Fundamentals

- Methodologies
  - Stages
  - OSSTMM
  - OWASP
  - NIST Cybersecurity Framework 1.1
  - NCSC CAF
- Ethics
  - Rules of Engagement (ROE)
  - Hacker Categories
  - Scopes
- Principles of Security
  - CIA Triad
  - Principles of Privileges
  - Security Models
  - Threat Modelling & Incident Response

Methodologies

Stages

Penetration tests can have a wide variety of objectives and targets within scope. Because of this, no penetration test is the same, and there are no one-case fits all as to how a penetration tester should approach it.

The steps a penetration tester takes during an engagement is known as the methodology. A practical methodology is a smart one, where the steps taken are relevant to the situation at hand. For example, having a methodology that you would use to test the security of a web application is not practical when you have to test the security of a network.

Before discussing some different industry-standard methodologies, we should note that all of them have a general theme of the following stages:

Stage	Description
Information Gathering	This stage involves collecting as much publically accessible information about a target/organisation as possible, for example, OSINT and research. **Note**: This does not involve scanning any systems.
Enumeration/Scanning	This stage involves discovering applications and services running on the systems. For example, finding a web server that may be potentially vulnerable.
Exploitation	This stage involves leveraging vulnerabilities discovered on a system or application. This stage can involve the use of public exploits or exploiting application logic.
Privilege Escalation	Once you have successfully exploited a system or application (known as a foothold), this stage is the attempt to expand your access to a system. You can escalate horizontally and vertically, where horizontally is accessing another account of the same permission group (i.e. another user), whereas vertically is that of another permission group (i.e. an administrator).
Post-exploitation	This stage involves a few sub-stages: 1. What other hosts can be targeted (pivoting) 2. What additional information can we gather from the host now that we are a privileged user 3. Covering your tracks 4. Reporting

OSSTMM

The Open Source Security Testing Methodology Manual provides a detailed framework of testing strategies for systems, software, applications, communications and the human aspect of cybersecurity.

The methodology focuses primarily on how these systems, applications communicate, so it includes a methodology for:

1. **Telecommunications (phones, VoIP, etc.)**
2. Wired Networks
3. Wireless communications

Advantages	Disadvantages
Covers various testing strategies in-depth.	The framework is difficult to understand, very detailed, and tends to use unique definitions.
Includes testing strategies for specific targets (I.e. telecommunications and networking)	*Intentionally left blank.*
The framework is flexible depending upon the organisation's needs.	*Intentionally left blank.*
The framework is meant to set a standard for systems and applications, meaning that a universal methodology can be used in a penetration testing scenario.	*Intentionally left blank.*

OWASP

The "Open Web Application Security Project" framework is a community-driven and frequently updated framework used solely to test the security of web applications and services.

The foundation regularly writes reports stating the top ten security vulnerabilities a web application may have, the testing approach, and remediation.

Advantages	Disadvantages
Easy to pick up and understand.	It may not be clear what type of vulnerability a web application has (they can often overlap).
Actively maintained and is frequently updated.	OWASP does not make suggestions to any specific software development life cycles.
It covers all stages of an engagement: from testing to reporting and remediation.	The framework doesn't hold any accreditation such as CHECK.
Specialises in web applications and services.	*Intentionally left blank.*

NIST Cybersecurity Framework 1.1

The NIST Cybersecurity Framework is a popular framework used to improve an organisations cybersecurity standards and manage the risk of cyber threats. This framework is a bit of an honourable mention because of its popularity and detail.

The framework provides guidelines on security controls & benchmarks for success for organisations from critical infrastructure (power plants, etc.) all through to commercial. There is a limited section on a standard guideline for the methodology a penetration tester should take.

Advantages	Disadvantages
The NIST Framework is estimated to be used by 50% of American organisations by 2020.	NIST has many iterations of frameworks, so it may be difficult to decide which one applies to your organisation.
The framework is extremely detailed in setting standards to help organisations mitigate the threat posed by cyber threats.	The NIST framework has weak auditing policies, making it difficult to determine how a breach occurred.
The framework is very frequently updated.	The framework does not consider cloud computing, which is quickly becoming increasingly popular for organisations.
NIST provides accreditation for organisations that use this framework.	*Intentionally left blank.*
The NIST framework is designed to be implemented alongside other frameworks.	*Intentionally left blank.*

Methodologies

NCSC CAF

The Cyber Assessment Framework (CAF) is an extensive framework of fourteen principles used to assess the risk of various cyber threats and an organisation's defences against these.

The framework applies to organisations considered to perform "vitally important services and activities" such as critical infrastructure, banking, and the likes. The framework mainly focuses on and assesses the following topics:

- Data security
- System security
- Identity and access control
- Resiliency
- Monitoring
- Response and recovery planning

Advantages	Disadvantages
This framework is backed by a government cybersecurity agency.	The framework is still new in the industry, meaning that organisations haven't had much time to make the necessary changes to be suitable for it.
This framework provides accreditation.	The framework is based on principles and ideas and isn't as direct as having rules like some other frameworks.
This framework covers fourteen principles which range from security to response.	Intentionally left blank.

Ethics

Rules of Engagement (ROE)

The ROE is a document that is created at the initial stages of a penetration testing engagement. This document consists of three main sections (explained in the table below), which are ultimately responsible for deciding how the engagement is carried out. The SANS institute has a great example of this document which you can view online here.

Section	Description
Permission	This section of the document gives explicit permission for the engagement to be carried out. This permission is essential to legally protect individuals and organisations for the activities they carry out.
Test Scope	This section of the document will annotate specific targets to which the engagement should apply. For example, the penetration test may only apply to certain servers or applications but not the entire network.
Rules	The rules section will define exactly the techniques that are permitted during the engagement. For example, the rules may specifically state that techniques such as phishing attacks are prohibited, but MITM (Man-in-the-Middle) attacks are okay.

Hacker Categories

Hat Category	Description	Example
White Hat	These hackers are considered the "good people". They remain within the law and use their skills to benefit others.	For example, a penetration tester performing an authorised engagement on a company.
Grey Hat	These people use their skills to benefit others often; however, they do not respect/follow the law or ethical standards at all times.	For example, someone taking down a scamming site.
Black Hat	These people are criminals and often seek to damage organisations or gain some form of financial benefit at the cost of others.	For example, ransomware authors infect devices with malicious code and hold data for ransom.

Scopes

Black-Box Testing

This testing process is a high-level process where the tester is not given any information about the inner workings of the application or service.

The tester acts as a regular user testing the functionality and interaction of the application or piece of software. This testing can involve interacting with the interface, i.e. buttons, and testing to see whether the intended result is returned. No knowledge of programming or understanding of the programme is necessary for this type of testing.

Black-Box testing significantly increases the amount of time spent during the information gathering and enumeration phase to understand the attack surface of the target.

Grey-Box Testing

This testing process is the most popular for things such as penetration testing. It is a combination of both black-box and white-box testing processes. The tester will have some limited knowledge of the internal components of the application or piece of software. Still, it will be interacting with the application as if it were a black-box scenario and then using their knowledge of the application to try and resolve issues as they find them.

With Grey-Box testing, the limited knowledge given saves time, and is often chosen for extremely well-hardened attack surfaces.

White-Box Testing

This testing process is a low-level process usually done by a software developer who knows programming and application logic. The tester will be testing the internal components of the application or piece of software and, for example, ensuring that specific functions work correctly and within a reasonable amount of time.

The tester will have full knowledge of the application and its expected behaviour and is much more time consuming than black-box testing. The full knowledge in a White-Box testing scenario provides a testing approach that guarantees the entire attack surface can be validated.

Principles of Security

CIA Triad

The CIA triad is an information security model that is used in consideration throughout creating a security policy. This model has an extensive background, ranging from being used in 1998.

This history is because the security of information (information security) does not start and/or end with cybersecurity, but instead, applies to scenarios like filing, record storage, etc.

Consisting of three sections: **C**onfidentiality, **I**ntegrity and **A**vailability (**CIA**), this model has quickly become an industry standard today. This model should help determine the value of data that it applies to, and in turn, the attention it needs from the business.

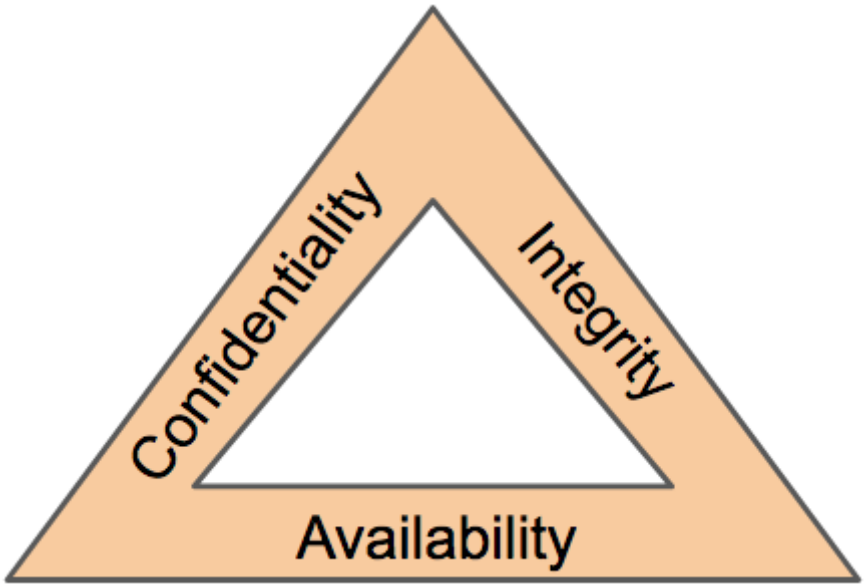

The CIA triad is unlike a traditional model where you have individual sections; instead, it is a continuous cycle. Whilst the three elements to the CIA triad can arguably overlap, if even just one element is not met, then the other two are rendered useless (similar to the fire triangle). If a security policy does not answer these three sections, it is seldom an effective security policy.

Whilst the three elements to the CIA triad are arguably self-explanatory, let's explore these and contextualise them into cybersecurity.

Confidentiality

This element is the protection of data from unauthorized access and misuse. Organisations will always have some form of sensitive data stored on their systems. To provide confidentiality is to protect this data from parties that it is not intended for.

There are many real-world examples for this, for example, employee records and accounting documents will be considered sensitive. Confidentiality will be provided in the sense that only HR administrators will access employee records, where vetting and tight access controls are in place. Accounting records are less valuable (and therefore less sensitive), so not as stringent access controls would be in place for these documents. Or, for example, governments using a sensitivity classification rating system (top-secret, classified, unclassified)

Integrity

The CIA triad element of integrity is the condition where information is kept accurate and consistent unless authorized changes are made. It is possible for the information to change because of careless access and use, errors in the information system, or unauthorized access and use. In the CIA triad, integrity is maintained when the information remains unchanged during storage, transmission, and usage not involving modification to the information. Steps must be taken to ensure data cannot be altered by unauthorised people (for example, in a breach of confidentiality).

Many defences to ensure integrity can be put in place. Access control and rigorous authentication can help prevent authorized users from making unauthorized changes. Hash verifications and digital signatures can help ensure that transactions are authentic and that files have not been modified or corrupted.

Availability

In order for data to be useful, it must be available and accessible by the user.

The main concern in the CIA triad is that the information should be available when authorised users need to access it.

Availability is very often a key benchmark for an organisation. For example, having 99.99% uptime on their websites or systems (this is laid out in Service Level Agreements). When a system is unavailable, it often results in damage to an organisations reputation and loss of finances. Availability is achieved through a combination of many elements, including:

- Having reliable and well-tested hardware for their information technology servers (i.e. reputable servers)
- Having redundant technology and services in the case of failure of the primary
- Implementing well-versed security protocols to protect technology and services from attack

Principles of Privileges

It is vital to administrate and correctly define the various levels of access to an information technology system individuals require.

The levels of access given to individuals are determined on two primary factors:

- The individual's role/function within the organisation
- The sensitivity of the information being stored on the system

Two key concepts are used to assign and manage the access rights of individuals, two key concepts are used: Privileged Identity Management (PIM) and Privileged Access Management (or PAM for short).

Initially, these two concepts can seem to overlap; however, they are different from one another. PIM is used to translate a user's role within an organisation into an access role on a system. Whereas PAM is the management of the privileges a system's access role has, amongst other things.

What is essential when discussing privilege and access controls is the principle of least privilege. Simply, users should be given the minimum amount of privileges, and only those that are absolutely necessary for them to perform their duties. Other people should be able to trust what people write to.

As we previously mentioned, PAM incorporates more than assigning access. It also encompasses enforcing security policies such as password management, auditing policies and reducing the attack surface a system faces.

Security Models

Before discussing security models further, let's recall the three elements of the CIA triad: Confidentiality, Integrity and Availability. We've previously outlined what these elements are and their importance. However, there is a formal way of achieving this.

According to a security model, any system or piece of technology storing information is called an information system, which is how we will reference systems and devices in this task.

Let's explore some popular and effective security models used to achieve the three elements of the CIA triad.

The Bell-La Padula Model

The Bell-La Padula Model is used to achieve confidentiality. This model has a few assumptions, such as an organisation's hierarchical structure it is used in, where everyone's responsibilities/roles are well-defined.

The model works by granting access to pieces of data (called objects) on a strictly need to know basis. This model uses the rule "no write down, no read up".

Advantages	**Disadvantages**
Policies in this model can be replicated to real-life organisations hierarchies (and vice versa)	Even though a user may not have access to an object, they will know about its existence -- so it's not confidential in that aspect.
Simple to implement and understand, and has been proven to be successful.	The model relies on a large amount of trust within the organisation.

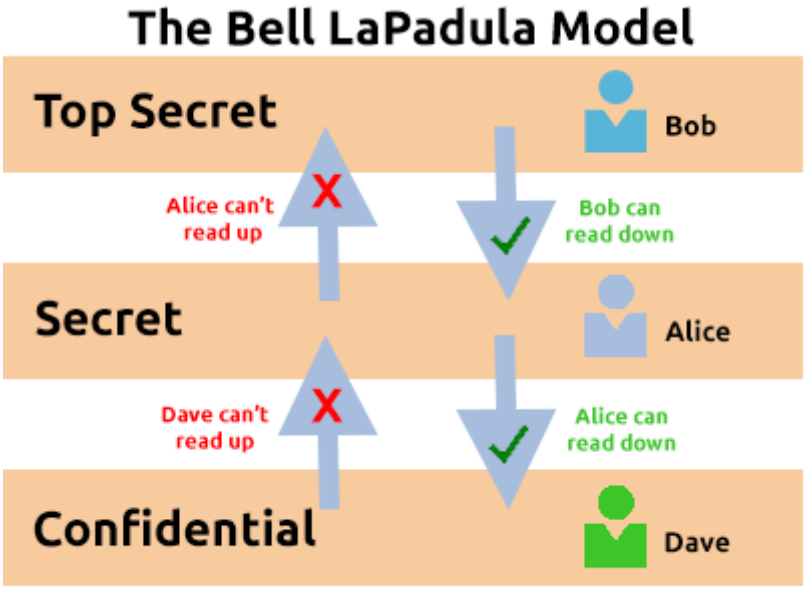

The Bell LaPadula Model is popular within organisations such as governmental and military. This is because members of the organisations are presumed to have already gone through a process called vetting. Vetting is a screening process where applicant's backgrounds are examined to establish the risk they pose to the organisation. Therefore, applicants who are successfully vetted are assumed to be trustworthy - which is where this model fits in.

Biba Model

The Biba model is arguably the equivalent of the Bell-La Padula model but for the integrity of the CIA triad.

This model applies the rule to objects (data) and subjects (users) that can be summarised as "no write up, no read down". This rule means that subjects can create or write content to objects at or below their level but **can only** read the contents of objects above the subject's level.

Let's compare some advantages and disadvantages of this model in the table below:

Advantages	Disadvantages
This model is simple to implement.	There will be many levels of access and objects. Things can be easily overlooked when applying security controls.
Resolves the limitations of the Bell-La Padula model by addressing both confidentiality and data integrity.	Often results in delays within a business. For example, a doctor would not be able to read the notes made by a nurse in a hospital with this model.

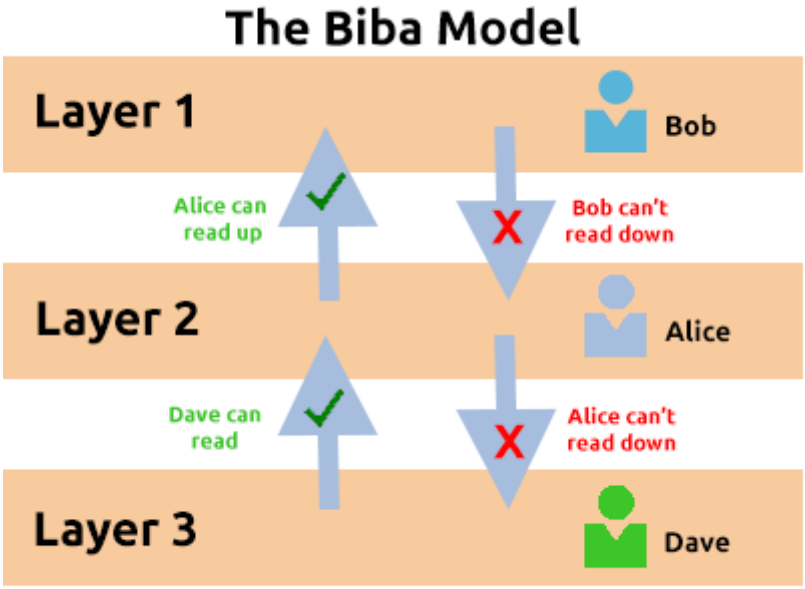

The Biba model is used in organisations or situations where integrity is more important than confidentiality. For example, in software development, developers may only have access to the code that is necessary for their job. They may not need access to critical pieces of information such as databases, etc.

Threat Modelling & Incident Response

Threat modelling is the process of reviewing, improving, and testing the security protocols in place in an organisation's information technology infrastructure and services.

A critical stage of the threat modelling process is identifying likely threats that an application or system may face, the vulnerabilities a system or application may be vulnerable to.

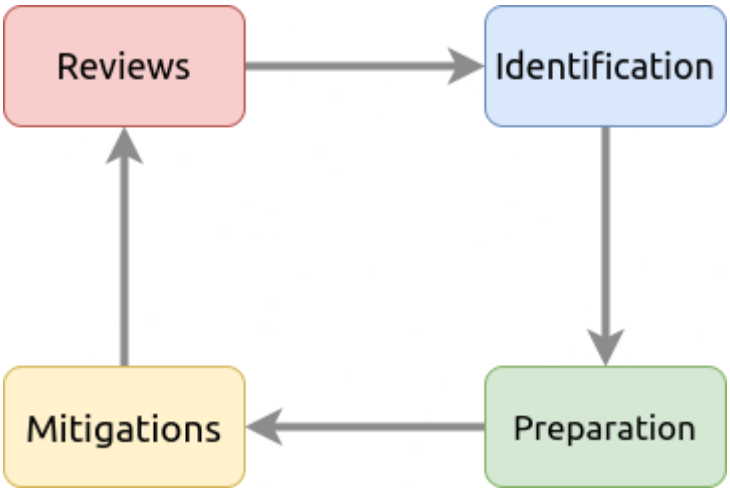

The threat modelling process is very similar to a risk assessment made in workplaces for employees and customers. The principles all return to:

- Preparation
- Identification
- Mitigations
- Review

It is, however, a complex process that needs constant review and discussion with a dedicated team. An effective threat model includes:

- Threat intelligence
- Asset identification
- Mitigation capabilities
- Risk assessment

To help with this, there are frameworks such as **STRIDE** (**S**poofing identity, **T**ampering with data, **R** epudiation threats, **I**nformation disclosure, **D**enial of Service and **E**levation of privileges) and

PASTA (**P**rocess for **A**ttack **S**imulation and **T**hreat **A**nalysis) infosec never tasted so good!. Let's detail STRIDE below. STRIDE, authored by two Microsoft security researchers in 1999 is still very relevant today. STRIDE includes six main principles, which I have detailed in the table below:

Principle	**Description**
Spoofing	This principle requires you to authenticate requests and users accessing a system. Spoofing involves a malicious party falsely identifying itself as another. Access keys (such as API keys) or signatures via encryption helps remediate this threat.
Tampering	By providing anti-tampering measures to a system or application, you help provide integrity to the data. Data that is accessed must be kept integral and accurate. For example, shops use seals on food products.
Repudiation	This principle dictates the use of services such as logging of activity for a system or application to track.
Information Disclosure	Applications or services that handle information of multiple users need to be appropriately configured to only show information relevant to the owner is shown.
Denial of Service	Applications and services use up system resources, these two things should have measures in place so that abuse of the application/service won't result in bringing the whole system down.
Elevation of Privilege	This is the worst-case scenario for an application or service. It means that a user was able to escalate their authorization to that of a higher level i.e. an administrator. This scenario often leads to further exploitation or information disclosure.

A breach of security is known as an incident. And despite all rigorous threat models and secure system designs, incidents do happen. Actions taken to resolve and remediate the threat are known as Incident Response (IR) and are a whole career path in cybersecurity.

Incidents are classified using a rating of urgency and impact. Urgency will be determined by the type of attack faced, where the impact will be determined by the affected system and what impact that has on business operations.

Urgency \ Impact	High	Medium	Low
High	1	2	3
Medium	2	3	4
Low	3	4	5

An incident is responded to by a **C**omputer **S**ecurity **I**ncident **R**esponse **T**eam (**CSIRT**) which is prearranged group of employees with technical knowledge about the systems and/or current incident. To successfully solve an incident, these steps are often referred to as the six phases of Incident Response that takes place, listed in the table below:

Action	**Description**
Preparation	Do we have the resources and plans in place to deal with the security incident?
Identification	Has the threat and the threat actor been correctly identified in order for us to respond to?
Containment	Can the threat/security incident be contained to prevent other systems or users from being impacted?
Eradication	Remove the active threat.
Recovery	Perform a full review of the impacted systems to return to business as usual operations.
	What can be learnt from the incident? I.e. if it was due to a phishing email, employees should be trained better to detect phishing emails.