability process. The Identity correlates an authentication factor with a subject:
  - Typing a username
  - Swiping a Smart Card
  - Waving a Token Device
  - Speaking a Phrase
  - Positioning Face, Hand or Finger for a camera or scanning device

7.2 Authentication

Authentication verifies the Identity of a Subject, thus Identification and Authentication are always a two step process, one useless without the other.

8. Passwords

This is the most common authentication technique. Passwords are poor security mechanisms for the following reasons:

- Users typically use passwords they can easily remember
- Random generated passwords are difficult to remember so the Subject tends to write them down
- Passwords are easily shared, written down, forgotten
- Passwords are easily stolen through observation, recording, playback, social engineering and security database theft
- Passwords often transmitted in clear or shrouded in simple to break encryption
- Short passwords can be discovered quickly by brute force attacks.

8.1 Password Selection

Passwords are broken into two groups:

- Static
  - Always remain the same
- Dynamic
  - One-time passwords, single-use passwords
  - Cognitive password
  - What is your date of birth?
  - What is your first pet's name?
  - What is your mother's maiden name?

Password policies should at a minimum force the subject to:

- Change the password regularly, minimum and maximum age
- Password characters should be dictated by the object during creation.
- Not all letters
- No number or letter sequences
- Does not contain the Identification name
- Minimum length
- Mix of letters and numbers, upper and lower case
- No password reuse

8.2 Password Security

Password theft methods include:

- Network Traffic Analysis
- Password file access
- Brute-force attacks
- Dictionary attacks
- Social Engineering

9. Biometrics

Biometrics refers to methods for uniquely recognising humans based upon one or more intrinsic physical or behavioural traits. For IT Security this typically falls into the following types:

- Fingerprints
- Face scans
- Iris Scans
  - Coloured area around pupil
- Retina scans
  - Pattern of blood vessels in back of eye
  - Most unacceptable by subjects as it can determine medical conditions in Subject (pregnancy, blood pressure) and it also blows air into the subjects eye
- Palm scans (Palm Topography)
- Hand Geometry
- Signature dynamics
  - Recognition of how a subject signs a set of characters
- Keystroke patterns (keystroke dynamics)
  - Flight time
  - Dwell time

9.1 Biometric Factor Ratings

- Errors
  - Type 1
    - Valid subject is not authenticated
    - False Rejection Rate (FRR)
      - The probability that the system fails to detects a match between the input pattern and a matching template in the database
    - It measures the percent of valid inputs which are incorrectly rejected
  - Type 2
    - Invalid subject authenticated
    - False Acceptance Rate (FAR)
      - The probability that the system incorrectly matches the input pattern to a non-matching template in the database
      - It measures the percent of invalid inputs which are incorrectly accepted

Point of intersection between FRR and FAR is known as the Crossover Error Rate (CER), it is the rate at which both accept and reject errors are equal. The lower the CER rate the more accurate is the system.

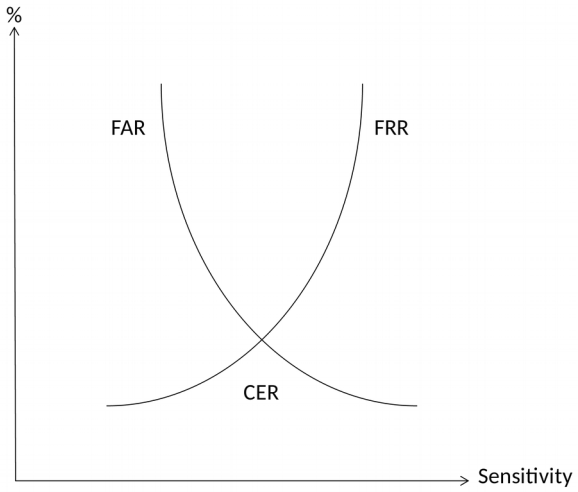

9.2 Biometric Registration

Biometric information for a subject is stored in a reference profile or template. Over time however the characteristics stored may not match the changes in the Subject due to ageing. Once subjects are enrolled refreshment of the biometric reference will need to be performed at regular intervals.

9.3 Biometric usage, acceptance and cost

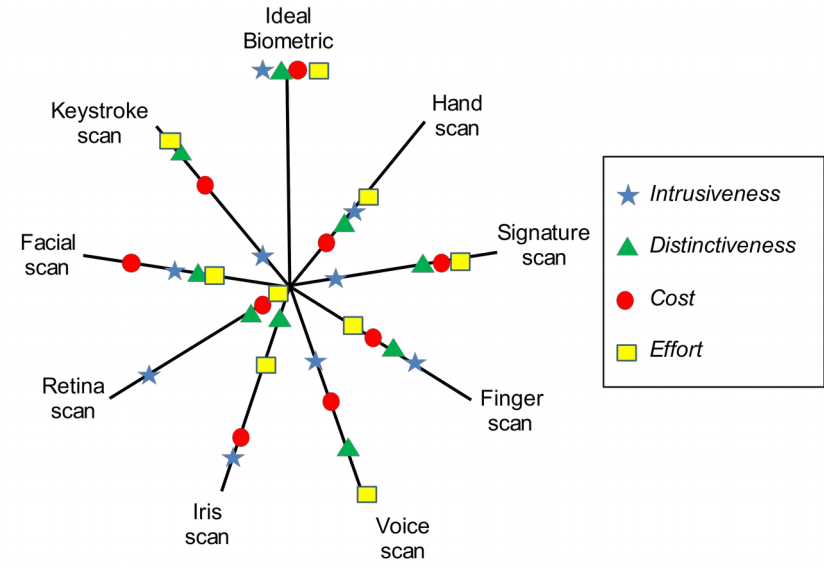

This Zephyr analysis chart shows the relation between ideal biometrics and most popular biometric technologies: iris, retina, voice, face, fingerprint, hand, keystroke and signature.

10. Tokens

A security token (or sometimes a hardware token, hard token, authentication token, USB token, cryptographic token, or key fob) may be a physical device that an authorised user of computer services is given to ease authentication. The term may also refer to software tokens.

There are four types of token:

- Static Tokens
  - Swipe card, disk, USB RAM Key or a physical key
- Synchronous dynamic password tokens
  - Device that generates new passwords at fixed time intervals
  - Subject enters generated password with PIN and passphrase/password
- Asynchronous dynamic password tokens
  - Device that generates new passwords on the occurrence of an event
  - Press a key on the token and the server for example, advances next password
  - Subject enters generated password with PIN and passphrase/password
- Challenge-response tokens
  - Passwords are generated by the token in response to instructions from the object

11. Tickets

Single Sign On (SSO), this means is a mechanism where multiple applications use one place to authenticate. From a user's point of view this means that he or she does not have to log into every single application when he or she moves between applications. A very common example of this will be Google, a single login permits access to Gmail, Google Calendar and other Google applications. Google uses Security Assertion Markup Language (SAML) Single Sign-On (SSO) service.

11.1 Needham-Schroeder Symmetric Key Protocol

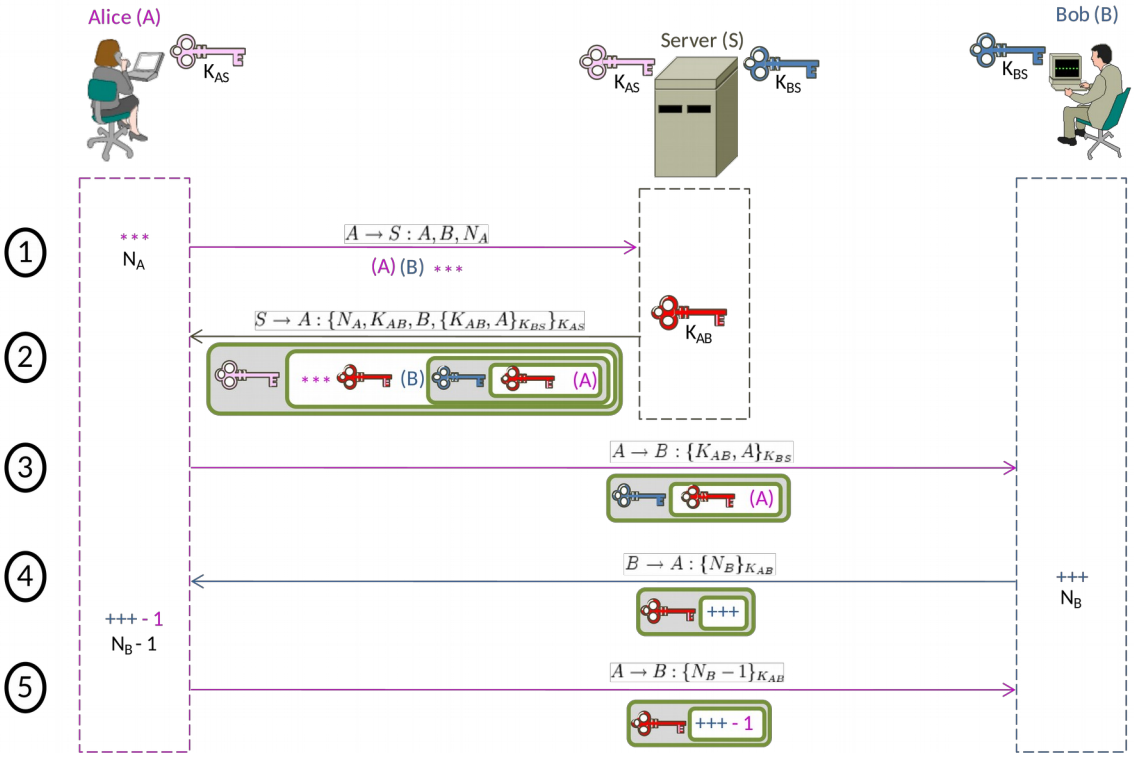

The Needham-Schroeder Symmetric Key Protocol is based on a symmetric encryption algorithm. It forms the basis for the Kerberos protocol. This protocol aims to establish a session key between two parties on a network, typically to protect further communication.

Here, Alice (A) initiates the communication to Bob (B). S is a server trusted by both parties. In the communication:

- A and B are identities of Alice and Bob respectively
- K_{AS} is a symmetric key known only to A and S
- K_{BS} is a symmetric key known only to B and S
- N_A and N_B are nonces generated by A and B respectively
- K_{AB} is a symmetric, generated key, which will be the session key of the session between A and B

Referring to the diagram:

- Alice sends a message to the server identifying herself and Bob, telling the server she wants to communicate with Bob.
- The server generates K_{AB} and sends back to Alice a copy encrypted under K_{BS} for Alice to forward to Bob and also a copy for Alice. Since Alice may be requesting keys for several different people, the nonce assures Alice that the message is fresh and that the server is replying to that particular message and the inclusion of Bob's name tells Alice who she is to share this key with.
- Alice forwards the key to Bob who can decrypt it with the key he shares with the server, thus authenticating the data.
- Bob sends Alice a nonce encrypted under K_{AB} to show that he has the key.
- Alice performs a simple operation on the nonce, re-encrypts it and sends it back verifying that she is still alive and that she holds the key.

11.2 Kerberos

Kerberos is a computer network authentication protocol, which allows nodes communicating over a non-secure network to prove their identity to one another in a secure manner.

Kerberos builds on symmetric key cryptography and requires a trusted third party. Extensions to Kerberos can provide for the use of public-key cryptography during certain phases of authentication.

Kerberos uses as its basis the symmetric Needham-Schroeder protocol. It makes use of a trusted third party, termed a Key Distribution Centre (KDC), which consists of two logically separate parts:

- Authentication Server (AS)
- Ticket Granting Server (TGS)

Kerberos works on the basis of "tickets" which serve to prove the identity of users.

The KDC maintains a database of secret keys, each entity on the network, whether a client or a server, shares a secret key known only to itself and to the KDC. Knowledge of this key serves to prove an entity's identity. For communication between two entities, the KDC generates a session key which they can use to secure their interactions. The security of the protocol relies heavily on participants maintaining loosely synchronised time and on short-lived assertions of authenticity called Kerberos tickets.

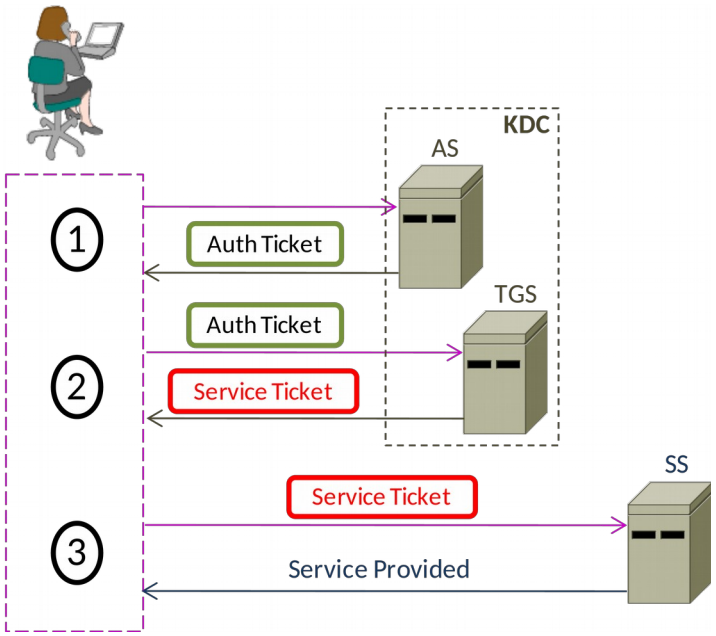

1) The client authenticates itself to the Authentication Server (AS) and receives a time stamped ticket.
2) It then contacts the Ticket Granting Server (TGS), and using the ticket it demonstrates its identity and asks for a service. If the client is eligible for the service, then the TGS sends another ticket to client.
3) The client then contacts the Service Server (SS), and using this ticket it proves that it has been approved to receive the service.

11.2.1 Kerberos Drawbacks

- Single point of failure
  - It requires continuous availability of a central server. When the Kerberos server is down, no one can log in. This can be mitigated by using multiple Kerberos servers and fallback authentication mechanisms
- Kerberos requires the clocks of the involved hosts to be synchronised
  - The tickets have a time availability period and if the host clock is not synchronised with the Kerberos server clock, the authentication will fail
  - The default configuration requires that clock times are no more than 10 minutes apart. In practice Network Time Protocol daemons are usually used to keep the host clocks synchronised
- The administration protocol is not standardised and differs between server implementations
  - Password changes are described in RFC 3244
- Since the secret keys for all users are stored on the central server, a compromise of that server will compromise all users' secret keys
- A compromised client will compromise the user password

12. Access Control Techniques

12.1 Discretionary Access Controls (DAC)

This is a kind of access control defined by the Trusted Computer System Evaluation Criteria (TCSEC) as "*a means of restricting access to objects based on the identity of subjects and/or groups to which they belong*". The controls are discretionary in the sense that a subject with certain access permission is capable of passing that permission (perhaps indirectly) on to any other subject (unless restrained by mandatory access control)".

- In other words access control is at the discretion of the owner
- DACs are often implemented using Access Control Lists (ACL) on objects
- DACs do not offer centralised management because owners can alter ACLs on their objects

12.2 Non-discretionary Access Controls

These are rules based systems in which a set of rules, restrictions or filters determine what can and cannot occur on the system, like granting subject access, performing an action on an object, or accessing a resource.

12.2.1 Mandatory Access Controls (MAC)

- The TCSEC, defines MAC as "*a means of restricting access to objects based on the sensitivity (as represented by a label) of the information contained in the objects and the formal authorisation (i.e., clearance) of subjects to access information of such sensitivity*".
- MAC has been closely associated with multi-level secure (MLS) systems. These are application of a computer system to process information with different sensitivities (i.e., at different security levels – Top Secret, Secret, Confidential, Restricted/Sensitive but Unclassified (SBU) and Unclassified), permit simultaneous access by users with different security clearances and needs-to-know, and prevent users from obtaining access to information for which they lack authorisation.
- Early implementations of MAC such as Honeywell's SCOMP, USAF SACDIN, NSA Blacker, and Boeing's MLS LAN focused on MLS to protect military-oriented security classification levels with robust enforcement.

12.2.2 Role-based Access Control (RBAC)

- Within an organisation, roles are created for various job functions. The permissions to perform certain operations are assigned to specific roles. Members of staff (or other system users) are assigned particular roles, and through those role assignments acquire the permissions to perform particular system functions.
- Since users are not assigned permissions directly, but only acquire them through their role (or roles), management of individual user rights becomes a matter of simply assigning appropriate roles to the user. This simplifies common operations, such as adding a user, or changing a user's department.
- RBAC differs from access control lists (ACLs) used in traditional discretionary access control systems in that it assigns permissions to specific operations with meaning in the organisation, rather than to low level data objects.
    - For example, an access control list could be used to grant or deny write access to a particular system file, but it would not dictate how that file could be changed.
    - In an RBAC-based system, an operation might be to create a 'credit account' transaction in a financial application or to populate a 'blood sugar level test' record in a medical application. The assignment of permission to perform a particular operation is meaningful, because the operations are granular with meaning within the application.
- RBAC is attractive to organisations with a high rate of turnover.

12.2.3 Lattice-based Access Control (LBAC)

- LBAC is a complex access control based on the interaction between any combination of objects (such as resources, computers, and applications) and subjects (such as individuals, groups or organisations).
- In this type of label-based mandatory access control model, a lattice is used to define the levels of security that an object may have and that a subject may have access to. The subject is only allowed to access an object if the security level of the subject is greater than or equal to that of the object.
- Mathematically, the security level access may also be expressed in terms of the lattice (a partial order set) where each object and subject have a greatest lower bound (meet) and least upper bound (join) of access rights. For example, if two subjects A and B need access to an object, the security level is defined as the meet of the levels of A and B. In another example, if two objects X and Y are combined, they form another object Z, which is assigned the security level formed by the join of the levels of X and Y.

13. Access Control methodologies and Implementation

13.1 Centralised Access Control

13.1.1 Advantages

- Managed by small team or individual
- Administrative overhead is low
- Single changes impact the complete system

13.1.2 Disadvantages

- Single point of failure
- If elements cannot access centralised access control system then subjects cannot access objects

13.1.3 Remote Access Dial-in User Service (RADIUS)

- This is a networking protocol that provides centralised Authentication, Authorisation, and Accounting (AAA) management for computers to connect and use a network service
- Developed by Livingston Enterprises, Inc., in 1991 as an access server authentication and accounting protocol and later brought into the IETF standards
- Because of the broad support and the ubiquitous nature of the RADIUS protocol it is often used by ISPs and enterprises to manage access to the Internet or internal networks, wireless networks, and integrated e-mail services
- RADIUS is a client/server protocol that runs in the application layer, using UDP as transport
- RADIUS serves three functions:
    - it authenticates users or devices before granting them access to a network
    - it authorises those users or devices for certain network services
    - it accounts for usage of those services

13.1.4 Terminal Access Controller Access Control System (TACACS)

- Another remote authentication protocol that is used to communicate with an authentication server commonly used in UNIX networks
- It uses TCP for transport
- TACACS+
    - TACACS+ is based on TACACS, but, in spite of its name, it is an entirely new protocol which is incompatible with any previous version of TACACS
    - Whereas RADIUS combines authentication and authorisation in a user profile, TACACS+ separates the two operations

13.1.5 Diameter

- It is a successor to RADIUS however a lot of the features of Diameter have been included in upgrades of RADIUS
- Uses Reliable transport protocols TCP or SCTP instead of UDP

13.2 Decentralised Access Control

- Advantages
  - No single point of failure
- Disadvantages
  - Large administrative overhead
  - Maintaining homogeneity becomes difficult

A domain is a realm of trust created were a collection of subjects and objects share a common security policy. Between these domains a security bridge called a trust can be established to allow subjects in one to access objects in the other.

14. Access Control Administration

14.1 Responsibilities

- User Account Management
- Activity Tracking
- Access rights and permission management

14.2 User Accounts

- User
  - Subject that has access to Objects to perform some action or work task
- Owner
  - Subject with the final responsibility for classification and labelling an Object
- Custodian
  - Subject that has been assigned responsibility of properly storing and protecting Objects

14.2.1 Enrolment

The function of creating and amending user accounts should be protected and secured through organisation security policies. The initial creation function is called enrolment.

User Accounts cannot be created without HR department request on new-hire or promotion.

- Formal request from HR department
  - User details
  - Security classification
- Users manager and security manager should verify and approve the assignment
- User should be trained on the organisations security policies
- User should sign a document agreeing to comply with the policies
- Document could have language that causes the employee to be subject to disciplinary action or dismissal should they knowingly breech the policy

14.2.2 Account Maintenance

Accounts will require maintenance through the life of the user, promotions, job changes will require amendments to the account. Again such changes should have an associated approval from the user's manager and security manager.

When employee's leave the organisation it is critical that the accounts are deleted, deleted or revoked.

14.2.3 Account, Log and Journal Monitoring

User Accounts, event logs and system journals help piece together the state of affairs for a server at any referenced point along the timeline of its operation. They capture events, changes, messages and other data that build up a history of activity on a system.

14.2.4 Access rights and permissions

It is important that subjects are only given access rights and permissions to those parts of an object that is necessary for them to perform their job function.

14.2.5 Principle of Least Privilege

This principle refers to the concept that all Subjects at all times should run with as few privileges as possible, and also launch applications with as few privileges as possible. Subjects should be restricted from accessing objects that they do not need access to in the course of their work.

14.2.6 Creeping Privileges

This is the idea that a long serving user accumulates privileges over time as roles change. This will eventually result in the user gaining excessive privileges.

14.2.7 Separation of duties (SoD)

This is the concept of having more than one person required to complete a task. It is alternatively called segregation of duties or, in the political realm, separation of powers.

Separation of Duties Control Matrix

	Control Group	Systems Analyst	Application Programmer	Help Desk and Support Manager	End User	Data Entry	Computer Operator	Database Administrator	Network Administrator	Systems Administrator	Security Administrator	Systems Programmer	Quality Assurance
Control Group		X	X	X		X	X	X	X	X		X	
Systems Analyst	X			X	X		X				X		X
Application Programmer	X			X	X	X	X	X	X	X	X	X	X
Help Desk and Support Manager	X	X	X		X	X		X	X	X		X	
End User		X	X	X			X	X	X			X	X
Data Entry	X		X	X			X	X	X	X	X	X	
Computer Operator	X	X	X		X	X		X	X	X	X	X	
Database Administrator	X		X	X	X	X	X		X	X		X	
Network Administrator	X		X	X	X	X	X	X					
System Administrator	X		X	X		X	X	X				X	
Security Administrator		X	X			X	X					X	
Systems Programmer	X		X	X	X	X	X	X		X	X		X
Quality Assurance		X	X		X							X	

X – Combining these functions may create a potential control weakness

15. Monitoring

This refers to a program applied to the activities of subjects by which they are held accountable for their actions while authenticated on a system. Monitoring allows management to detect malicious actions by Subjects while also monitoring attempts at intrusion or simply object system failures.

Tools used to extract useful information from log files are called Data Mining tools. Such data mining can be done in almost real time for intrusion detection and such is called a Intrusion Detection System (IDS).

15.1 Intrusion Detection System (IDS)

An Intrusion detection system (IDS) is software and/or hardware designed to detect unwanted attempts at accessing, manipulating, and/or disabling computer systems, mainly through a network, such as the Internet. These attempts may take the form of attacks, as examples, by crackers, malware and/or disgruntled employees. An IDS cannot directly detect attacks within properly encrypted traffic.

An intrusion detection system is used to detect several types of malicious behaviours that can compromise the security and trust of a computer system. This includes network attacks against vulnerable services, data driven attacks on applications, host based attacks such as privilege escalation, unauthorised logins and access to sensitive files, and malware (viruses, trojan horses, and worms).

An IDS can be composed of several components:

1. Sensors which generate security events
2. Console to monitor events and alerts and control the sensors
3. Central Engine that records events logged by the sensors in a database and uses a system of rules to generate alerts from security events received

There are several ways to categorise an IDS depending on the type and location of the sensors and the methodology used by the engine to generate alerts. In many simple IDS implementations all three components are combined in a single device or appliance.

15.1.1 Host-based IDS (HIDS)

A HIDS monitors all or parts of the dynamic behaviour and the state of a computer system. Much as a NIDS will dynamically inspect network packets, a HIDS might detect which program accesses what resources and discover that, for example, a word-processor has suddenly and inexplicably started modifying the system password database. Similarly a HIDS might look at the state of a system, its stored information, whether in RAM, in the file system, log files or elsewhere; and check that the contents of these appear as expected.

One can think of a HIDS as an agent that monitors whether anything or anyone, whether internal or external, has circumvented the system's security policy.

15.1.2 Network-based IDS (NIDS)

A NIDS reads all the incoming packets and tries to find suspicious patterns known as signatures or rules. If, for example, a large number of TCP connection requests to a very large number of different ports are observed, one could assume that there is someone conducting a port scan of some or all of the computer(s) in the network. It also (mostly) tries to detect incoming shellcodes[1] in the same manner that an ordinary intrusion detection system does.

A NIDS is not limited to inspecting incoming network traffic only. Often valuable information about an ongoing intrusion can be learned from outgoing or local traffic as well. Some attacks might even be staged from the inside of the monitored network or network segment, and are therefore not regarded as incoming traffic at all.

15.1.3 Knowledge (Signature) Based Detection

Network traffic is examined for preconfigured and predetermined attack patterns known as signatures. Many attacks today have distinct signatures. In good security practice, a collection of these signatures must be constantly updated to mitigate emerging threats.

15.1.4 Behaviour (Statistical anomaly) Based Detection

A behaviour based IDS establishes a performance baseline based on normal network traffic evaluations. It will then sample current network traffic activity to this baseline in order to detect whether or not it is within baseline parameters. If the sampled traffic is outside baseline parameters an alarm will be triggered.

[1] a shellcode is a small piece of code used as the payload in the exploitation of a software vulnerability. It is called "shellcode" because it typically starts a command shell from which the attacker can control the compromised machine.

16. IDS Related Tools

16.1 Honey Pot

A honey pot is a trap set to detect, deflect, or in some manner counteract attempts at unauthorised use of information systems. Generally it consists of a computer, data, or a network site that appears to be part of a network, but is actually isolated, (un)protected, and monitored, and which seems to contain information or a resource of value to attackers.

A honey pot that masquerades as an open proxy to monitor and record those using the system is a sugarcane. Honey pots should have no production value, and hence should not see any legitimate traffic or activity. Whatever they capture can then be surmised as malicious or unauthorised. One practical implication of this is honey pots that thwart spam by masquerading as the type of systems abused by spammers. They categorise trapped material 100% accurately.

16.1.1 Enticement

A Honey Pot placed with open security vulnerabilities and services with known exploits is enticement. The opportunity for illegal or unauthorised activity is provided but the perpetrator makes his/her own decision to perform the exploit.

16.1.2 Entrapment

This is where the honey pot actively solicits subjects to access it and then the owner charges them with unauthorised intrusion. This type of activity is illegal.

16.1.3 Vulnerability scanner

A vulnerability scanner is a program designed to search for and map systems for weaknesses in an application, computer or network.

1) The scanner will first look for active IP addresses, open ports, OS's and any applications running.
2) It may at this point create a report or move to the next step.
3) Try to determine the patch level of the OS or applications. In this process the scanner can cause an exploit of the vulnerability such as crash the OS or application.
4) The final phase the scanner may attempt to exploit the vulnerability. Scanners may either be malicious or friendly. Friendly scanners usually stop at step 2 and occasionally step 3 but never go to step 4.

16.1.4 Types of vulnerability scanners

- Port Scanner
  - Application designed to probe a network host for open ports
- Network Scanner
  - Application used to retrieve user names, and info on groups, shares and services of networked computers. This type of program scans networks for vulnerabilities in the security of that network
- Web Application Security Scanner
  - Application that communicates with a web application through the web front-end in order to identify potential security vulnerabilities in the web application and architectural weaknesses
- Computer worm
  - A computer worm is a self-replicating computer program. It uses a network to send copies of itself to other nodes (computers on the network) and it may do so without any user intervention. Unlike a virus, it does not need to attach itself to an existing program. Worms almost always cause at least some harm to the network, if only by consuming bandwidth, whereas viruses almost always corrupt or devour files on a targeted computer

16.2 Penetration Testing

Penetration testing is a method of evaluating the security of a computer system or network by simulating an attack from a malicious source, known as a Black Hat Hacker, or Cracker.

The process involves an active analysis of the system for any potential vulnerabilities that may result from poor or improper system configuration, known and/or unknown hardware or software flaws, or operational weaknesses in process or technical countermeasures. This analysis is carried out from the position of a potential attacker, and can involve active exploitation of security vulnerabilities. Any security issues that are found will be presented to the system owner together with an assessment of their impact and often with a proposal for mitigation or a technical solution. The intent of a penetration test is to determine feasibility of an attack and the amount of business impact of a successful exploit, if discovered. It is a component of a full security audit.

16.3 Padded cell

A padded cell is like a honey pot but is used for intruder isolation. When the IDS detects an intruder he/she is transferred to the padded cell. The padded cell has the look of an actual system but with fake programs and data, a simulated environment of sorts.

17. Methods of Attack

17.1 Brute force and dictionary attacks

Used to find or break passwords for user accounts.

17.1.1 Brute Force Attack

A Brute force attack is a crude form of attack where every possible password is tried. This method is unlikely to be practical unless the password is relatively short. It also depends on there being no limit to the number of attempts.

17.1.2 Dictionary Attack

A dictionary attack is a technique for guessing passwords by trying to determine by searching likely possibilities. This is made possible as users often choose weak passwords or simply leave the default password. Users often choose weak passwords like:

- "password", "passcode", "admin" and their derivatives
- a row of letters from the qwerty keyboard -- qwerty itself, asdf, or qwertyuiop)
- the user's name or login name
- the name of their significant other, a friend, relative or pet
- their birthplace or date of birth, or a friend's, or a relative's
- their car license plate number, or a friend's, or a relative's
- their office number, residence number or most commonly, their mobile number
- a name of a celebrity they like
- a simple modification of one of the preceding, such as suffixing a digit, particularly 1, or reversing the order of the letters
- a swear word

Dictionary attack programs can usually determine passwords where it is a single word found in dictionaries, given and family names, any too short password (usually thought to be 6 or 7 characters or less), or any password meeting a too restrictive and so predictable, pattern (eg, alternating vowels and consonants). Repeated research over some 40 years has demonstrated that around 40% of user-chosen passwords are readily guessable by sophisticated cracking programs armed with dictionaries and, perhaps, the user's personal information.

17.2 Denial of Service (DoS) attacks

These are an attempt to make a computer resource unavailable to its intended users. Although the means to carry out, motives for, and targets of a DoS attack may vary, it generally consists of the concerted efforts of a person or people to prevent an Internet site or service from functioning efficiently or at all, temporarily or indefinitely. Perpetrators of DoS attacks typically target sites or services hosted on high-profile web servers such as banks, credit card payment gateways, and even root name-servers.

One common method of attack involves saturating the target (victim) machine with external communications requests, such that it cannot respond to legitimate traffic, or responds so slowly as to be rendered effectively unavailable. In general terms, DoS attacks are implemented by either forcing the targeted computer(s) to reset, or consuming its resources so that it can no longer provide its intended service or obstructing the communication media between the intended users and the victim so that they can no longer communicate adequately.

Denial-of-service attacks are considered violations of the IAB's Internet proper use policy, and also violate the acceptable use policies of virtually all Internet Service Providers. They also commonly constitute violations of the laws of individual nations.

17.2.1 SYN Flood attacks

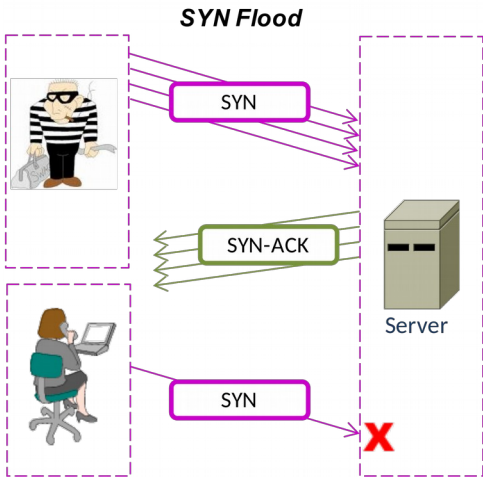

A SYN flood is a form of DOS attack in which an attacker sends a succession of SYN requests to a target's system.

When a client attempts to start a TCP connection to a server, the client and server exchange a series of messages which normally runs like this:

- The client requests a connection by sending a SYN (synchronise) message to the server
- The server acknowledges this request by sending SYN-ACK back to the client
- The client responds with an ACK, and the connection is established

This is called the TCP three-way handshake, and is the foundation for every connection established using the TCP protocol.

The SYN flood is a well known type of attack and is generally not effective against modern networks. It works if a server allocates resources after receiving a SYN, but before it has received the ACK.

There are two methods, but both involve the server not receiving the ACK. A malicious client can skip sending this last ACK message. Or by spoofing the source IP address in the SYN, it makes the server send the SYN-ACK to the falsified IP address, and thus never receive the ACK. In both cases the server will wait for the acknowledgement for some time, as simple network congestion could also be the cause of the missing ACK.

If these half-open connections bind resources on the server, it may be possible to take up all these resources by flooding the server with SYN messages. Once all resources set aside for half-open connections are reserved, no new connections (legitimate or not) can be made, resulting in denial of service. Some systems may malfunction badly or even crash if other operating system functions are starved of resources this way.

17.2.2 Smurf attack

The Smurf attack is a way of generating significant computer network traffic on a victim network. This is a type of DOS attack that floods a target system via spoofed broadcast ping messages.

In such an attack, a perpetrator sends a large amount of ICMP echo request (ping) traffic to IP broadcast addresses, all of which have a spoofed source IP address of the intended victim. If the routing device delivering traffic to those broadcast addresses delivers the IP broadcast to all hosts (for example via a layer 2 broadcast), most hosts on that IP network will take the ICMP echo request and reply to it with an echo reply, multiplying the traffic by the number of hosts responding. On a multi-access broadcast network, hundreds of machines might reply to each packet.

17.3 Spoofing

A spoofing attack is a situation in which one person or program successfully masquerades as another by falsifying data and thereby gaining an illegitimate advantage.

17.3.1 Man in the middle attacks

The man-in-the-middle attack, in which an attacker spoofs Alice into believing the attacker is Bob, and spoofs Bob into believing the attacker is Alice, thus gaining access to all messages in both directions without the trouble of any cryptanalytic effort.

The attacker must monitor the packets sent from Alice to Bob and then guess the sequence number of the packets. Then the attacker knocks out Alice with a SYN attack and injects his own packets, claiming to have the address of Alice. Alice's firewall can defend against some spoof attacks when it has been configured with knowledge of all the IP addresses connected to each of its interfaces. It can then detect a spoofed packet if it arrives at an interface that is not known to be connected to the IP address.

Many carelessly designed protocols are subject to spoof attacks, including many of those used on the Internet.

17.3.2 Spamming

Spam is the abuse of electronic messaging systems (including most broadcast media, digital delivery systems) to send unsolicited bulk messages indiscriminately. While the most widely recognised form of spam is e-mail spam, the term is applied to similar abuses in other media: instant messaging spam, Usenet newsgroup spam, Web search engine spam, spam in blogs, wiki spam, online classified ads spam, mobile phone messaging spam, Internet forum spam, junk fax transmissions, social networking spam, and file sharing network spam.

Spamming remains economically viable because advertisers have no operating costs beyond the management of their mailing lists, and it is difficult to hold senders accountable for their mass mailings. Because the barrier to entry is so low, spammers are numerous, and the volume of unsolicited mail has become very high. The costs, such as lost productivity and fraud, are borne by the public and by Internet service providers, which have been forced to add extra capacity to cope with the deluge. Spamming is widely reviled, and has been the subject of legislation in many jurisdictions.

17.3.3 Sniffers

A sniffer attack is a snooping activity carried out by malicious users gaining access to network traffic, gathering it and either extract login information of users or use the packets in a form of replay attack.