DIRECTORATUL NAȚIONAL
DE SECURITATE CIBERNETICĂ

ISACA

Keep your Information System Safe (KISS)

Practical Steps for Implementation

Best Practices and Legal Considerations

Authors:

Theodor Octavian Adam

Florin Andrei

Larisa Găbudeanu

Vasile Victor Rotaru

Editor and guest writer:

Alexandru Mircea Rotaru

SITECH

Craiova, 2021

Publisher: Sitech Publishing, Craiova, 2021

ISBN 978-606-11-7931-2

Published in October 2021 at https://cert.ro/vezi/document/keep-your-information-system-safe

Summary

Foreword

The amazing volume of innovation in technology in the past years has brought us amazing opportunities for communication, access to information, development in business environments and, without any doubt, value in our private lives. At the same time, challenges to appropriately embrace all the changes and protect business and personal data have increased exponentially.

Last year (2020) broke all records when we speak about data breaches and numbers of cyber-attacks on companies, government, and individuals. And the concern does not stem only from the numbers of incidents but also from the sophistication of threats.

And when you learn that 86.2% of organizations were affected by at least a successful attack[1] or cybercrime is up 600% due to the COVID-19 pandemic[2] or each 1.12 seconds a new successful cyber incident happens[3] the need for raising awareness, educating people, increasing the level of cyber-maturity within organizations is becoming an imperative.

But all of these can be only supported with highly qualified people, extensive hands-on experience in dealing with current cyber challenges and continuous investments in people. The knowledge and skills of specialists ought to cover broader areas than exclusively strong academic and technical ones.

And beyond the enormous amount of literature on emerging technology, the role of cyber community members is to share the know-how, be curious and vigilant.

The current volume encompasses the hard work of 4 security professionals carried out over half a year, who committed themselves and with enthusiasm and

[1] https://cyber-edge.com/cdr/ , last accessed on 18 September 2021.

[2] https://purplesec.us/resources/cyber-security-statistics/ , last accessed on 18 September 2021.

[3] https://eng.umd.edu/news/story/study-hackers-attack-every-39-seconds, last accessed on 18 September 2021.

perseverance and assumed a very high responsibility of sharing their experience through periodic writing.

The aim of this book is not only to put together some principles, methodologies and any other relevant theoretical aspects on critical cybersecurity topics (offensive security, security incident handling, healthcare and supply chain security challenges, etc.) but also to provide an end-to-end overview of such selected topics and provide guidance on practical aspects for implementation, based on their own experience and perspective. Moreover, the content makes for accessible reading independent of the level of seniority a specialist might have. If you are a beginner, then this is a perfect way to gain an insight into the topics and an invitation to start searching for additional materials depending on your needs and curiosity. In case you are a senior in the area, then you might find the information helpful as a means to compare and validate your approach and vision but at the same time it might give you new perspectives on specific topics.

An aspect which cannot go unnoticed is the legal aspects of many subjects covered in this volume. And this is extremely valuable information as combining technology and legal perspective is not at all an easy feat.

Congratulations to all team members for starting and continuing the initiative, for your commitment, enthusiasm and sometimes your endless energy and determination. It is a fantastic journey you embarked on and I am honored to support you.

Kudos to the team!

Gratiela Magdalinoiu

President of ISACA Romania

1. Introduction

This book aims to assist organizations when setting-up their offensive security process (or parts thereof) or when analyzing how to improve the existing offensive security process. The purpose is to provide the audience with practical approaches that can be implemented by organizations (small, medium or large) by taking into account standards and best practices in this field.

The three main angles from which security is analyzed throughout this book are Red Teaming, penetration testing and Blue Team.

Red teaming is detailed in terms of approach to be taken when analyzing an organization in this manner, both from the Red Team's and the organization's point of view.

Penetration testing has become a very important tool for new IT systems and for day-to-day analysis of the IT landscape. The chapters addressing this process include useful steps and approaches both from the organization's and the penetration tester's perspective.

The Blue Team process within the organization is outlined in terms of roles and responsibilities, framework for operating and includes also a SOC team creation tips and tricks.

On top of these three main angles, we have included practical considerations from a legal and privacy perspective.

The data governance chapters and sub-chapters address the entire life-cycle of personal data within the offensive security processes. Emphasis is placed on the collection of data, internal use of data and disclosure of data to third parties (e.g. CRISTs, vendors, entities from the same industry, authorities). In this respect suggestions of methodologies and steps are included, together with relevant examples.

Legal requirements outlined in the specific chapters and in sub-chapters refer mainly to contractual documentation (and processes) with third parties, especially those involved in the Red Teaming and penetration exercise.

Furthermore, this book includes the analysis concerning forensic steps to be taken in case of incidents (including gathering of evidence for potential litigation), handling whistleblowing requests and sharing incident data (or data gathered from offensive security) to third parties.

Given the increasing focus on supply chain attacks, practical considerations for supply chain security and privacy management are included in a dedicated chapter and throughout all other relevant chapters.

In view of providing further guidance on the above topics, the book includes a series of chapters dedicated to specific areas encountered in real life.

As it is an area of focus in the recent years from multiple perspectives, the healthcare sector is analyzed in specific chapters, from handing patient data as per the day-to-day activity, eHealth implementation to use of medical devices within the healthcare organization.

Specific emphasis is placed on the use of geographic information system (GIS) by organizations (including for information security purposes), on security steps to be taken during software development life cycle (SDLC) and, given the recent trend of remote work, on endpoint security.

As the role of management is essential in information security, practical steps and mechanisms for presenting the current information security status and future projects/proposals to management have also been detailed.

2. Process Management for Red Team and Threat Intelligence Activities in an Organization

Because cyber threats are continuously evolving, we need to get an accurate picture of our organization's defenses. We need to clearly understand how effective our organization's security is, by considering the current threat landscape, in order to continuously improve the organization's security effectiveness, which involves people, process, and technology. Only a holistic approach can address this challenge; it is obviously broader than pen-testing, and it should be a continuous process with the following main objectives:

- To prepare the organization for targeted attacks (especially multi-step, multi-vector attacks like Advanced Persistent Threat (APT));
- To assess the organization's prevention strategy and program effectiveness;
- To identify and mitigate any kind of vulnerabilities in the organization's infrastructure;
- To minimize the organization's public digital footprint (and consequently, its digital attack surface);
- To enhance the organization's security team ability to detect and respond to real-world incidents.

The steps an attacker must complete to carry out a successful attack are described in Lockheed Martin's Cyber Kill Chain model. This model consists of seven sequential steps:

- Reconnaissance
- Weaponization
- Delivery
- Exploitation
- Installation
- Command and Control
- Actions on Objectives

To disrupt an attack, one or more of these steps must be broken for the entire chain to fail. We need to understand the adversary's behavior before we can apply effective defensive tactics and techniques.

2.1 RECONNAISSANCE

The first step of any cybersecurity attack is collecting information about the victim, also known as reconnaissance. There are two different phases of reconnaissance:

- Passive reconnaissance;
- Active reconnaissance;

During the passive reconnaissance phase, the attacker will use indirect methods to gather information about his target(s) from publicly available sources, like:

- Whois
- Google
- Shodan
- Company websites
- Job Listings

A company is a sum of its employees, procedures, protocols, standards, and technology. Today, each employee has their own digital footprint. Whether it's in their "personal digital life" or their "professional digital life", each employee performs some online public digital actions across the public web (uploading a resume to a website, publishing an article on LinkedIn or other website, commenting on a blog post or tweet, etc.). All these little bits of information an individual or a company leaves on the public web represents their respective "digital footprint". Some of that footprint could be from online public digital actions dating back years ago, when cyber security awareness was virtually non-existent. Any connections between these digital footprints and the company can be used as an enabler (attack vector) against the company. Having a digital footprint is normal – they are very difficult to avoid. We can, however, become aware about what it looks like and actively manage it.

SpiderFoot is a reconnaissance tool that your security professionals might consider using to assess the organization's digital footprint on the public web. It automatically queries over 100 public data sources (OSINT) to gather intelligence on IP addresses, domain names, e-mail addresses, and names to build up an understanding of all the entities involved and how they relate to each other. OSINT (Open Source Intelligence) is data available in the public domain. This includes DNS, Whois, Web pages, passive DNS, spam blacklists, file meta data, and threat intelligence lists, as well as services like SHODAN and HaveIBeenPwned?.

No matter what tool you will choose to use, you should always keep in mind that there is a dark side to intelligence gathering: anything that can be found by security professionals can also be found (and used) by threat actors. Having a clear strategy and framework in place for intelligence gathering is essential — simply looking for anything that could be interesting or useful will inevitably lead to burnout. Also, every organization runs on limited resources, and that includes time, processing power, funds, and everything else needed to do intelligence gathering. By not making deliberate choices on where to invest those resources, you are creating waste in your company. Even the smallest amount of aimlessly invested resources can prove disastrous in detecting and mitigating threats, as every second gained can mean the difference between a minor cybersecurity incident and a PR and legal nightmare. Unfortunately, with how massive the internet has become, it's very easy to go on wild goose chases during intelligence gathering if there's no strategy in place, Therefore, a very careful and responsible attitude is mandatory for intelligence gathering, as it can change the risk landscape of the assessed environment.

Done right, the data returned from an intelligence gathering scan will reveal a lot of information, providing insight into possible data leaks, vulnerabilities, or other sensitive information that can be leveraged during a penetration test, Red Team exercise, or for threat intelligence. All this information is very helpful in creating a good understanding of where you might be exposed.

Defending against passive reconnaissance means limiting the level of details you expose. To know where your weak points are, you need to understand your complete digital footprint and view your organization from a hacker's perspective and try to limit your exposure whenever possible. For instance, this can start with your organization's websites and web/mobile applications. Your organization should consider removing specific error messages from its public servers. In addition, as mentioned previously, employees are a critical component of your cyber security protection, and are often the weakest link. There is no replacement for employee cyber security education, which emphasizes various tactics used by adversaries, and details the consequences of doing anything that exposes their data. Limiting the information your organization puts on job postings and on its websites helps to reduce its digital footprint. You should also take into account your business partners when doing any digital footprint assessment.

With GDPR, now the basis for European data protection law, the Right to be Forgotten and the pre-GDPR obligation to delete unnecessary data (for which there is no legal basis of storing), can help remove inaccurate or out of date data. Furthermore, the data minimization principle, which entails limiting data exposure, as well as anonymisation or Pseudonymisation of data can help minimize the organization's digital footprint. So, don't hesitate to use these. Always ask yourself: does this person need access to this data? And do we really need to keep this data for a specific purpose?

The rule of thumb to minimize any digital footprint is to think before posting. This is one of the easiest ways to keep someone safe online and reduce their digital footprint. Don't share details you don't want the world to see, and don't post something you don't want out in the world forever. Another rule of thumb is to always keep in mind how attackers can use your data to cause harm to you, the people close to you, and your organization(s). Information is power, so don't give ill-intentioned strangers the chance to have power over you.

Once an attacker has collected as much public information as possible about his target(s), he will move on to active reconnaissance. Active reconnaissance involves some level of interaction with your organization.

During this phase, the attacker will actively probe your network looking for open ports and services. Their goal is to discover exploitable communication channels and find various ways to intrude the target system. The tools used for active reconnaissance include: nmap, OpenVAS, Nikto, netcat, and Metasploit.

Network reconnaissance is a crucial part of any hacking operation. Any information that a hacker can learn about the target environment can help in identifying potential attack vectors and targeting exploits to potential vulnerabilities. Vulnerability scanners are very loud and obvious, so attackers will usually limit their scope or scan slowly over a longer period of time to avoid detection.

An attacker may choose to obfuscate his scan. A common obfuscation method is to spoof many source IP addresses along with the real source IP for a port scan. The target machine will likely log the scan, but it will be extremely difficult for the network admin to determine from which IP address the port scan actually originated.

During the COVID-19 pandemic, many companies decided to allow remote working through remote services such as VPNs, Citrix, and other access mechanisms that allow users to connect to internal enterprise network resources from external locations. This increases the risks because adversaries may use remote services to access and/or persist within your network.

For active reconnaissance, your first protection measure is at the network infrastructure level, by ensuring that all unused ports and services are disabled. A stateful firewall with IPS capabilities placed on a network perimeter is likely one of the best prevention measures for any intrusion. The firewall should be configured to allow only the necessary traffic and should log multiple connection attempts from the same source and/or IP address. This limits the number of entry points an attacker can use to get into your system. Using strong two-factor or multi-factor authentication for remote service accounts will mitigate an adversary's ability to leverage stolen credentials and will disable or block remotely available services that may be unnecessary.

The main goal of the reconnaissance phase is to find weaknesses that can be exploited. Once the attacker has found at least one weakness, they can move on to the next step; they cannot digitally infiltrate your organization’s systems without a weakness to exploit. This is why your security professionals must find and fix all the weaknesses in your systems as soon as possible, so that an attacker doesn’t have the chance to find and exploit them. This is the reasoning behind the defense in depth concept.

2.2 WEAPONIZATION

Once an attacker has found a weakness, their next step is to find or create an attack plan to exploit that vulnerability. The weapon of choice will depend on the information collected during the reconnaissance step.

Some commonly used weapons during this phase are tools like Metasploit or Exploit-DB. These are repositories for known exploits. The Veil framework is commonly used to generate Metasploit payloads that bypass common anti-virus solutions. TheFratRat is an easy tool to generate backdoors and to post exploitation attacks. The Social-Engineer Toolkit (SET) might be used if the attacker decides to deliver the malware through a social engineering campaign. These are only a few options among a wide range that an attacker can use to build his weapon. Your organization also needs to consider that the attacker might choose to craft his weapon using a different pattern than those already available.

Unfortunately, the vast majority of today's breaches occur because the basics are still not fully covered. These include unpatched servers and computers, outdated antivirus, installed plugins, and many others. An attacker can and will exploit any of these paths. For this reason, patch management, along with up-to-date antivirus programs and disabled plugins and macros continue to be the best defensive measures against the weaponization phase.

This phase is all about what the attacker can use as a weapon. Reducing the exposure by keeping the operating systems and antivirus up to date are critical because an attacker can't exploit a vulnerability that doesn't exist to begin with. An IDS/IPS tuned to look for exploit attempts will help you see what vulnerabilities attackers have tried to exploit, so that your organization can mitigate them and thus shut down potential paths for future attacks.

During the weaponization phase, the attacker chooses which weapon(s) to use, but has not delivered it yet. How the attack is delivered is as critical as the weapon chosen. This brings us to the next phase.

2.3 DELIVERY

At this point, the attacker has selected the weapon based on their earlier reconnaissance. Now, they will try to use one or multiple paths to deliver their weapon(s). The delivery path varies by the kind of attack, but the most common examples are: e-mail, websites, social media, and USB flash drives.

Through e-mail, an attacker might embed a malware into an attached file and might disguise (“phish") the e-mail to make it look like it’s coming from a partner/supplier they found during the reconnaissance phase. That way, an employee is most likely to open it.

Attackers can choose to infect a website frequently used by members of the organization, or use social media. However, when using a social network as a vector for targeted attack, the attacker must have some level of interaction with this environment.

Attackers can also use offline pathways as much as online ones, particularly in the form of remote storage devices, such as USB flash drives. For instance, during a conference, an attacker can easily obtain one or more USB flash drives with the conference materials, infect them with a rootkit, and then return them. Also, there’s the even worse and less daring way of leaving an infected USB flash drive somewhere in a public area, around employees, hoping that the temptation for them to plug it into their computer will be big enough to make it happen.

The single best security measure against the delivery of the attack is user awareness. This includes security training for both employees and security staff about threats and good security practices. Careless or uninformed staff are the second most likely cause of a serious security breach, behind Phishing/Malware (which still requires a human error to activate), hence why employee awareness is important in keeping the organization safe.

To limit the delivery paths an attacker can use you can implement the following measures:

- SPF (Sender Policy Framework) is an email validation protocol designed to detect and block email spoofing. SPF is a “proposed standard” that helps protect email users from potential spammers.

- DKIM (DomainKeys Identified Mail) lets an organization (or handler of the message) take responsibility for a message that is in transit.

- DMARC (Domain-Based Message Authentication Reporting and Conformance) is an added authentication method that uses both SPF and DKIM to verify whether or not an email was actually sent by the owner of the “Friendly-From” domain that the user sees.

- Installing Web Application Firewall (WAF) or at least web filtering to prevent a user to access known bad websites.
- DNS filtering to prevent DNS lookup attempts.
- Disabling USB ports on the computers and not giving administrative rights to users.

Nevertheless, your organization needs to do SSL inspection for all of its delivery channels. Only using a full SSL Inspection or Deep SSL Inspection can your organization do antivirus scanning, web filtering, email filtering, etc. to filter out malicious content from network traffic. New generations of firewalls also support a second type of SSL inspection, called the SSL certificate inspection. An attacker will almost always use encrypted connections to avoid being caught. If you are not doing full SSL deep packet inspections, you have no chance to detect any communication attempts going through any sort of encrypted communication tunnel.

If the attacker succeeds to deliver the weapon, then he will move to the next phase.

2.4 EXPLOITATION

The Exploitation phase begins when the attacker has delivered his weapon(s) of choice to the victim(s) and the attack has been executed; at this point, all of your prevention measures to keep the weapon outside of your environment have failed, and the only thing left to do is wait for the moment when the attacker will pull the trigger and 'detonate' the attack.

Once the attacker has been able to execute the exploit, the protective measures you can rely on are very limited. If you are lucky and figure out what exploit the attacker used, you can still use Data Execution Prevention (DEP) and the anti-exploit technology embedded in antivirus.

Data Execution Prevention (DEP) is a security feature within an operating system that prevents applications from executing code from a non-executable memory location. It was designed to secure information systems against memory-based code exploits and is available on Windows, Linux and Mac OS.

You should always expect that the attacker will also attempt to compromise additional systems and/or accounts by gaining admin-type rights, which is why you should re-focus your prevention measures on finding and then protecting the areas of your organization that have yet to be attacked. An appropriate countermeasure might be deploying an automatic rights escalation and de-escalation to streamline the process and contain an infection in case of an intrusion, and thus limit its fallout on your systems.

Although there's no guarantee that sandboxing will stop zero-day threats, it offers an additional security layer by separating the threats from the rest of the network because it runs on a separate system. Once threats get quarantined, cybersecurity experts can study them to identify patterns, helping to prevent future attacks and identify other vulnerabilities. It also allows IT to test malicious code in an isolated testing environment to understand how it works within a system as well as more rapidly detect similar malware attacks.

All in all, an exploit is not a one hit operation. The goal of the exploit is to gain better access on our resources. This leads us to the next phase.

2.5 INSTALLATION

From the attacker's perspective, gaining better access allows them to control the victim, even after the system has been patched or restarted. The attempts to gain better access are not limited to one system; the attacker might choose to take their time and restart internally the cyber kill chain by starting a fresh reconnaissance phase after they get full access to one of your systems.

During the installation phase, an attacker aims to create persistence and to get better access, with the ultimate goal of accomplishing the mission. Persistence consists of techniques that adversaries use to keep access to systems across restarts, changed credentials, and other interruptions that could cut off the attackers' access.

Traditional detection sensors based on traffic analysis and comparisons with known attack signatures do not detect effectively subtle attacks like APTs, or unknown threats like zero-day.

Endpoint Detection and Response (EDR) focuses on detecting attackers that evaded the prevention layer of an Endpoint Protection Platform (EPP) solution and that are now active in the target environment. The EDR can detect when an attack has taken place, take immediate action on the endpoint to prevent the attack from spreading, and provide real-time forensic information to help investigate and respond to the attack. The EDR is today considered an essential part of the EPP.

Once you determine that a system was infected, you can start the process to restore that system to a known stable state.

2.6 COMMAND AND CONTROL

At this stage, the system is compromised and under the attacker's control. If they completed the previous steps correctly, their access persists even after patching the vulnerability or restarting the system. The infected system can be used either to carry out the attack mission or it can wait for future instructions from its command-and-control server.

Your defending tactics should center on detecting unusual activity and limiting what the attacker can control. Network segmentation and application segmentation will make any lateral movements the attacker performs harder and detecting their presence easier. Detection tools can look for changed patterns of network usage, such as increased amounts of information sent or downloaded to/from an external server, or changed behaviors on the internal servers, like an unusual spike in CPU or memory utilization, or at unusual hours.

Lateral movement represents the set of steps that an attacker who gained a foothold in a trusted environment takes to expand their level of access, to move to additional trusted assets, and to further advance towards their ultimate target.

You can use the Breach and Attack Simulation (BAS) tool to simulate real attacks against your data center, so you can review the results and take action. This is the best way to stay ahead of the attacker. However, keep in mind that "the best" does not equal "perfect" in this case, so make sure your organization stays vigilant even after performing a BAS.

2.7 ACTIONS ON OBJECTIVES

The system is now infected and the attacker is in full control. Now, they can do what they need to achieve their objectives. Their motivation may be financial, political, or simply moving laterally to target a more important system in the network; regardless, figuring out the attacker's motivation will help you predict how they might act moving forward.

Lateral movement is a common step an attacker takes once they gain the access into a system. At this point, they might begin a newer reconnaissance phase to gain information about the internal network.

The pandemic forced many companies to suddenly move some or even all of employees outside of their "traditional" network. Given the increased risks coming from working remotely, such as unauthorized access to assets, organizations with remote workers should transition to a zero trust security model as soon as possible to improve the organization's security.

Zero trust security is an IT security model that requires strict identity verification for every person and device trying to access resources on a private network, regardless of whether they are located within or outside the network perimeter. A zero trust network assumes that there are attackers both within and outside of the network, so no users or machines should be automatically trusted, and gaining access to resources on the network requires verification for everyone.

A key principle of zero trust security is least-privilege access: users only get as much access as they need. This minimizes each user's exposure to sensitive parts of the network and can prevent privilege escalation, which is a vital part of the cyber kill chain.

Zero trust networks also utilize micro-segmentation: breaking up security perimeters into small zones to maintain separate access for separate parts of the network. This principle of micro-segmentation, included within the design of an application, ensures that the privacy by default principle gets implemented.

Multi-factor authentication (MFA) is another principle of zero trust security. MFA simply means requiring more than one piece of evidence to authenticate a user; just entering a password (i.e. only one authentication factor) is not enough to gain access.

A very helpful knowledge base of adversary tactics and techniques is MITRE ATT&CK. This knowledge base can be used as a foundation for developing specific threat models and methodologies. It provides recommendations for detecting and mitigating known threats. You can find it at https://attack.mitre.org/.

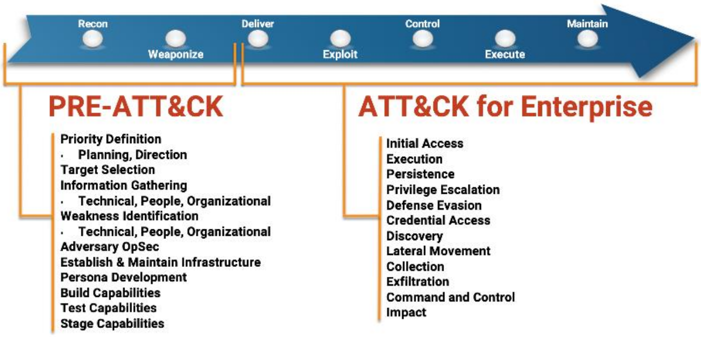

PRE-ATT&CK

Priority Definition
- Planning, Direction

Target Selection
Information Gathering
- Technical, People, Organizational

Weakness Identification
- Technical, People, Organizational

Adversary OpSec
Establish & Maintain Infrastructure
Persona Development
Build Capabilities
Test Capabilities
Stage Capabilities

ATT&CK for Enterprise

Initial Access
Execution
Persistence
Privilege Escalation
Defense Evasion
Credential Access
Discovery
Lateral Movement
Collection
Exfiltration
Command and Control
Impact

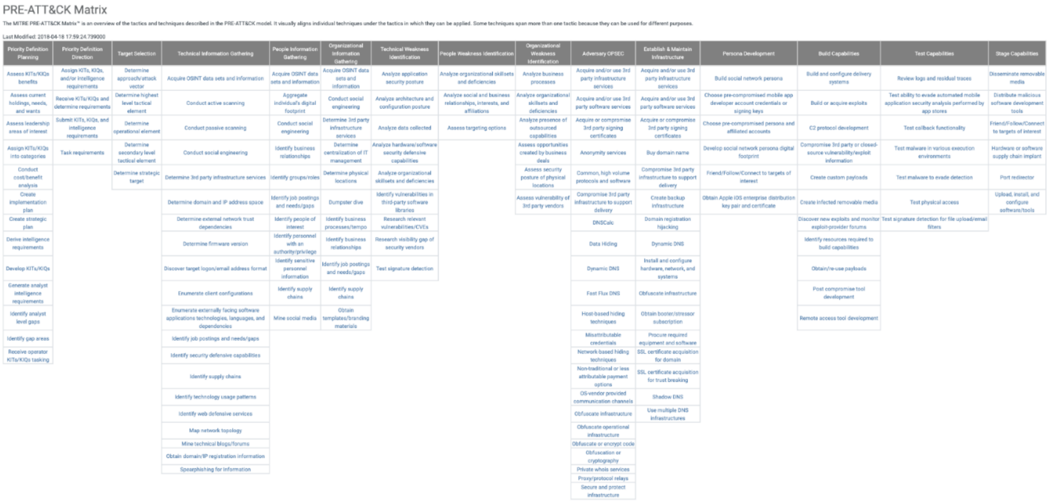

PRE-ATT&CK Matrix

The MITRE PRE-ATT&CK Matrix™ is an overview of the tactics and techniques described in the PRE-ATT&CK model. It visually aligns individual techniques under the tactics in which they can be applied. Some techniques span more than one tactic because they can be used for different purposes.

Last Modified: 2018-04-18 17:59:24.739000

Priority Definition Planning	Priority Definition Direction	Target Selection	Technical Information Gathering	People Information Gathering	Organizational Information Gathering	Technical Weakness Identification	People Weakness Identification	Organizational Weakness Identification	Adversary OPSEC	Establish & Maintain Infrastructure	Persona Development	Build Capabilities	Test Capabilities	Stage Capabilities
Assess KITs/KIQs benefits	Assign KITs, KIQs, and/or intelligence requirements	Determine approach/attack vector	Acquire OSINT data sets and information	Acquire OSINT data sets and information	Acquire OSINT data sets and information	Analyze application security posture	Analyze organizational skillsets and deficiencies	Analyze business processes	Acquire and/or use 3rd party infrastructure services	Acquire and/or use 3rd party infrastructure services	Build social network persona	Build and configure delivery systems	Review logs and residual traces	Disseminate removable media
Assess current holdings, needs, and wants	Receive KITs/KIQs and determine requirements	Determine highest level tactical element	Conduct active scanning	Aggregate individual's digital footprint	Conduct social engineering	Analyze architecture and configuration posture	Analyze social and business relationships, interests, and affiliations	Analyze organizational skillsets and deficiencies	Acquire and/or use 3rd party software services	Acquire and/or use 3rd party software services	Choose pre-compromised mobile app developer account credentials or signing keys	Build or acquire exploits	Test ability to evade automated mobile application security analysis performed by app stores	Distribute malicious software development tools
Assess leadership areas of interest	Submit KITs, KIQs, and intelligence requirements	Determine operational element	Conduct passive scanning	Conduct social engineering	Determine 3rd party infrastructure services	Analyze data collected	Assess targeting options	Analyze presence of outsourced capabilities	Acquire or compromise 3rd party signing certificates	Acquire or compromise 3rd party signing certificates	Choose pre-compromised persona and affiliated accounts	C2 protocol development	Test callback functionality	Friend/Follow/Connect to targets of interest
Assign KITs/KIQs into categories	Task requirements	Determine secondary level tactical element	Conduct social engineering	Identify business relationships	Determine centralization of IT management	Analyze hardware/software security defensive capabilities		Assess opportunities created by business deals	Anonymity services	Buy domain name	Develop social network persona digital footprint	Compromise 3rd party or closed-source vulnerability/exploit information	Test malware in various execution environments	Hardware or software supply chain implant
Conduct cost/benefit analysis		Determine strategic target	Determine 3rd party infrastructure services	Identify groups/roles	Determine physical locations	Analyze organizational skillsets and deficiencies		Assess security posture of physical locations	Common, high volume protocols and software	Compromise 3rd party infrastructure to support delivery	Friend/Follow/Connect to targets of interest	Create custom payloads	Test malware to evade detection	Port redirector
Create implementation plan			Determine domain and IP address space	Identify job postings and needs/gaps	Dumpster dive	Identify vulnerabilities in third-party software libraries		Assess vulnerability of 3rd party vendors	Compromise 3rd party infrastructure to support delivery	Create backup infrastructure	Obtain Apple iOS enterprise distribution key pair and certificate	Create infected removable media	Test physical access	Upload, install, and configure software/tools
Create strategic plan			Determine external network trust dependencies	Identify people of interest	Identify business processes/tempo	Research relevant vulnerabilities/CVEs			DNSCalc	Domain registration hijacking		Discover new exploits and monitor exploit-provider forums	Test signature detection for file upload/email filters	
Derive intelligence requirements			Determine firmware version	Identify personnel with an authority/privilege	Identify business relationships	Research visibility gap of security vendors			Data Hiding	Dynamic DNS		Identify resources required to build capabilities		
Develop KITs/KIQs			Discover target logon/email address format	Identify sensitive personnel information	Identify job postings and needs/gaps	Test signature detection			Dynamic DNS	Install and configure hardware, network, and systems		Obtain/re-use payloads		
Generate analyst intelligence requirements			Enumerate client configurations	Identify supply chains	Identify supply chains				Fast Flux DNS	Obfuscate infrastructure		Post compromise tool development		
Identify analyst level gaps			Enumerate externally facing software applications technologies, languages, and dependencies	Mine social media	Obtain templates/branding materials				Host-based hiding techniques	Obtain booter/stressor subscription		Remote access tool development		
Identify gap areas			Identify job postings and needs/gaps						Misattributable credentials	Procure required equipment and software				
Receive operator KITs/KIQs tasking			Identify security defensive capabilities						Network-based hiding techniques	SSL certificate acquisition for domain				
			Identify supply chains						Non-traditional or less attributable payment options	SSL certificate acquisition for trust breaking				
			Identify technology usage patterns						OS-vendor provided communication channels	Shadow DNS				
			Identify web defensive services						Obfuscate infrastructure	Use multiple DNS infrastructures				
			Map network topology						Obfuscate operational infrastructure					
			Mine technical blogs/forums						Obfuscate or encrypt code					
			Obtain domain/IP registration information						Obfuscation or cryptography					
			Spearphishing for Information						Private whois services					
									Proxy/protocol relays					
									Secure and protect infrastructure					

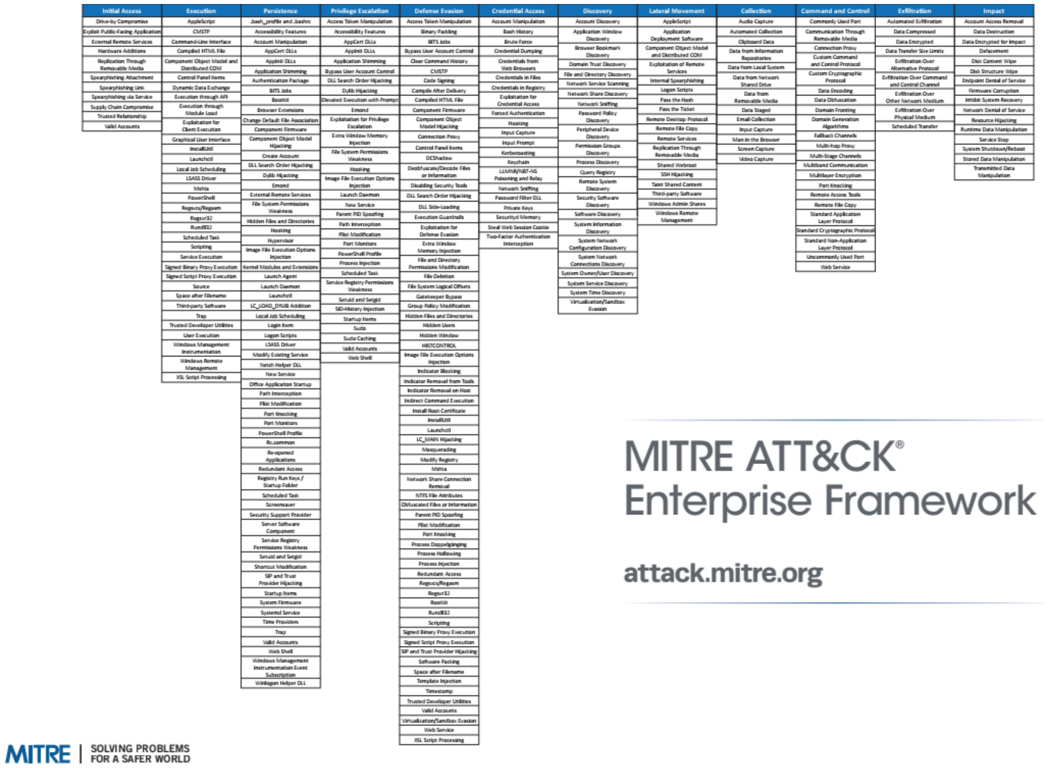

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Collection	Command and Control	Exfiltration	Impact
Drive-by Compromise	AppleScript	.bash_profile and .bashrc	Access Token Manipulation	Access Token Manipulation	Account Manipulation	Account Discovery	AppleScript	Audio Capture	Commonly Used Port	Automated Exfiltration	Account Access Removal
Exploit Public-Facing Application	CMSTP	Accessibility Features	Accessibility Features	Binary Padding	Bash History	Application Window Discovery	Application Deployment Software	Automated Collection	Communication Through Removable Media	Data Compressed	Data Destruction
External Remote Services	Command-Line Interface	Account Manipulation	AppCert DLLs	BITS Jobs	Brute Force	Browser Bookmark Discovery	Component Object Model and Distributed COM	Clipboard Data	Connection Proxy	Data Encrypted	Data Encrypted for Impact
Hardware Additions	Compiled HTML File	AppCert DLLs	AppInit DLLs	Bypass User Account Control	Credential Dumping	Domain Trust Discovery	Exploitation of Remote Services	Data from Information Repositories	Custom Command and Control Protocol	Data Transfer Size Limits	Defacement
Replication Through Removable Media	Component Object Model and Distributed COM	AppInit DLLs	Application Shimming	Clear Command History	Credentials from Web Browsers	File and Directory Discovery	Internal Spearphishing	Data from Local System	Custom Cryptographic Protocol	Exfiltration Over Alternative Protocol	Disk Content Wipe
Spearphishing Attachment	Control Panel Items	Application Shimming	Bypass User Account Control	CMSTP	Credentials in Files	Network Service Scanning	Logon Scripts	Data from Network Shared Drive	Data Encoding	Exfiltration Over Command and Control Channel	Disk Structure Wipe
Spearphishing Link	Dynamic Data Exchange	Authentication Package	DLL Search Order Hijacking	Code Signing	Credentials in Registry	Network Share Discovery	Pass the Hash	Data from Removable Media	Data Obfuscation	Exfiltration Over Other Network Medium	Endpoint Denial of Service
Spearphishing via Service	Execution through API	BITS Jobs	Dylib Hijacking	Compile After Delivery	Exploitation for Credential Access	Network Sniffing	Pass the Ticket	Data Staged	Domain Fronting	Exfiltration Over Physical Medium	Firmware Corruption
Supply Chain Compromise	Execution through Module Load	Bootkit	Elevated Execution with Prompt	Compiled HTML File	Forced Authentication	Password Policy Discovery	Remote Desktop Protocol	Email Collection	Domain Generation Algorithms	Scheduled Transfer	Inhibit System Recovery
Trusted Relationship	Exploitation for Client Execution	Browser Extensions	Emond	Component Firmware	Hooking	Peripheral Device Discovery	Remote File Copy	Input Capture	Fallback Channels		Network Denial of Service
Valid Accounts	Graphical User Interface	Change Default File Association	Exploitation for Privilege Escalation	Component Object Model Hijacking	Input Capture	Permission Groups Discovery	Remote Services	Man in the Browser	Multi-hop Proxy		Resource Hijacking
	InstallUtil	Component Firmware	Extra Window Memory Injection	Connection Proxy	Input Prompt	Process Discovery	Replication Through Removable Media	Screen Capture	Multi-Stage Channels		Runtime Data Manipulation
	Launchctl	Component Object Model Hijacking	File System Permissions Weakness	Control Panel Items	Kerberoasting	Query Registry	Shared Webroot	Video Capture	Multiband Communication		Service Stop
	Local Job Scheduling	Create Account	Hooking	DCShadow	Keychain	Remote System Discovery	SSH Hijacking		Multilayer Encryption		System Shutdown/Reboot
	LSASS Driver	DLL Search Order Hijacking	Image File Execution Options Injection	Deobfuscate/Decode Files or Information	LLMNR/NBT-NS Poisoning and Relay	Security Software Discovery	Taint Shared Content		Port Knocking		Stored Data Manipulation
	Mshta	Dylib Hijacking	Launch Daemon	Disabling Security Tools	Network Sniffing	Software Discovery	Third-party Software		Remote Access Tools		Transmitted Data Manipulation
	PowerShell	Emond	New Service	DLL Search Order Hijacking	Password Filter DLL	System Information Discovery	Windows Admin Shares		Remote File Copy		
	Regsvcs/Regasm	External Remote Services	Parent PID Spoofing	DLL Side-Loading	Private Keys	System Network Configuration Discovery	Windows Remote Management		Standard Application Layer Protocol		
	Regsvr32	File System Permissions Weakness	Path Interception	Execution Guardrails	Securityd Memory	System Network Connections Discovery			Standard Cryptographic Protocol		
	Rundll32	Hidden Files and Directories	Plist Modification	Exploitation for Defense Evasion	Steal Web Session Cookie	System Owner/User Discovery			Standard Non-Application Layer Protocol		
	Scheduled Task	Hooking	Port Monitors	Extra Window Memory Injection	Two-Factor Authentication Interception	System Service Discovery			Uncommonly Used Port		
	Scripting	Hypervisor	PowerShell Profile	File and Directory Permissions Modification		System Time Discovery			Web Service		
	Service Execution	Image File Execution Options Injection	Process Injection	File Deletion		Virtualization/Sandbox Evasion					
	Signed Binary Proxy Execution	Kernel Modules and Extensions	Scheduled Task	File System Logical Offsets							
	Signed Script Proxy Execution	Launch Agent	Service Registry Permissions Weakness	Gatekeeper Bypass							
	Source	Launch Daemon	Setuid and Setgid	Group Policy Modification							
	Space after Filename	Launchctl	SID-History Injection	Hidden Files and Directories							
	Third-party Software	LC_LOAD_DYLIB Addition	Startup Items	Hidden Users							
	Trap	Local Job Scheduling	Sudo	Hidden Window							
	Trusted Developer Utilities	Login Item	Sudo Caching	HISTCONTROL							
	User Execution	Logon Scripts	Valid Accounts	Image File Execution Options Injection							
	Windows Management Instrumentation	LSASS Driver	Web Shell	Indicator Blocking							
	Windows Remote Management	Modify Existing Service		Indicator Removal from Tools							
	XSL Script Processing	Netsh Helper DLL		Indicator Removal on Host							
		New Service		Indirect Command Execution							
		Office Application Startup		Install Root Certificate							
		Path Interception		InstallUtil							
		Plist Modification		Launchctl							
		Port Knocking		LC_MAIN Hijacking							
		Port Monitors		Masquerading							
		PowerShell Profile		Modify Registry							
		Rc.common		Mshta							
		Re-opened Applications		Network Share Connection Removal							
		Redundant Access		NTFS File Attributes							
		Registry Run Keys / Startup Folder		Obfuscated Files or Information							
		Scheduled Task		Parent PID Spoofing							
		Screensaver		Plist Modification							
		Security Support Provider		Port Knocking							
		Server Software Component		Process Doppelgänging							
		Service Registry Permissions Weakness		Process Hollowing							
		Setuid and Setgid		Process Injection							
		Shortcut Modification		Redundant Access							
		SIP and Trust Provider Hijacking		Regsvcs/Regasm							
		Startup Items		Regsvr32							
		System Firmware		Rootkit							
		Systemd Service		Rundll32							
		Time Providers		Scripting							
		Trap		Signed Binary Proxy Execution							
		Valid Accounts		Signed Script Proxy Execution							
		Web Shell		SIP and Trust Provider Hijacking							
		Windows Management Instrumentation Event Subscription		Software Packing							
		Winlogon Helper DLL		Space after Filename							
				Template Injection							
				Timestomp							
				Trusted Developer Utilities							
				Valid Accounts							
				Virtualization/Sandbox Evasion							
				Web Service							
				XSL Script Processing							

MITRE ATT&CK®
Enterprise Framework

attack.mitre.org

MITRE | SOLVING PROBLEMS FOR A SAFER WORLD

Figure 1: MITRE ATT&CK Framework

As you can see, preventive measures alone are not enough to deal with adversaries; an organization needs to create a consistent prevention, detection, and response program. Establishing an offensive security program can help improve the security posture of your organization and identify weaknesses in its prevention, detection, and response to security incidents. You need to make sure that your security is effective and that your security staff has the appropriate skills to deal effectively with incidents. Though penetration testing is valuable, it is also limited in scope and cannot cover the protection needs of the organization, which is why your organization's defense system should include more than just penetration testing. Being focused on compliance and penetration testing only creates a false sense of security. Nevertheless, this holistic approach brings challenges from several perspectives.

Getting buy-in from top management will be a significant challenge you will face when setting up an offensive security program. A possible solution might be to propose a lightweight offensive penetration test on the applications/parts of the network chosen after a risk assessment prioritizing exercise. This helps identify if and where more investments would be useful for the organization. Starting from scratch might seem rather intimidating, but it's also a great opportunity. The results of this initial assessment provide the building blocks to improve the organization's security and can lead to an initial proposal for improving the 'quick wins' - preferably high risk items - based on the company's budget for initial security investment. The most likely scenario will involve starting with a one-person show (with specific external providers used for very specific tasks), and growing organically into a team by demonstrating its value and business impact. Without demonstrating the added value and the business impact of your work in security to the top management, you only have theoretical chances that you will transform it in a success story.

Your main goal is to continuously protect your organization. The only proactive pre-compromise tool available is Red Teaming. Red Teaming as a whole is a goal-based adversarial testing process. A Red Team assessment is ultimately a simulated attack effort that targets a defined set of goals; it uses the same tools, techniques, and methods that a real hacker would. In addition, by reference to penetration testing, Red Teaming is interactive, as it entails identifying attacks and back-and-forth interaction with the attackers, similar to real life scenarios. You want to find out if the organization can prevent those attacks from succeeding in the first place; detect those attacks if they do succeed; and respond to them, returning the organization into a state of normal operations. Only through a Red Team / Blue Team exercise can you find all these. Red Teaming is the sharpest weapon available to fight against threats and your organization needs to use it in order to consistently evaluate how good your security is, and to improve the organization's security as well as the security staff training.

Because today's attacks are complex, you need to build models that aim to anticipate the attacker's behaviors. Organizations today likely lack the ability and resources to defend against all threats. A zero-day threat can be a success factor of an attack because all organizations rely on signature-based detection mechanisms. Attempting to handle unknown threats without a systematic plan will fail. Incident handlers and response teams must have a methodology in place that enables them to respond to unknown or unidentified threats, thus protecting the critical assets and data that businesses rely on.

Understanding the problem is half the solution; therefore, you need a way to model threats against your environment. If you can understand all the different ways in which your organization can be attacked, you can design effective countermeasures. It is as simple as that: *the better you understand the threats, the better you can defend your environment*.

“Problems are nothing but wake-up calls for creativity” – Gerhard Gschwandtner

So, you need to put your creativity at work to describe and depict potential attacks, and build countermeasures to protect your environment. Threat modeling and drawing the “threat picture” using the attack/threat tree will help describe and depict potential threats.

Important benefits of threat modeling:

- Spots design flaws that traditional testing methods and code reviews might overlook;
- Evaluates new forms of attack that might not otherwise be considered;
- Models threats against the existing infrastructure and evaluate the potential to create damage;
- Evaluates which of the current countermeasures are likely to succeed or fail;
- Helps design proper remediation countermeasures in order to reduce threats;

Simply put, there are two sides: attackers and defenders. *The side that learns the fastest wins*. The only way to win this game is to find and fix what makes your organization vulnerable *sooner* than the attackers.

Even though this chapter focuses on cyber security, you should consider the entire attack surface, which is not always entirely digital. Some attack vectors might be non-technical. In hybrid attacks, attackers frequently leverage physical threat vectors in order to bypass digital controls. Therefore, you should not ignore the physical threat vectors of a potential attack that can start in the parking lot, or an open location where an employee finds and picks up or gets an USB stick, or when an intruder gets into the building together with one of the employees, by smuggling the access control.

The attack tree is a tool to explore vulnerabilities in a system, be it physical, digital or both (technical and non-technical). It is particularly suitable for analyzing a system’s security against malicious attackers. It puts the security expert in the shoes of an attacker to gain new insight into vulnerabilities of the system. It is a structured process to anticipate cyber-attacks and reveal the attack surface according to the attack goal analyzed.

The goal of building an attack tree is to explore attacks on a system and expose vulnerabilities. Therefore, the root (first node) of an attack tree is a goal an attacker would have (e.g. access customer data or disrupt the flow of business). After the root is set, the rest of the tree should be created by refining each node until the action in the node becomes trivial.

An attack tree consists of the following components:

- Root node – the goal of the attack and the starting point of the attack tree;
- OR node – a node of which ONLY ONE of its child nodes needs to be successful;
- AND node – a node of which ALL of its child nodes need to be successful;
- LEAF node – the activity performed by the attacker;

Nodes between the leaf nodes and the root node depict intermediate states or attacker sub-goals. The attacker may gain benefits at any level of the tree.

The following line of questioning can be used to start the most used form of an attack tree:

1) What would be the goal of an attack to our environment?
2) What areas can be attacked to reach the goal?
3) What attacks can be performed on these areas?
4) What steps are required to execute such an attack?
5) What alternative approaches are there to the step?

Let’s take an example in order to simplify the understanding of how to build an attack tree. You would like to analyze how an attacker can gain unauthorized physical access to one of our buildings. So, the goal of the attacked will be “Gaining Unauthorized Physical Access to Building”. In order to do this, there are some options for the attacker, represented in Figure 1.

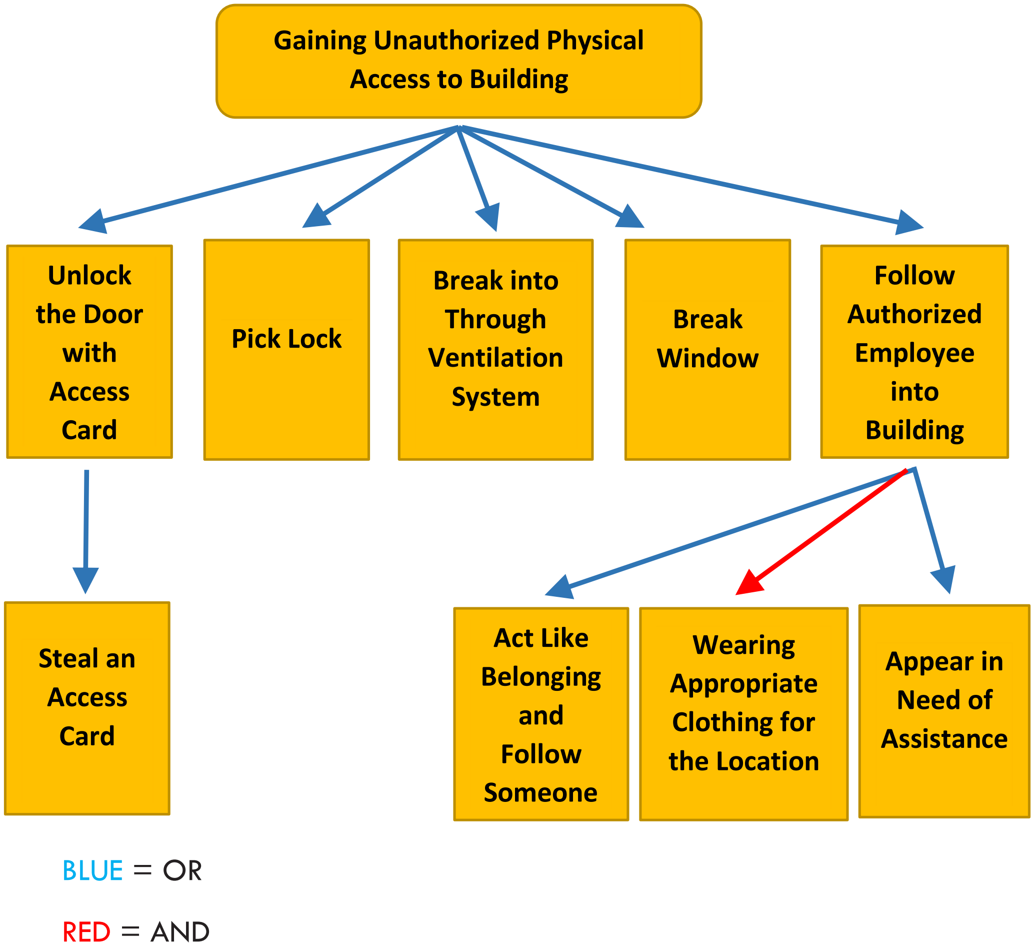
Gaining Unauthorized Physical Access to Building
Unlock the Door with Access Card
Pick Lock
Break into Through Ventilation System
Break Window
Follow Authorized Employee into Building
Steal an Access Card
Act Like Belonging and Follow Someone
Wearing Appropriate Clothing for the Location
Appear in Need of Assistance
BLUE = OR
RED = AND

Figure 2: Attack Tree Example - Gain Unauthorized Physical Access to Building

Now, what if you were to put the attacker's goal at the top of the kill chain and connect the seven phases of the chain to it? Will it look like an attack tree? Of course it will! And what if you were to also use the MITRE PRE-ATT&CK and MITRE ATT&AK matrices to depict the attacker's modus operandi? You'd have a priceless knowledge base you can rely on to build specific attack trees related to your environment. There is no reason not to use it. And don't forget to consider the physical threats along with digital threats when you build the attack trees.

2.8 OODA Loop

In certain circumstances, you might find difficult making decisions when trying to identify threats and define them as goals to build attack trees. The OODA loop might prove helpful in such situations. The OODA (Observe, Orient Decide, Act) loop is a method for dealing with uncertainty; it is a learning system and a decision-making framework that helps with orienting to situations and acting faster than the adversary can adapt. The OODA loop can be applied in threat hunting, as well as in other information security areas and in business. When the environment is volatile, and uncertainty is high, managers know that the same uncertainty facing them is facing their competitors too. Influencing and shaping the environment's uncertainty can provide them with the tools needed to create a competitive advantage.

Observe

To observe effectively, you need to have good situational awareness. Ideally, you should act and make decisions in a way that sets up your opponent and makes them vulnerable to having their rhythm broken. Therefore, you need to develop a clear understanding of your operational environment and context before making any new decision. For threat intelligence or threat hunting, observations include your situation, your opponent's situation and the environment more broadly - the physical, mental, and moral dimensions. In other words, this is the data collection phase – you should just aggregate what's available.

Orient

Orienting is the most important part of the OODA loop. The orienting phase aims to find mismatches: errors in your previous judgment or in the judgment of others. As a general rule, bad news is the best kind because, as long as you can catch it in time, you can turn it to your advantage.

For threat hunting, the goal of orienting should be identifying a set of possible threats that are relevant to the operational environment being analyzed, how they would present themselves in various possible scenarios, and what could be done to mitigate them. The information collected during the Observe phase is now analyzed, evaluated, and prioritized. If you feel uncertain, make sure you're devoting more time and resources to orienting.

The success of the overall OODA loop relies on creating models or concepts during this phase and trying to validate them before operation, in order to assure that you have the confidence that your models or concepts will work before you actually need to use them.

Decide

Invariably, this phase produces a hypothesis: the decision-maker predicts what the best course of action will be based on their understanding of the situation. Decision-makers should now be well-positioned to decide on the appropriate response. When deciding, you're essentially moving forward with your best hypothesis about which model(s) or concept(s) will work. To find out if your hypothesis is correct, you have to test it.

Act

This step is about testing the hypothesis generated in the decision phase. Action is how you find out if your models or concepts are correct. If they are, you win the battle; if they aren't, you need to start the OODA Loop again using your newly observed data.

Ideally, you should have multiple actions or experiments going on at the same time so that you can quickly discover the best model or concept for a particular situation.

Because the OODA loop is, after all, a loop, Act is never the last step of the process. Rather, the information gained from testing the hypothesis will be used during the next cycle of the OODA loop.

The power of the OODA loop comes from its simplicity. The one who runs successful consecutive OODA loops the fastest will win. Remember that "Orient" shapes the way you "Observe", the way you "Decide", and the way you "Act".

2.9 RED TEAMING

Now, more than ever before, robust Red Teams are needed to challenge emerging operational concepts and current security practices, in order to ultimately discover weaknesses before real adversaries do. Successful Red Teaming leads to robust decision-making, ameliorates risk, and helps the organization prepare for the unexpected.

How do we know that our investments in security are doing a good job? Red Teaming can give the answer to this question. Red Teaming is the only way we have available to measure how well our defenses will hold up to a real-world attack.

Red Teaming is a function that can compare and test approaches and plans, by considering a range of hypotheses or alternative outcomes in order to help the organization mitigate against potential attacks. One of the Red Team's key roles is to challenge your basic assumptions and to provide a different perspective on the assessment process and alternative views on adversaries.

Regardless of their individual skillsets, all Red Team members should have a full understanding of the problem analyzed and should ensure they are familiar with the relevant systems and processes your organization uses. The team should contain critical and creative thinkers who can approach the problem from different perspectives and can deal with complex systems and challenging constructs.

When forming a Red Team the following attributes should be considered and included as appropriate:

- The ability to see things from alternative perspectives;
- Imagination, a particularly desirable attribute, enabling freedom of thought;
- Self-awareness. 'Know thy enemy but not yourself, wallow in defeat every time' (Sun Tzu);
- Understanding of the operational environment, its critical variables and the decision-making process;
- Familiarity with cyberwar gaming and experimentation best practices;
- The confidence to challenge conventional or established Blue thinking;
- The ability to communicate effectively;

- Strong leadership;
- Effective facilitation;

Red Teaming is not easy; establishing an effective team and applying sound processes are challenges in themselves but are essential for the Red Team's success. The Red Team must:

- Have a clear objective;
- Be independent from Blue, but be close to the decision-making process and have adequate interaction with Blue;
- Contain critical and creative thinkers with relevant expertise;
- Have the full support of the top management;
- Help to detect possible deception and denial strategies by an adversary;
- Assist security team in understanding how much confidence to place in information and judgments derived from it;

Trust is crucial for Red Teaming. In the wrong hands, the information they handle can be deadly for your business.

Key Responsibilities of a Red Team

- Test the effectiveness of the organization's security programs and the performance of the internal security team;
- Improve the organization's ability to respond to real-world threats and incidents;
- Assess the internal security standards and practices in order to take needed steps to maximize the performance of the Blue Team;
- Identify and mitigate sophisticated security flaws before an attacker does;
- Use appropriate tactics to discover exploitable vulnerabilities, get access to the target, steal sensitive information, and, at the end, take proper measures to fix and improve the overall organization's security;
- Compiles detailed security assessment reports of discovered vulnerabilities and the measures taken to mitigate them;

- Checks the overall organization's security and creates a strategy to fix and enhance it;
- Educates Blue Team members and senior management to maintain and improve the organization's security;
- Complies with all laws, regulations, policies, programs and Rules of Engagement;

The following mind map depicts the main activities of a Red Team engagement:

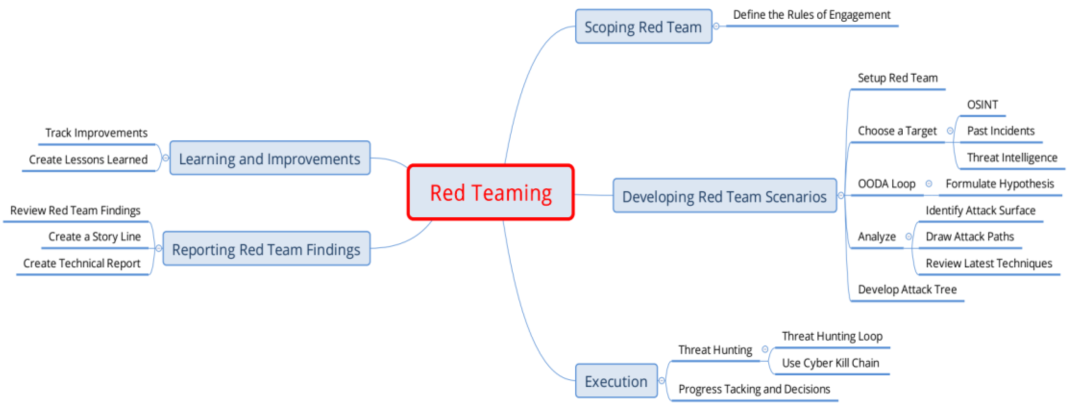

Figure 3: Red Teaming Mind Map

All organizations face a major decision when setting-up a Red Team: setting-up a Red Team with internal resources versus outsourcing the Red Team to get an independent perspective. Both approaches have pluses and minuses. However, a Red Team's effectiveness depends on their mastery of the needed skills. In most cases, you cannot find all required skills within your internal team and you will build a Red Team using both internal and external resources. There are two areas of expertise you should consider – technical and tactical. Technical means specialized expertise to build and run the tools needed during the Execution phase. Tactical means specialized expertise to threat model and to develop Red Team scenarios. By default, each Red Team member must be able to communicate effectively with the business.

Scoping the Red Team

A Red Team engagement aims to simulate a real-world attack by expanding the initial scope of the engagement to include the entire organization (in case of gaining access to the target). Therefore, the scope of a Red Team cannot be limited to specific systems. It should be scenario driven, with specific goals, based on real security threats. Those scenarios can be developed with to the input of other security teams, such as the Threat Intelligence team.

The Rules of Engagement for the Red Team should include:

- A list of goals to be achieved by the Red Team during the exercise

Some examples of goals might be:

- ▪ Obtaining physical access to a server room;
- ▪ Gaining access to an environment holding sensitive data;
- ▪ Taking control of a mobile device;
- ▪ Compromising the account credentials of a top manager;

- Specific techniques that are excluded from the engagement (if applicable);
- Specific areas or assets excluded from the scope and Red Teaming exercise (if applicable);
- The official testing period;
- References to the applicable legislation, policies, code of conduct & ethics, etc. (some service providers have their own set of guidelines for performing assessments; they also require proper permission to be obtained before performing any assessment);
- Communication and collaboration rules:

- ▪ Functional and operational escalation points the Red Team can use and in what circumstances;
- ▪ When to share the knowledge with the Blue Team members (if the organization’s intent is to test its response to a security event without prior warning, the Red Team will be allowed to share knowledge with the Blue Team at the end of the exercise);

- Incident response rules:
  - In case of critical vulnerabilities and exploits found during the engagement;
  - In case of emergency;

A letter of authorization should be prepared and provided to the Red Team for all on-site activities performed during the testing period. The opening of the Red Team engagement can be done only through a written authorization given by the organization's representative.

Developing Red Team Scenarios

Once the Red Team receives the written authorization, it will start the black-box assessment. The initial work done in black-box assessment is information gathering. The goal is to gather data on the target organization, and it is critical to the operation because this information is the basis for the development of the early plans for the attack. This should be regarded as initial reconnaissance to identify what potential targets are the most vulnerable, then to analyze the intelligence information and define the potential attack surface and develop the attack tree.

Execution

This phase involves the execution of the attack on the identified targets based on the attack plan and scenarios that are formulated in the previous phase. The attack execution phase should be closely monitored. All steps taken by the Red Team and their observations should be continuously captured in the Exercise Log Report with the level of details as defined by the Rules of Engagement.

Reporting Red Team Findings

At the end of the exercise, the first responsibility of the Red Team is, where possible, to remediate immediate issues found during the exercise, as well as eradicate any left-over attack tools and artifacts.

Then, the Red Team will prepare the Exercise Report, capturing all aspects of the exercise where the attack was detected, observed, reacted upon, tracked, contained and eradicated, or lost sight of. This reconciliation should be used to identify security controls, either missing or in need of improvement, that would otherwise have prevented or detected the attack.

For learning purposes, a joint post-attack exercise can be organized to enable a step-by-step replay of the attack for the Blue Team members learning benefit. Also, this post-attack exercise can be organized in the actual live environment, to demonstrate failed controls in real-time.

At the very end, the Red Team will prepare the Final Report of the exercise. The Final Report will include security strengths, comprehensive analysis of organizational capability, with recommendations for remediation and enhancements. This report will also contain the methodology, the evidence of goals achieved, the details of the attack paths undertaken and the concessions, if any, used.

In addition to such report, once the Red Team has completed their exercise, an in-depth debrief should occur with the Blue Team. During this debrief, the Red Team should describe the conclusions they have arrived at, successes, failures, as well as preventative measures and security controls they recommend.

Learning and Improvements

This phase should be under the lead of the Blue Team. The Blue Team should prepare the lessons learned of the exercise and implement the recommendations and enhancement proposed by the Red Team. At the end of the implementation a new overall security check would be required to assure that the risks are under control.

2.10 Conclusions

As already mentioned, Red Teaming is the only way we have available to measure how well our defenses will hold up to a real-world attack. Only through Red Teaming can you get the peace of mind you need to operate safely and maintain the security of our organization and its data.

3. Penetration Testing and Red Teaming – Blueprints for Your Fortress

The first section of this chapter discusses the stages an organization should go through when preparing for a penetration testing activity, along with practical approaches to take - especially in terms of prioritizing the applications to be tested based on risk assessments and mechanisms for addressing change management in terms of recurrence of testing.

The second section covers the main differences between Red Teaming and penetration testing, so that organizations can identify the best suited actions to take at a certain point in time or on specific IT assets. Furthermore, the section outlines the main principles to follow in Red Teaming exercise. The third section, outlines the main tools used in both Red Teaming and penetration testing exercises.

3.1 Penetration Testing – The Organization's Perspective

This section begins with describing the penetration testing concept, emphasizing on the situations in which it is useful. It continues with recommendations on establishing roles and responsibilities in your organization for this type of exercises. Next, the section outlines the strategy for prioritizing IT assets for which penetration testing will be performed, based on risk assessment results and the CIA (Confidentiality, Integrity, Availability) rating.

Following the risk assessment and the CIA rating exercise, the organization can establish scheduling for penetration testing and create a mechanism for addressing change management scenarios. Practical recommendations are also included in this section.

3.2 What is Penetration Testing?

This chapter focuses on one of the important processes at your disposal that help secure our applications and infrastructure, ensuring smooth and secure operations and, hopefully, peace of mind: penetration testing. As web technologies have become increasingly commonplace, so, too, has discussion around this topic and, sure enough, demand for professionals with working knowledge of the process.

So, what is, exactly, penetration testing? For those of you more familiarized with the broader IT concepts, you could think of it as the functional testing of security features of an application, infrastructure component, etc. The security requirements for new developments are similar to how, in the design phase of a typical infrastructure component or application, the project team gathers the functional requirements (i.e. what the application is intended for and what are its functions) before deployment. The project team needs to identify, at least at a high level, the security requirements that the new application needs to fulfill.

At the end of development/implementation, these will need to be tested to ensure they have been implemented and working correctly. Penetration testing is the process that ensures security measures have been implemented according to specifications and data, and that your whole network is ultimately secured from malicious activity.

This is not a one-off process. One cannot assume that if done once and passed with no notable issues (as there is no application 100% secured, just as there is no such thing as 0 risk), the application will be as secure in 3 years' time. The threat landscape is constantly evolving, new techniques of attack are constantly developed, and new vulnerabilities are discovered every day in all applications, frameworks, and even in security products.

Thus, maintaining a sound level of security for your data requires regular testing. Penetration testing should be a continuous process in which all applications - both new and existing - are included. Penetration testing should follow every application throughout its life-cycle.

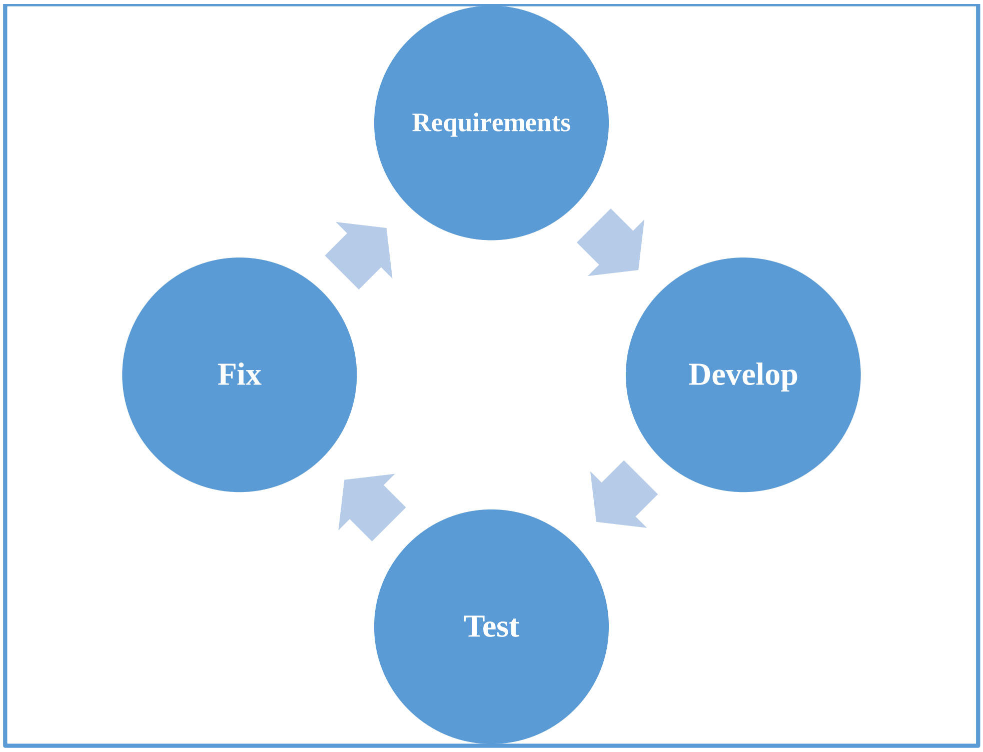

Figure 4: Penetration Test Process – High Level Overview

Then, there are the questions of what applications to test, when to test them, and what kind of testing fits which situation. The following sections will address each of them, along with many others, with the end goal of providing at least a baseline understanding of the penetration testing concept and how it can be implemented within an organization and embedded within its processes.

Understanding penetration testing requires a broader description of the whole process of ensuring application security – from requirements generation to the actual testing. These processes are interlinked, which is why they need to be treated together in a single framework, rather than discussing just about the actual process of testing – which is, in itself, a rather technical discussion; a detailed look at penetration testing is beyond the scope of this chapter.

Roles and Responsibilities

Just as with pretty much anything else, implementing a penetration testing process should start off with defining the roles and responsibilities. These most likely already exist in some form in the organization, as they would be part of a much broader implementation of an IT management framework, Information Security Management Framework, or IT Risk management framework.

CISO (Chief Information Security Officer) – overall responsible (and accountable) for defining the security strategy for the organization. Depending on the size of your organization, the CISO might have multiple teams under their command.

Security Manager – or, depending on size, Security Managers. The Security Manager is responsible with implementing the security strategy – developing and implementing projects aligned with the goals as defined in the security strategy. They might lead one or multiple teams of security professionals with different areas of expertise.

Application Owner – The person from the organization that is accountable for the use, maintenance, and security of a particular asset (application, server, database, etc.).

Penetration Test Engineer – technical person, part of the internal security team or a vendor's team, with knowledge and skills to perform security testing for various assets.

It is very important that these roles are defined within the organization - particularly the Application Owner; they are the person from the business that holds accountability of the application's good functioning and security.

As with all other aspects of security, getting buy-in from the business leadership is crucial in testing, but even more so in remedying any vulnerabilities found, as most of the time this will mean extra resources (either internal or external, usually time and) for development/configuration and retesting. Ultimately, these extra resources will impact the company bottom line, but so will any attacks that exploit the threats that your security team has yet to discover - usually more so than the investments in remedying vulnerabilities. This is why the owner needs to be educated on security matters and to understand the security gaps and the risks associated with not addressing these gaps, as well as the risks their organization faces. That way, they would become able to support the remediation efforts.

The overall penetration testing program relies on a risk-based approach, as you will not always be able to choose to address all identified issues, nor will you test all applications with same frequency. But more on that later in the chapter.

Defining the strategy

Having a risk management framework in place within the organization would help with implementing a penetration testing program; while it is not absolutely mandatory, it would definitely help, as you could map the pen test process to the framework and use the same risk rating system as with the rest of risk assessments being performed. This way, it would make much more sense and be more relevant to the business, especially when discussing about remediation efforts and costs versus risks.

3.3 Risk Management – A Quick Overview

There are plenty of risk management frameworks out there to choose from and implement. Those include the NIST SP 800-37 - Risk Management Framework, ISO27005 Risk Management Framework, and the ISACA Risk IT framework. Any of these would make a good choice for implementing risk management into your organization.

What is risk?

Risk is defined, in simple terms, as the possibility of something bad happening. Even though there are definitions out there that say that the existence of risk involves opportunities as well, the chapter will focus on the former meaning. As such, in the IT risk world, risk is defined as:

Risk = P x I where

P = probability of something happening (e.g. a threat occurring)

I = impact of that particular event.

This is what a graphical representation of the risk would look like:

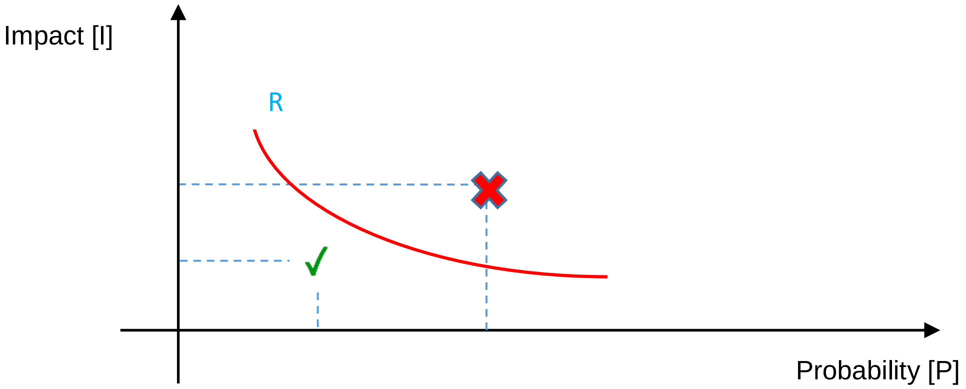

Figure 5: The Definition of Risk

By IT risks, we mean any event that could affect one or more of the three main characteristics of any IT system from a security perspective: confidentiality, integrity, and availability (CIA). Confidentiality is the property of something being secret, integrity refers to accurate, complete and unaltered data, and availability represents a system's ability to operate uninterrupted, thus offering uninterrupted access to data.

The risk management framework has 4 main components: risk identification, risk response, addressing the risk and monitoring the risk.

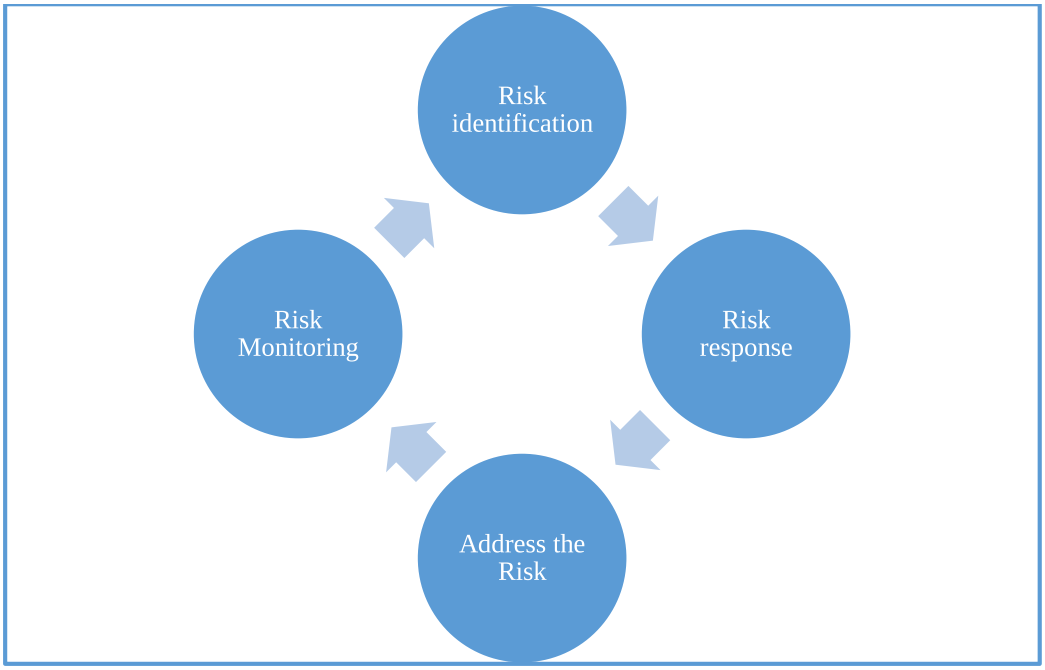

Figure 6: The Risk Management Process

A form of risk definition usually exists at organizational level and serves to identify and evaluate risk levels. Probably, the most prevalent one is the use of a risk matrix that will define the organization's view on risk. It could look like the one below:

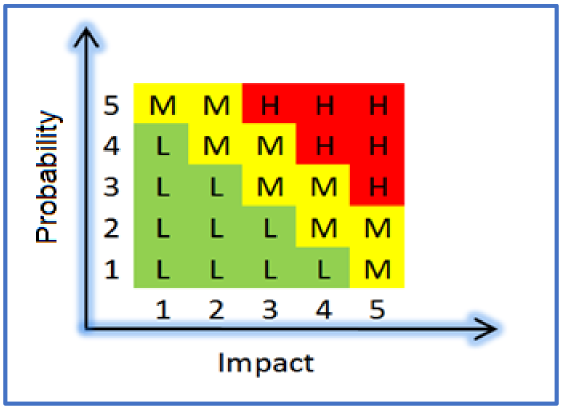

	1	2	3	4	5
5	M	M	H	H	H
4	L	M	M	H	H
3	L	L	M	M	H
2	L	L	L	M	M
1	L	L	L	L	M

Figure 7: Risk Matrix

Rating the Applications

The next important question comes down to what exactly to test. Not all applications have the same function, process the same kind of data, nor are accessible to the same audience. Some applications might be internal, some external. Some of them might be using confidential data, PII (personally identifiable data), or health data, while others might contain just some marketing data or even data that would be readily available publicly as well. This complicates things and might generate a lot of confusion as to what to test exactly, when to test, or how often to do it. This is where rating applications comes into play.

Rating the applications could very well be done as part of the overall IT risk management program – applications could be rated, for example, in terms of:

- Confidentiality – what is the level of secrecy of the data and, subsequently, how losing the data stored in those applications would impact the business?
- Integrity – how would altering the stored data in any way impact the business?
- Availability – what are the availability levels of these applications and, of course, how would losing the access to these applications impact the business?

Rating the confidentiality, integrity, and availability levels of an application could be done using simple notations indicating the level of importance of that data or application. For example, using levels from 1 to 4 - or from 1 to 10 if you want to be more granular, to indicate how important the application is when it comes to any of the three characteristics of the C-I-A triad. These levels could very well match the data classification scheme at the organizational level, as defined within the risk framework, or you could use a stand-alone rating system solely for this purpose.

However, using the company-wide data classification will ensure that everyone will understand their meaning when they see the ratings. An example of a data classification scheme that takes into consideration the confidentiality, integrity, and availability of data could look like this:

Area	Classification Level	Classification	Data property
Confidentiality	C-1	Public	Public data
	C-2	Internal	Data for internal use. Some damage if data lost/made public
	C-3	Confidential	Important internal data. Significant damage if lost/made public
	C-4	Secret	Data that, if made public, would significantly affect the business

Area	Classification Level	Classification	Data property
Integrity	I-1	Public	Data is publicly accessible. Alteration would not bring any damage or impact would be very limited
	I-2	Internal	Internal data. Alteration might have limited impact on people/departments using it.
	I-3	Restricted	Important data. Alteration of this data would cause damage to organization – financial, reputational, etc.
	I-4	Critical	Alteration of this data would bring serious damage to the organization.

Area	Classification Level	Classification	Data property
Availability	A-1	Non-critical	Low or very limited impact should the asset be unavailable
	A-2	Sensitive	Unavailability of this data would cause some impact to the organization
	A-3	Vital	Loss of availability would lead to significant damage for the organization
	A-4	Critical	Serious adverse effects for the organization should the data/asset be unavailable.

Applying this rating/classification system to all your applications through the process of impact assessment would lead to having a full view of what data you have and process, as well as to better understanding of the overall risk the organization faces(e.g. running a plethora of C4 applications makes it clear that the organization handles critical data, and that exposing or losing said data would have significant adverse effects on the business; possible outcomes include compliance or legal implications, bad publicity, or downright losing the business). Ideally, of course, the application ratings would be noted in a CMDB (configuration management database) and monitored.

As it is usually the case, applications will evolve with the business; thus, a recurrent review cycle is needed. A re-evaluation of the rating is deemed necessary once every 2-3 years. A reasonable interval needs to be chosen so that the re-evaluation interval is long enough that it does not put a significant burden on your staff while yielding the same results, but short enough to timely capture any significant changes. Usually, a re-evaluation once every 2-3 years would be an optimal interval for most organizations.

Some organizations, depending on their nature or size, might not have a fully developed risk management framework or a sound data classification scheme implemented across the board. In this case, one alternative option of rating your applications while also deciding how and when to test them is the OWASP ASVS scoring system. OWASP stands for Open Web Application Security Project. Per their website, OWASP "is a nonprofit foundation that works to improve the security of software. Through community-led open source software projects, hundreds of local chapters worldwide, tens of thousands of members, and leading educational and training conferences, the OWASP Foundation is the source for developers and technologists to secure the web."[4] The OWASP Foundation has a range of projects covering application security: including OWASP Top 10 – which covers the top web application security risks - and Mobile and Web application testing security guides.

The OWASP ASVS (Application Security Verification Standard) "provides a basis for testing web application technical security controls and also provides developers with a list of requirements for secure development."[5]

The OWASP ASVS is a standard intended to help developers build secure applications but also to "allow security service vendors, security tools vendors and consumers to align their requirements and offerings." The ASVS has multiple uses and may be looked at as a guide for security architecture, as an alternative to off-the-shelf secure coding checklists, and as a guide for testing. While not part of the chapter's scope, we encourage a deeper inspection of the whole standard; however, we will mention that the ASVS provides a rating system that may be applied for your applications. There are three security verification levels defined within the standard, each with varying degrees of complexity:

- ASVS Level 1 is for low assurance levels, and is completely penetration testable
- ASVS Level 2 is for applications that contain sensitive data, which requires protection and is the recommended level for most apps
- ASVS Level 3 is for the most critical applications - applications that perform high value transactions, contain sensitive medical data, or any application that requires the highest level of trust.

The standard provides a list of requirements but also gives guidance on testing and verification.

[4] https://owasp.org/, last accessed on 10 August 2021.

[5] https://owasp.org/www-project-application-security-verification-standard/, last accessed on 10 August 2021.

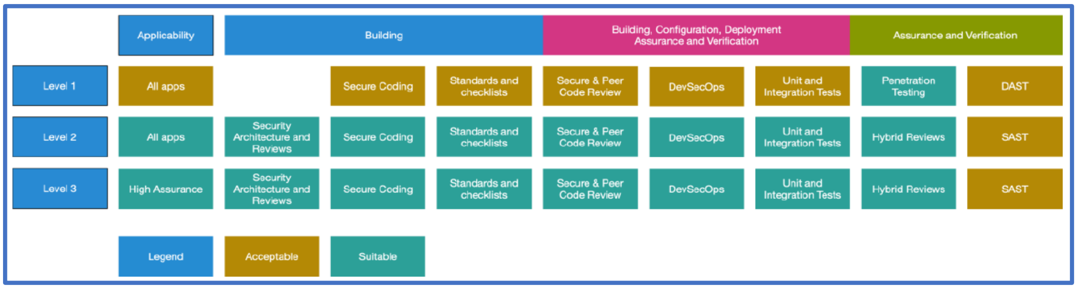

	Applicability	Building			Building, Configuration, Deployment Assurance and Verification			Assurance and Verification	
Level 1	All apps		Secure Coding	Standards and checklists	Secure & Peer Code Review	DevSecOps	Unit and Integration Tests	Penetration Testing	DAST
Level 2	All apps	Security Architecture and Reviews	Secure Coding	Standards and checklists	Secure & Peer Code Review	DevSecOps	Unit and Integration Tests	Hybrid Reviews	SAST
Level 3	High Assurance	Security Architecture and Reviews	Secure Coding	Standards and checklists	Secure & Peer Code Review	DevSecOps	Unit and Integration Tests	Hybrid Reviews	SAST
	Legend	Acceptable	Suitable						

Figure 8: OWASP ASVS Levels

Using the same approach, the applications may be given a rating/level, based on the data they process and level of risk they pose to the organization. The ASVS rating system covers a broader set of requirements and controls, but, at a very minimum, one could use it as guidance for rating the applications.

3.4 Defining Testing Schedules based on Application Ratings

Establishing the penetration testing program within your organization requires constantly evaluating and rating the applications it uses, regardless of the method used. Next, you need to identify how and when each of the applications should be tested. Some applications are more important than others. Some process personal client data, while other applications may contain only internal data or marketing data. Not all have the same level of importance to the business and, consequently, losing any of their individual C-I-A triad components will not have the same impact on the business as would losing those of another application.

Once all applications have been rated, either using the impact assessment method or ASVS scoring system presented or any other method, you can then just map these to a testing interval that is reasonable for the business: just as in the case of the assessment interval, it will be long enough to ensure there is no overburden on resources and that there is no unnecessary testing being performed, but short enough to capture potential changes. At the same time, new attack methods are continuously identified and published, as well as vulnerabilities for underlying operating systems, development frameworks, databases, etc. Given the complex nature of IT systems and the speed at which new vulnerabilities are identified the testing interval should not be too big. Compared with the assessment interval, where a reassessment once every 2-3 years would be enough to capture any significant/major changes, the actual testing should be done every 1-2 years, depending on the application, the data it processes, and its ratings.

Consider the impact assessment and scoring system presented. We will take each of the four possible levels (C1-C4) and map these to a testing interval. While creating the testing schedule, you also need to consider what type of testing should be done for each level. Taking all of this into account, a potential result could look something like this:

C1 – application contains public data or data that is readily available from other sources as well. Losing this data would not have any impact on the business. For these applications, a testing interval once every 2-3 years could be applied. It should, however, be no longer than the assessment interval. By matching testing with the assessment interval, you can ensure you capture any significant changes and test them. For example, if an application changes rating at next assessment and becomes a C2, then you can also test it immediately in the next period as it would have been tested anyway.

C2 – these applications might contain internal company data or any other data that is not readily available to the public. The impact is higher in case of security breaches, compromise, and lost data. These applications need a tighter testing interval- perhaps 1-2 years. In terms of testing methods, while for C1 applications you would go with just a dynamic black box testing, for C2 the option of dynamic testing and perhaps grey box or white box testing, where the testers have knowledge of the application design and can perform authenticated tests would be a choice. By performing authenticated tests and having knowledge of the architecture, more vulnerabilities could be identified that would not have been visible in the case of a black box testing.

C3 – applications rated as C3 contain internal sensitive information, perhaps also some combination of personal data, financial data, etc. These applications should be tested at least on an annual basis. One option is to use static testing, which may be done using automated code review tools.

C4 – the most critical applications for your business in terms of data they process. Depending on the area of activity, they may contain sensitive personal data, health data or financial data. They may contain sensitive business logic or trade secrets. Needless to say, should these applications be compromised, they would majorly impact the business. To match their importance, the interval for testing should be, perhaps, once every 6 months, but not rarer than yearly. All testing methods should be used: white box testing, dynamic testing as well as static testing – both automated and manual reviews, where possible.

The end result of how a testing schedule would look could be as shown:

Application rating	Assessment interval	Testing interval	Testing method
C1	2-3 years	2-3 years	Dynamic, black box
C2	2-3 years	1-2 years	Dynamic, white box
C3	2-3 years	1 year	Dynamic, Static
C4	2-3 years	6 months/1 year	All test methods

3.5 Handling Change Management

So far, the testing schedule has been defined considering that the applications remain the same; however, the business environment is changing constantly and, as a result, so are the applications.

Throughout their lifetime, applications evolve: new functionalities are added, some merge into one bigger application, or they may just collect additional data as part of the same type of processing – due to business needs or as demanded by regulatory requirements, etc. Regardless, applications do change, and this change, in turn, should be taken into account when evaluating their impact and establishing the testing interval.

Your framework needs to account for major changes as well. The change management process needs to account for major changes to applications and/or trigger either a reassessment of the application or a test before new functionalities are promoted into production. New functionalities such as processing additional personal data or adding a module for health data should definitely be evaluated and incorporated in the application rating accordingly. Testing ensures that the new features do not have significant vulnerabilities that could be exploited. In short, test every application whenever it undergoes changes to avoid missing any new vulnerabilities that come with the new features and/or with removing existing ones.

3.6 Adjusting the Testing Schedule – Balancing Business Needs and Security Requirements:

There are a lot of moving parts involved in the testing schedule, and they need to be managed in such a way that always maintains alignment with the changes/new projects schedule. The teams overseeing the security testing schedule should collaborate constantly with the teams overseeing new developments. A discussion on coming testing and new projects/functionalities would ideally take place once every 3 months for alignment. This way, the security team can reprioritize the testing based on the IT planning.

For example, one application is due to be tested in 2 months' time. But the development schedules estimates that significant new functionality to the application will be added and estimated to go in production in 3-4 months' time. The security team can use this information to postpone the testing of the application for an additional 2 months to ensure they cover the new functionality as well. In this case, even though the application test would be delayed for 2 months, and perhaps surpass the testing interval established, it would cover more ground and would be more relevant for the new version of the application.

There is no established formula on how and when to do the testing; this is something that the organization will optimize by doing. The example above shows that testing may be delayed with good reasoning, or, in the same manner, done faster. Sound reasoning and constant collaboration between the security team and rest of business is essential. As is almost always the case, there will be resource constraints, priority changes, etc. As long as the communication channels are open with the security team and decisions follow a risk-based approach that will ensure both efficient use of resources and security of your data and applications, the business has only to gain.

3.7 Red Teaming – Methods and Practices – The Organization's Perspective

The next section describes Red Teaming and reviews some of practices that should be considered during a Red Team engagement. It begins with illustrating the distinction between penetration testing and Red Teaming exercises. Next, it outlines the main principles to use in the Red Teaming exercise, from the perspective of the Red Team and from that of the organization, along with the main technical tools and techniques to consider in this case.

Differences between Penetration Testing and Red Teaming

The Red Team is brought in to assess the level of your security posture and/or the effectiveness of your security controls. Basically, the Red Team simulates an attack on your systems in a controlled fashion (preferably), in order to identify potential gaps.

Although you might be tempted to say that penetration testing and Red Teaming are the same (they are often used interchangeably), there are differences between the two:

Scope: Penetration testing is performed in order to identify potential cybersecurity vulnerabilities – application layer flaws, network or system level control gaps, or even physical security vulnerabilities. It's looking at your environment through the eyes of an attacker. The pen test is planned, usually known about, and the target is agreed upon prior to the assessment. Red Teaming, on the other hand, while at a high level has the same aim (identifying and exploiting vulnerabilities in an effort to gain access to systems), it differs in its approach. Red Team operations do not entail agreeing necessarily on a target – their job is to assess the level of security and try to obtain access to systems in any way possible. The approach entails, of course, "attacking" the organization from multiple angles, simultaneously.

Team members: The Red Teaming effort usually entails more people, time and resources compared to a pen test, which is usually done by 1-2 people in a limited time-frame.

Tools and techniques: The Red Team will probably use the same tools and techniques you would see a pen tester use. The difference lies mainly in the time, people, and resources available, not knowing the actual target, and, in general, in the approach.

A simple analogy for this is: “pen testers are pirates — ready to rampage and pillage wherever and whenever they can. In this analogy, Red Teamers would be more like ninjas, stealthily planning multi-faceted, controlled, focused attacks”[6].

3.8 Main Principles in Red Teaming

In terms of desirable practices for the Red Teaming effort, without getting into the technical aspects of testing, there are a few worth mentioning:

a) Plan in Detail

The team should take the time to detail the approach as much as possible: outline the approach, define the roles in the team, and create a general attack plan. The relevant persons from the organization that are aware of the assessment should ideally have common discussions with the Red Team and take into consideration their valuable input.

The discussions at this stage, of course, will not go into detail about the Red Team strategy and targets, but focus on the perimeter of the organization and on specific limitations to the scope of the Red Teaming exercise.

b) Document the Approach

The Red Team documents the target, the roles assigned to each team member, and their general approach (basically, document the results of the planning phase). Even though the target might not be specifically defined, it is a good idea to document the approach and the discussions with the organization about the attack methods used, the depth of the test and how far the Red Team would be allowed to go (e.g. prove that access may be obtained to a database versus actually extract/delete data from that database).

This documentation is useful for the lessons learnt part, in order for the organization to track the areas of improving its security.

c) Diversify

The Red Team should consider using a variety of tools and techniques in their assessment and not limit themselves to just one attack method or tool. They should use a variety of vulnerability scanners, penetration testing tools, web application attacks,

[6] https://www.redteamsecure.com/blog/penetration-testing-vs-red-teaming, last accessed on 10 August 2021.

frameworks, social engineering techniques, etc. Basically, the Red Team tries as much as possible to simulate a real-life attack. An attacker would usually exploit the same vulnerability in multiple ways and would try out as many techniques as possible to gain access. The same needs to happen with the Red Team.

This is aimed at testing the reaction of the organization's Blue Team under attack circumstances.

d) Document Findings

At the end of the engagement, the Red Team should document all findings in a detailed report:

What are the findings?

What were the vulnerabilities identified and how have they been exploited?

What is the risk level for the organization, given the deep understating the Red Team has gained of the organization and their environment?

What are the actions they need to take to remediate the findings and address the risks?

The Red Team's report should contain all of this information so that the organization can take a qualified decision on the next steps.

This step is closely tied to the continuous feedback given by the Red Team to the organization throughout the Red Teaming exercise. On the side of the organization, a similar report is usually prepared to identify the improvements that can be made in terms of detection, assessment, and response to potential attacks.

3.9 Tools and Platforms to use

Now let's look at some of the tools that might be used during a penetration test or a Red Team engagement. Generally, there is a combination of automated tools and expert analysis in each of these cases.

a) Vulnerability Scanners

These are tools that scan your systems (network devices, servers, workstations, dedicated appliances) and provide you with details related to vulnerabilities present on these systems. They often list their findings/the vulnerabilities using CVE identifiers that provide information on known weaknesses.

b) Proxies

These are tools that can be installed locally and intercept traffic between the host and destination. These may be used while communicating with web applications to modify responses to server requests or alter the inputs or requests from client to server. They can be very useful in testing for a variety of web application vulnerabilities.

c) Web Application Scanning Tools

These are tools that have been built for the sole purpose of testing web applications for vulnerabilities. They usually come with test templates out of the box but they may also be configured to perform specific tests.

d) Dedicated Pen Test Platforms/Frameworks

These are systems or collections of tools dedicated to scanning, identifying, and exploiting vulnerabilities. These may be in the form of full operating systems containing a suite of tools (e.g. Kali Linux) or dedicated tools (e.g. Metasploit)

e) Packet Sniffers

Also known as packet analyzers or network analyzers, these are pieces of hardware or software that are used to monitor network traffic. These may be used for traffic sniffing and analysis.

f) Password Crackers

Tools dedicated to cracking passwords. May come in use when, during the test, encrypted passwords are found.

g) Scripting and Programming Languages

Generally, the testing would not be based solely on automated test performed by dedicated tools. Manual testing is an important part of the process and team members would be acquainted with scripting languages – perl, bash, powershell or programming languages – HTML, javascript, Java, SQL, etc. Testers can make use of these to craft malicious input to applications, alter responses sent to server requests or build custom scripts to perform certain tasks, as needed.

3.10 Conclusions

When an organization is considering including penetration testing in its existing processes, this has to be correlated with the risk management process. This is essential, firstly, in order to analyze and address the IT systems with the highest risk (thus, the impact on the organization's business) and, secondly, to manage cost-benefit properly after having a holistic risk overview for a specific IT system or process.

This chapter describes a risk-based view of the Penetration testing and of the Red Teaming strategy and practices; the difference between the two and their principles and methods are explained.

4. The Penetration Testing Process - Hunting for Vulnerabilities

The penetration testing process, at a high level, is quite straightforward: the organization tests its assets, in an effort to identify weaknesses in the way they are set up or developed. The details of how to do this are, however, more complicated, as penetration testing requires a unique set of skills and knowledge that not everybody possesses - not to mention a rather special frame of mind: one needs to think like an attacker. While not being extremely technical, this chapter will present the penetration process, what it looks like, what it entails, what are its steps, and what are some of the tools available out there to get the job done.

The term “assets”, as opposed to “applications”, illustrates the penetration testing process’ breadth of scope more accurately, because attackers do not necessarily focus only on applications, and attacks from nefarious parties might come in different forms and target various components of the organization: applications, network components, physical locations, and even people. Some of these attacks are purely technical, targeting perhaps applications or network components, while other represent a combination of tools and techniques leveraging both the technical aspect and the human factor.

The pen testing process consists of a few stages, with each stage giving input to the next one. However, before going into details of the actual testing process. This chapter will outline some of the types of tests that are usually employed. The grouping of tests may be done based on knowledge of the tester on the target or, in the case of applications more specifically, based on the types of methods being used to test. Let’s take them one by one and understand what the defining characteristics are for each of them.

4.1 Knowledge of the Target

Penetration testing comes with some variety of approaches to implementation. Each approach has its benefits and drawbacks, of course. Here are the main types of testing based on knowledge of the targets:

- Black box testing
- Grey box testing
- White box testing

Black Box Testing

Black box testing means the person testing has minimal knowledge of the target details. They will have no knowledge of how their targets have been developed, configured, and deployed, nor do they have any knowledge of any internal processes or procedures. For example, in the case of a web application testing, there will be no knowledge on how the application was developed, no knowledge of the code, nor how it was deployed in the network. Or, perhaps, in the case of testing physical location, the tester will have no knowledge of the access mechanism employed, what controls are in place, nor of any of the processes supporting the physical access to buildings.

The biggest benefit in black box testing is that it basically simulates an attack when the attacker has no knowledge of the target. It is probably the closest thing you have to a real attack, and gives the opportunity to identify vulnerabilities just as an attacker would. You basically get the view that they would when targeting the organization (depending on the skills of the tester, of course). The drawback is that the results might not necessarily reveal all weaknesses.

Grey Box Testing

In grey box testing, some information related to the target is revealed to the tester. Partial knowledge of the internal workings of the target are known. To use the previous examples, in the case of a web application, parts of the code might be revealed to the tester, or how it is deployed in the infrastructure. In the case of a physical access testing, the tester might be shown the details of the physical controls deployed – what is being used, how it is being used. Basically, the idea of grey box testing is that when performing the test, you have partial knowledge of the target.

In grey box testing one could potentially reveal more vulnerabilities than in black box testing. As it is usually the case, pen testing endeavors have a limited time frame in which they are conducted – typically 1-2 weeks. This is not enough time for a black box type of approach to identify all weaknesses. But more information on the target would significantly help concentrate on areas of interest and speed up the process of vulnerability identification.

White Box Testing

White box testing is essentially the complete opposite of black box. While in black box testing there is no knowledge of the targets being testing, in white box testing you have complete knowledge of how the target operates, how it has been deployed but also complete knowledge of processes and procedures supporting the operation of it. In web application pen tests, there is full knowledge of the code and also of the infrastructure supporting the application. In physical pen-tests, there is full knowledge of physical controls used and processes and procedures supporting the physical access control in locations – what systems are used to monitor and grant access, what the process of obtaining access is, etc.

Based on all the information at hand, the tester can develop a more focused attack strategy. They may also identify a whole plethora of vulnerabilities that would have otherwise been overlooked. Just to give some more examples, in vulnerability scanning, this might mean an authenticated scan is used as opposed to unauthenticated scans used in black box testing, which usually leads to identification of much more vulnerabilities. Or, in case of web app scanning, besides knowing the code, the tester might also be granted access to the application to be able to test all the menus and pages of the application.

While this type of testing does give the pen tester all the information related to their targets and has the potential to reveal much more vulnerabilities, it does expose the organization to the risk of testers concentrating on obvious weaknesses, thus not behaving like real attackers and, potentially, miss vulnerabilities that an attacker with less knowledge of the system would identify.

These are the three main types of tests that would be encountered and they may be used for any type of testing – application testing, physical access tests or testing various infrastructure devices and controls such as firewalls, email gateways, proxies, VPN concentrators, etc. There are other types of tests as well but these are rather more focused on application penetration test. Nevertheless, we will cover them as it is important to understand their specifics.

4.2 Testing Approach

There are two major categories of testing in this area and they refer to how the application is being tested from a code analysis perspective – if the code is executed or not. The two categories are:

- Static testing
- Dynamic testing

Static Testing

Static testing, also referred to as SAST, entails testing the application code. This means that the application is not being executed in a runtime environment, but rather testers are looking at how it was developed in an effort to identify any vulnerabilities in the code. This can be done either manually or may be achieved with the help of various tools that help analyze the code. Given the amount of code an application might have, the use of static test tools is usually employed. There are various code analyzers available online that can help with the testing.

Dynamic Testing

Dynamic testing, or DAST, usually means testing the application in a run-time environment. Application is deployed and running and testers try to manipulate input to the application and analyze the responses in an effort to identify vulnerabilities. Dynamic application testing is, basically, a simulation of an attack on the application. As is the case with SAST, dynamic testing of an application could be done either manually but also by using a variety of tools to generate and manipulate input for applications. The next chapter will cover some of these, when discussing about the testing process.

In most testing scenarios you will find that a combination of automated and manual testing is used. Usually, testers will use different tools to automate part of the testing and help them identify areas that are more prone to vulnerabilities. Some of these applications already have built in testing procedures that they simply run against the application. Once areas of interest are identified, then the testers hone in on those areas and try to exploit these vulnerabilities by manual means which could mean writing bits of code/scripts to get the application to respond in a certain way.

4.3 The Process

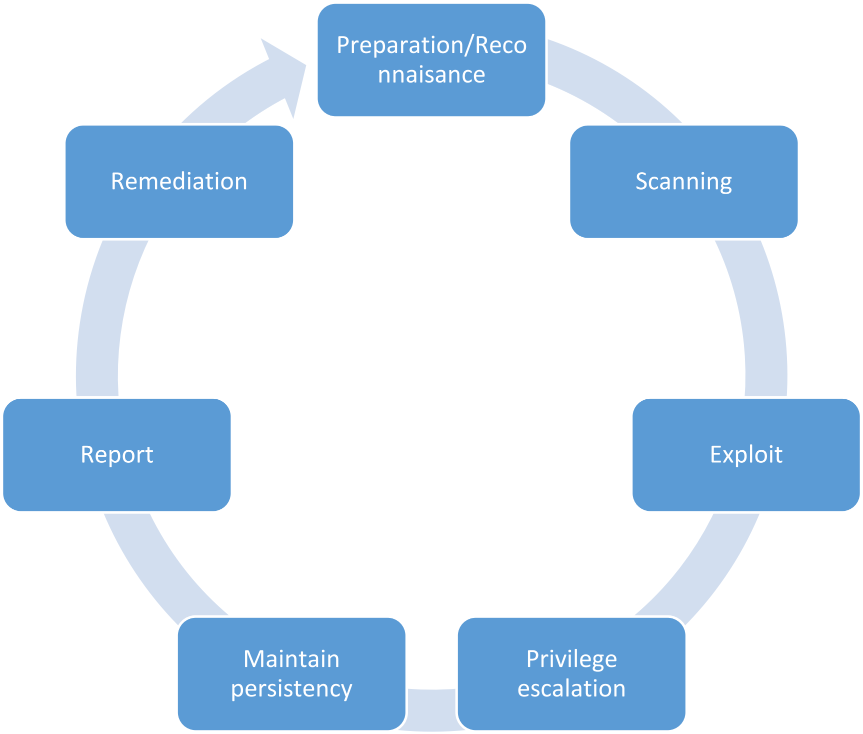

Figure 9: The Penetration Testing Process

There are a few well known and usually employed frameworks or methodologies when developing the actual penetration testing process and how you would want to do things with this respect. These frameworks have been developed by independent organizations or standards bodies, each with its own perspective on things.

OWASP

The OWASP framework addresses mostly web application scanning and security. The Open Web Application Security Project is an open project maintained by the community that helps organizations to identify vulnerabilities in their applications as well as secure them. It addresses web application and mobile applications security. There are numerous resources that could be used, such as the top 10 web and mobile application vulnerabilities, lists of controls as well as a testing framework – the Application Security Verification Standard.

NIST

The National Institute of Standards and Technology is a great resource when it comes to information security and risk management practices. They have created a breadth of standards and guidelines covering topics such as risk management, security controls and testing methodologies. In particular for this topic, the NIST special publications (SP) 800-53A and 800-115 may be of interest as they touch upon the building of an assessment program (800-53A) but also lay out instructions/technical guideline son how to test.

OSSTM

The Open Source Security Testing Methodology Manual provides a framework for testing and has tools useful for analysis and measurement of the test results – especially when organizations choose to become OSSTMM certified.

PTES

The Penetration Testing Execution Standard consists of 7 main sections covering the whole process from inception – reasoning behind a test – to reporting and presenting results to different stakeholders, including suggestions of what to include in your test reports.

Lockheed Martin's Kill Chains

Originally used as a military concept and more recently adapted to information security, it is intended to describe the structure of an attack[7]. It describes 7 attack phases which begin with reconnaissance and end with the "actions on objectives" phase, in which the attacker carries out his intended goal. The framework may be used as guideline for penetration testing as well as for implementing controls to prevent successful attacks.

Of course, various frameworks and organizations will have their own take on the penetration test process, even though there is a number of common elements to all of the above presented frameworks. The penetration test process has the following major stages: pre-engagement, preparation, scanning, exploitation, reporting and retesting. All these have a sub-set of steps that are to be fulfilled at each stage in the process. Each step produces results that are then fed into the next all the way up to final report preparation.

[7] https://en.wikipedia.org/wiki/Kill_chain, last accessed on 10 August 2021.

1) Pre-Engagement

What to test?

The pre-engagement phase is not necessarily part of the actual testing process, but it is closely related to it. At this initial stage, the security team, together with stakeholders – business owners, risk management, application owners, etc. – need to decide what exactly should be tested. This stage could very well be a one-off effort at the beginning of the year when, ideally, a testing schedule would be created for the whole year. In order to be able to do that, the security team will need input from various teams and decide, together with them, what is to be tested and when. Input for this endeavor might be – risk ratings of systems, importance to the business, scheduled changes or new developments (presumably we have a calendar for that as well). As it is most always the case, there will not be enough resources to test everything so a decision needs to be taken on what exactly will be tested. Just as an example, final list of assets to undergo a penetration test might be comprised of only high risk applications, or maybe by externally facing applications and devices (firewalls, WAFs, etc.).

Why and How to Test?

The rationale behind the testing will, ideally, be documented. This could take the form of a document signed by all stakeholders, or a testing schedule approved by all parties involved. This serves to identify and address vulnerabilities within the organization's infrastructure in an effort to improve the overall security posture.

How to test is equally important – will this be done by internal teams? By external third parties? What kind of testing will be required? Will it be a static or dynamic test? All these need to be decided before the actual test begins and, of course, should be documented.

From a tester's perspective, this stage is equally important. Knowing what they will test, when they will test will allow them to properly organize and plan for the testing period. Communication of the aforementioned test plan with the penetration testing team is mandatory. Having knowledge of the testing target, the complexity of the test and desired interval would allow them to properly assign resources and prepare the test.

As already mentioned, one can either to go through this exercise with an internal team or with an external provider of such services. While things may be easier in the case of the internal team, in the sense that all that would be needed is the communication of the test plan and agreement on the schedule, for the external provider things are a bit more complicated. One will need to go through the vendor selection process, negotiating terms, contract signing, etc. – which, depending on company processes, might take quite a bit of time. The provider will have to ensure they meet all the requirements and are legally permitted to engage in such activities.

It is of paramount importance for the tester to be backed up by contracts and formal agreements before actually starting the test. A test should never be started until all these are completed.

Once the pre-engagement stage is complete, everything has been decided and you have a tester or a testing team available, you're good to go for the actual penetration test. Next, you need-to-know what the test entails at every step of the way.

2) Preparation

In the preparation phase, the penetration tester will gather as much information as possible regarding the asset being evaluated. Of course, depending on the test type chosen, some information might already be available for them. Nonetheless, they will undergo a few actions the gather information for themselves.

3) Footprinting

This initial stage in which information about a computer system is gathered is known as footprinting or reconnaissance. Depending on whether the testing is done on an internal or externally facing app, testers will have different tools at their disposal. Still, most of them may be used in both cases. Testing can make use of simple tools such as ping and trace route commands to gather more information on the assets but also deploy other tools and techniques as well.

Other techniques that may be employed during this phase are:

- Network enumeration
- Port scans
- DNS queries
- WHOIS queries

- Operating system identification
- Google searches – for externally facing applications
- Spidering – in case of web applications – internal or external

All these techniques will give the tester more information on the asset under testing. Port scans may reveal certain open ports that could be exploited – either because of vulnerable services that have known vulnerabilities and, possibly, even exploits or because there are services/protocols in place that are just plain simply not secure – for example Telnet. The same goes for Operation system identification – you might think it is not much, but actually knowing the OS running, coupled with other results from footprinting, could give you a very detailed "map" of the vulnerabilities that might be present on the host. Not to mention if the OS is an older one. Spidering, or the use of a web crawler, can reveal pages on a web application that are not usually accessed by typical users but could be useful in building up an attack against the website – thinking here of possible error pages that do not filter out information and give much more details than they should (such as what technologies are used in the backend), or just giving our information that just should not be public.

4) Scanning

In the scanning phase, you take it one step further and perform vulnerability scans on the identified systems. You use the information gathered in the reconnaissance phase, such as IP addresses, hostnames, etc. and launch scans against those targets. This is done automatically through the use of vulnerability scanners, pointing them towards hosts, servers, network equipment, etc., depending on what the target is. This is also true for web applications. However, to some extent testing for vulnerabilities in web applications may be done manually as well – testing input fields in web pages, for example, by trying to pass input to the application that will determine it to process data in an unwanted way.

There are multiple types of scanners available online:

- Vulnerability scanners – these are usually dedicated tools that scan your target with the purpose of identifying known vulnerabilities present on it. They present their results in relation to the CVE identifier for the particular vulnerabilities they identify. These are complex tools that can actually do a plethora of tasks, such as scanning ports, identifying Operating systems and also test for known vulnerabilities.

- Port scanners – a somewhat simpler version of the vulnerability scanner, the port scanner does what its name implies – scans 1 or more ports to identify if these are open.

- Fuzzing – although not practically a scanner type of application, we consider fuzzing as part of the Scanning phase. Fuzzing is a technique of providing incorrect or unexpected input to an application/ a particular field and observe the responses the application returns to unexpected inputs. Usually, you are looking for error messages, crashes. Fuzzing may be done with dedicated tools, specially built for this or by manual means. They are typically used to detect potential buffer overflows, memory leaks, undefined behaviors, errors, etc. – all which pose a security risk for the applications.

5) Exploitation

The Scanning phase will yield various result to the tester – from operating system used to potential known vulnerabilities present, open ports and fields with issues and errors. Next phase of the process is to gather all this information and try to exploit the (potential) vulnerabilities identified. This may be done by using known attack techniques that have been proven to work in a particular case or the penetration testers can develop their own approach to exploitation.

Without going into the technical aspects, these are some of the techniques that might be employed:

- Code execution in applications

Fields that have been identified as having issues will be further explored. The penetration tester will generate pieces of code and will pass them to the application to run them. If the code was poorly written, then this additional chunk of code might get executed by the application, thus performing unwanted operations. Running additional code may be very dangerous, especially if the test is done in production (which we do not recommend), so extreme care must be taken when choosing what to run. For purpose of proof of concept, a function that would display something on the screen rather than executing operations on the database would be enough to showcase the presence of vulnerabilities.

- Malware delivery

The penetration tester might choose to deliver a malware either to people or to the system. Delivering malware to people might be done through phishing emails that could be distributed to all employees. This would test also the level of awareness within the organization. Another option of delivering malware would be through upload fields in applications. This would test if the input fields have been limited to specific expected types of files and also if the uploaded files may be run in the back. Basically, the test is to see if they could pass on malicious content to the organization and if that gets run on not. Of course, when we say "deliver malware" we are not suggesting to deliver actual malware that might wreak havoc in the organization. Rather, the penetration tester would make use of known test files, such as the EICAR malware test files, or craft their own that would perform a certain benign activity just to prove the existence of a vulnerability.

6) Privilege Escalation

This step is the peak of the penetration testing endeavor. Privilege escalation is actually "the act of exploiting a bug, design flaw or configuration oversight in an operating system or software application to gain elevated access to resources that are normally protected from an application or user."[8] The purpose of the penetration test process is to identify and prove the existence of vulnerabilities and obtain privileged access to systems. While the proof of the existence of vulnerabilities is definitely useful, obtaining privileged access to systems is the king for a penetration testing endeavor as it showcases that it is actually possible for an ill-intended party to do so (at least from my point of view).

The result of this stage is either an application that has more privileges than intended by the developers or a user with admin access that could perform unauthorized actions on systems/applications.

Privilege escalation could be split in two types/variants:

- Vertical privilege escalation – this would be the actual "privilege escalation". This involves a user/application with a lower level of access obtaining a higher level of access that was not intended for them. Such

[8] Privilege Escalation Wikipedia article - https://en.wikipedia.org/wiki/Privilege_escalation, last accessed on 10 August 2021.

an example would be a normal user gaining administrative access on his workstation or a normal user gaining admin access on the server.

- Horizontal privilege escalation – this entails the access of content and functions that is normally reserved for other users of the same level. One such example would be a user accessing another user's account in a web application.

An important note should be made here. Should a penetration tester find a really serious vulnerability and gain access to important accounts (such as domain admin accounts, admin accounts in applications, etc.) by exploiting it, this should be reported immediately to the organization, instead of through the final penetration test report (more on this shortly). The existence of such vulnerabilities with potential catastrophic outcomes for the organization need to be addressed immediately. Of course, the final report may contain details about this finding as well, but fixing such a gaping problem needs immediate attention. The organization can address this finding while the penetration tester moves on to testing other areas.

7) Report

Each penetration testing endeavor should end with a detailed report. The report, ideally, would be split into two parts:

- Management summary

- Detailed technical analysis

Management Summary

The management summary part should a 1 – 2 pages maximum report outlining in general, the findings, presenting the overall risk rating and a breakdown of vulnerabilities and their associated risk ratings together with a general opinion on the security stance of the organization based on the scope of testing and findings. The organization management is interested on the impact to the bottom line, so leave any irrelevant technical details out of this part.

Detailed Technical Analysis

This part of the report is presenting in depth the whole testing process. For each vulnerability identified, it should present the details of what tests were performed, as well as the steps performed to identify it. Moreover, print screens showcasing the finding should also be included. Give as much details as possible to that the technical teams can actually reproduce your findings for confirmation. Besides the technical details, each vulnerability should also have its associated risk rating mentioned. Ideally, each vulnerability would also have recommendations for remediation made by the penetration tester.

Determining the risk level for each vulnerability might be quite difficult, especially if the penetration tester is a third-party and might not be accustomed with the risk framework and risk rating system used internally; however, they can always do an initial estimation based on known best practices. The penetration tester may make use of several resources when rating the vulnerabilities identified:

- The CVE scoring system
- OWASP Top Ten web application vulnerabilities
- MITRE ATT&CK framework
- Own experience

8) Recommendations

The report must contain the penetration tester or team's recommendations on how to remediate each of the identified vulnerabilities. Recommendations may span from general, such as "patch the operating system" or "apply patch x for the application" to more specific items, such as details on how the configuration should be modified or what areas of the applications should be modified such that the vulnerability is addressed. Recommendations for remediation should be as specific as possible. This will depend, of course, on the knowledge and expertise of the tester/testing team.

4.4 When to Test

The penetration test process is cyclical; this accounts for how applications evolve over time, technology changes at a rapid pace, and new vulnerabilities are discovered every day. When adding the ever-evolving attack tools, tactics, and techniques the attackers use to the mix, penetration testing clearly is not just a one-off effort, but rather is cyclical and repetitive. Applications and the organization's infrastructure should undergo initial tests, retests as well as constant testing at specific intervals (established by the organization as per existing risk rating of the IT systems).

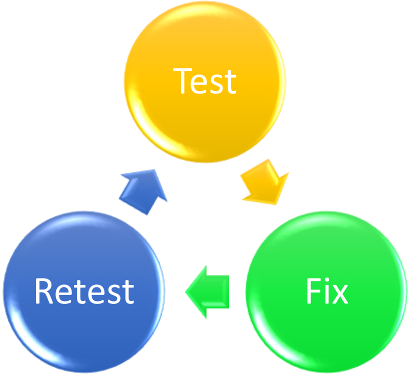

Figure 10: Penetration Testing Process

Initial Testing

The organization wants to implement a penetration test process. Where should it start? The first step is to establish what to test. Moreover, the retesting interval should also be defined. More on this in the previous chapter, but suffice to say that testing and prioritization of what needs to be tested may be done based on risk scores of applications and elements of their infrastructure. You and the organization can discuss the testing process, prioritization, and timeline, based on your experience.

For initial testing, there are two main areas of focus:

- Penetration tests for existing applications
- Penetration tests for new applications

Penetration Tests for Existing Applications

In the case of existing applications, testing can split further this into two types:

- Testing the application itself
- Testing the application in operational context

Testing the application itself is pretty much straightforward. Based on the risk rating (or any other method of prioritization the organization chooses), the organization selects the application that needs to be tested. The organization selects who will be testing, the type of test that the organization wants to perform, and the penetration test process will be the one described in this chapter.

For existing applications, however, there is another type of testing that should be performed – the application in operational context. This can very well be part of the same testing endeavor but must not be overlooked. The tests should cover the application's uses of various departments, any integrations it may have with other applications, essentially covering all the data flows to and especially from the application. Is the application sending data to another application (integration) or are there any processes that heavily rely on its outputs? These would also need to be tests to assess the real impact on the organization. Testing just the application itself would reveal important information, but it's just one part of the problem.

Penetration Tests for New Applications

In the case of new applications, the rule of thumb is always test before going into production with the application. The process usually follows the same steps outlined in this chapter: selection of tester/testing team, type of test, etc. The penetration test will need to happen after development has been finished and before the application goes into production to ensure that potential vulnerabilities are identified and addressed prior to it becoming operational.

Retesting

One important aspect of the whole process is the retesting of vulnerabilities. As previously mentioned, new applications need to have their vulnerabilities addressed before going live. The process is simple: the penetration tester issues an initial report outlining the findings, risk ratings, and recommendations, as already covered. The report is shared with the development team and internal teams. They will need to address the vulnerabilities before declaring the application production-ready.

Once vulnerabilities are fixed/addressed, a retest of the initial vulnerability presented in the report will need to be done to validate it has been addressed. This process of develop and retest, ideally, should be done until there are no vulnerabilities left open – or at least the ones that remain open are low, in which case the owner of the application may choose to accept those temporarily. At the very least, critical, high, and medium vulnerabilities identified should be addressed.

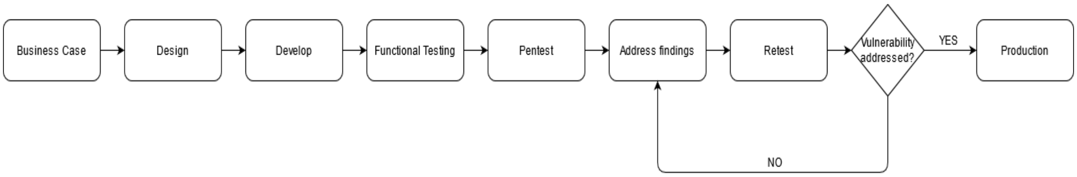

Figure 11: Retesting of Vulnerabilities

In the case of existing applications, retesting findings usually takes the same route – although there is a significant difference here: the application is already in use and in production, so there is no dependency in the retest to validate the fixing of vulnerabilities before it being available to users. Once the penetration test is performed and findings presented, the addressing of findings will usually take the form of an action list/projects list for the internal teams (developers, infrastructure, network, etc.). They will work on creating fixes/applying patches, etc. However, once each point is addressed, a retest will be needed to confirm that the initial vulnerability is no longer present. While, in the case of new applications, fixing the vulnerabilities would be done immediately, as there is also pressure to promote the app to production, for existing applications, the plan may span multiple months, depending on existing operational issues and tasks, resources available and importance of findings. Given the constraints, the list of findings should be prioritized and the most critical vulnerabilities addressed first and as soon as possible.

4.5 Where to Test

We have covered the types of testing, the details of the penetration test process as well as why and when to test. In the end, there is one important aspect to discuss and that is the testing environment – or where you test.

There is one rule: ideally, because of the potential disruptive nature of a penetration test engagement, these should never take place in the production environment as daily operations may be impacted which, in turn, may translate to loss of money or reputation for the organization (at the very least). Testing should be conducted in a test environment that is configured to replicate the live environment as much as possible. Ideally, it would be a 1-to-1 copy of the production environment, but this might be extremely costly to set up and maintain, so a close enough replica that would allow you to deploy the application being developed (if it is new) or that would already host the existing applications and simulate all connections it may have will do. The important thing is that you have a test environment where all functional and non-functional tests, including penetration testing, would be performed.

Just to mention a few of the characteristics for test environments:

- Replicates production environment – including operating systems used, development frameworks, databases
- Is isolated from production environment
- Typically, should not contain live, production data

4.6 Conclusions

This chapter describes what the penetration process entails, its steps and some of the tools available for the specific activities, from the angle of the person or organization that is responsible for execution. The knowledge about the target makes a difference in choosing the right approach for the testing. The reason for the pen testing mission has to be understood and the actual process must be properly prepared. Once executed, writing the reporting (including analyzing the findings and making recommendations) needs to deliver actionable information, addressing the vulnerabilities found, offering iterative retesting after remediation. The outcome should be an improved state of security.

5. The Blue Team's Role in Security

Though it wasn't officially called "Blue Team" to begin with, the concept of the Blue Team has been used in the past decade more often within organizations, generally to complement the Red Team. This chapter focuses on the concept of Blue Teaming, especially on practical methodologies and their implementation within the organization. The interaction between the Blue and the Red Team follows. Thus, this chapter is aimed at outlining the main directions that any organization can take in terms of preventing, identifying, and mitigating security incidents, together with practical examples for each step of the Blue Team framework.

The Blue Team encompasses all personnel within an organization tasked with identifying or responding to a security incident in the context of ensuring the security defense of the organization. The Blue Team has to undergo rigorous preparation in order to be able to identify and respond to a security incident, and help implement security measures as a result.

The role of the Blue Team can be placed in the organization depending on which specific steps from the sections below are implemented and on the type of IT systems/data that needs protection.[9] The Blue Team should collaborate closely with departments that can assist with the incident identification mechanisms (such as the operations team: networking, infrastructure, software development) and with incident investigation and management (compliance, legal, data protection officer, risk management). The roles of each department should be clearly established before an incident occurs and periodically revised and adjusted to the specifics of the organization and based on lessons learned.

To fully benefit from a Red Teaming exercise, the Red Team and the Blue Team should interact at certain points of the process, go through it and discuss the manner

[9] Susan Lincke, "*Security Planning: An Applied Approach*", Springer, 2016.

in which both teams have viewed the events, with the end goal of producing a report that the Blue Team can use for future improvement.

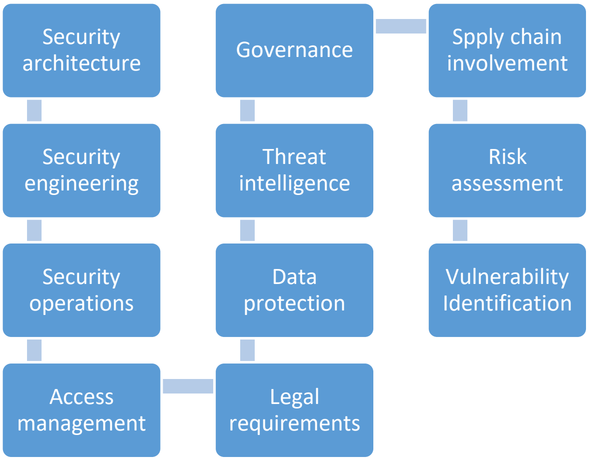

Figure 12: Elements of Blue Team approach

5.1 Methodologies and Practices for Implementing Blue Team

There are specific steps that can be implemented in order to assist the Blue Team in detecting potential threats or attacks, but also on the analysis and response side. Some steps are manual and entail experienced security professionals within or outside of the organization. In addition, certain steps can be automated either through configuration or specific automatic tools (especially ones based on machine learning).

In order to properly identify threats/attacks, the following three main elements are relevant: data (collected from the network, devices and servers), baselines (such as network baselines), and threat intelligence.

There is a series of best practice methodologies that address certain points that should be covered by the Blue Team,[10] including the ones mentioned below, NIST SP 800-61 - Computer Security Incident Handling Guide, NIST SP 800-86 - Guide to Integrating Forensic Techniques into Incident Response, NISTIR 7622 - Notional Supply

[10] https://www.bsigroup.com/en-GB/Cyber-Security/Managing-your-IT-and-cyber-security-incidents/Standards-for-managing-IT-security-incidents/ , last accessed on 14 March 2021.

Chain Risk Management Practices for Federal Information Systems and ISO/IEC 27035 information security incident management standard.

The European Central Bank published the TIBER-EU Framework, which includes guidelines for the two perspectives (red and Blue Teams): How to Implement the European Framework for Threat Intelligence-based Ethical Red Teaming and TIBER-EU White Team Guidance.

This chapter outlines a framework of the main points to be covered by the Blue Team, with reference to the relationship between Blue Team and Red Team.

Further, the actions of the Blue Team have to be viewed in the wider legal context of the organization, as there may be specific regulatory requirements in terms of identification, investigation and notification of authorities when an incident occurs, including in the banking sector, energy sector, for organizations under the NIS Directive and obligations under the data protection legislation. These legal obligations have to be integrated into the manner in which the Blue Team operates throughout the entire cycle of its activity.

The main steps the Blue Team can take[11], by reference to the activity of the Red Team, are detailed in the above figure. There are two different types of actions the Blue Team may take that have to be correlated and calibrated in order to ensure swift and proper response to any potential incidents.

On the one hand, the Blue Team has to build internally the hardening of systems, gathering of threat intelligence and implementation of proper security measures based on internal analysis. This is the passive phase of the Blue Team activity, which has to be properly performed before occurrence of an attack.

On the other hand, there are the actions that the Blue Team has to take in order to identify, investigate and respond to an attack. This is an active phase of the Blue Team activity that entails interaction with the threat actor.

The important aspect is that the two types of actions both have a cyclic approach, as there are lessons learned from the active phase that have to be implemented in both the active and passive phase and there are new threat/security measures identified in the passive phase that, once implemented, also improve the active phase response of the Blue Team.[12]

[11] Diogenes, Yuri, Ozkaya, Erdal, “Cybersecurity – Attack and Defense Strategies: Infrastructure security with Red Team and Blue Team tactics (English Edition)”.

[12] Alan J White, “Blue Team Field Manual (BTFM) (RTFM) Paperback”.

Generally, within Red Teaming exercises, it is assumed that the Blue Team has already gone through the passive phase and is currently engaged in the active phase.

This section focuses on the passive phase, with the active phase being detailed in the next section.

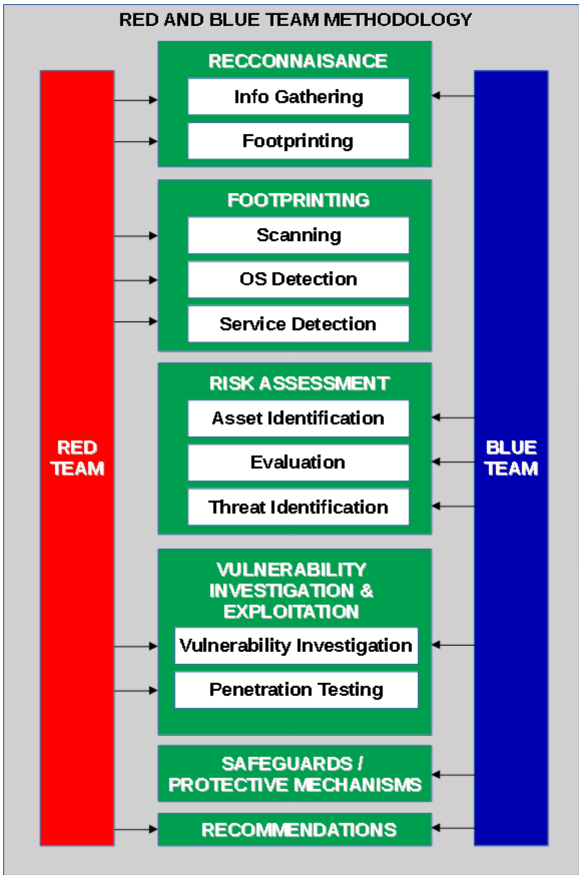

Figure 13: Red and Blue Team Methodology

Reconnaissance: During the reconnaissance phase, the Red Team is trying to identify as much information as possible about the organization and about potential vulnerabilities in a passive manner, by accessing various publicly available information and information gathered by other threat actors previously. The same type of information is targeted by the Blue Team in order to have an overview of the information that can be known by threat actors: from information made publicly available by its employees in social media to information sold by threat actors that targeted the organization previously. This type of information helps the Blue Team to figure out the manner in which threat actors are likely to try to attack the organization and focus on diminishing the chances of success for these vulnerabilities first.

Foot Printing and Risk Assessment: For the foot printing phase, the Red Team is performing scanning activities in order to enhance its knowledge about the organization and potential vulnerabilities that can be exploited.

The corresponding phase in the Blue Team focuses on risk assessment and threat hunting. This entails that the Blue Team takes into account the knowledge it gathered through threat intelligence, best practices, overview of the threat landscape and previous attacks on the organizations and identifies use cases concerning potential threats and evaluates their likelihood. This exercise is especially useful in order to identify the risks on which the organization should first focus and attempt to mitigate them through implementation of technical or organizational controls.

Prioritization should also take into account the data classification / CIA rating for the data stored in the IT systems in scope. If certain IT systems are maintained by third parties, an integration exercise for a coherent approach in terms of Blue Team tasks with these third parties has to be in place.

Thus, the two actions (of the Red and Blue Teams) have similar objectives, even if these are from different viewpoints.

Vulnerability Identification:

There are various manners in which vulnerabilities can be identified by the Blue Team, including vulnerability scanning[13], penetration testing, analysis of threat intelligence, for the IT environment managed by the organization or by entities in its supply chain.

Given the changing IT landscape and the changing attack types, such types of analyses should be performed periodically, with a periodicity established based on the threat landscape, risk rating, types of data in the IT systems and internal/external resources available.

The difficulty stems from identifying vulnerabilities in the supply chain. On the one hand, this can be identified through continuous monitoring of threat intelligence or monitoring/auditing of vendors. On the other hand, given the wide range of vendors, sub-contractors of vendors and tools/code of third parties used directly or

[13] https://github.com/rabobank-cdc/DeTTECT/wiki , last accessed on 14 March 2021.

indirectly by the organization, it is difficult to monitor or review all aspects of the IT systems and services offered by these entities.

In case of more regulated sectors, such as banking, insurance, energy or entities falling under the NIS Directive, additional steps should be taken to ensure vendors undertake obligations about implementing security requirements. It may be argued that certain steps also stem from GDPR requirements on state-of-the art security measures being implemented, as detailed under article 32 of the GDPR.

Safeguards and Protective Measures:

The main point to start from in this case is the design of a security architecture and implementation of secure coding practices. In case of existing IT landscape and legacy IT systems and networks, these can be adjusted in time and prioritizing on the most vulnerable parts, as these have been identified through risks assessment or vulnerability identification.

There are several approaches that can be taken in order to achieve this objective. In this section we are outlining the zero trust architecture, the ten design principles for defensible architecture mentioned in literature and the main directions for security controls.

Zero trust architecture has been gaining ground in terms of approach towards defendable systems and is based on the following main principles outlined by the NIST publication[14]:

- All data sources and computing services are considered resources: this ensures that each data location is properly hardened or has implemented adequate security measures.
- All communication is secured regardless of network location: all communication within the organization or external communications are properly secured.
- Access to individual enterprise resources is granted on a per-session basis: this is in line with the need-to-know principle and goes further in ensuring that the authentication and authorization takes place for each session.
- Access to resources is determined by dynamic policy: this entails that the state of client identity, client device and requested application/service are monitored in order to identify any abnormalities that can be an

[14] NIST SP 800-207, "Zero trust architecture".

indication of compromise. This analysis may include other behavioral and environmental attributes.

- The enterprise monitors and measures the integrity and security posture of all owned and associated assets: this entails that, on the one hand, there is a clear overview of existing software and hardware used within the organization and also that the relevant CIA ratings and security controls in place are constantly monitored for adequacy and identification of needed updates/changes.

The below steps should be implemented as a repeated set of steps in order to improve the security measures and security architecture, with periodic evaluation for improvement.

Proper training of relevant staff is essential for the Blue Team role. This training has to be in line with recent trends in attacks, investigation techniques and Blue Team management framework. Thus, they should address both the technical side and the governance side.[15]

The above are in line with the **10 design principles for defensible architecture:**

1. Assign the least privilege possible – limiting user access to the data/IT systems they need for the work tasks.

2. Separate responsibilities – this ensures that each department/person concentrates on specific tasks that work well together for identification, analysis and remediation of incidents.

3. Trust cautiously – implementing zero-trust approach within the organization and outside the organization.

4. Simplest solution possible – given the fast pace in which new threats appear, in order to be able to act swiftly and effectively, the organization should approach risks and vulnerabilities on a risk-based approach. Further, complex systems without a specific purpose for the complexity may be difficult to manage on the long run.

5. Audit sensitive events – a prioritization has to be established in terms of events (e.g. alerts, logs) in order to allocate appropriate resources.

[15] Luis Tello-Oquendo et al., "A Structured Approach to Guide the Development of Incident Management Capability for Security and Privacy".

6. Fail securely and use secure defaults – ensure that proper mechanisms are in place for data protection and business continuity in case an incident occurs.

7. Never rely upon obscurity – as discussed in specialty literature in cryptography, for all security aspects, the obscurity of IT systems or tools used should not account as a factor in the security approach (except, of course, for passwords, keys, etc.). Obscurity may bring an additional layer that the attacker has to overcome, but additional steps should be taken to ensure proper security measures are in place.

8. Implement defense in depth – in correlation with lack of obscurity, multiple layers of identification of incidents and prevention of incidents should be implemented. This ensures that it takes threat actor a longer period of time and additional skills to enter and compromise the system. In this manner, the organization has a higher probability of identifying the incident and/or prevent/remediate it before damages to IT systems occurs.

9. Never invent security technology – organizations should review the security products/services landscape in order to choose existing tools – either open-source or not. Creation of tools from scratch may prove expensive and time-consuming. Cooperation and use of tested solutions are the best approach in terms of security.

10. Find the weakest link – this entails thinking like an attacker and Red Teaming exercises are useful in this respect. An IT system or an infrastructure is as safe as its weakest link.

The controls implemented for protection of data and of IT systems have as main goal the detection and prevention of incidents. Periodical assessment of the efficiency of such controls has to be conducted, either in the risk assessment and/or separately.

Controls can cover a wide range of aspects. A list of proposed controls that can be tailored on the specific infrastructure of the organization can be found, for instance, in NIST's SP 800-53. The main points that should be covered are: checking on domain expirations, including email filters, threshold, and spam rules, implementing two-factor authentication, denying long relay request, application whitelisting, segmentation, managing keys securely, proper configuration and patch management and securing group policy settings. Further, through the architecture and implementation of IT systems, the aim of the organization should be to diminish the attack surface.

Further, security awareness training is essential as well,[16] given that the human factor has great influence on the success of certain types of attacks, especially based on phishing.

Recommendations: Recommendations are part of the Red/Blue Team exercise and involve the exchange of information between the teams in order to increase the security of the organization based on the vulnerabilities identified or exploited by the Red Team. As a best practice, it is recommended that this occurs throughout the Red/Blue Team exercise and not just at the end of the exercise. A step-by-step debriefing between the two teams can help the Blue Team with valuable information about each step of the exercise, instead of learning just a summary of the vulnerability exploited by the Red Team at the end of the exercise. For example, the manner in which the Red Team performs reconnaissance – e.g. the use of certain OSINT tools now known by the Blue Team – can help the Blue Team in real life scenarios by increasing the knowledgebase of the entire team.

5.2 Incident Handling by the Blue Team

For the purpose of incident handling, the Blue Team has to take a cyclic approach in order to improve each stage with each iteration of this cycle, either in Red Teaming exercises or in real-life attacks.

As in the case of the passive phase, in this active phase, the Blue Team has to work with various departments within the organization and with external entities. A non-exhaustive list is provided below. For the below steps of the Blue Team activity, it can interact with one or more of the below. For instance, for incident detection, Blue Team establishes the rules for detection based on discussions with other teams and, afterwards, receives the incident alerts from the incident response team. Further, depending on the structure of the organizations and the internal resources available, the below activities may be performed by multiple departments or by the same department. In the below, we refer to Blue Team as the team performing certain activities and also to the activity of the below departments.

[16] Joel Brynielsson, Ulrik Franke, and Stefan Varga, “Cyber Situational Awareness Testing”.

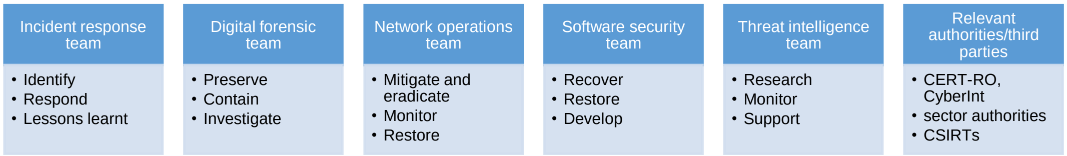

Incident response team	Digital forensic team	Network operations team	Software security team	Threat intelligence team	Relevant authorities/third parties
• Identify • Respond • Lessons learnt	• Preserve • Contain • Investigate	• Mitigate and eradicate • Monitor • Restore	• Recover • Restore • Develop	• Research • Monitor • Support	• CERT-RO, CyberInt • sector authorities • CSIRTs

Figure 14: Departments and Entities involved in Incident Handling

In terms of steps to be taken by the Blue Team, we have outlined below the main four phases of the incident handling: scoping (gathering of relevant information to ensure proper incident identification), identification and assessment (steps to be taken to identify incidents and to analyze them), remediation (identifying and applying initial response and remediation steps, as well as remediation steps to be considered for medium term) and lessons learned (identifying aspects that can be changed in the business process, IT environment or incident handling process to prevent similar incidents happening in the future or, at least, earlier detection of such incidents).

a. Scoping

The first step in incident handling entails determining main areas and mechanisms for identification of incidents.[17] Of course, in order to cover all aspects of security, physical, infrastructure, networking and IT system security angles should be covered. This means that the organization has to ensure it has an updated list of IT systems and infrastructure, which can be obtained through an internal framework outlining roles and responsibilities of departments to keep the list updated. Shadow IT generally exists, but the aim is to minimize or eliminate this through awareness and proper data collection and aggregation. This ties in with the passive phase actions taken, especially in terms of risk analysis, but also with further threat intelligence gathered constantly on the potential threat specific to the architecture and sector of the organization.[18] Nevertheless, for any detection solution chosen, after a proof of concept is made, analysis on the scaling up of the solution has to be analyzed, based

[17] https://blog.cyberint.com/threat-hunting-with-the-mitre-attck-framework , https://www.siriussecurity.nl/blog/2019/5/8/mapping-your-blue-team-to-mitre-attack , last accessed on 14 March 2021.

[18] https://digitalguardian.com/blog/threat-hunting-mitres-attck-framework-part-1 , last accessed on 14 March 2021.

on the existing infrastructure (and the types of IT systems and layout of networks) any extension plans.

In addition to the data from the IT systems stored/managed by the organization, the IT systems stored/managed by third parties should also be analyzed. The analysis of data can entail sending raw data from the IT systems maintained by third parties to the incident handling team or sending just the incident data. There are pros and cons for each situation. On the one hand, if only incident data is sent, this means that the organization relies on the third-party for incident identification. Further, analysis of only a limited amount of data about the organization's IT systems may lead to not detecting all incidents. On the other hand, if the third-party if managing IT systems for multiple entities, it can gather sufficient data to identify incidents more accurately. In case incident handling is left to the third-party (or its sub-contractors, as part of the supply chain), best practice entails that the organization establishes with the third-party the standards the latter follows to ensure proper incident handling framework. In certain instances, periodical monitoring or auditing may be useful to review that this framework is properly implemented.

With the extensive use of cloud systems, cloud providers, as provider, have to be involved in a constant cooperation with the organization in terms of threat identification. In this case, generally, cloud providers allow penetration tests to be performed on their systems, provided prior approval is obtained from them. Further, in case vulnerabilities are identified in their systems or in the supply chain used by their system, cloud providers have taken an active approach and, aside from notifying their clients about the incidents/vulnerabilities, have also issued appropriate mitigation plans and tools in a timely manner.

The lessons learned from previous incidents, aside from integration into the risk assessment, should also be included in the incident identification framework. The creation of this identification framework is also closely tied with threat hunting. Threat hunting exercises can be useful in identifying potential incidents or future incidents and should take into account the outcome of the risk assessment and a prioritization of threat intelligence gathered.

Moreover, analysis of historical data may prove useful, as new threat models have been developed since the first time that data was assessed. This entails intrusions not detected during the first analysis may be detected during a subsequent analysis. This is especially useful when the intrusions are still active in the organization environment and steps can be taken to address them.

The aim of such extensive analysis is to identify and document relations between events that have occurred in relation to the organization or to other organizations in order to better understand the threat landscape. Generally, more data can lead to useful information, if properly analyzed. Determining relevant rules or pattern identification can prove sometimes difficult in practice. Approaches that can be taken by companies include using third parties for the analysis, establishing rules by experience or using machine learning approaches.

In this case, legal requirements for log retention and sharing should be observed. A balance has to be made between the legal requirements to delete data and the practical need for data for analysis. Further, both internally and externally, the organization has to implement the need-to-know principle, allowing access to data only for individuals that need to access such data to fulfil their job tasks and balancing need for data with legal requirements to delete data. Access to data – need-to-know basis. When sending response process instructions to operations team, disclose only needed information.

Further, the organization can explore the use of honeypots for both defense (to deflect threat actors from the actual IT systems/infrastructure used by the organization) and for threat information gathering (obtaining information about existing threat that are targeting the organization). The honeypots are generally placed within the perimeter (behind the firewalls), but, for specific reasons, they can be placed before the firewalls. Honeypots can be implemented and analyzed third parties or by the internal teams.

In case of a Red Teaming exercise, during the various stages of the exercise and especially after the reconnaissance and recommendations phases, the Red Team can provide useful information in this respect.[19] If scoping suggestions are made during the scoping phase of the Blue Team, the Red Teaming exercise can prove useful for the Blue Team by experiencing in practice new types of detection.

The data quality is essential for accurate incident detection. The data quality should be monitored and periodically adjusted to reflect the relevant data for detection. The main characteristics of data include:

- Accuracy – integrity of data and lack of errors in data collection and transmission.
- Completeness – the analyzed data should be complete in terms of timing and data sources.

[19] David Mugisha, "Cyber Security: Improving Cyber Defense Through Coherent Joint Red Team and Blue Team", Journal of Defense Modeling & Simulation, 2019.

- Consistency – the data received from various entry points should not be contradictory.
- Timeliness – data should be up-to-date and proper historical data should be available.

The goal for maintaining data quality in incident analysis is to reduce the time it takes to fixing, validating and correlating data, but it is also useful in order to be able to rely on data from multiple sources during the analysis and enhance use of automation in the incident identification process.

Aside from the practical commercial aspects, use of an external SOC/CSIRT, there are certain legal points to consider[20] (such as, commercial secret sharing with third parties, ensuring timing responses and notifications from the third parties, ensuring third parties comply with legal requirements applicable to the organization in terms of incident handling).

Generally, the data anonymisation cannot be achieved in such situations, as the SOC and CSIRT teams need to have access to the entire infrastructure relevant for the role they are playing in the Blue Team. Nevertheless, data minimization and need-to-know principles should be implemented in this type of exercise. In addition, the retention period and deletion should be implemented.

The scoping phase is closely tied with business continuity and disaster recovery, as it represents the trigger for setting in motion the steps necessary to keep the activity of the organization going until the incident is properly investigated, contained and remediation are in frameworks, including ones from NIST and ISO (ISO 22301 Business continuity management systems requirements, ISO 22313 Business continuity management systems guidance).

b. Discovery and Assessment

For the discovery step, the organization has to ensure that proper data is collected from the IT systems (network, hosts, servers) in order for the identification methodologies to be applied to this set of collected data.[21] The assessment of the data collected is essential in determining the main direction for further investigation,

[20] The Forrester Wave™: "Enterprise Detection And Response", Q1 2020, https://reprints.forrester.com/#/assets/2/482/RES146957/reports, last accessed on 14 March 2021.

[21] Don Murdoch, "Blue Team Handbook: Incident Response Edition: A condensed field guide for the Cyber Security Incident Responder. 2nd Edition".

data collection and assessment.[22] This step generally is implemented as a cyclical step, as it may take several iterations until the relevant data is obtained and the relevant types of threat / attacks are searched for.

- **Hypothesis driven investigations** – when the Blue Team has indications that threat actors may target specific organizations or specific vulnerabilities. This can be achieved through threat intelligence gathering or through sharing of existing attacks within the sector/national/international community.

- **Known indicators of compromise (IOC) or indicators of attack (IOA)** – given the current types of attacks, the Blue Team identifies within the organization's infrastructure the preliminary steps in the cyber kill chain that were used in previous attacks. Depending on the step in the cyber kill chain, this may give the Blue Team a heads-up that can help in reducing the footprint of the attack or its consequences. Without knowing IOCs or IOAs from other previous attacks and scanning the infrastructure for these, attacks can go undetected for large periods of time.

- **Advanced analytics and machine learning** – these can be used to identify potential anomalies in the fingerprinting of a system, device or in the traffic of a network or towards/from a server/database. This step entails choosing the appropriate algorithm for each situation. In certain cases clustering may help identify the areas of concern and in other fuzzy logic may prove more useful. This can be decided and implemented on a case-by-case basis, as the Blue Team considers most fit for delivery of swift and useful information in order to allow time for the Blue Team to react to the potential threat/incident.

According to Bianco's Pyramid of Pain[23], in order to slow down or stop a threat actors, relevant information should be obtained about the attack approach and mechanisms, preferably early on the cyber kill chain. Thus, with a low value in terms of detection, the organization may consider binary hash values, IP addresses and domain names. These are easy to change by the threat actor once they are compromised and are also generally shared within the security community and are embedded in certain security solutions. The more valuable information that can be gathered is network/host artefacts, tools and TTPs (Tactics, Techniques and Procedures). These are very particular for the attack and/or for the threat actor and

[22] https://www.us-cert.gov/CISA-Cyber-Incident-Scoring-System , last accessed on 14 March 2021.

[23] http://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html , last accessed on 14 March 2021.

represent valuable information that allows organizations to stop the attack at its early stages. General tools that can be used for the Blue Team responses to the cyber kill chain include the following, mapped to the response steps.

For detection, there are various methods that can be used, as detailed above and may include: web analytics, NIDS, HIDS, audit logs, SIEM, vigilant users.

In order to deny an attack throughout the cyber kill chain, various tools can be used, including NIPS, proxy filter, firewall ACL, patches, outbound ACL.

Similarly, for disrupting the cyber kill chain may be performed by various tools, including NIPS, DLP, DEP, Inline AV.

Deceiving is usually achieved through DNS redirection or honeypots. Degrading the attack is usually performed through queuing, tarpit, limitation quality of service.

Containment may be achieved through various methods: trust zones, app-aware firewall, EPP, inter-zone NIPS, firewall ACL. This generally relates to the implementation of zero trust architecture.

In both Red Teaming exercises and real life scenarios,[24] once information about the environment being breached has been detected, there are a series of steps to be taken in a specific order:

- Alert the appropriate persons within the organization, as per internal procedures. This may include the Blue Team, incident response, legal and risk departments, board of directors. Generally, a roles and responsibilities matrix with thresholds sets-out the specific situations in which each department is notified and the input needed from each department. In addition, role play exercises should be performed prior to incidents in order for all individuals from each department to be aware of their role and for all department to work swiftly together.

- From an organization perspective, it is essential to implement at this stage the internal procedures in place for notification of relevant authorities in case of incidents. This type of roles has to already be included in internal procedures or instructions and rehearsed by the team members prior to an incident occurring. Generally, there are multiple authorities (and, in certain cases, affected individuals/client) that have to be notified in a particular case, but viewed from different angles.

[24] Don Murdoch, "Blue Team Handbook: SOC, SIEM, and Threat Hunting (V1.02): A Condensed Guide for the Security Operations Team and Threat Hunter Paperback".

- The relevant persons within the Blue Team decide the containment and context analysis to be performed. At this stage, external parties may be called-in to assist on specific investigation points.

- Throughout the process, the identification of relevant evidence concerning the incident and proper preservation thereof is necessary.

- After analysis is completed or, even, in some cases, after partial analysis is completed, a remediation plan is discussed. Further on the remediation steps in the next phase of the incident handling.

During this phase, the role of third parties (such as vendors) involved in the IT systems affected by the incident is essential. This entails prior contractual provisions are in place to outline the involvement of third parties: allocation of third-party team members, access to third-party logs or documentation, response times for questions and root cause reports, mitigation steps to be taken by third parties. This type of additional services may have an impact on the workload of the third-party and on the cost of the contract.

The communication of threat information to other entities within the same sector of activity or other communities should be performed in an anonymous manner, without giving valuable commercial/architecture/personal data information. One common information model used widely is the STIX model.

The secondary aim of the discovery phase is to limit the false positives, while analyzing the events in a more complex context in order to be able to identify patterns for threat/attacks.

For this phase, specific metrics can be created in order to improve the process, such as estimated time to detection or estimated time to recovery. This correlates to metrics of the Red Team, such as mean time to compromise, mean time to escalation, mean time to detection.

c. Response Process Development

The response development should identify the steps to be taken to repair the affected systems, to eradicate the part of the intrusion that is still in the network/IT system (e.g. reinstalling applications/OS, using back-up version of application before indications of compromise existed in it, eradicating viruses).

Building on the defense in depth principle, additional security measures identified during the incident assessment can help in case of future incidents in ensuring that the parts of the defense in depth structure that had been damaged are identified and repaired swiftly and, further, that the actual data/IT systems are not compromised.

The remediation has to be performed based on the CIA rating of the application/IT system/network involved in the incident and based on the level at which the incident occurred (e.g. at the level of the firewall, within the email server, within the application server through privilege escalation from an employee laptop).

In certain cases, correlation with other affected entities or instructions from authorities or from vendors may also be needed.

In terms of implementation of remedial steps, the relevant IT/business owners of the process/IT system should be involved in the process and should agree on the budget and timeline. In certain cases, the buy-in of the board of directors may be needed.

The remediation steps may change before they are implemented or afterwards. Thus, as in the case of controls, they should be periodically reviewed to ensure that they represent the most efficient and effective manner to prevent the occurrence of future similar incidents.

Follow-ups should be performed in order to ensure that remediation steps have been implemented properly both internally and, if needed, by third parties.

d. Reporting and Lessons Learned

The reporting of incident analysis results should be made internally, within the security team, before the relevant internal stakeholders (including the board of directors, risk management), but also before authorities, if this is required under the law. Further, from a public relations perspective, constant reporting to the public may also be useful.

The lessons learned step after an incident occurrence is useful in order to identify internal processes within the organization that should be improved either for identifying incidents, increasing security prevention measures, improving employee awareness, updating incident handling procedures or updating agreements with third parties providing IT services to the organization.

In the case of a Red Teaming exercise, this is the phase of recommendations that entails discussions about the steps taken by the Red Team.

The knowledge base for incident response can also be improved by the lessons learned in the incident handling exercise.

5.3 Conclusions

Blue Team entails a cyclic activity to be performed continuously by the organization. The actions to be taken start by setting up the prevention measures and mechanisms. They continue with identification and assessment of the vulnerabilities identified, together with relevant response and recovery steps. The last action includes lessons learnt in order to improve the Blue Team process.

6. Main Steps in Setting-Up a SOC Team

Day by day, security threats are evolving in complexity and diversity. Right now, no matter what business you are in, no matter what size your organization has, no matter where you are located, your business is at risk. And this risk increases day by day. We already are aware that, in our age, information security prevention is no longer optional, and that security, and cyber defense in particular, should be high priorities.

The key to cyber defense today is to develop an organizational structure that continuously evolves in order to counter advanced attacks. This organizational structure is called a Security Operations Center (SOC), and relies on skilled security specialists, technology, and processes. The SOC's goals are to prevent; monitor and detect; investigate and respond to all types of cyber threats.

A Security Operations Center (SOC) is a functional unit specifically built to be the first line of defense for your organization. Therefore, its area of operation (AO) must be very well defined and should cover all equipment and applications used by the organization. This is the only way to assure a proper and effective protection against threats. SOCs are led by a SOC Manager or a Chief Security Officer.

SOCs perform the following key tasks:

- Proactive Network Maintenance and Monitoring
- Assessment of the Network Security
- Incident Analysis
- Threat Intelligence

- Information Assurance
- Information Security Compliance
- Security Governance
- SOC Support

One of the key elements of a SOC is a Security Information and Event Management (SIEM), a software solution that aggregates and correlates data from different security feeds across the monitored infrastructure. A SIEM performs proactive monitoring and analysis, and provides event correlation, alerting, and data visualization. Through these functions, it helps the organization detect and mitigate threats. Even organizations with limited resources can implement a SIEM, as there are many open source solutions available.

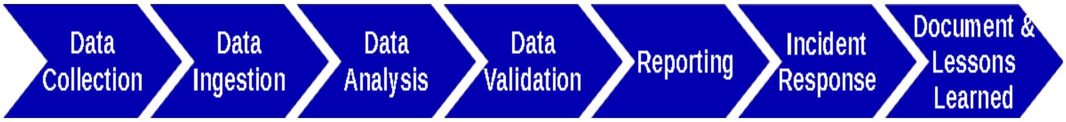

Figure 15: Typical workflow for Security Operations Center

- **Data Collection** – Log data are collected from various sources / security feeds (devices, applications, etc.) on the network and sent to the SIEM;
- **Data Ingestion** – Collected data together with threat and contextual data are ingested into SIEM to easily standardize casing of all fields parsed and produce security alerts;
- **Data Analysis** – Alerts generated by SIEM are reviewed and evaluated on their urgency and relevancy;
- **Data Validation** – Alerts triage is performed and incidents are validated;
- **Reporting** – Validated incidents are escalated to response team through the ticketing system;
- **Incident Response** – Incident Response Team reviews incidents and performs incident response activities;

- **Document & Lessons Learned** – Document incident for audit purposes and lessons learned;

Without proper visibility and control over the entire infrastructure, blind spots can form in the network security posture. These blind spots represent a weakness in your defenses and may be subject to exploitation or entry points for ill-intended parties. This is why the SOC's goal is to gain a complete view of the business's threat landscape, including third-party services and traffic flowing between these assets.

Your SOC team is crucial regardless of what technology and applications you might implement as part of your SOC. Their level of expertise will determine how fast they will detect a threat and how fast they will identify, build, and apply a response. Considering the sensitivity of their activity, all SOC team members must be well trained, as their knowhow is critical to identifying and responding quickly to new threats, which is vital to your organization.

One of the main challenges to create a SOC is staffing. Finding skilled people and keeping them on board is a challenge for most SOC Managers. In a business context like information security, which demands high level of expertise, the competition goes beyond a company's profile. Let's face it: a company down the street in a different industry is still your competitor for talented specialists. The evolving threats push you to continuously improve your security team's professional skills in order to keep them competitive, fighting against threat actors. This is why training your security staff is a win-win business, and can become a differentiator that serves two purposes: getting better performance from the workers you have and showing those workers that you value them enough to invest in them. This creates employee retention as well. If you want more stability, which means less turnover, you need to offer them something that each of them values: training and professional growth.

The roles and responsibilities you should consider for your SOC are:

- **Security Analyst** - Reviews the latest (SIEM) alerts to determine relevancy and urgency of identified events. Creates new trouble tickets for alerts that signal an incident and require Incident Response review. Runs vulnerability scans and reviews vulnerability assessment reports. Manages and configures security monitoring tools. Having a former white hat hacker experience is a big plus for this role.

- **Threat Hunter** - Reviews asset discovery and vulnerability assessment data. Explores ways to identify stealthy threats that may have found their way inside your network, without detection, using the latest threat intelligence. Conducts penetration tests on production systems to validate resiliency and discover areas of weakness to fix. Recommends how to optimize security monitoring tools based on threat hunting discoveries. This is a senior security analyst role, which requires expertise in threat hunting.

- **Security Engineer** - Maintains tools used, recommends new tools, and applies security updates for those tools. Designs and builds a security infrastructure and network security for an organization. Oversees the security architecture build over different systems.

- **Security Manager** - Supervises the SOC team's activity. Recruits, hires, trains, and assesses the staff. Manages the escalation process and reviews incident reports. Acts as Incident Response Manager when required. Develops and executes crisis communication plan to CISO and other stakeholders. Runs compliance reports and supports the audit process. Measures SOC performance metrics and communicates the value of security operations to business leaders.

- Other roles you should consider for your SOC include:

- **Penetration tester** - also known as "ethical hacker", is a highly skilled security specialist that uses different tools and techniques, attempting to breach computer and network security systems.

- **Compliance officer** - ensures that a company complies with its outside regulatory and legal requirements, as well as internal policies and regulations.

There are two main responsibilities involved with the SOC team:

- Maintaining security monitoring and analyzing your security on an ongoing basis. Detecting, analyzing, and responding to security incidents using a combination of people, processes and technology.

- Proactively investigating suspicious activities to keep your infrastructure secure by ensuring that potential security incidents are correctly detected, identified, analyzed, investigated, and escalated.

Setting-up a SOC requires the following steps:

- **Define a SOC strategy for your organization**: Defining a Mission and Vision for your SOC, along with defining the SOC objectives, will create, in a few sentences, the same understanding both within the organization and for external parties about the SOC implementation you lead;

- **Define, approve and implement the organizational structure for SOC**: For a SOC to become operational, the designed organizational structure must be approved and implemented. Sometimes, this process can take months. Thus, your SOC might start off working using an interim structure, to enable work to progress;

- **Hire and appoint staff**: The new structure must be filled with competent and skilled staff. It is likely that only a few positions will be filled with current staff; thus, additional staff members must be hired. You should consider that potential recruits often lack the required competences, and you will need to give them some time and appropriate training to assure their professional growth;

- **Preparation of facilities:** Facilities must be prepared by taking into account physical security and appropriate access rights - at least the security monitoring room should be protected from unauthorized physical access;

- **Development and implementation of detailed processes and procedures**: Development and implementation of the SOC processes, policies and procedures; IT processes and procedures; information security policy; security controls; and procedures within SOC;

- **Implementation of technology for the automation of processes:** Installing, configuring, documenting, and testing technologies for automation of processes within SOC;

- **Define a training plan for different staff roles**: A yearly training plan should be defined for each role of the SOC team in order to assure the continuous improvement of their skills;

- **Execute training for different staff roles according to the training plan**: Apply the training plan as defined. Further training can be carried out as planned. Skills gaps can be identified using crisis drills and blue-Red Teaming exercises;

- **Signing of relevant agreements with the constituency, stakeholders and partners**: Ensure expectations and authorities of the SOC are well-defined and recognized from the start, especially by those in the SOC's management chain;
- **Test run of SOC services and tuning of results**: Once the processes and technologies have been implemented, it is important to run tests for at least couple of days in order to identify any deficiencies in processes and technologies. Tuning actions should then be carried out for appropriate adjustments of the implemented processes and technologies;
- **SOC Go Live:** Now your implementation of SOC is ready. You can launch your SOC into production and celebrate;

Conclusions

Building a SOC is a challenging endeavor. Successful implementation requires careful definition and planning. Do a few things well rather than many things poorly. The main focus of the SOC team should be on prevention by enforcing security policy and controls, as well as assessing and mitigating risks.

You need to ensure strong quality control for everything delivered by the SOC. You need to gain trust and credibility for your SOC implementation and SOC team as well.

Define KPIs and measure them on monthly basis in order to ensure the SOC services are delivered effectively, thus proving that your investment in SOC means money well spent.

7. Security Incident Handling – Legal Considerations

This chapter explores the various legal aspects that have to be taken into account when a security incident occurs, focusing on the evidence gathering and supply chain management perspectives. Other related aspects of concern covered here include information disclosure to the public, information disclosure to the data subjects, notification and cooperation with relevant authorities (including CERT-RO), and the attribution angles.

7.1 Supply Chain Management

In terms of incident notification, the contractual provisions with the provider and the operational steps between the organization and its IT solutions providers (referred to as provider in this chapter) must occur as soon as possible. At the end of these steps, organizations will have established a maximum timeline by which an incident notification must reach them.

Of course, replicating this type of obligation and operational aspect throughout the supply chain will ensure that the organization is aware of all security incidents.

The organization can then use this information to take the necessary steps internally to mitigate incident consequences and to prevent similar incidents from happening in the future. Furthermore, the organization can also use the information to fulfil all their obligations towards the relevant authorities and affected individuals/entities, per existing legal requirements, as well as to reduce the negative effects of the security incidents on the affected parties.

Incident notification

In order to provide the organization with enough time to analyze the incident and decide upon all the legal and operational steps to take, the Provider needs to notify the organization of any incidents as early as possible; generally, the contract includes a maximum period, which tends to be around 24-72 hours. Another option is to have the Provider notify the organization immediately upon learning of the incident. Nevertheless, a clear timeline is critical in such situations because every second gained in addressing a threat reduces its overall fallout, resulting in a lower level of damages the organization has to incur in its aftermath.

In terms defining what an 'incident' is, in this notification context, it can be viewed as a 'potential incident', respectively, a risk that an incident may have occurred, without clear confirmation of this occurrence yet. Thus, the best approach to ensure proper compliance with legislation and swift identification of incidents (or, as mentioned under the GDPR, data breaches), is to have the Provider notify the company when they suspect an incident has happened, even if its occurrence has not been confirmed clearly.

Organizations can negotiate with their providers to have an obligation to promptly notify the organization about any identified vulnerabilities in the IT systems, as well as the Indicator of Compromise (IOC).

Assistance throughout the incident handling process

Including certain provisions concerning assistance from the provider (and its sub-contractors) will ensure that any incidents get investigated swiftly and that remedial steps get implemented quickly. This type of clause is generally heavily negotiated.

On the one hand, it is important for the organization to have its providers on standby in case of incidents (especially if caused by any of the providers themselves), in order to investigate the incident, to ensuring timely reporting toward authorities, and to swiftly remedy the root cause of the incident.

On the other hand, the provider has to have predictability when allocating resources and costs associated with a given contract. Having experts on stand-by 24/7 can be unfeasible from an operational and commercial perspective for certain providers. For this reason, providers often suggest limitations on their involvement and the time allotted for such tasks, and charging additional fees for such services.

At the same time, the organization may have certain data or metadata it retrieves. For example, in cloud services, certain types of logs are only kept by the cloud service provider. For this reason, it is essential to have specific contractual requirements for the provider to disclose data to the organization when incidents that require such information to be analyzed for mitigation purposes occur.

Interaction with authorities

Certain types of incidents, such as incidents occurring in certain sectors need to be notified to relevant authorities (e.g. data protection authority, banking regulatory authority). In this case, per the relevant legislation, authorities usually accept initial details on an incident with subsequent submissions completing the picture with more details and evidence, as they get discovered. Subsequent reports can include, for instance, details about the root cause, or mitigation measures to stop the consequences of the attack or future similar attacks.

Additionally, the respective authority may request more information on certain points relating to the incident, may perform an on-site audit of the situation, and/or may request that the organization implements certain controls to address the incident. These aspects should be reflected in the contractual clauses with all providers, in order for said providers to assist on these points as well.

Implementing controls

Every organization must clearly define the scope of their mitigation controls, which aim to prevent future incidents from occurring. Certain mitigation controls may be required by law (or mentioned as guidelines by relevant authorities), which is why no organization should go without them. However, a provider may request additional fees for such actions, depending on their complexity and their utility for its other clients, so having too many controls of this type can become financially unsustainable. Hence, beyond the legal requirements, every organization should tailor its mitigation controls to the reasonably likely and highly damaging risks specific to their operations.

Some argue that the provider should implement such controls, as they are closely related to the software they provide to clients regulated by such specific legislation. Of course, this is closely related to tailoring the IT system to the client's needs. For instance, when the IT system is aimed at a specific sector, such as banking, this may be argued easily. For IT systems created per the client's instructions or off-the-shelf IT systems, it may more difficult to argue. For data protection aspects concerning privacy by design, it may be argued that the IT system should, from the outset, respect all privacy by design requirements without the need for specific requests from the provider's clients in this respect.

For this reason, the main aspects negotiated here (depending on legal requirements and needs of the organization) are the costs for assistance, the extent of the assistance, and the timing for response/implementation.

Confidentiality of data

Another point to consider is that concerning the confidentiality of the data obtained during the incident analysis, attribution, vulnerability identification, and mitigation. This confidentiality and any disclosure of such information has to be on a need-to-know basis (only to the individuals that need to have access to this in order to perform a specific task). Further, the information should be deleted once it served the purpose for which it was disclosed.

7.2 Forensic Analysis and Preservation of Evidence

Forensics is an important part of incident handling. On the one hand, it can assist with identifying the root causes and the steps taken by attackers, which turns it into a valuable "lessons learned" tool. On the other hand, it can be used as evidence before the courts of law in case of litigation concerning the incident.

When creating forensic copies of data, one should have in mind the following principles: the forensic copy should be admissible (comply with any legal requirements in terms of evidence gathering and preservation), authentic (ensured through the best practices used during the forensic collection phase and through the chain of custody implemented properly), complete (the entire context needed to analyze the incident and the need for reaching a conclusion concerning the incident), and reliable (based on the forensic collection and preservation process used, which implements best practices in this respect). For this latter point, one should remember that, when performing forensic collection, the steps taken, when reconstructed, should lead to the same outcome.

For performing this type of activity, there are various tools that can be used in order to comply with best practices. We are mentioning below a few of them that can be further explored depending on the needs of the organization:

- SANS - SIFT Workstation - https://www.sans.org/tools/sift-workstation/
- Autopsy - http://www.sleuthkit.org/
- FTK Toolkit - https://accessdata.com/product-download
- Caine - https://www.caine-live.net/page5/page5.html

Preparing for incidents

Having a forensic expert on stand-by is crucial in preparing for incidents, as they are able to preserve evidence that can be used in the future, before authorities or before a court of law. There are two options in this case: an in-house forensic expert or an external one. From a practical perspective, unless the organization requires frequent preservation of evidence, it may be useful to have a framework agreement for an external forensic expert.

For this situation, a confidentiality agreement and a data processing agreement should be in place with the external forensic expert, as they will have access to confidential information (which would most likely include personal data).

Handling of forensic and investigation work in parallel

From an operational perspective, once the incident investigation commences, the incident handling team should be able to investigate the incident without impacting on the evidence gathering process and without destroying potential evidence.

The best approach begins with the evidence gathering process, once the incident occurrence is confirmed or probable. This should abide by the best practices in the field, including in terms of various IT assets/devices to be copied and the types of data to be collected (including volatile data, such as data within the virtual machine).

Applying best practices in forensic collection

In order to ensure proper evidence gathering that can be used before the courts of law later on, one should use scientifically derived and proven methods for preserving, collecting, validating, analyzing, interpreting, documenting and presenting digital evidence. This allows events to potentially be accurately reconstructed in the future.[25]

Prioritizing data gathering must rely on best practices, such as starting with the most volatile evidence/data and working towards the more persistent evidence/data.

Generally, one should not shutdown or reboot the IT system before collecting evidence/data, as the evidence/data may be lost or altered. The same recommendation applies when copying or preserving a program on the IT systems-otherwise, the data/evidence can get altered.

In case of complex IT ecosystems and/or complex incidents, it is essential to work fast, which can mean, for instance, prioritize data collection based on most relevant IT systems and parallel copying sessions for multiple servers, devices, etc.

Nevertheless, the data collected should be proportional to the purpose for which it is collected. Thus, no more than the data needed for the incident investigation, reporting to authorities and for potential legal disputes on the matter should be collected.

Further, any forensic collection should be made in close correlation with the Security Operations Center (SOC) team actions (either internal or external) and investigation phases for the incident. In addition, if specific actions have to be taken, per legal requirements, quicker than the forensic data collection can be finalized, alternative solutions should be analyzed in order to ensure both actions get completed successfully. Given the number of parties involved in the process when such correlations occur, there must be periodical drills that ensure everyone involved knows what to do in case of an incident.

Additionally, during the forensic copying of the data, data handling and corruption of original data should be minimized. This can be ensured by using best practices and a methodology for forensic data collection.

In terms of legal compliance, the forensic collection scope and limits should be clear for the organization and for the forensic professional. It is important to collect only the data needed, but sufficient data for further analysis, especially in terms of context and the IT system's state. As a court of law may request details on the collection process, a specific clause should be included in the forensic services

[25] Gary Palmer, "A Road Map for Digital Forensic Research", DFRWS 16, 2001, http://www.dfrws.org/2001/dfrws-rm-final.pdf , last accessed on 8 June 2021.

agreement concerning the forensic professional testifying before a court of law, if needed.

Preserving data

After data collection, data preservation is also important. The organization can keep the forensic copy internally, with proper security and chain of custody rules in place which ensure that no tempering occurs. Alternatively, it can be kept with a third-party at their location, with the same principles in place.

In general, it is preferred to have the original intact and not tempered with, with Forensics making available copies of the IT system (e.g. bit-by-bit, a snapshot at a given time). If this is not possible, the forensic copy should be prepared based on the forensic best practices and kept securely as per a well-documented chain of custody processes. Further, any subsequent analysis should not be performed directly on the forensic copy, but on secondary copies thereof. These secondary copies, of course, should be created based on forensic best practices, as the initial forensic copy was created.

Further, it is recommended, if possible, to maintain also the original (e.g. in case of laptops) or an auditable copy (i.e. forensic copy as per best practices) in order for a court appointed expert to be able to re-perform the incident analysis.

For this forensic copy, as per data protection requirements, a retention period should be established, with the data being deleted afterwards.

Data sharing

In terms of sharing the forensic copy or data from the investigation, this can be shared with specific provider for certain aspects. With the provider that provided the services/IT system under investigation, it may be useful to share certain data or, even, a forensic copy of the relevant data in order for this provider to analyze the data and identify the root cause of the incident, for instance. Another situation of data sharing might be towards a security incident investigation company, in order for this company to provide information about the incident after investigation. For any such scenario, the organization has to apply the data protection requirements, starting with proper data processing agreements in place, minimization of data being disclosed, and deletion of data once it is not needed by the provider.

7.3 Role of Relevant Departments within the Organization

The organization should have in place procedures that outline the role of each department in case of a security incident. For instance, the operational departments can be guided by another department – e.g. the security investigation department. It is essential to identify and involve all relevant departments from the outset – e.g. legal, data protection, risk. The internal procedure should also ensure that the key employees needed for the incident investigation and remediation steps are on stand-by or easily reachable.

Further, management should be informed once sufficient information to qualify an event as potential incident and should be kept informed about subsequent information gathered, incident reports prepared.

All notifications towards third parties should also be sent as soon as sufficient information has been gathered. This might be the case for individuals affected by the incident, companies using the services offered by the organization and subject to the incident, etc.

In addition, once an initial report about the incident consequences and root cause are prepared, this should be shared with key departments for they input on mitigation measures, containment measures, notification of authorities, etc., depending on the specifics of the department and of the organization. Mitigation steps are to be decided as per the usual internal rules of the organization – e.g. security team, management team.

At this stage, depending on the type of incident, the organization can consider if witness statements can be useful for future litigation. If yes, these should be obtained as quickly as possible from the relevant individuals.

After the incident has been mitigated to prevent similar incident from occurring in the future, it may be useful to re-check the affected IT systems, with external auditing or penetration testing, in order to ensure that the mitigating controls have been implemented properly and that no additional vulnerabilities were generated by the incident.

7.4 Conclusions

This chapter includes practical steps to be taken firstly internally by the organization when an incident takes place. The forensic steps should be taken in parallel with the incident investigation in order to gather information as soon as possible and not to interfere with the investigation.

The role of third parties in the forensic process is essential and should be established before an incident occurs.

The lessons learnt principle should be implemented for all the security incident handling.

8. Supply chain management – practical considerations

In the last couple of year, there has been significant emphasis on supply chain management, as there has been a growing number of successful attacks that started by targeting a service provider in order to reach its clients (e.g. for exfiltration of data, for affecting the integrity of the data, for deploying ransomware within that client organization).

Thus, recently, there have been a series of best practice guides for handling third parties providing services to an organization or holding (or processing) the data of an organization. This includes NIST[26] with best practices and SP 800-161, Cyber Supply Chain Risk Management Practices for Systems and Organizations and CISA, with best practices on this topic.[27] These best practices and guidance working papers focus on technical and organizational aspects of handling supply chain management. In this section, we are focusing on practical aspects in order to put in place appropriate processes in relation to the supply chain.

Further, from a legal perspective, the requirements in terms of third parties and sub-contractor have also increased in the past decade. In certain cases, this impact is clearly detailed, as is the case of outsourcing in the financial services sector or in data protection legislation (e.g. the obligations undertaken by the co-contractor of the organization have to be replicated throughout the supply chain to all of its sub-

[26] NIST, "Cyber Supply Chain Best Practices", https://csrc.nist.gov/CSRC/media/Projects/Supply-Chain-Risk-Management/documents/briefings/Workshop-Brief-on-Cyber-Supply-Chain-Best-Practices.pdf , last accessed on 21 July 2021. NIST, SP 800-161, "Cyber Supply Chain Risk Management Practices for Systems and Organizations", https://csrc.nist.gov/publications/detail/sp/800-161/rev-1/draft , last accessed on 21 July 2021.

[27] CISA, "ICT Supply Chain Risk Management Toolkit", https://www.cisa.gov/ict-supply-chain-toolkit , last accessed on 21 July 2021. CISA and NIST, "Defending Against Software Supply Chain Attacks", https://www.cisa.gov/publication/software-supply-chain-attacks , last accessed on 21 July 2021.

contractors and so on). In other cases, the impact is less clearly detailed, as is the case of the NIS Directive (including national implementation legislation), which expressly emphasize the obligations undertaken by the co-contractor of the organization.

In this section we are focusing on a couple of first steps that can be taken in this direction and which can be implemented by organizations having a low or medium level of maturity in terms of third-party management.

It is important to use resources efficiently within an organization. To this end, a prioritization should be performed on the approach to be taken to ensure security throughout the supply chain. Below we have included a couple of preliminary characteristics of co-contractors or services/products provided by them that can be used for establishing a risk-based third-party framework. Such an approach is aimed at using existing resources in an efficient manner and at addressing the risks with the highest impact or probability, depending on the risk appetite of the organization.

Based on the impact on data (and personal data), co-contractors/sub-contractors can be classified as follows:

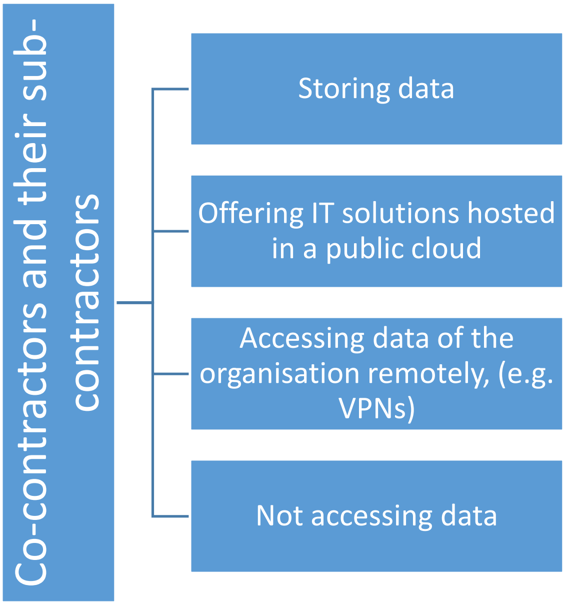
Co-contractors and their sub-contractors
Storing data
Offering IT solutions hosted in a public cloud
Accessing data of the organisation remotely, (e.g. VPNs)
Not accessing data

Figure 16: Third parties classified based on their access to data (and personal data)

Further, regardless if data access/data management/data storing is involved, third parties can be classified based on the impact of their services/products on the security landscape of the organization, as follows:

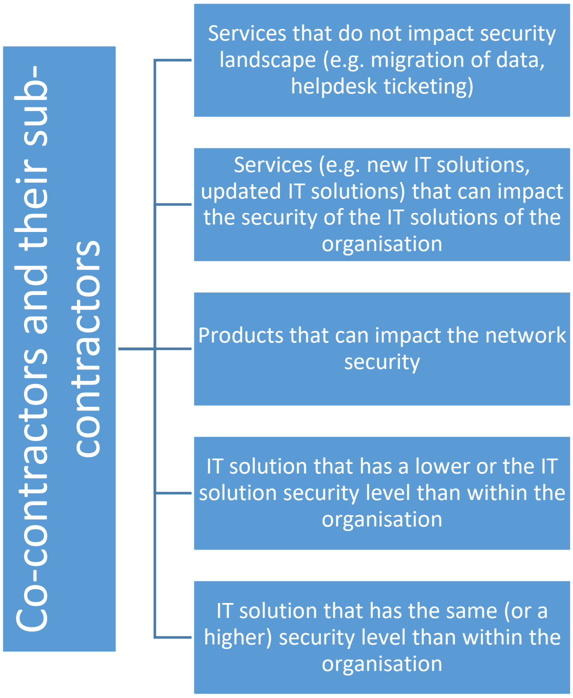
Co-contractors and their sub-contractors
Services that do not impact security landscape (e.g. migration of data, helpdesk ticketing)
Services (e.g. new IT solutions, updated IT solutions) that can impact the security of the IT solutions of the organisation
Products that can impact the network security
IT solution that has a lower or the IT solution security level than within the organisation
IT solution that has the same (or a higher) security level than within the organisation

Figure 17: Third parties classified based on their impact on the security landscape

Regardless of the classification above, there are certain steps that should be taken in relation to third parties. These can be considered the supply chain life-cycle. The specific actions for each step can be established by each organization based on specific guidelines (such as the ones above) and applicable thereto.

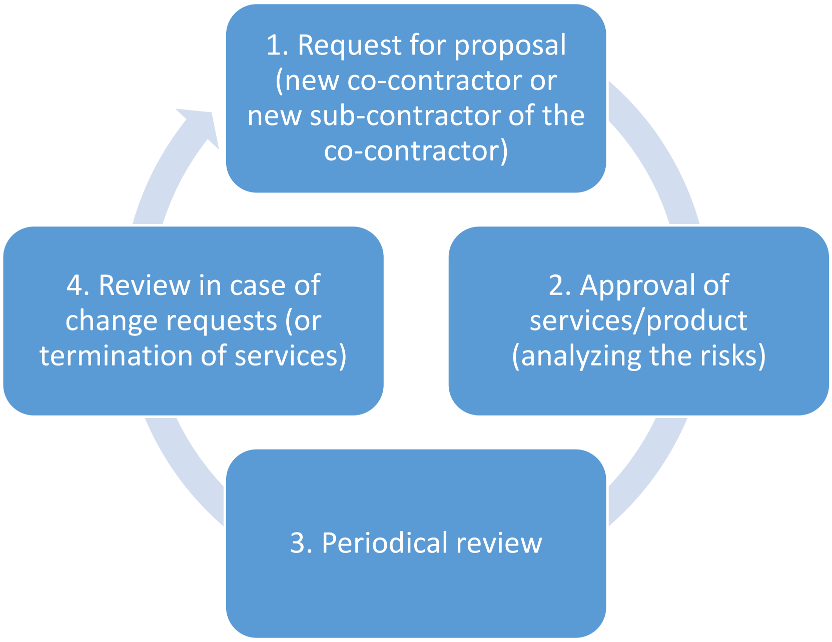
1. Request for proposal (new co-contractor or new sub-contractor of the co-contractor)
2. Approval of services/product (analyzing the risks)
3. Periodical review
4. Review in case of change requests (or termination of services)

Figure 18: Supply chain life-cycle

Below, we are detailing specific aspects to consider for each of the above steps. Throughout the below sections, we refer to the co-contractor of the organization as the vendor.

8.1 Request for Proposal

When an organization wishes to obtain a specific service or product, the first step is to prepare a request for proposal procedure, whether there is a single potential vendor or there are multiple potential vendors.

In this phase, the organization can request a baseline of security measures (and data protection requirements) to be confirmed by the entities submitting the proposals. Generally, lack of full compliance with such baseline entails the exclusion of the respective vendor from the request for proposal process. This assessment should be performed by individuals within the organization that have expertise in the security, respectively, data protection field.

The baseline requirements can be tailored based on the types of services provided by the vendor, focusing on the aspects outlined below.

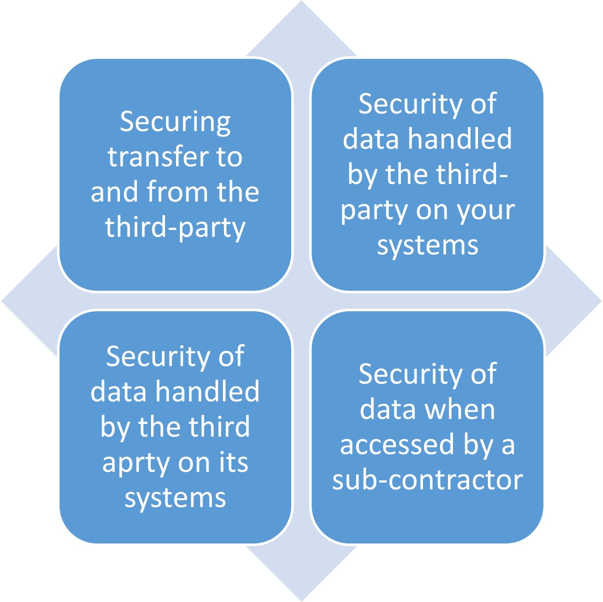
Securing transfer to and from the third-party
Security of data handled by the third-party on your systems
Security of data handled by the third aprty on its systems
Security of data when accessed by a sub-contractor

Figure 19: Security approaches based on services provided by third parties

The same process is applicable in case of sub-contractors. Generally, for custom types of services and products (e.g. not for cloud services, off-the-shelf software), the organization can request to be notified of any sub-contractors being contemplated before these are contracted by the vendor.

8.2 Risk Analysis

The selected vendor, aside from the brief baseline confirmation/analysis during the request for proposal phase, should be subject to a more detailed analysis in terms of security (and data protection), correlated with the matters detailed in the contract being negotiated with this vendor.

From a legal and privacy perspective the following main aspects have to be taken into account at the outset and during the use of third parties for various services:

- **Know your IT (and third-party) landscape** – it is important to understand the impact of the services/products on the organization and on the data/information held by the organization.

- **Understand the risk** – once the IT landscape is clear, the risk assessment can commence, taking into account the services/product stand-alone and within the IT ecosystem of the organization. In this respect, there a series of risk assessment frameworks that can be used or adapted, such as Octave, NIST RMF, etc.

- **Manage the risk** – for the identified risks, proper measures (controls) can be put in place as a response to the risk (e.g. to reduce the risk level). Such measures can be implemented by the organization or by the vendor. They can be organizational (e.g. procedure, obligation undertaken in a contract) or technical (e.g. changes to the It solution, changes to the infrastructure).

- **Establish workflows with the vendor** – the organization will require assistance in certain situations from the vendor and its sub-contractors. These should be discussed and established from the outset, periodically refreshed in order for all parties to know their role in the process. The organization may consider establishing workflows for:

    - security incident identification and investigation,

    - assistance in case of investigations from authorities or complaints from clients of the organization,

    - reporting obligations of the organization to relevant authorities,

    - auditing of vendor by an independent auditor or by the organization itself

From a technical and operational perspective, the security (and data protection) baseline can be analyzed taking into account the below:

- Relying on third-party's external auditors – an organization should generally prefer to have an analysis performed by external auditors chosen by it, as this entails knowing exactly the scope, the documents provided by the third-party, etc.

- Using auditors chosen by the organization – this is the preferred approach. It may be used as a periodical review and not necessarily at the outset of the involvement with the vendor.

- Relying on third-party's internal auditors or assessment – this should generally be avoided, as this does not present a level of independence needed.

- Automating the process – for certain reporting that the organization may want to monitor throughout the relationship with the vendor, the organization may set-up indicators (e.g. KRI, KPI) in an automated manner in order to efficiently gather and analyze the data.

- Allocating proper resources – these actions of analysis, monitoring, etc. require in the organization employees with proper experience.

- Sharing threat intelligence and having in place proper notification and investigation processes – for prevention of incidents, swift response in case of a security incident or identification of a vulnerability is essential.

Less experienced vendors may require training and guidance in terms of incident identification and incident handling. More experienced ones may establish a process for sharing of information about threats.

Often, a software or infrastructure component used by the vendor (or its sub-contractors) may include a vulnerability that is used by attackers. In this case, proactive analysis of potential vulnerabilities by each entity in the supply chain, the constant communication throughout the supply chain and swift reaction in case such vulnerability is essential.

- Controls within the organization – following the analysis, certain measures may need to be implemented within the organization. Proper monitoring thereof should be in place.

- Controls within the co-contractor's IT infrastructure (or its sub-contractors) – in addition, the analysis may reveal the need for certain measures to be implemented by the vendor (or its sub-contractors). This is the reason why it is essential to have an analysis before signing the agreement with the vendor and before deploying the solution of the vendor. Certain changes may need to be made to the vendor's solution and certain obligations may need to be included in the contract with the vendor.

Proper monitoring of the implementation of such measures should be made by the organization.

- Verify security within the IT infrastructure the IT solution is placed after deployment – it is essential not view the services/product offered by the vendor in context, respectively, in the IT ecosystem it is going to be deployed in. Thus, during this risk analysis phase, this should be the angles from which the assessment is performed.

8.3 Periodical Review

The above analysis should be performed periodically. The option for analysis can range, as detailed above, from independent auditors, internal auditors of the organization, and auditors of the vendor or independent certifications of the vendor, depending on the classifications mentioned at the beginning of the section, the previously identified risks and the resources of the organization. These should be also had in mind in terms of frequency of analysis. For higher risk vendors, more frequent reviews should be considered, whereas, for less high risk vendors, less frequent reviews can be chosen.

8.4 Review in Case of Change Requests

Throughout the lifetime of service provision, various changes may occur, such as new modules, changes in functionality of existing modules, exclusion of certain modules. These situations should be identified within the organization and can serve as triggers for a re-analysis. Changes, either new elements or exclusion of new elements, can both bring additional risks and the need for additional measures.

8.5 Conclusions

In the previous years, emphasis has been placed on the supply chain in terms of security incidents and prevention thereof. Guidance has been issued given this increase in cyber-criminal activity in terms of evaluating and monitoring suppliers of services and of products, including by NIST and CISA.

This chapter has described practical aspects that can be implemented in an organization or in the relation of the organization with suppliers in order to implement such guidance, including from a contractual and process establishment perspective.

9. Data Protection and Legal Aspects in the Context of Red Team and Penetration Testing Activities

Services related to offensive security involve access of third parties to the architecture of the network/applications of the company and also to data (personal data, confidential data, trade secrets). This entails the need to ensure proper steps and procedures are followed when handling such data, in order to prevent any negative consequences on the operations of the company. In such cases, as the services providers have access to the data usually handled by company employees, it is recommended to reflect best practices in terms of confidential/personal data handling in such interaction. Thus, such service providers have access to the internal IT systems of the company similarly to employees working on/with those IT systems and, consequently, measures similar to those imposed on employees are recommended to be reflected in the relation with the service providers. This entails the need to focus on preventive measures and ensure contractual documentation allows for remedial measures to be taken by the company.

A risk assessment of the planned actions for the offensive security exercise, reveals the sensitive areas or assets for which the company might decide to have contractual provisions.

For this reason, it is recommended to analyze at the outset the actions taken by a company internally or through external service providers in terms of offensive security from a data protection and legal perspective. Depending on the type of actions taken and the manner of implementation, specific aspects may be contemplated to be included in the implementation and/or specific clauses may be included in the agreement with the service providers. The purpose of such clauses is to set out the landscape for the exercise and act as a basis for the training made to individuals on the assignment.

Generally, such contractual provisions are useful to provide clarity on actions to be taken by each entity in cases where a data breach occurs, either in a voluntary or involuntary manner or in cases when a security incident occurs on an IT system on which penetration testing occurred.

This chapter first provides insight into the scope of offensive security service agreements, including practical suggestions on how to structure the scope and correlation with the other clauses in the agreement, especially from a liability and confidentiality obligation perspective.

Subsequently, we detail the data protection and confidentiality practical points to be analyzed and reflected in the agreement and/or implementation. These are sensitive aspects, as it is advisable for these to reflect legal requirements and also are closely linked to the business operations. The use of third-party tools is also analyzed from these perspectives.

The scope and the confidentiality clauses reflect also important point in terms of establishing the limitations of the exercise and to provide awareness of consequences in case of going beyond these limits or not fulfilling requirements properly. These situations and implications are detailed as well, together with conditions for liability in order to emphasize the usefulness of including clear contractual provisions from the outset and enforcing such contractual provisions throughout the exercise, including through training of employees of the parties.

Throughout the chapter, we reference penetration agreements, but the matters can be applied also for Red Teaming agreements, provided that they are relevant for the scope of the Red Teaming agreement.

The below analysis points and recommendations relate to preventive steps to be taken in terms of compliance, data protection and security of data, but also remedial steps in case of breach of contractual obligations or tort liability.

9.1 Scope of Offensive Security Agreements

The scope of the penetration test or Red Team exercise has to be established before its start and may be changed subsequently (for the performance of a subsequent penetration test), when unanticipated use cases/aspects become relevant. In terms of Red Team exercise, the scoping might be more general, referring to the purpose/end result, rather than detailed - as is the case of penetration testing. The below points referring to penetration testing are also relevant for Red Team exercise.

The scope is set through a risk-based approach performed by the company or together by the company and the penetration testers. The analysis can be based on importance of the IT system in the ecosystem of the company. Also, the risk analysis has to be aligned to the expectance of threats on such IT systems based on the current threat landscape. Thus, the scope can be determined together by the company and the penetration testing service provider. In case of penetration testing, the analysis can be performed based on internal risk assessment methodologies, service provider methodologies or legal requirements that have to be monitored continuously. This approach is not applicable in case of Red Team exercises, as these are focused more on the goal of the exercise, leaving the methodologies choices up to the Red Team.

The scope of the penetration testing has to also be defined by reference to the type of IT system being analyzed. Thus, the penetration testing can focus on the software (mobile application, web application, server and client side, together with any related middleware), on the network of the company (or part of the network) or implementation of the infrastructure (including servers, operating systems and firmware on hardware).

Further, the scope of work can include the methodology to be used for the testing phase and for the threat rating. On the one hand, this is relevant in order to ensure that all the aspects that have been agreed to be tested are covered in the testing exercise. On the other hand, this is relevant in order to ensure that all legal obligations concerning IT system testing are covered (including use cases, threat testing, legal obligations directly incumbent on the company or indirectly, when the company provides services for other entities subject to such legal requirements). This is applicable for certain sectors having specific legislation in terms of security of IT assets.

The specification of scope for penetration testing or for Red Teaming is relevant in defining the limits of the IT system accessing, and for protecting both parties when access without a right has to be determined. This aspect is described below, together with the proof of concept.

Such scoping exercise is relevant also in cases where penetration testing reports are used by auditors in their assessment of the IT system.

In terms of penetration testing lifecycle, it is advisable to consider from the outset if the service provider will perform any re-testing after certain findings are mitigated in order to include it in the penetration agreement. Alternatively, the re-testing can be performed by another service provider.

From a contractual perspective, there are two approaches that can be taken: either a single agreement reflects all aspects of the penetration service exercise (including confidentiality/non-disclosure aspects, commercial points, scope of work), which is usually used for one-off assignments, or a master services agreement is concluded with the service provider (containing confidentiality/non-disclosure aspects, some commercial points), whereas the scope of work and certain commercial points are included in statement of works for each individual penetration testing exercise.

Thus, in order for the scope of work to be clear for and binding on both parties (the company and the service provider), it is recommended for it to be included in the penetration services agreement/statement of work, in purchase orders or circulated through the manner of communication established between the parties for sending instructions.

In terms of obligations being binding on the service provider, it is recommended to have all essential requirements included in a document signed by both parties or mentioned in the agreement as being binding in terms of obligations. Generally, if specific requirements are mentioned in the RFP phase (e.g. number of certified persons performing the testing), in order for these to be legally clear for the parties, it is recommended to include them in the contractual documentation as well. This ensures an easy-to-follow framework for the service provision, making it easier for any member of the team working on the exercise to know the overall aim of the exercise and the steps to be followed. Usually, agreements mention that they supersede any prior agreements and information provided in RFP phase is not an obligation of the service provider to act in a certain manner.

9.2 The Concept of Personal Data in Offensive Security

Personal data represents data that can lead to the identification of an individual or makes an individual identifiable when correlated with other data available to the entity trying to identify the individual. Thus, in an organization environment, this includes data and any pseudonymised data in IT systems. As testing and development environments (by applying the data minimization and need-to-know principles) generally should not hold production data, these environments should contain anonymized or synthetic data. This is an implementation of the minimization principle and of the need-to-know principle. On the one hand, only personal data specifically needed for a data processing purpose should be used. In this case, generally as long as the IT systems and IT architecture is similar to the production one, actual production data is not needed. Only persons that need to have access to personal data should have access to it. In this case, in most cases, access to actual personal data is not needed, as the interest is to test the various technical and organizational aspects of the IT infrastructure and IT systems.

However, the protection of personal data has to be viewed in context. Even if the Red Team exercise or penetration test is performed on testing environments or on production environments for a limited period of time, the protection of personal data has to be ensured also for the future. Vulnerabilities identified in testing environments exist also on production. Thus, offensive security exercises also contribute to a higher level of protection of personal data.

In terms of specifics concerning the protection of personal data, offensive security exercises should take into account the following. Depending on the specifics of the exercise, certain steps can be taken in terms of the below points. For the Red Team exercises, the main aspects to be agreed with the service provider are the location of data extracted from the IT systems and general processing during and upon completion of the Red Team exercise for the data extracted/collected by the Red Team, whereas the Blue Team has to ensure security measures are in place in the company IT infrastructure.

- **Accessing data:** The amount of data and types of data to be accessed by the service provider are to be assessed on a case-by-case basis and they are to be closely correlated with the scope of the agreement. Section 9.3 below outlines several steps for protection of confidential data that can be applied also for access to personal data.

- **Transfer of data:** Any transfer of data has to be analyzed in terms of compliance with legal requirements. This is especially necessary when data is transferred to servers pertaining to the service provider or to third parties. In case of public cloud storage, data protection analysis has to be made in terms of the cloud service provider.

- **Scope of data processing:** The data processing scope should reflect the data processing activities performed in order to fulfil the scope of the offensive security agreement. This entails the identification of each type of data processing activities and the types of personal data that need to be processed. Section 9.3 below outlines several steps for protection of confidential data that can be applied also for protection of personal data in terms of scope of processing and data minimization of data disclosed.

- **Sub-contractors:** Sub-contractors of the offensive security service provider are most likely data processors under the data protection legislation. When sub-contractors are used, it is recommended for the company to identify which obligations from the service agreement should be replicated in agreements with sub-contractors. Emphasis is placed on transfer of data, confidentiality requirements and limitation of liability. Use of open-source solutions is useful from a practical and cost perspective. However, in such cases, an analysis has to be made on where data is stored and transferred.

- **Data retention:** The data collected during the offensive security exercise is to be held by the offensive security service provider and the company only for the amount of time needed. For example, the service provider should hold the personal data only until it delivers the report to the company, while, afterwards, it should pseudonymised the data (if it is needed to prove the performance of the agreement), delete it or anonymize it (if it is not needed for other purposes). In turn, the company can decide to hold the personal data, for instance, only until it identifies how to remedy the vulnerability identified in the offensive security exercise.

- **Security details:** In order to prioritize their implementation in accordance with legal requirements, the specific security technical and organizational matters can be implemented on a risk-based approach and taking into account the level of access of the service provider, the type of personal data which it has access to and the location for storing data during the performance of the exercise.

- **Anonymisation/Pseudonymisation:** In certain instances of penetration testing or of Red Team testing anonymized or pseudonymised data can be used. This is the case when production-like environments are used for the exercise. As a simplified definition, anonymized data entails that no individual can be identified or identifiable from the anonymized data, whereas pseudonymised data entails that an individual cannot be identified or is not identifiable without additional data kept separately from the pseudonymised data.

- **Data breach notification:** A process should be established for notifying the company in case of data breach occurring on the side of the service provider or its sub-contractors. A data breach can occur in the context of a penetration testing in case the actions taken by the penetration testers lead to loss of confidentiality, integrity or availability of personal data that was not in the scope of the penetration agreement. For example, if data is accidentally exposed to the public or if a service becomes unavailable to customers of the company because of actions taken by the penetration testers. Alternatively, identification of a data breach may also be relevant when a penetration tester identifies an already existing data breach in the IT systems of the company. This type of obligation is to be cascaded throughout the supply chain and should also take into consideration any other legislation concerning data breach notification.

The above aspects can be detailed in a data processing agreement concluded with the service provider. Generally, the service provider can be qualified as a data processor and, in such situations, the conclusion of a data processing agreement is mandatory. In case of sub-contractors being used, the service provider should undertake to replicate its obligations under the data processing agreement in any agreement with its sub-contractors.

9.3 Practical Aspects Concerning Confidentiality Obligations and Conflicts of Interest

In terms of protection of company know how and confidential information, there are several legal aspects that have to be included in the penetration agreement.

Trade secrets are defined by law[28] as information that satisfies all of the following conditions:

- It is a secret that is not known among the circles that normally deal with this kind of information;
- It has commercial value because it is a secret; and
- The company took reasonable steps given the specific circumstances to keep it secret.

Whereas the knowledge of certain information among professionals in a specific sector or known to the public can be determined to some extent, the commercial value of such information may be difficult to prove by a company. This entails proving a directly or indirectly liaison between the secrecy of information and the product/service sold by the company.

The first step in order to ensure protection of trade secrets is to have a confidentiality agreement in the offensive security agreement in this respect. This agreement can only cover the information that fulfils all of the above conditions cumulatively. In order to reflect legal requirements, it is recommended for the clause to include at least the following aspects:

- The types of information that is confidential, for instance, any information provided throughout the performance of the agreement.
- The situations in which confidential information may be provided to the service provider, such as, instructions, written documentation, oral presentations, during the performance of the agreement when accessing the company IT systems.

[28] Directive (EU) 2016/943 of the European Parliament and of the Council of 8 June 2016 on the protection of undisclosed know-how and business information (trade secrets) against their unlawful acquisition, use and disclosure.

- The exclusions from the confidentiality agreement, such as data publicly available, data obtained through other sources, data developed independently, data already known by the service provider. The exclusion concerning information already known publicly (or in certain circles of professionals) stems from the legal provision outlining the conditions of trade secret qualification.

- The disclosure right, which includes situations in which confidential data may be disclosed to third parties. This generally covers situations of legal obligations, such as disclosure to auditors, authorities or courts of law.

- Details of the persons to whom the confidentiality obligation applies, which may include the service provider and all employees of the service provider. For the latter, a commitment can be undertaken by the service provider to ensure that its employees comply with this requirement and specific training on this can be made.

The confidentiality agreement is generally applicable for intentional disclosure of confidential data. In case of unintentional disclosure, the liability clause becomes applicable.

In terms of timing, in order to produce effects, the confidentiality agreement should be signed prior to any confidential information is disclosed to the service provider. In certain situations, this entails that it should be signed during the request for proposals phase, as, in this phase, by asking questions about specific infrastructure products and configurations, confidential information about the company IT systems is disclosed.

Among the situations in which trade secrets are lawfully obtained is the observation, study, disassembly or testing of a product that has been made available to the public. It could be argued that, in this case, the offensive security testers testing a web application or mobile application have rights to use any such results of their analysis. Of course, this can be argued for any black-box situation, as in grey-box or white-box scenarios, they have access to more information than what is only made available to the public. In order to avoid such situations, the agreement for the services provided can include a limitation on using any information resulting from the activity performed under the offensive security agreement.

Thus, as mentioned above, in order to be able to protect trade secrets, the company has to ensure that it mentioned expressly what constitutes trade secrets (confidential data) and has to take active steps in protecting the unauthorized disclosure of such trade secrets. In case there are sub-contractors involved in the services provided, the company has to ensure all such sub-contractors make the same guarantees and undertake the same obligations as the service provider.

The condition of active steps taken by the company entails that these have to be properly documented and may include, aside from the contractual clauses mentioned above, technical and organizational steps, including the examples below.

Active steps in this context can be established on a case-by-case basis, depending on the medium on which the confidential information is stored, the manner in which it is communicated to the service provider and the instructions received by the service provider at the outset of the service provision and throughout the service provision:

- **Reiteration of confidentiality:** Reiterate the confidentiality level of document/information in writing or verbally when the document/information is provided.

- **Access limitation:** The service provider can make use of enterprise computers with enterprise **accounts**. Alternatively, the service provider uses a VPN connection for certain tests, when possible. Access is limited based on the scope of the testing exercise and alerts should be triggered when the scope is exceeded. Further, the testing should end when a proof of concept is completed, without the need to go further than that with data exfiltration, for instance. For Red Team exercises, as it entails creative means of entering the IT systems and extraction of data and, consequently, not a play-by-play establishment of the tests performed during the exercise, the access management mentioned above is usually not applicable in this case. Further, it is important for the company not to forget to remove access once the penetration test is completed. Even in cases when the same service provider is used for re-testing or for other penetration testing exercises.

- **Logging:** It is recommended for the exercise to be covered by logging both on the side of the company and of the service provider. This is useful in order to determine the facts when there are queries in this respect from either party.

- **Real-time monitoring:** The real-time monitoring is useful in case the internal team of the company is aware of the testing exercise in order to identify any situation in which the scope of the exercise is exceeded. In such cases, time is of the essence.

- **Verification of deletion of data:** At the end of the testing process, as per the retention period and deletion obligation under data protection legislation, all data should be deleted from the IT systems of the service provider. A statement from the service provider in this respect can be obtained (including reference to deletion of all copies of data from any storage) to ensure the process was completed successfully.
- **Integrity check:** For certain situations, such as storage spaces, web application source code, integrity checks can be performed after the exercise is completed in order to ensure that these have not been changed during the exercise.
- **Vulnerabilities/incidents identified:** Contractual obligations can provide that, during the exercise, if vulnerabilities outside of the scope of work or prior security breaches (that took place on the IT systems) are identified by the service provider, these should be notified to the company. Further, any such details should fall under the confidential data definition and prohibition of disclosure to third parties, the press or third parties should exist.

The relevant controls can be put in place on a case-by-case basis, depending on the specifics of the organization and on the type of confidential data disclosed to the service provider. For instance, in case of Red Team exercises, some of the above may not be feasible in practice due to the nature of the exercise.

The main idea that should be found throughout the confidentiality agreement/clauses is for the service parties to mention the scope of data needed for the exercise and, consequently, for the provider to have an idea of the type of data it needs to complete the exercise. In this manner, it is ensured that the provider obtains only the data needed for the performance of the service agreement and not to use the data for any personal/commercial purposes of the service provider, its employees or of third parties. This entails that the service provider guarantees that it trained its employees and has controls in place to prevent any breach of the obligations undertaken by the service provider in the agreement.

Confidentiality should also be analyzed in terms of the destination of the data. For instance, when a cloud solution is used for the tests, company data may be transferred in the cloud storage for that particular tool. In this scenario, analysis should be made on the appropriateness to share the data to this third-party and on assurance that data is deleted after the testing is completed. Further, in case the service provider uses machine learning tools for conducting its testing, service provider has to ensure that no confidential information is included in the machine learning tools.

In terms of using tools for which the service provider has license or other rights of use or use of third-party tools, the service provider has to guarantee on the one hand that it has the right to use the tool for the offensive security exercise and, on the other hand, that confidential data (including personal data) is not transferred or stored in such tools. Or, if data is stored, it is deleted appropriately at the completion of the exercise.

Additional aspects relate to the confidentiality of the offensive security report. Generally, a company would like to use this report as basis for auditors to perform their report, in case of litigation, queries from authorities, for required notifications to authorities and for requests from clients. Of course, for each of these instances, before the penetration testing report is issued, the company has to identify the future disclosure needs for such report. In case the Red Team also provides a report on specific vulnerabilities it identified, the same is applicable.

In certain cases, clients of the company may request to review the content of the penetration testing. This is the case especially when the company provides IT services to its clients, such as cloud services, any hosting services or web application hosting services. In certain cases, the clients might request the penetration testing, as they have legal obligations in terms of security assessments and/or performance of penetration testing for the entire IT system it uses.

The use of reports for other purposes than internal review may be subject to additional restrictions imposed by the service provider. Generally, service providers allow for disclosure of their report to auditors, authorities or potential purchasers of the company, but they do not provide reliance on the report. Thus, third parties can become aware of the report, but they cannot rely on the report and, consequently, the service provider is not liable towards these third parties for the content of their report or the manner in which they conducted the exercise.

There are certain situations in which the company might need to prove level of security and/or the fact that it makes regular penetration tests on the IT systems it uses or on the IT systems that are integrated with its products (e.g. cloud services uses for storing certain data within the application flow).

Generally, the service provider provides a partial/truncated version of the report in certain situations. This approach is taken because on the one hand, as per legal requirements, certain details of identified vulnerabilities should be known by a limited number of individuals (based on the need-to-know basis) and, on the other hand, that disclosing such sensitive information to third parties (e.g. potential clients in the RFP process, to the public on its website) may lead to exploiting the unresolved vulnerabilities or use of resolved vulnerabilities as a starting point for attacks (as part of the reconnaissance process).

Thus, in such cases, a general description of the identified vulnerability together with risk rating should be sufficient, provided there are no specific legal requirements for additional details.

In terms of reproducing the penetration report, as this is protected by copyright held by the penetration service provider, a specific right to reproduce it in other documents or publications is required. This applies also when reproducing parts of the report on its website or towards authorities/third parties.

Concerning the limiting of conflict of interest, it should be avoided for an individual that previously worked for the company or was involved as external provider in the development of IT systems or auditing activities to participate in offensive security testing.

Further situations of potential conflict of interest that may lead to inefficiency in service provision or suspicion of this by authorities and third parties involve not using the penetration service provider also for SOC or auditing services. This ties in also with regulatory prohibition in specific legislation, such as the auditing legislation.

One additional point to consider in relation to confidentiality and also liability is the performance of penetration testing exercises on software/infrastructure of third parties (e.g. vendors of the company whose application is used by the company and is integrated with the company's IT systems, cloud services used by the company for storage or in another form – IaaS, PaaS). In such cases, the company does not have the right to approve a penetration test on the systems it uses. The license for use does not cover such types of uses. This is mainly because such types of services are used by multiple clients and the availability level for all clients has to remain within the agreed levels. A penetration test might lead to perturbation of the activities for other clients than the company.

Additionally, such penetration testing might lead to accessing of confidential information pertaining to the service provider or to other clients. This leads to the breach of agreements between the company and the service provider or between the service provider and the other clients.

In cases where testing of infrastructure/software pertaining to other entities is needed, prior discussions and approvals from the third-party with respect to the scope of the agreement and the confidentiality of the information uncovered is needed.

Certain service providers have anticipated this needed of their clients (e.g. cloud service providers) and have provided on their website a notification mechanism for intention to perform penetration tests on their infrastructure. After a notification is submitted outlining the exact parameters for the penetration test, this is analyzed and approved or rejected by the service provider. This analysis is needed in order for the service provider to have an overview of the tests to be performed and ensure that these cannot damage/obtain unnecessary access to its IT systems. This approval mechanism can be detailed in the service agreement, depending on the needs of the company, and certain use cases can be provided from the outset.

Lack of an agreement between the company and the service provider concerning penetration tests or not respecting such understanding can lead to liability of the company or of the penetration testing service provider, as detailed in the following sections.

9.4 Prevention of Potential Consequences in Offensive Security

There are certain legal aspects that should be taken into consideration when access to IT systems is involved. In this section we outline some of these potential consequences with the aim of the readers to implement prevention mechanisms, including training, in order to avoid such types of consequences. The below represent a wide range of examples of potential situations that can occur, with the actual potential consequences depending on a case-by-case basis. It is worth noting that, in most cases detailed below, intent is a requirement for legal implications for the individual performing the action. Negligence generally does not lead to such legal implications.

Further, the below is aimed at providing awareness about legal implications of specific actions that may take place in the context of professional services provision, first describing certain legal concepts and the conditions that have to be fulfilled in order for these to become applicable.

Actions taken by service providers outside of the agreed scope of work may lead to certain consequences in accordance with legislation, including criminal law implications. As mentioned, in this section, we outline the main points to take into consideration, together with practical aspects in order to prevent any perpetration of criminal offences and criminal law liability during offensive security testing. The section focuses on the provisions of the Budapest Cybercrime Convention of 2001 and, as an example, on the Romanian Criminal Code. As countries generally have independent decision power in terms of the criminal law they adopt at the national level, such provisions may differ slightly from country to country, even if they implement the same international convention.

The first important concept to be clarified in this context is the IT system concept. The Budapest Cybercrime Convention defines IT systems as any device or a group of interconnected or related devices, one or more of which, pursuant to a program, performs automatic processing of data.[29] Case law has extended the concept of IT systems in order to cover any types of software, web application, database, microservices, computer, API, email server, etc.

The IT system access can be analyzed at different levels, depending on the method of accessing: can be located at the back-end application level (lack of SQL injection protection for databases, source code or other vulnerabilities), at the transport level (for example, in terms of encryption of passwords or lack of proper protection of ports for accessing a web application) or at the front-end application level (for example, the manner in which tokens are stored on the user's device for the login process).

Computer data means any representation of facts, information or concepts in a form suitable for processing in a computer system, including a program suitable to cause a computer system to perform a function.[30] This entails that any type of data, regardless of format (e.g. database, metadata, logs, files), location of storage or whether in transit or at rest falls under this definition.

"Perpetrator"[31] is defined as the person (individual or legal entity) that directly committed the actions that constitute a criminal offence.

"Accomplice"[32] is defined as the person (individual or legal entity) that, with intent, helps the perpetrator or makes it easier for the perpetrator to commit the criminal offence. An accomplice can either act before the perpetration occurs (e.g. leaving a back door access to an IT system for access in the future) or assistance is given to the perpetrator at the moment the perpetration takes place.

[29] Article 1 of the Budapest Cybercrime Convention.

[30] Article 1 of the Budapest Cybercrime Convention.

[31] See article 46 of the Romanian criminal code.

[32] See article 48 of the Romanian criminal code.

We continue with more legal concepts that are required in order to understand the conditions for criminal offence perpetration and in order to be able to set out workflows and awareness trainings that prevent such situations. This overview of legal concepts is followed by specific examples that reflect potential consequences to take into consideration in case of offensive security exercises.

In criminal law, it is important to determine if a criminal offence was perpetrated with intention or not. This is relevant, as some actions are considered criminal offences only if they are perpetrated with intention, whereas, other actions are considered criminal offences when perpetrated with intention or negligence. First, as an example, let us understand these two concepts under Romanian law and, then, apply them to use cases for each type of criminal offence covered by this section. Their definition might differ slightly in each country, but, generally, the concepts include the below conditions.

"Intention"[33] is defined as the type of fault in perpetrating a criminal offence whereby:

- The perpetrator foresees the consequences of his/her/its actions, having the purpose of such consequences occurring (direct intent); or
- The perpetrator foresees the consequences of his/her/its actions and, even if not having as purpose such consequences, he/she/it accepts them (indirect intent).

"Negligence"[34] is defined as the type of fault in perpetrating a criminal offence whereby:

- The perpetrator foresees the consequences of his/her/its actions, but considers these will not occur (foreseeable negligence); or
- The perpetrator does not foresee the consequences of his/her/its actions, even if he/she/it should have (unforeseeable negligence).[35]

The above subjective nature of taking actions can be analyzed by reference to the main individual performing the action or to individual assisting him/her. Under criminal law, individuals related to the perpetrating of a criminal offence can also be

[33] See article 16 of the Romanian criminal code.

[34] See article 16 of the Romanian criminal code.

[35] Chandler, Jennifer A., "Negligence Liability for Breaches of Data Security. Banking and Finance Law Review", Forthcoming. https://ssrn.com/abstract=998305 , last accessed on 28 February 2020.

held liable (e.g. instigator, accomplice). Below we cover the cases of perpetrator and accomplice, which are to be encountered in the use cases described in this section.

Further, as an example, under Romanian law, legal entities[36] can become liable from a criminal law perspective, provided the conditions under the Romanian criminal law are fulfilled. Generally, a private entity can be held liable for criminal offences perpetrated by individuals related to its business activity or on its behalf or for its benefit. Thus, a legal entity can be a perpetrator or an accomplice, as defined below. There are legislations in certain countries that have a similar approach towards legal entities. However, there are also countries that do not provide for the liability of legal entities, but only of individuals.

After this brief background on the main concepts in criminal law, we are continuing the section with use cases that may be encountered in offensive security exercises and manner of preventing that they are potentially interpreted from a criminal law perspective. The use cases mentioned below are not meant to be exhaustive, but are the main ones that can be encountered in cyber offences in general. This is useful for both companies and offensive security service providers in order to identify examples to be included in awareness trainings and to be had in mind when interacting with IT systems.

The main criminal law implications derive from accessing IT systems/confidential data that should not be accessed under the service agreement, extracting data from the IT systems of the company without right and intercepting communication towards/from the IT systems of the company.

The most relevant criminal offence is the accessing of an IT system (including data therein) without a right. This criminal offence is meant at protecting the social relations related to the confidentiality of the data in the IT systems.[37] This data can belong to a company or to individuals (employees, customers, co-contractors of the company). Thus, this provision protects both the ownership of the entity owning the IT systems and the ownership of the data held on these IT systems. This ties in with the data protection legislation protecting the data in the IT systems.

The scope of this criminal offence refers to the components of the IT systems, from the hardware infrastructure to the databases and applications found within the IT systems, together with the data found in these components (including logs, metadata, data found in databases, unstructured data found on servers, cookies).

[36] See article 135 of the Romanian criminal code.

[37] I.C. Spiridon, “Reflecții cu privire la legislație română în domeniul criminalității informatice”, Dreptul, no. 6/2008.

Further, the mere analysis of metadata, browsing history, types of cookies stored on the device, the types of apps installed on the device/their version, location, status of various sensors placed on the device and traffic data may also constitute unlawful access to an IT system, if there is no legal right or consent of the user to access such data.

The illegal accessing of an IT system includes various degrees of access: the authentication (entering the system as a user thereof), bypassing the authentication system (through various means, such as, brute force, social engineering), reading content in the IT system, copying/deleting data from the IT system or using the IT system to perpetrate other criminal offences. The legal provisions for this criminal offence include an aggravated version if the perpetrator surpasses certain security measures to enter the IT system (as, for instance, in the device monitoring context is found in some types of deep packet inspection). The legal doctrine is divided in terms of the need for such an aggravated version, as the impact on the rights of individuals is the same. The distinction between the usual type of perpetration and the aggravated criminal offence may be useful in correlation with obligations to ensure security measures are in place in an IT system. Thus, for the civil part of the litigation, the entity that did not ensure proper security measures for the IT system can be held liable for a part of the damages.

The transfer of data from the IT system also constitutes a criminal offence. This usually involves the prior access to the IT system. The transfer is unauthorized, in the sense that the perpetrator either does not have any legal or contractual right to transfer the data or the perpetrator exceeds its right to transfer data as part of its usual business activity or transfers the data to another location than in the course of its usual business activity. This entails the breach of the confidentiality and potentially integrity and availability of data (in case data is modified, deleted or there is a denial of service).

The criminal offence refers to the transfer of data outside of a given IT system (outside of, for example its databases, its storing spaces in case of data at rest or its infrastructure in case of data in transit).

There are certain specific situations in which there can be a violation of privacy. This criminal offence is rather new in Romanian legislation and case law is scares on the topic. Generally, it protects the social relation of one's life from illegitimate intrusions from others, either through the taking of pictures or from the listening of private conversations by others. This is applicable in situations where the illegally accessed IT system involves interaction with users (clients) or with employees of the company. In other jurisdictions this type of criminal offence may not be regulated. Thus, specific analysis of the relevant jurisdictions is necessary.

9.5 Use of Third-party Tools in Offensive Security

There are certain situations in which, either the company or a service provider does not hold on premise licensed offensive security tools or self-built offensive security tools. In case of use of third-party tools, these entities may opt to use third-party tools that are hosted on the servers of such third parties.

This is generally the case of SMEs or certain cyber security start-ups as it is less expensive and swifter to set-up, whereas established international groups usually opt for on premise offensive tools or tools that they have built internally. Of course, in addition, there is a degree of privacy and security risks associated with using third-party tools, as it involves another entity in the supply chain which gains access to confidential/personal data.

In case of using third-party tool in offensive security, the main aspects to be considered by the company relate to liability, confidentiality and data protection, as detailed in this section.

Nevertheless, in case a service provider offers tools that can be used in offensive security, such service provider has to analyze the legal requirements that can be applicable to it and the active steps that it should take to fulfil such legal requirements. The main consequence that triggers this is the possibility of such tool being used for illegal activities.

Firstly, it is recommended to prohibit the use of the tools for illegal activities in the agreement that allows the use of the tools. However, it may be considered that this is not enough to ensure lack of liability of the entity providing the offensive security tools. One solution might be to periodically verify the manner in which the tools are used and the existence of approval from the entities against which the tools are used. This monitoring might be useful in order for the entity providing the tools not be considered an accomplice to any criminal offences perpetrated by its clients, as it may be argued that the entity is acting with indirect intent.

Another factor to consider when providing such tools to entities for their use is that holding tools that can be used for criminal offenses on IT systems with the intent to perpetrate a criminal offence may be a criminal offence itself, depending on the applicable legislation. This is the case under Romanian law and similar criminal offences may be included by other jurisdictions. Of course, the intent has to be proven through any means available.

The provision of such penetration testing as a service (as detailed above) also entails the storing or confidential data that results from the offensive security tools used. This entails that the entities using the tools should guarantee that they have rights to store.

9.6 Specific Commercial Points to Consider

Aside from the legal and commercial aspects mentioned above as relevant for analysis when concluded an agreement, there are certain commercial points concerning liability that have to be considered as well. The below details represent awareness of the various aspects to consider when drafting an agreement and should be reflected in the agreement on a case-by-case basis, depending on the specifics of the situation.

Clauses concerning liability of the parties in terms of manner of performance of the agreement (including specific requirements and limitations under the agreement) are essential in case of offensive security services.

Liability can stem from breach of contractual obligations (either intentionally or not), breach of legal provisions (in case legal requirements are incumbent on the service provider or in case of criminal offences) or tort liability (in case of actions that create a prejudice for the company).

Contractual liability entails that the service provider, intentionally or not, has not performed an action or has performed an action without respecting the contractual or legal requirements. This type of liability relates only to specific obligations undertaken under the contract.

Breach of legal provisions entail the existence of a specific legal obligation applicable directly or indirectly to the service provider. The breach of legal provisions can have an impact in terms of sanctions applied by public authorities and of damages to be paid by the service provider if it is determined that it is its fault for the damages incurred.

Tort liability entails that an action of the service provider is a direct link to damages incurred by the company or other third parties, such as the customers of the company. Generally, the employer (in this case the service provider) is liable for the actions of its employees during their work activities. Tort liability can occur in any circumstances and does not relate to the contractual relation between entities.

In view of transferring risk, the company can opt for an insurance policy to cover potential liability occurring from the actions of the service provider, with the insurance being provided by the service provider to the company.

Below, we have outlined a couple of use cases that can be debatable from a liability perspective and potential approaches from a contractual drafting perspective. This is useful in view of illustrating the various aspects to consider when concluding an agreement.

- The service provider performed actions in addition to the scope of work. Generally, this should not result in liability of the service provider, provided other contractual obligations or tort damages are generated by such actions.

- The service provider did not perform the actions under the scope of work (either did not provide all actions or did not provide them properly). This can generate contractual liability.

- The service provider delivered an incomplete advice in terms of the scope of work needed in order to cover the legal requirements for offensive security testing. This can be rather difficult to establish, as it always requires the scope of work concerning consultancy to have been included in the service agreement. Further, there is the question of legal interpretation for the duty of care and negligence of the service provider in providing such guidance. The legal interpretation depends on a case-by-case basis and may be influenced on whether a lawyer was involved in the analysis or not.

- The service provider unintentionally created damages to the company (or to third parties) or generated a security incident through their performance of the testing. Examples in this respect include: accidentally running malware found on the company's IT system, accidentally, accidentally created a DoS for the IT systems of the company or its clients. In certain cases, the service provider can even be held liable for unintentional damages. However, this depends on the contractual provisions (whether such liability was undertaken by the service provider under the contract) and on legal obligations of the service provider under the applicable law.

- The service provider intentionally created damages to the company (or to third parties) or generated a security incident/eased the access of attackers through their performance of the testing (including extracting confidential data). This situation is generally covered by the service agreement, as it is closely correlated with the scope of work.

- The report did not reflect properly the impact or the probability of a risk, as this is usually perceived in the industry based on past events. The adequate risk assessment and risk rating depending on industry and best practices in a specific field is generally essential for companies to prioritize investments in security controls. This is closely tied to the proper performance of actions as required by the service agreement and legislation requiring a specific level of security.

- Damages generated by the automatic tool used during the testing exercise, including any machine learning or similar tools. One should analyze the amount of knowledge of the entity that used the automatic tool or lack of knowledge in order to determine liability.

- Leaving eavesdropping tools or similar tools in the IT systems of the company for future monitoring or extraction of information from these IT systems. This is generally performed with intent and falls under contractual liability general, but also may involve tort liability and legal liability (especially in case of criminal law angles).

Under Romanian law, certain types of damages cannot be excluded, such as the following:

- Material damages caused through intent or gross negligence.

- Damages concerning physical integrity, emotional integrity or health of an individual.

Depending on the applicable jurisdiction, similar or other exclusions may be applicable.

As the actual offensive security testing is performed by individuals, either employees of the service provider or external consultants thereof can be involved in the process. Thus, it is worth having a guarantee from the service provider that the employees/external providers are aware of the obligations undertaken under the contract or, alternatively, the employees/external providers can acknowledge and agree to such obligations directly.

An interesting situation from a legal perspective is the limitation of liability for damages generated by the company on the IT system of the service provider (e.g. malware that gets transferred to the IT systems of the service provider). This can be limited to situations in which this damage was created with intent.

In terms of limiting the liability of the service provider, there are certain points that can be included in the offensive security contract:

- Amount in damages covered – the amount can be negotiated between the parties.
- Limitation of actions for which the service provider is responsible – some of the above actions can be limited, including intent of actions.
- Limitation of damages covered by the liability clause – generally, direct damages resulting directly from the actions of the service provider are covered by service providers. Also, some indirect damages – loss of profit, costs with lawyers, litigation costs are sometimes included.

Liability clauses in an agreement are mainly a commercial point that has to be discussed and agreed between the parties. Nevertheless, it is essential to have a full picture of the potential implications of actions taken by each party, in order to tailor accurately the liability clauses.

9.7 Conclusions

In the context of evaluating the efficiency of the security of the organization, including the penetration tests and the Red Team activity, there are certain legal and data protection aspects to consider.

The purpose of the contracts – the necessity of clarity concerning the services provided, the importance of defining the steps to be followed and the limitations for their provision.

Accessing the personal data – when a legal requirement concerning personal data has to be implemented within internal processes of the service providers and of the organization opting for offensive security.

Identifying incidents concerning security and personal data – addressing this situation in the contractual documentation is necessary in order to clarify the responsibilities of each entity involved and the individuals that have to be informed of such situations.

Confidentiality obligation – the chapter details specific measures to be taken to ensure confidentiality and the implications of confidentiality in the context of offensive security.

Avoiding conflicts of interests – when choosing the service provider for the offensive security conflicts of interests should be avoided.

10. Governance of Data Related to Offensive Security and Blue Team

Governance in terms of data related to offensive security and Blue Team entails setting-up internal policies and procedures for implementation of proper mechanisms and legal requirements concerning:

- Gathering of personal data and sharing such data with third parties that are either private entities or authorities.
- Managing whistleblowing filings concerning data processing and information security
- Proper collection and preservation of evidence in case of security incidents

Thus, a company may collect, process and transfer data when choosing and implementing proper security measures and during the phase of security incident analysis.

The below sections outline the main legal requirements and legal risks to be taken into account, together with specific methodologies that can be implemented by the company.

10.1 Governance of Personal Data in Interaction with Third Parties and Authorities

In terms of international cooperation concerning cyber security, EU member states have been cooperating among themselves and with the US since 2010 on a joint approach on certain points, such as establishing standard good practices, incident handling procedures and raising awareness of cyber threats.[38] As outlined in this section, these types of initiatives can be further detailed in order to ensure proper and consistent implementation between stakeholders involved in the data sharing.

Further, cyber security has been on the agenda of the EU in the past decade. In 2017 the strategy of the EU in this respect has been mentioned by way of several impact assessments, draft legislation and communications. These documents essentially mention four ways forward: (i) more cooperation among member states and on an international level in terms of preventive steps to be taken, (ii) immediate cooperation among members states and on an international level in case of cyber-attacks, (iii) raising awareness on cybercrime matters, and (iv) enhanced research and development in the cyber security field and ensuring a greater number of specialists in this field.[39] These four ways forward, implemented in a cyclical and comprehensive manner by multiple stakeholders at an international level lead, in time, to an increase in the maturity level of information security in companies.

The cooperation on cyber-attacks and, consequently, on preventive steps against cyber-attacks has proven over the last decades as essential in growing the level of information security. This becomes more useful as the number of entities participating in the cooperation grows in terms of sectors of activity and territorial location. This constant cooperation from an early stage of cyber-attack identification can decrease the time until characteristics for identification of cyber-attack are established and efficient preventive measures are identified. Constant cooperation, in time, leads to an increase into the research on cybercrime and cyber security and, subsequently, in the increased need for cyber security specialist individuals.

Nevertheless, in order to implement an efficient cooperation and ensure the implementation of preventive steps, the actual end-users have to be aware on the

[38] European Commission, "Cyber security: EU and US strengthen transatlantic cooperation in face of mounting global cyber-security and cyber-crime threats", 2011, http://europa.eu/rapid/press-release_MEMO-11-246_en.htm , last accessed on 22 December 2020.

[39] Joint Communication to the European Parliament and the Council, *Resilience, Deterrence and Defence: Building strong cybersecurity for the EU*, September 2017. Commission Recommendations (EU) 2017/1584 of 13 September 2017 on coordinated response to large-scale cybersecurity incidents and crises.

one hand about the cybercrime techniques and, on the other hand, about the preventive steps to be taken to ensure information security.

Data protection requirements generally stem from legal provisions, which either mention a principle for data processing or detailed provisions to be implemented. In any situation, even if data protection entail compliance with legal requirements, in order to proper implement this in a company, a data governance framework has to be implemented. There are various approaches in this respect, from the NIST Privacy Framework to ISO 27701 on privacy information management.

A data governance model used by a company should cover certain main steps and actions to be taken by the company. We detail below a data governance model that includes steps for both data protection and information security:

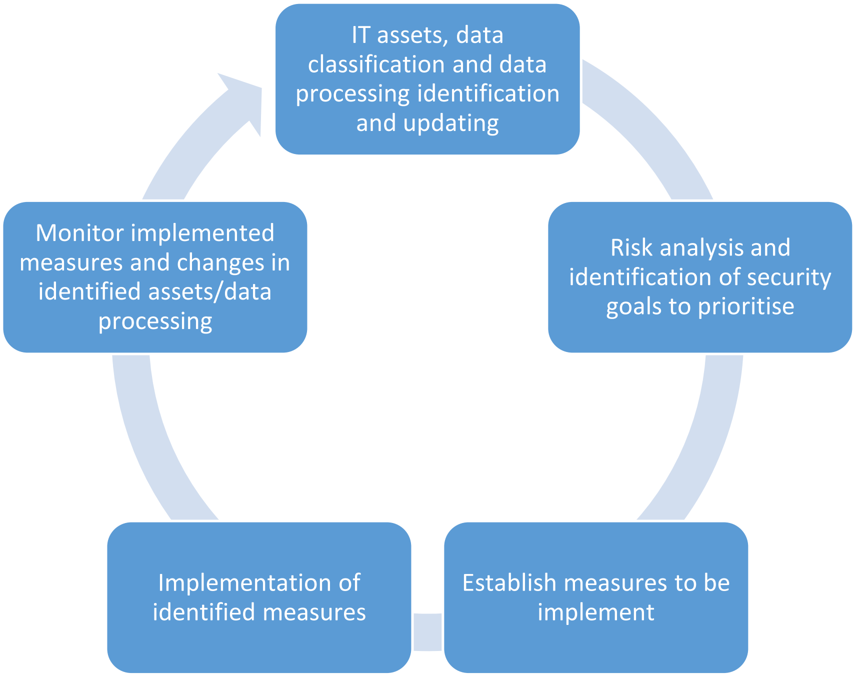
IT assets, data classification and data processing identification and updating
Risk analysis and identification of security goals to prioritise
Establish measures to be implement
Implementation of identified measures
Monitor implemented measures and changes in identified assets/data processing

Figure 20: Data governance model for data protection and information security

Within the above data governance context, in terms of transfer to third parties, we outline below a specific data governance model that can be integrated in the overall data governance framework of the company, while maintaining the specifics in case of third-party involvement. The approach we propose includes the following main steps in the data governance and can be applied for scenarios referring to data shared by the company and data received by the company. Further, each step outlined in this model contains data protection, legal and security requirements.

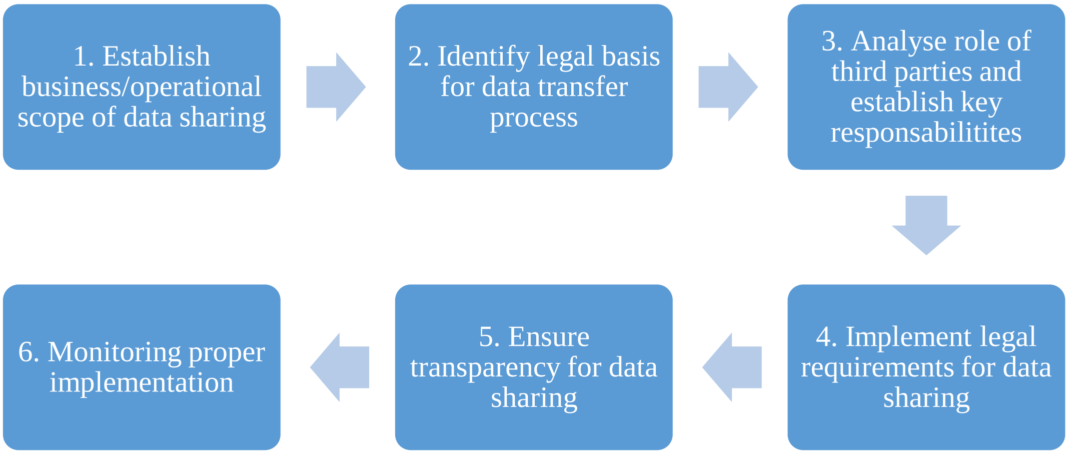
1. Establish business/operational scope of data sharing
2. Identify legal basis for data transfer process
3. Analyse role of third parties and establish key responsabilitites
4. Implement legal requirements for data sharing
5. Ensure transparency for data sharing
6. Monitoring proper implementation

Figure 21: Steps concerning third parties

The first step mentioned above starts with the business/operational need identified in this case by the IT/Information security department. Steps 2, 3 and 4 involve the cooperation of the data protection, legal, IT and information security departments in order to ensure proper compliance with legal requirements and sharing of data useful for the identified business/operational scope, with each department outlining requirements impacting the sharing process which are relevant from their perspective. Step 5 is ensured by the data protection department together with the business department having contact with the individuals whose personal data is being shared.

This section outlines the below aspects for each step of the data governance, with practical examples and aspects to be analyzed.

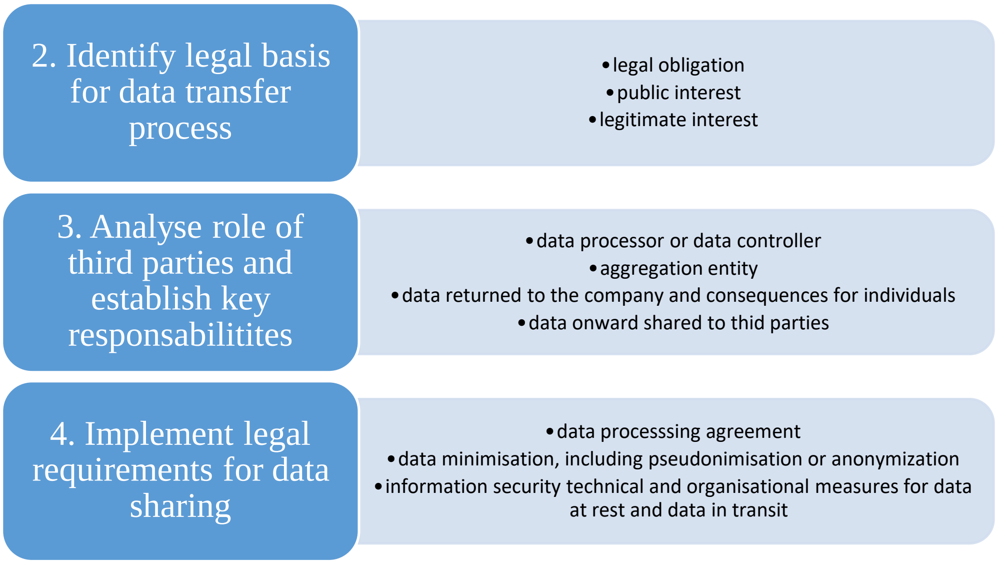
2. Identify legal basis for data transfer process
• legal obligation
• public interest
• legitimate interest
3. Analyse role of third parties and establish key responsabilitites
• data processor or data controller
• aggregation entity
• data returned to the company and consequences for individuals
• data onward shared to thid parties
4. Implement legal requirements for data sharing
• data processsing agreement
• data minimisation, including pseudonimisation or anonymization
• information security technical and organisational measures for data at rest and data in transit

Figure 22: Details of the steps concerning third parties

The sharing of data for prevention purposes may be performed to multiple types of third parties.[40] The data can be shared within the same group of companies, with service providers that aggregate data in order to determine trends in cyber-attacks, with an organization for a sector in order to gather information about sector-wide potential attacks or with other third parties. Further, data may be sent to the entire supply chain in order to avoid supply chain attacks or to a particular entity providing forensic services (as the private entity that collected the data does not have necessary skills in this respect). In certain cases, due to the potential global coverage of a potential attack, sharing of data on a publicly accessible website may be contemplated.[41] Each case may have different implications in terms of legitimate interest for data sharing.

In this section we cover the main aspects to be analyzed by the company in case of sharing of data about security incidents or about vulnerabilities with third parties.

[40] The Global Cyber Security Capacity Center, "Computer Security Incident Response Teams(CSIRTs): An Overview", 2014, https://www.sbs.ox.ac.uk/cybersecurity-capacity/system/files/CSIRTs.pdf , last accessed on 2 February 2020.

[41] Hewling, Moniphia, "Cyber Intelligence: A Framework for the Sharing of Data", International Conference on Cyber Warfare and Security, 2018.

In terms of legal obligations concerning sharing of data on security incidents[42] for IT systems, the NIS Directive is the main source of such obligations.[43] However, data protection legislation may also be considered relevant, as it entails state-of-the-art security measures to be implemented for the confidentiality, integrity and availability of data and for the resilience of the organization.

The NIS Directive[44] goes further and establishes national Computer Security Incident Response Teams (CSIRTs) and, on the other hand, a network of CSIRTs. [45] The CSIRTs within a country can be also private entities.

As an example, under Romanian law,[46] notifications of security incidents include generally details on the incident, impact of the incident and preliminary measures adopted. The law expressly mentions that no data bringing negative consequences on the rights and liberties of individuals/third parties involved in the incident should be provided.[47]

On the one hand, the notification obligations[48] are useful in terms of sharing information and correlated these in order to identify patterns and prevent future attacks.

On the other hand, information granted concerning security incidents/data breaches should be limited in terms of access and content, in order to be in compliance with applicable legal requirements. For instance, confidential information should not be included in the notification sent to the data subjects, but limit the data disclosed to the minimum requirement under the law.[49] Such information may be used by the

[42] Article 14 of the NIS Directive.

[43] Erich Schweighofer, Vinzenz Heussler, Peter Kieseberg, "Privacy by design data exchange between CSIRTs", Annual Privacy Forum, 2017.

[44] Articles 9 and 12 of the NIS Directive.

[45] The White House, "Fact Sheet: Cyber Threat Intelligence Integration Center", 2015, https://www.whitehouse.gov/the-press-office/2015/02/25/fact-sheet-cyber-threat-intelligence-integration-center , last accessed on 1 February 2020.

[46] Article 27 of law 362/2018 implementing the NIS Directive.

[47] Luis Tello-Oquendo et al, "A Structured Approach to Guide the Development of Incident: Management Capability for Security and Privacy", https://pdfs.semanticscholar.org/023e/d70a52d6c8396e463188be7ddd88544869ec.pdf , last accessed on 2 February 2020.

[48] Hong, Seung-Hun and Alazab, Mamoun, "Cybercrime and Data Breach: Privacy Protection through the Regulation of Voluntary Notification", 2017. Prepared for the Korea Legislation Research Institute (KLRI), 2017 Legal Scholar Roundtable, How Law Operates in the Wired Society, Seoul, Korea, 2017. https://ssrn.com/abstract=3042174 , last accessed on 28 February 2020.

[49] Dähn, Marie-Christine and Pernice, Ingolf and Pohle, Jörg and Goldman, Zachary and Nemitz, Paul Friedrich and Christakis, Theodore and Milch, Randal S. and Kent, Gail and Wetzling, Thorsten Manuel and von Lewinski, Kai and Djeffal, Christian and Herpig, Sven and Krüger, Philipp S. and Grafenstein, Maximilian and Barker, Tyson and Rubinstein, Ira and

perpetrator or other individuals in order to identify vulnerabilities in the system of the company that incurred the security incident (additional information may be obtained through data subject access requests).[50] Further, information about the perpetrator may be contained in the notification and this may not need to be available to data subjects, but only to the authorities assisting in the investigation and monitoring of the security incident.

Nevertheless, in some cases, it may prove useful to share information even about an unsuccessful attack.[51] Although not expressly mentioned as a legal obligation in practice, the data gathered during the attack or in case of active security measures may be useful for prevention of further attacks.

As a general note, in case of transfer between data controllers who act independently, each data controller has the obligation to fulfil its own legal requirements as data controller. Thus, each data controller has to ensure that the sent data is transferred based on a legal basis and in accordance with legal requirements.[52] In addition, each data controller has to ensure that the data is received based on a legal basis and in accordance with legal requirements. Thus, a data controller cannot assume that the data it receives is in accordance with legal requirements without verifications in this respect.

For sharing of data between such independent data controllers, usually, in the agreements between them, there are specific clauses that ensure that the data sender has taken specific steps to comply with data protection requirements upon collection of the personal data and, usually, that it has also informed the data subjects about the transfer to the data receiver.[53] Sometimes, the data sender undertakes to perform

Lenssen, Klaus and Gitter, Rotraud, "Privacy and Cyber Security on the Books and on the Ground" (August 1, 2018). Edited volume. Berlin: Alexander von Humboldt Institute for Internet and Society. ISBN: 978-3-9820242-1-9. https://ssrn.com/abstract=3250354 , last accessed on 28 February 2020.

[50] ECJ, College van burgemeester en wethouders van Rotterdam/M.E.E. Rijkeboer, C-553/07, 7 May 2009, concerning the right of access of the data subject. Bucharest Court of Appeal, decision no. 158/2019.

[51] ENISA, "Standards and tools for exchange and processing of actionable information", 2014. https://www.enisa.europa.eu/publications/standards-and-tools-for-exchange-and-processing-of-actionable-information/at_download/fullReport , last accessed on 28 February 2020. Maria Bada et al., "Computer Security Incident Response Teams (CSIRTs) An Overview, Cybersecurity Capacity Portal", https://www.sbs.ox.ac.uk/cybersecurity-capacity/content/computer-security-incident-response-teams-csirts-overview , last accessed on 28 February 2020. ENISA, "A air for sharing - encouraging information exchange between CERTs.", 2011. https://www.enisa.europa.eu/publications/legal-information-sharing-1/at_download/fullReport , last accessed on 28 February 2020.

[52] NIS Cooperation Group, "Reference document on security measures for Operators of Essential Services, CG Publication 01/2018".

[53] Bucharest Tribunal, decision no. 182/2019 concerning the legal obligation of fiscal authorities concerning the public disclosure of receivables. Bucharest Tribunal, decision no. 4925/2019 concerning having a protocol between prosecution unit and land book registry as basis for data transfer.

the information obligation on behalf of the data receiver, as it has direct contact with the data subject.

In the case of joint controllers,[54] matters relating to transparency about data sharing should be discussed and agreed between the joint controllers. For the sharing of threat information, as the data is useful for the common purpose of all entities involved, it may be argued that these act as joint controllers. In such case, the liability of each of them should be detailed in the data sharing agreement between them, including any limitations of future uses of received data. Thus, matters relating to liability of each data controller (independent or joint controllers) can be detailed in the contractual documentation.

In case of threat data sharing, depending on the structuring of the data sharing, the entities participating in the system may be joint controllers or independent controllers. This is relevant from a liability perspective.

The main data protection risks are related to transparency and legal basis for transferring. Lack of these aspects may result in sanctioning with fines for one or both of the data controllers.[55]

The data gathered for monitoring of device activity for preventing cyber-attacks and fraud should not be used for subsequent / other purposes. This ties in with the net neutrality discussions over the last decade.[56] These mainly refer to ISPs (especially given their deep packet inspection capabilities), but are also applicable to other companies that have access to large amounts of data about individuals (as is the case of applications that scan and monitor devices for threats), especially in case such data is shared with third parties or aggregated in a single central database.[57] The implementation of the need-to-know and data minimization

[54] Hongxin Hu, “Detecting and resolving privacy conflicts for collaborative data sharing in online social networks”, ACSAC '11 Proceedings of the 27th Annual Computer Security Applications Conference, pages 103-112.

[55] James C. Cooper, “Anonymity, Autonomy, and the Collection of Personal Data: Measuring the Privacy Impact of Google's 2012 Privacy Policy Change”, 2017, https://ssrn.com/abstract=2909148 , last accessed on 17 December 2020. George H. Pike, “Google, YouTube, Copyright, and Privacy”, 2007, https://ssrn.com/abstract=1636395 , last accessed on 17 December 2020. Ira Rubinstein & Nathan Good, “Privacy by Design: A Counterfactual Analysis of Google and Facebook Privacy Incidents”, 2013, https://ssrn.com/abstract=2128146 , last accessed on 17 December 2020. Trautman, Lawrence J., “How Google Perceives Customer Privacy, Cyber, E-Commerce, Political and Regulatory Compliance Risks”, 2017, https://ssrn.com/abstract=3067298, last accessed on 17 December 2020.

[56] Bert-Jaap Koops, Jasper Paul Sluijs, “Network Neutrality and Privacy According to Art. 8 ECHR”, European Journal of Law and Technology, Volume 3, No 2, 2012.

[57] European Data Protection Supervisor, “Opinion of the European Data Protection Supervisor on net neutrality, traffic management and the protection of privacy and personal data”, 2011, http://www.edps.europa.eu/EDPSWEB/webdav/site/mySite/shared/Documents/Consultation/Opinions/2011/11-10-07_Net_neutrality_EN.pdf , last accessed on 28 February 2020.

principles under data protection legislation is essential in this respect, both in terms of the aggregated database containing the data and also by each company that has access to such data. This also assists in ensuring compliance with article 8 of the ECHR (European Court of Human Rights).

Transparency aspects

The transparency principle entails that the data controller informs the data subject in a clear, concise and easy to understand manner of the data processing and data sharing. In case consent is required for the data sharing, the information is performed prior to obtaining the consent of the data subject. This is relevant also for the criminal law analysis in terms of the conditions for validity of the consent exemption to be fulfilled.

Interesting in case of data sharing between two data controllers is the manner in which the receiving data controller performs its information obligation towards the data subject. Whereas the data sender may have collected the data directly from the data subject, the data receiver generally has obtained indirectly the personal data (and may not be in direct contact with the data subject). Usually, in practice, as per an agreement between the data sender and data receiver, the data sender provides the needed information notice on behalf also of the data receiver.

Data protection authorities in member states[58] have commenced to express their opinion that a complete list of data receivers has to be provided to individuals (irrespective of whether consent is needed or not for the data sharing). This impact also the manner in which data subject requests are dealt with, as it implies a cooperation between the data sender and all the data receivers to correlate and take into account any request from the data subject. Further, in case of additions to the list of data receivers, the updated list has to be brought to the attention of the respective data subject.

On the transparency aspect, one matter mentioned by CNIL[59] (Commission Nationale de l'Informatique et des Libertés) was that the information notice was difficult to read and understand by individuals, due to the manner in which it was presented, but also due to lack of clarity about the data processed by the large

[58] CNIL, "Guidance concerning the sharing of data with business partners", 2018, https://www.cnil.fr/fr/transmission-des-donnees-des-partenaires-des-fins-de-prospection-electronique-quels-sont-les , last accessed on 21 December 2020.

[59] CNIL, CNIL imposes financial penalty against Google LLC, 2019, https://www.cnil.fr/en/cnils-restricted-committee-imposes-financial-penalty-50-million-euros-against-google-llc , last accessed on 3 December 2020.

number of services provided by Google (around 20), which combined data about users among themselves in various manners.

The decision of CNIL was mainly on the lack of proper consent obtained from individuals, as this consent was not properly informed, specific or unambiguous. Thus, the sharing of data among the various entities involved in the advertising does not have a proper legal basis under data protection legislation.

In a similar context, the Dutch data protection authority investigated the aggregation of data obtained through its various services and products (search engine, web browser, email client, video streaming, and online maps). The authority concluded that the consent obtained for the sharing of data was ambiguous and not sufficiently informed. [60] Further, the necessity for aggregation of data under the legitimate interest legal basis was not sufficiently substantiated.

Thus, this transparency aspect relates to the reasonable expectation of the data subjects about the transfer of their data based on the information they have been provided. This should be ensured for clients of private entities. However, it may be debated for data pertaining to perpetrators (potential perpetrators), as detailed in the previous section.

The purpose limitation for the data transfer should also be clearly stated in the privacy policy and in the contractual documentation concluded with the entities that receive the data. The purpose for transfer should be compatible with the purpose for which the data was initially collected. In this case, for activities of prevention of potential cyber-attacks. Thus, data may not be used for any other purpose, especially for any segmentation of clients and marketing purposes. [61]

Data sharing can take several forms, depending on the receiver of data and number of stakeholders involved, respectively, reciprocal exchange of data between two entities, entity(ies) sending data to third-party(ies), several companies putting together information they hold, one-off disclosure of data to third parties.

Provided that data is anonymized when shared with third parties, it may be argued that only non-personal data is being transferred. In terms of non-personal

[60] Dutch Data Protection Authority, “Investigation into the combining of personal data by Google, Report of Definitive Findings”, 2013, https://autoriteitpersoonsgegevens.nl/sites/default/files/downloads/mijn_privacy/en_rap_2013-google-privacypolicy.pdf , last accessed on 5 December 2020.

[61] Saberlotodo Internet, S.L. - Judgment of June 6, 2012 - Spanish National Court of Appeal, https://www.iberley.es/jurisprudencia/sentencia-administrativo-an-sala-contencioso-sec-1-rec-594-2009-06-06-2012-13777081 , last accessed on 5 December 2020.

data to be shared between companies, EU has recently adopted a regulation[62] in this respect and issued a guidance[63] for the interaction of the legislation concerning data protection and the one concerning non-personal data. The guidance includes reference to cases where there is a mixed dataset (including both personal and non-personal data), which usually occurs when profiling activities are also intended by the companies, including big data,[64] artificial intelligence or internet of things network. In such case, there are two approaches:

- If the personal and non-personal data can be divided, GDPR is applicable for the personal data part of the dataset and the above regulation is applicable for the non-personal data.
- If the personal and non-personal data are inextricably linked, data protection legislation is applicable to all data. This situation can occur if it would be impossible, technically not feasible or economically inefficient to separate the two types of data. Also, there may be cases where separation of the dataset can decrease the value of the dataset or it may be difficult to clearly differentiate between the two types of data. Nevertheless, the separation of these two types of data is not mandatory and is left as a choice of the companies holding the data.

The below requirements are analyzed from the perspective of sharing data and their specifics in this scenario. As a general note, the ECJ (European Court of Justice) held that the rights of data subjects override, as a rule, the economic interests of companies.[65] However, if the data is transferred only for prevention of future attacks, this would not be included in the concept of economic interests of a company.

In case of threat data sharing, the main types of personal data found in the files shared may pertain to (i) employees, (ii) clients, (iii) individuals related to employees/clients, (iv) the perpetrator or (v) to the individual holding the IT systems used during the attack. For these categories of data subjects, it is difficult to implement the information obligation for the data processing, due to lack of proper details of

[62] Regulation (EU) 2018/1807 of the European Parliament and of the Council of 14 November 2018 on a framework for the free flow of non-personal data in the European Union.

[63] European Commission Communication, “Guidance on the Regulation on a framework for the free flow of non-personal data in the European Union”, 29 May 2019, https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=COM:2019:250:FIN&from=EN , last accessed on 5 December 2019.

[64] Sumithra, R. and Parameswari, R., Security, “Privacy Issues and Challenges in Big Data and Cloud Security: A Survey”. International Journal of Advanced Studies of Scientific Research, Volume 3, No. 10, 2018. https://ssrn.com/abstract=3319251 , last accessed on 29 February 2020.

[65] ECJ, Case C131/12, Google.

the data sharing until it takes place. For the employees, client and related data subjects, a general statement may be included in the information notice they are provided with at the outset of the relation with the entity. This should reflect also a description of the data sharing system and the possibility of the data to be transferred to entities in countries that do not have an adequate level of protection of personal data. For the perpetrator, it depends on certain specifics of the situation. For instance, if the information is also sent to the criminal investigation bodies, the perpetrator should not be notified of the data that is included in the case file and the content of the case file is not to be disclosed.[66] This is relevant if the attack was successful or not. If the data is not sent as part of a criminal file, it may be argued that the prior notification of data processing should be made to the perpetrator. In US legal doctrine, it was mentioned that entities could include a file on the desktop of the honeypot (a visible location) with the data processing details and that this is sufficient in terms of brining to the attention of the perpetrator the information notice.

For the entities receiving the data, it may prove impossible or disproportionate to provide such information notice and, consequently, they may try to invoke this exemption from the information obligation.

Thus, the transparency requirement may prove tricky in terms of bringing to the attention of the individuals whose data is being processed the data sharing activity.

Legal basis for transfer

There are a number of types of legal basis for sharing data mentioned under data protection legislation. For the sharing of data mentioned in this section, the below potential legal basis are analyzed, in order to identify whether there was a right for the access and transfer of data. This impacts on the criminal law angles. In the case of intermediaries, legitimate interest and consent are frequently used. The applicability of the main types of legal basis are analyzed in the following sections.

Legal obligation: The legal obligation may relate to the sender of the data or to the receiver of the data. This is a situation in which it is clear which data and to whom it has to be sent. This legal basis is complemented by the legitimate interest in cases where there is a general legal obligation (e.g. to perform reporting to an authority in a centralized manner for a group of companies), if the data to be processed for the reporting and the need to transfer such data to the other members

[66] Stephanie von Maltzan, “No Contradiction Between Cyber-Security and Data Protection? Designing a Data Protecton Compliant Incident Response System”, European Journal of Law and Technology, volume 10, No 1, 2019.

of the group is not expressly mentioned under the legislation. In this case, it depends on the interpretation of ensuring state-of-the-art security measures and notification of security incidents. However, as these two legal obligations are rather general, further clarification is required in order to consider this legal basis applicable for all data sharing situations detailed above. Currently, for any data transfer not covered by the legal obligation, the legitimate interest analysis is performed prior to data sharing.

Legitimate interest:[67] The legitimate interest may pertain to the data sender or to the data receiver. However, this has to not have negative consequences on the data subject's rights and liberties. One interesting case involving the sharing of data between the members of a gas station association in Sweden[68] involved the sharing of CCTV of vehicles that left the gas station without paying. The aim was to prevent future similar actions of the identified vehicles in gas stations. This approach was considered excessive by the Swedish data protection authority due to the large-scale processing and creation of blacklists. Interpretation of the applicability of the legitimate interest as a legal basis for processing, as it has to be assessed on a case-by-case basis. Legitimate interest with the intent to ensure prevention of systemic attacks in a particular sector and, thus, comply with proper security measures in order to face existing attacks.

Consent: Consent obtained for data sharing has to fulfil (as in other cases when consent is needed) certain conditions. German courts have mentioned that it is "the authority of the individual to decide for himself, on the basis of the idea of self-determination"[69]. The consent would be difficult to implement as legal basis. For data pertaining to clients, this entails the deletion of data once the consent is withdrawn (from the IT systems of the private entity that collected the data and from the IT systems of the subsequent receivers of personal data). Further, a mechanism for managing consent has to be implemented. For data pertaining to the perpetrator, this legal basis is not practical in terms of obtaining the consent and ensuring withdrawal thereof.

Public interest: It may be argued to some extent, that, as NIS Directive includes the basis for sharing of data on types of attacks, it can be a basis for sharing of

[67] ECJ, C-13/16, Rigas on the interpretation of the legitimate interest concept.

[68] Nymity, "Deciphering legitimate interests under the GDPR", https://info.nymity.com/deciphering-legitimate-interests-under-the-gdpr , page 24, last accessed on 5 December 2019.

[69] German Constitutional Court, BVerfGe 65,1; 1983.

minimum personal data related to such attacks in order to ensure the public interest covered by this directive.

Relevant for the legal basis for sharing data are cases when data is collected for one purpose and transferred for another (subsequent) purpose. In such cases, a compatibility test had to be performed between the initial and subsequent purposes. The purpose of the processing should match or be similar to the purpose for which the data had been collected. This ties in with the reasonable expectation of individuals in terms of the processing of their data. Such aspects have been detailed by Working Party Article 29 in the context of behavioral advertising[70] and in relation to purpose limitation.[71]

Thus, any processing following collection should be checked for the purpose compatibility test.[72] A different purpose does not necessarily entail that it is incompatible to the initial purpose.

The compatibility test focuses on the following points:

- Similarity between the initial purpose and subsequent purpose.
- Reasonable expectation of the individual with respect to the subsequent purpose. This depends on the first point on similarity of purposes and on the information that was provided to the individual at collection time.
- The types of data processed and the consequences of the subsequent purpose on the individual needed are similar to the initial purpose.

Further, the compatibility test is performed on a case-by-case basis. Therefore, for each type of data sharing, if other legal basis is not applicable, the compatibility test has to be performed for each type of personal data shared.

The usual example given in this respect is the improvement of a mobile application. The subsequent purpose helps in the service provisioning that represents the initial purpose. Individuals are expecting that the mobile application will be improved from a technical and functionality perspective. This can be supported further by an information provided to individuals about such processing for data actually needed for technical and functionality improvements. This is applicable if the data processed is limited to the data needed for the technical and functionality

[70] Working Party Article 29, "Opinion 2/2010 on online behavioral advertising", 2010.

[71] Working Party Article 29, "Opinion 3/2013 on purpose limitation", 2013.

[72] Sabah S. Al-Fedaghi, "Beyond purpose-based privacy access control", ADC '07 Proceedings of the eighteenth conference on Australasian database - Volume 63, Pages 23-32.

improvements (there are some examples in the below sections in which more data than needed is transferred).[73] Data processed for this subsequent purpose impacts only the technical improvements of the mobile application, without any impact on the type of service or manner of providing the service to the individual.

An example of data used for subsequent purposes is mentioned in a decision concerning a medical clinic that collected the email address of individual with the purpose of establishing appointments and making surveys about the satisfaction of the individuals with the medical services provided by the medical clinic. The medical clinic subsequently used the email addresses to send commercial communication (containing various offers of medical services) to the individuals. In this case, the court clearly stated that such subsequent use of the data is not in compliance with the initial purpose for collecting the data and is not in compliance with data protection legislation (as it also required the consent of the individual for receiving marketing materials).

The public interest legal basis and vital interest legal basis may also be considered. However, there applicability is rather narrow.

The legal basis has to also be analyzed in case of transfer of data to a country without an adequate level of protection of personal data. This may occur in case of aggregation of data by a private entity at a global level or in case potential threat information is posted on a publicly accessible website.[74] In such cases, personal data may be transferred if it is expressly required for public interest or vital interest purposes. In addition, the transfer can occur in case standard model clauses are signed between the entities that participate in the data sharing exercise.

[73] ENISA. 2013. Detect, "SHARE, Protect – Solutions for Improving Threat Data Exchange among CERTs". (Oct. 2013). https://www.enisa.europa.eu/activities/cert/support/information-sharing/detect-share-protect-solutions-for-improving-threat-data-exchange-among-certs/at_download/fullReport , last accessed on 28 February 2020. Václav Stupka et al., "Protection of personal data in security alert sharing platforms", 2017, https://dl.acm.org/doi/10.1145/3098954.3105822 , last accessed on 28 February 2020.

[74] Cedric M.J. Ryngaert and Nico A.N.M. van Eijk, "International cooperation by (European) security and intelligence services: reviewing the creation of a joint database in light of data protection guarantees", International Data Privacy Law, Volume 9, No. 1, 2019.

Data minimization principle

This legal requirement translates into the following three aspects that should be taken into account.

Only data needed for the processing purpose should be collected and processed, with no excessive data being collected, stored or processed[75]. This depends on the specificity of the purpose for processing or sharing data. Thus, only data needed for the processing should be collected. Further, if data has already been collected and is being stored, for subsequent processing[76] or new iterations of the initial processing a verification has to be performed before the data processing or data sharing takes place in terms of the amount of data to be shared. One example in the case law of the ECJ[77] refers to metadata collected about individuals. In this decision, the excessiveness of data collected for the purpose of providing a calling service, such as metadata on date, time, duration and type of a communication, identification of communication equipment and location thereof, the number called and an IP address for internet services, as such information could provide a very detailed profile about an individual.[78]

Data protection aspects and especially data minimization is relevant also in case of BYOD models, as, in such cases, generally, it may be argued that only the data pertaining to the work container can be extracted without the consent of the employee at the moment of extraction, as the internal policies and procedures generally provide for such situations from the outset.

One needs to analyze whether using aggregate data is sufficient for the purpose of the attack data modelling and threat prevention. In certain situations, the aggregation would lose the granularity that the data modelling requires. In other situations, there is a need to be able to revert to the individuals whose data was profiled in order to be able to prevent attacks and understand if multiple types of

[75] Wright, D., & Raab, C., “Privacy principles, risks and harms", 2014, International Review of Law, Computers & Technology, 28(3), 277–298.

[76] ECHR, Gardel vs. France, para 62, concerning the inadequate use of data.

[77] ECJ, Cases C293/12 and C594/12, Digital Rights Ireland, 2014.

[78] Michael Barbaro & Tom Zeller, “A Face Is Exposed for AOL Searcher No. 4417749”, NY TIMES, Aug. 9, 2006 details how anonymized search logs can be used to re-identify individuals. Christopher Kuner, *European Data Privacy Law and online business*, 2nd Ed. Oxford University Press 2007, page 91-95 – IP personal data

attacks are used by a particular perpetrator and, thus, the aggregation brings issues in achieving this purpose.[79]

The data minimization requirement is linked to the limitation of entities/individuals having access to the data and to the period for which such access is granted. Thus, it should be analyzed in each case of data sharing. This ties in with the requirement of any stakeholder (which should include intermediaries) to limit access to the personal data on a need-to-know basis. In cases where data is stored by the intermediaries for their users, one may argue that the obligation is applicable for the users and not the intermediary. In this respect, further legal clarification on the level of support the intermediaries should give to users for setting-up a default level of access management and, implicitly, security of data.

In practice, entities use third parties to set up the communication system for threats or potential threats. As mentioned above, the need-to-know principle can be implemented through a centralized system that stores all data shared between the parties. If the transfer is not performed through a centralized system, it would be much more difficult to manage from a regulatory perspective. Thus, a centralized system would be preferred. Examples of sharing formats on the market include STIX and TAXII.[80]

For the centralized system, if algorithms are in place in order to identify patterns of attacks or correlation between multiple attacks, the algorithm has to be analyzed from a data protection perspective. If automated decisions are taken based on the outcome of applying this algorithm, the specific requirements on automatic decisions have to be implemented.

The establishment of limitation of data sharing can also take into account (especially in terms of legitimate interest)[81] the number of individuals potentially affected by the data sharing, the number of individuals potentially affected by the identified security incident, percentage of IT system having the particular vulnerability, level of impact of the vulnerability use, including economical.

[79] Kaltheuner, F. and Bietti, E., "Data is power: Towards additional guidance on profiling and automated decision making in the GDPR", IRP&P.

[80] ENISA, "Technical Guidelines for the Implementation of Minimum Security Measures for Digital Service Providers", December 2016, https://www.enisa.europa.eu/publications/minimum-security-measures-for-digital-service-providers/at_download/fullReport , last accessed on 8 March 2020.

[81] Cf. Céline van Waesberge and Stéphanie De Smedt, "Cybersecurity and Data Breach Notification Obligations under the Current and Future Legislative Framework", 2016, EDPL .

The centralized system should be created with rules for deletion of data that becomes irrelevant or that does not need to be in the system. Further, when participants retire from the system, the data they send can be either deleted or not, depending on the legal basis for transfer and the data they received can be deleted or not, depending on the legal basis.

Mechanisms for new participants (and history of data they can see), making sure that perpetrators (whose data might be stored in the system) do not become participants themselves. If no mechanism is in place to prevent such abuses by participants, the participants can be considered accomplices to the perpetrator.

A framework for handling data subject requests and deletion of data should be established. Further, the receivers of data should inform the other participants of any data breach on their side concerning the data obtained from other participants. Each participant should be obliged through the participation agreement not to use the data for active defense mechanisms that are illegal.

Aggregation is performed also by network companies through the metadata obtained from multiple clients from all over the world. In this case, there are two levels of agreements to be set in place: one between the entities from which the data is collected and the network companies and one between[82] the network companies and the recipients of the data. For this data transfer the network company becomes data controller and it is liable for all of the matters mentioned above. It has to manage the participants in terms of data access.

The use of centralized SIEM (security information and event management) systems with data sharing between private entities[83] may be considered excessive, as a significant amount of data is not relevant for threat prevention, as it represents day-to-day activity of the private companies. Thus, in terms of legal basis for data processing, this cannot be substantiated, as it does not fall under a legal requirement and the interests of the private companies do not surpass the rights of the data subjects whose data is aggregated, profiled and shared with a large number of entities. This type of data sharing is likely to be considered intrusive by the data subjects. As a result, if a SIEM centralized system is used by multiple private entities, the data of each entity has to be kept separately.

[82] ENISA, "Incident Notification for DSPs in the Context of the NIS Directive", https://www.enisa.europa.eu/publications/incident-notification-for-dsps-in-the-context-of-the-nis-directive/at_download/fullReport , last accessed on 8 March 2020.

[83] Stephanie von Maltzan, "No contradicton between Cyber-Security and Data Protecton? Designing a Data Protecton compliant Incident Response System", European Journal of Law and Technology, volume10, Issue 1, 2019.

Of course, in terms of current legislation, the data sharing for significant security incidents is generally defined from private entities to public authorities and between public authorities. Unfortunately, even for these scenarios, there is no guidance on the types of personal data actually needed for the data processing purpose. Further guidance is needed and can be created on a scenario-based approach, depending on the type of attack. Further, the NIS Directive applies only to certain sectors and to certain private entities within those sectors.

The transfer of data concerning threats and potential threats may be made to third parties that are private entities, including entities within the same group of companies, to private organizations and to entities that have set-up the security operation center, that operate the SIEM (security information and event management) or that centralize the security incident data. These may include transfer of confidential data, that, as per legal requirements, either should not have been in the possession of the sender or that should not be disclosed by the sender.

Other legal implications of data transfers

In relation to the data pertaining to the perpetrator, this action may involve transfer of data without a right. In this case, the criminal offence of transfer of data obtained without a right may be applicable.[84]

Further, violation of private life may also be applicable for private data obtained without a right and for transfer of such data to third parties.

Access to beacons or to malware found in the files copied by the perpetration should not be given to other private entities, as this would constitute, aside from the criminal angles mentioned above, breach of other legal provisions in terms of surveillance.

The transfer of any data that was obtained illegally by the victim, by perpetrating the criminal offences related to entrance, change or disturbance of the IT system of the perpetrator (or other IT systems involved in the attack) constitutes in itself a criminal offence. Further, the setting-up of a system by private entities in a specific sector to share such information among them constitutes a criminal offence

[84] Irish Data Protection Commissioner, “Data Sharing in the Public Sector”, 2019. Gina Fisk et al., “Privacy Principles for Sharing Cyber Security Data”, 2015 IEEE Security and Privacy Workshops, https://ieeexplore.ieee.org/document/7163225, last accessed on 28 February 2020.

perpetrated by all members of the system, irrespective if they shared or not data, as they are accomplices to any criminal offence perpetrated in this respect.

Data in the files that were copied by the perpetrator (for instance, data used by a private entity to find them on Darknet), should not be given to any private entities (aside from a service provider that actively searches on behalf of the victim for these on Darknet), as this would constitute breaches in terms of data protection and criminal offences related to professional information.

10.2 Handling Whistleblowing Related to Data Security

In 2019, the EU has adopted a Directive[85] outlining principles for companies to set up whistleblowing hotlines and specific policies and procedures related to collecting information about potential breaches of the company of EU law in specific sectors (including protection of privacy and personal data, and security of network and information systems), together with procedures for investigating these and for protecting the whistleblower from any negative consequences.

We are outlining in this section the main points to consider when setting-up such hotline, as outlined in the Directive. For each specific country in the EU, further details may be provided in national legislation implementing the Directive. This type of hotline is also relevant in the context of offensive security, as, it may be the case that either employees of the company or service providers performing the offensive security exercise may wish to bring to the attention of the company certain aspects through this hotline.

The Directive is applicable for whistleblowers that are employees of a company or not, as detailed in article 4 of the Directive. The main takeaway in this section is that individuals working on behalf of the offensive security provider may also qualify as whistleblowers under this Directive. In addition, the persons to be subject to protection in the context of whistleblowing (especially against any retaliation) includes the whistleblowers themselves and, among others, the companies these work for.

The Directive covers the reporting of breaches of EU legislation. In terms of level of knowledge of a breach or a potential beach, the Directive mentions 'information on breaches' to mean "information, including reasonable suspicions, about actual or

[85] Directive (EU) 2019/1937 of the European Parliament and of the Council of 23 October 2019 on the protection of persons who report breaches of Union law.

potential breaches, which occurred or are very likely to occur in the organization in which the reporting person works or has worked or in another organization with which the reporting person is or was in contact through his or her work, and about attempts to conceal such breaches".[86]

Companies, as per the Directive, have to establish internal and external communication channels (hotlines) which includes anonymous reporting, together with mechanisms for analysis, including the following steps:

- Designation of impartial person/department for receiving information.
- Record keeping of the exact information received – e.g. recording of conversation, saving of emails.
- Establishment of an impartial team to investigate any allegations in a diligent manner.
- A reasonable timeframe for resolving the allegation and timely feedback on the progress of the investigation.
- The company to ensure confidentiality about the identity of the whistleblower during the investigation and afterwards, except for any legal obligations to make such disclosure, like disclosure to authorities or courts of law.

In certain circumstances outlined under the Directive (such as, prior disclosure to internal hotline without a proper response or imminent danger to public interest), public disclosure of information about breaches can fall under the Directive and ensure protection for the whistleblower.

Any disclosure in such circumstances relating to IT security is essential and it is recommended to analyze any vulnerability identified and disclosed through the hotline as soon as possible in order to avoid third parties or public disclosure of such vulnerabilities.

The company should not allow for retaliation to occur, including aspects such as suspension, dismissal, harm (including any reputation damage), discrimination. Further examples are included in article 19 of the Directive.

[86] Please see Article 5 of Directive (EU) 2019/1937 of the European Parliament and of the Council of 23 October 2019 on the protection of persons who report breaches of Union law.

10.3 Conclusions

When implementing offensive security, generally personal data is also processed. This chapter includes the main principles of data protection tailored for the offensive security actions.

This entails the process of collection of data within the organization (either by the employees of the organization or by the third-party service providers) and the transfer of data outside of the organization (e.g. to authorities, to other organizations).

11. Cyber resilience

Cyber Resilience became important because traditional security measures are no longer enough to protect the organization. In our days, we cannot be only reactive to cyber threats. We need to build for our organizations the capability to anticipate, prepare for and adapt to changing conditions and withstand and recover rapidly from disruptions.

Cyber Resilience is the measure of an organization's ability to continue to work normally while attempting to *anticipate*, *prevent* and *detect* any cyber threats, applying c*orrect*ions to the operational environment if needed, and *respond* to, and *recover* from those cyber threats. An organization is cyber resilient when it can *defend* against cyber threats, has an efficient *cyber security risk management* in place and can guarantee *business continuity* during and after cyber incidents.

The goal of cyber resilience is to maintain the entity's ability to deliver the intended outcome despite adverse cyber events.

In order to establish and maintain cyber resilience within organization, we need to involve everyone working or doing business with that organization. In this respect, we should consider prevention as highly important for cyber resilience as well as user training and awareness. Let's not forget that the goal of cyber resilience is to keep the organization operational even in unexpected circumstances.

In order to establish an effective cyber resilience we should have:

1. A well-defined strategy to drive properly and to improve continuously the cyber resilience implementation;
2. A clear understanding of what the organization's critical assets are;
3. A clear view of the organization's key threats and vulnerabilities;
4. An assessment of the organization's cyber resilience maturity;

5. The design of appropriate cyber resilience plans using best practices and guidance;

6. An appropriate balance of controls to prevent, detect and correct issues in the operational environment;

7. An appropriate incident response process to assure effective response to security incidents and to assure proper escalation to business continuity and disaster recovery plans if required.

8. A continual review and improvement process allowing the fine tuning of the cyber resilience implementation;

A strong cyber resilient program ensures continuity of operation with minimum impact to business despite any incident. As an operational capability, cyber resilience is an iterative process providing the means to anticipate, identify, protect, and detect an attack, and respond and recover from it needed. The following mindmap depicts this iterative process:

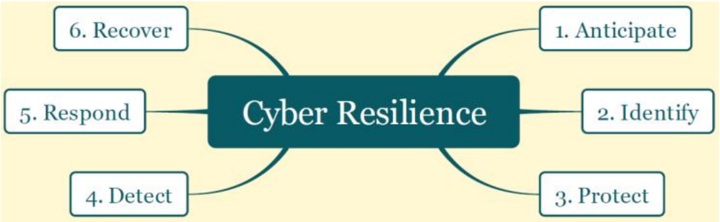

Figure 23: Cyber Resilience iterative process

The operational capability presented in the above figure is not enough to assure Cyber Resilience within our organization. We will consider them as operational objectives of Cyber Resilience, but we need more than the operational capability to build our Cyber Resilience Framework.

Moving further from the eight requirements defined above, we can conclude that for cyber resilience framework we need: strategy, design, transition, operation and continual improvement. If we look to the available frameworks, the one containing all these components is ITIL. The below figure presents the ITIL framework.

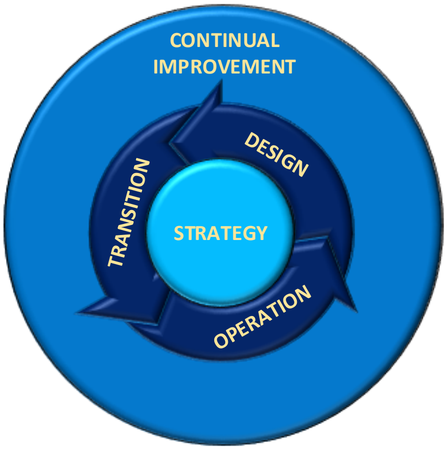

Figure 24: ITIL Framework

Many organizations have ITIL in place and are familiar with the ITIL framework. It became a standard for the IT organization worldwide and it is a strong and mature framework. Because it contains all the required elements we need to define a solid framework, it makes sense to use this framework also for cyber resilience. Our goal is to build a prioritized, scalable, and cost-effective path for our organization to be cyber-resilient.

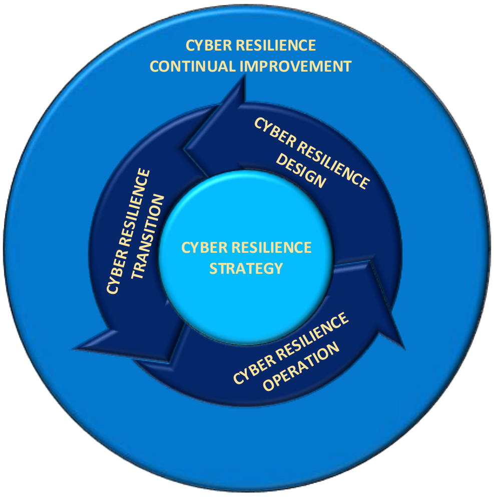
CYBER RESILIENCE CONTINUAL IMPROVEMENT
CYBER RESILIENCE DESIGN
CYBER RESILIENCE TRANSITION
CYBER RESILIENCE STRATEGY
CYBER RESILIENCE OPERATION

Figure 25: Cyber Resilience Framework

Let's depict the content of each component of the Cyber Resilience Framework presented in the above figure:

11.1 Cyber Resilience Strategy

The cyber resilience strategy must cover the entire product lifecycle as well as supporting business operations. Good cyber resilience is a complete, collaborative approach driven by the organization's board, but involving everyone in the organization focusing on people, suppliers, and resources. Effective cyber resilience requires an organization-wide risk-based strategy that proactively manages the vulnerabilities, threats, risks and impacts on its critical information and supporting assets.

A cyber resilience strategy cannot be effective if risk management is not the foundation. Cyber resilience controls are best determined when a comprehensive cyber risk management approach is adopted, which understands the enterprise strategy and associated cyber risk exposure in the ever-changing business landscape.

We need to ensure that cyber resilience activity is based on clearly understood objectives and supports the achievement of the organization's goals and it is aligned with the organization's strategy.

Both strategies have in common the organization's critical assets. So, we need to identify the organization's critical assets (what services, information and systems are the most important for the business) and what are the potential threats they might face.

The Cyber Resilience Strategy should outline:

- The importance of Cyber Resilience for the organization;
- The organization's vision and mission regarding Cyber Resilience;
- The organization's Cyber Resilience objectives;
- The organization's cyber risk appetite;
- The organization's stakeholders high-level requirements;
- The framework, and high-level approach to Cyber Resilience;

- The organization’s resilience targets and implementation plan;
- A narrative about how the cyber resilience program will be delivered, managed and funded;
- A roadmap on how to continuously improve Cyber Resilience Maturity within organization;

The following are some representative success factors for a strong Cyber Resilience implementation:

- A clear understanding of the cyber resilience program ownership;
- A clear definition of roles and responsibilities;
- The alignment between the Cyber Resilience Strategy and the Business Strategy;
- A correct and complete identification of the organization’s critical assets;
- A clear view of the organization’s key threats and vulnerabilities, particularly those targeting critical assets;
- An appropriate balance of controls to prevent, detect and correct security issues;
- An appropriate incident response process to assure effective response to security incidents;
- An internal audit process helping with the monitoring and measurement of the implementation progress, adequacy and effectiveness of the cyber resilience program;
- An assessment process for the organization’s cyber resilience maturity and design of appropriate plans to improve it using best practices and international standards and guidelines as guidance;
- Regular review and update of the cyber resilience strategy to ensure that organization can continue its business operation regardless the evolution of the cyber risk environment;

According to NIST 800-160 Volume 2[1] publication, “…any discussion of cyber resiliency is predicated on the assumption that adversaries will breach defenses and that, whether via breaches or via supply chain attacks, adversaries will establish a long-term presence in organizational systems. … The assumption of a sophisticated, well-resourced, and persistent adversary whose presence in systems can go undetected for extended periods is a key differentiator between cyber resiliency and other aspects of trustworthiness."

Remember that Cyber Resilience is about anticipating.

11.2 Cyber Resilience Design

We can consider that we achieved resilience when any of our critical assets, as were identified in the Cyber Resilience Strategy, is capable to return to its normal healthy range or one close to it regardless the security incident it faces. This capability can be achieved for the new information systems we build, but for all critical assets it cannot be achieved overnight because at least some of them are not aligned with cyber resilience practices and the alignment is not an easy task.

One of the Cyber Resilience Design's challenges is to rearchitect the current information systems (to rearchitect the existing critical assets) in order to apply cyber resilience practices. This may include redesigning and reimplementing or replacing existing cyber resources. For the new information systems we build, we just add those principles in the requirements and design phases, but for the existing information systems the changes to apply cyber resilience practices are not always applicable because of technical limitations and then we will need to replace components or even the entire technical solution.

Cyber Resilience practices are approaches that are applied to the architecture or design of business functions and cyber resources in order to achieve cyber resilience objectives.

According to MITRE Cyber Resiliency Framework PR 11-4436[2] publication, "3.2 Information Systems Security Engineering", "Some security engineering principles are specific to resilience (Stoneburner, et al., 2004):

"Principle 16. Implement layered security (Ensure no single point of vulnerability).

Principle 17. Design and operate an IT system to limit damage and to be resilient in response.

Principle 18. Provide assurance that the system is, and continues to be, resilient in the face of expected threats.

Principle 19. Limit or contain vulnerabilities.

Principle 20. Isolate public access systems from mission critical resources (e.g., data, processes, etc.).

Principle 21. Use boundary mechanisms to separate computing systems and network infrastructures.

Principle 22. Design and implement audit mechanisms to detect unauthorized use and to support incident investigations.

Principle 23. Develop and exercise contingency or disaster recovery procedures to ensure appropriate availability."

As defined by MITRE Cyber Resiliency Design Principles PR 17-0103[3], "2 Representative Cyber Resiliency Design Principles", the below figure shows representative cyber resiliency design principles:

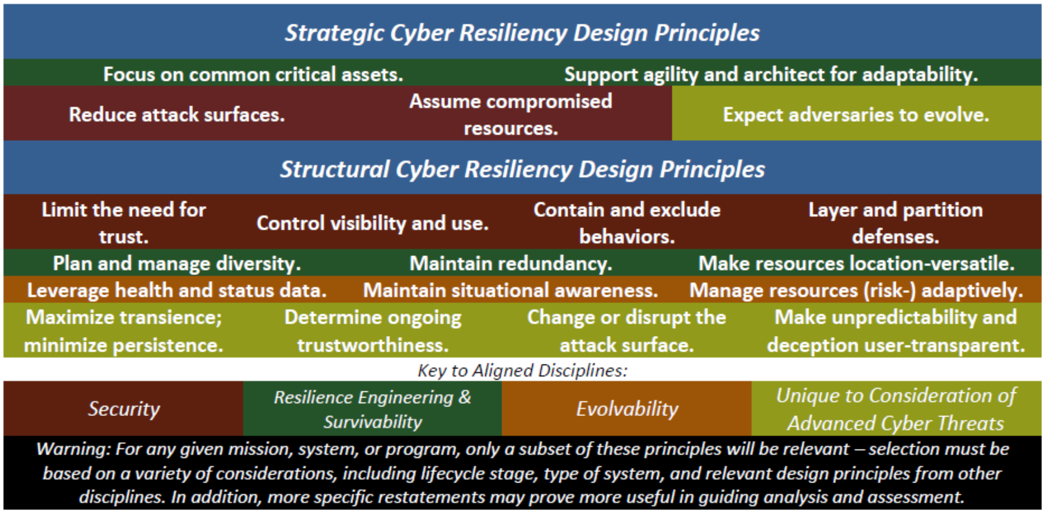

Strategic Cyber Resiliency Design Principles

Focus on common critical assets.

Support agility and architect for adaptability.

Reduce attack surfaces.

Assume compromised resources.

Expect adversaries to evolve.

Structural Cyber Resiliency Design Principles

Limit the need for trust.

Control visibility and use.

Contain and exclude behaviors.

Layer and partition defenses.

Plan and manage diversity.

Maintain redundancy.

Make resources location-versatile.

Leverage health and status data.

Maintain situational awareness.

Manage resources (risk-) adaptively.

Maximize transience; minimize persistence.

Determine ongoing trustworthiness.

Change or disrupt the attack surface.

Make unpredictability and deception user-transparent.

Key to Aligned Disciplines:

Security

Resilience Engineering & Survivability

Evolvability

Unique to Consideration of Advanced Cyber Threats

Warning: For any given mission, system, or program, only a subset of these principles will be relevant – selection must be based on a variety of considerations, including lifecycle stage, type of system, and relevant design principles from other disciplines. In addition, more specific restatements may prove more useful in guiding analysis and assessment.

Figure 26: Representative cyber resiliency design principles, MITRE PR-17-0103

11.3 Cyber Resilience Transition

The scope of this phase is to test the correct operation of the technology, controls and procedures. At this stage, we can refine the incident detection for our critical assets. At the end of this stage, we will move our implementation to production (to operational use).

At this stage we ensure that the risks introduced by the change of the operational environment through the technology and/or controls we implement are minimized through rigorous testing. Testing should be based on standard testing framework (ISO/IEC/IEEE 29119 for example for software testing). In transition to production, during testing, we need to ensure that we do not introduce vulnerabilities into the operational environment, vulnerabilities that can be exploited by hackers. As a minimum, the testing should include test against the latest OWASP top 10 risks (https://owasp.org/www-project-top-ten/).

Another important aspect of the transition phase is the users and IT staff training. Without appropriate training, users and IT staff will not have good knowledge to operate the system and will cause errors, security incidents and breaches. From cyber resilience point of view, the training session should include:

- Acceptable use policy;
- Data protection and secure data handling;
- Principles and procedures for secure information disposal;
- Secure operating procedures;

11.4 Cyber Resilience Operation

The goal of this stage is to operate the technologies, controls, and procedures, and to detect and manage cyber security and cyber resilience related events and incidents. This includes continual evaluation of the implemented controls in order to ensure they are effective and consistent.

As presented in the above figure, the operational objectives we can consider for this stage are the following:

- Anticipate potential threats against the operational environment. Prevention is the key factor and the ability to anticipate the next move of threat actors makes effective the measures taken to protect the organization.

- Identify potential threats against the operational environment. The ability to identify from earlier moments a potential threat.

- Protect critical infrastructure services. Limit or contain the impact of any potential threat.

- Detect strange events and suspected data breaches or data leaks before major damage occurs. This demands constant security monitoring.

- Respond to a detected security breach or failure. An end-to-end incident response plan to ensure business runs as usual in the face of a cyber-attack.

- Recover to restore any affected infrastructure, capabilities or services that were compromised during a cybersecurity incident. Making a timely return to normal efforts.

11.5 Cyber Resilience Continual Improvement

Continual Improvement ensures that cyber resilience continue to provide the protection as needed. After each event, including incidents there are experiences and lessons learned. All these experiences can lead to modifications of procedures, design, technology, strategy and will need to update the training for users.

Effectiveness of the implemented controls should be measured continuously in order to ensure that they operate as desired.

Risk assessment should be performed on both our organization and third parties interconnected or connecting to our organization.

Improvement ideas and suggestions should be part of the cyber resilience maturity increase roadmap and should be prioritized according to urgency for improvement needs of the cyber resilience program. There are two questions you should get answer to regarding cyber resilience maturity: What maturity level do you think is appropriate for our cyber resilience? Why is this the right level of maturity?

11.6 Conclusions

Cyber resilience should be a complete, collaborative approach driven by the board and involving every employee and business partner.

A cyber resilience strategy cannot be effective if risk management is not the foundation. Cyber resilience controls are best determined when a comprehensive cyber risk management approach is adopted.

The importance of aligning your cyber resilience risk management to the organization's enterprise risk framework cannot be ignored. Therefore, we need to embed our cyber risk governance within the existing organizational governance framework to ensure consistency in directing, monitoring, and evaluating cyber risk mitigation within the entire organization.

Let's not underestimate the potential risks regarding user training and third-party suppliers.

User training and awareness is as important as the implementation itself. Well trained and aware user will provide valuable support to protect organization's critical assets.

Poorly secured cyber suppliers are a huge vulnerability that can be easily exploited by cyber threats and expose the organization to significant damages as well as to regulatory and legal penalties.

12. Security in SDLC – Secure Software Development Lifecycle – SSDLC

The Software Development Lifecycle (SDLC) is a framework developed to design, build, test and deploy high quality software that meets or exceeds customer expectation and reaches completion within estimated time and budget.

The SDLC framework includes the following steps:

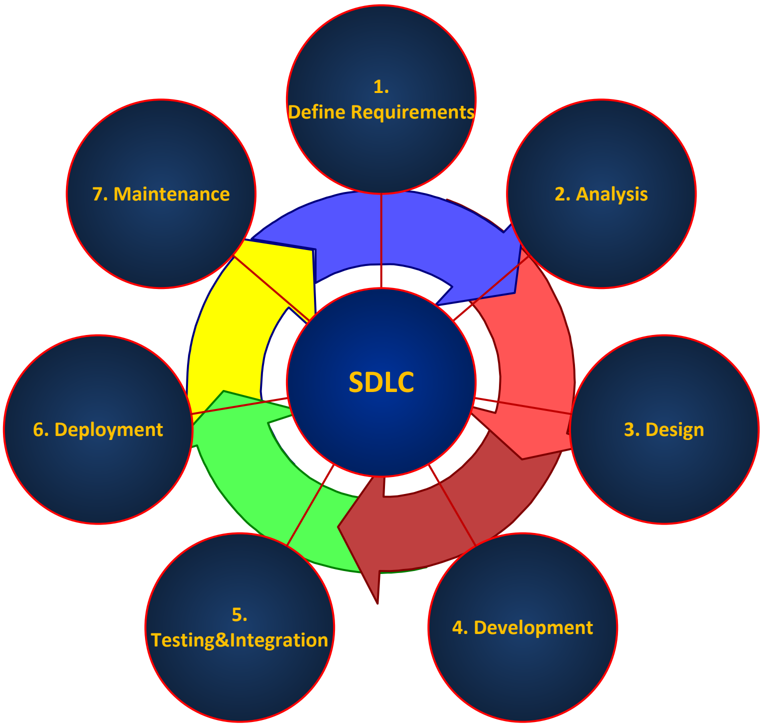

Figure 27: The SDLC Framework

The traditional SDLC steps are described below:

1) **Define Requirements** – The requirements definition phase shapes the major functions and features of the intended application or system. The following should be considered:

   a. Business requirements;

   b. Operational requirements;

   c. Security requirements;

   d. Architecture requirements;

The validation of the collected and defined requirements with the stakeholders of the intended application or system is very important.

2) **Analysis** – The requirements analysis will produce the specification for the intended application or system. Using the specification and the defined requirements it will be easy to shape the acceptance criteria you will use to assess the completeness of the application or system before deployment. This is the phase where the feasibility study is done.

3) **Design** – The requirements and specification from the previous phase will help in this stage for the creation of the conceptual design of the intended application or system. The workflows together with the detailed software architecture are created during this phase and the applied standards are defined.

4) **Development** – This is the phase where the design documentation developed in the previous phase is converted into an application or system that should meet the requirements and specification.

5) **Testing & Integration** - the testing team (or the quality assurance team) uses different frameworks to execute unit tests, functionality testing, system integration testing, interoperability testing as well as user acceptance testing in order to ensure that the code is clean (bug free) and the requirements are met.

6) **Deployment** – the tested application or system is moved to production. This phase includes the work necessary to deploy the final solution into the target production environment. Also, you need to create some documentation like: creating guides for installation, system operations, system administration, and end-user functionality. For complex software projects, you need to create a

detailed plan for implementing the application or the system across the organization.

7) **Maintenance** – this is the final stage of the software development lifecycle and includes maintenance and regular updates. Through maintenance, the deployed application or system is fine-tuned according to the user's feedback about its performance. Also, periodically, the application or system is enhanced and upgraded in order to comply with the end user needs.

As it was defined, SDLC is a framework focused on software delivery, but can be enriched with best practices in order to improve the quality of its results. **Secure SDLC** (SSDLC) is a collection of best practices enriching SDLC framework to make secure the software developed through SDLC.

Next, we will pass through the SDLC steps presented above and add to it the common best practices used to make SDLC secure, or to transform SDLC in SSDLC.

One of the most important steps for both SDLC and SSDLC is the **Requirements Definition phase**. In order to properly shape the security requirements, you need to perform a risk assessment and use it as reference, to define these security requirements. Moreover, it is very important to identify any security considerations for business requirements, operations requirements and architecture requirements. There are some security aspects you should consider, regardless of the size of your project:

- Access control. Identification and Authentication;
- Data security. Security of data in transit (communications protection). Security of data at rest (information protection and integrity);
- Media protection;
- Physical protection;
- System protection and integrity.

A good guideline for all these security aspects is NIST 800-171 Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations, https://csrc.nist.gov/publications/detail/sp/800-171/rev-2/final.

Other very important aspects you should consider from the Requirement Definition phase are the legal and regulatory constraints. For example laws like GDPR, Sarbanes-Oxley and HIPAA put constraints on data collection, processing and dissemination. This may affect how you will design your security for your application or system. Therefore, for each of the requirements topic you develop, you should define a dedicated topic of security constraints that apply and present them in the Security Requirements.

The **Analysis phase** it is also a validation step for the collected requirements and constraints during the Requirements Definition phase. If the feasibility study you perform during this phase reveals some compliance or regulatory issues because of a poor security requirements collection, the project should return to the Requirements Definition phase. The specification developed at this phase should embed security specification and the acceptance criteria you design should address all security requirements collected during Requirement Definition phase.

During the **Design phase**, you need to perform a Risk Analysis, or a Threat Modelling and at the end to review the design in order to ensure that all the security requirements and specifications are enough to protect against the threats identified in the studied threat or risk models. The solution architecture should embed all security requirements and specification. The overall technical solution should assure end-to-end security specification as defined.

At this phase, a Security Testing Plan should be developed based on the threats identified during Threat Modelling or Risk Analysis exercise.

For the **Development phase** the goal is to make sure that the written code is clean and well written. Your organization should adhere to rigid coding standards. If code runs or plays a part in gathering information for mission-critical applications, it is too important to leave it to chance and should be controlled by coding standards that are constantly kept up-to date. Obtaining the right standards and keeping them current with the latest best practices should be a top priority for organizations with a software development team.

For the **Testing and Integration phase** a security test plan should be defined prior starting the Testing and Integration phase. The security test plan can be derived directly from the results of threat modelling. Unit testing and Integration testing are the main testing tasks performed in this phase, but are not enough to ensure the security of the developed application or system.

A code review is required in order to ensure that there is no issue with the intended application or system and that coding standards are respected and applied correctly. Ensuring continuous code quality, both in the development and maintenance phases, reduces considerably the costs and risks of security and reliability issues in software as well. The code review can be either manual or automated using technologies such as static application security testing (SAST). These open-source components are usually checked using Software Composition Analysis (SCA) tools.

The key objectives of the code review are:

- The design goals are being met;
- The security objective is being met;
- The implementation is robust;
- The coding standards are respected and applied correctly;

Keep in mind that software is developed by humans, and humans make mistakes. The later the issues are discovered in the SDLC, the more difficult they are to correct and the more work that may need to be redone as a result. Static code analysis tools are not capable of detecting every potential vulnerability within an application because some vulnerabilities are only apparent at runtime, and static code analysis tools do not execute the code that they are examining. For the runtime vulnerabilities, penetration testing will reveal them and this should be done before moving the application or system to production.

The **Deployment phase** should start with a strong security testing before moving the application or system to production. A reliable security testing phase should contain not only vulnerability assessment or penetration testing scenarios, but also a testing phase of the incident response, especially if the application or system you are planning to test is mission-critical.

Platform security cannot be ignored, for while the application itself might be secure, the platform it operates on might have exploitable flaws. Therefore, platforms need to be made secure by taking appropriate measures like turning off unwanted services, running the machines on the least privilege principle, and making sure there are security safeguards such as IDS, and firewalls.

The overall security testing of the entire system should be performed in order to ensure that the developed application or system together with the platform it operates on are secured and can allow users to use it. Penetration testing will identify the vulnerabilities and will allow the implementation team to fix them in order to ensure a secure configuration within production environment.

Keep in mind that once an application or system is deployed in production, security vulnerabilities become exponentially costlier to fix. Regardless of the sophistication of the software and thorough testing there will always be glitches and bugs.

There are also on-going tasks during the **Maintenance phase** because security is an on-going process and updates, patches and enhancements to the application code are constantly required. It is a cycle that repeats itself, but security, even at the time of these modifications, must always be in focus. Performing penetration testing after implementing important updates for the application or supporting platform should be part of your Operational Assurance process.

Keep in mind that at the end of its lifecycle, a software or a platform must be disposed properly in order to maintain the same level of security for your operational environment. Sometimes, the disposal might be performed through a dedicated project where all aspects are considered in order to minimize the risks.

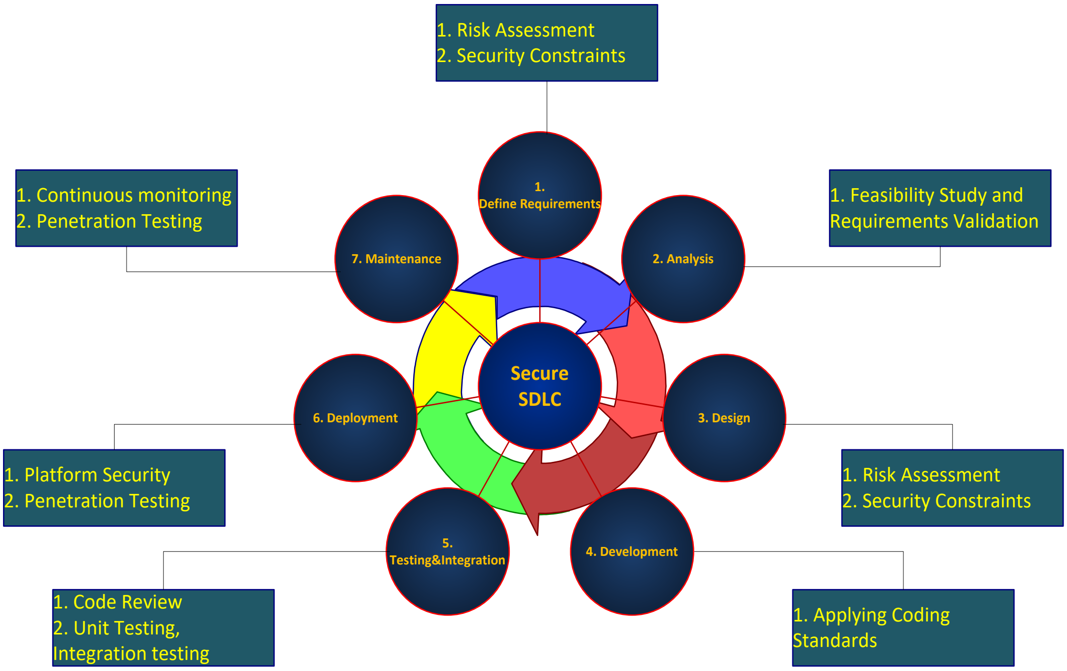

Figure 28: The SSDLC Framework

Above, the figure presents SDLC with the main best practices needed to make it Secure SDLC or SSDLC.

The continuously evolving threats demand organizations to improve their security posture and therefore, a framework like SDLC must be enriched with the best practices required to make it become SSDLC and deliver secure software.

As presented in this chapter, secure development lifecycle or secure SDLC helps developers and organizations plan, create, deploy and maintain secure software because security becomes part of the software development process in each of its stages and controls that the definition, analysis, design, development, testing, deployment and maintenance phases deliver secure results.

Mature implementations of SDLC already had in place the best practices to make it secure before being defined as Secure SDLC because in time, it became a necessity to embed security best practices in their Software Development Lifecycle and they understood that security must be everywhere. It should begin at project inception and be on the mind of every engineer during requirements analysis, design, coding, testing and deployment. This is the only way that security can be reliably improved in every application or system you build to deliver secure software and solutions.

The above SSDLC best practices were described in the context of a traditional waterfall development method, but they can naturally be extrapolated to iterative methods (agile/scrum, devops).

Conclusions

Embedding security actions and decisions into any SDLC method creates the right mindset for developers and engineers and lead to secure solution’s implementation and operation.

13. Securing Agile: Getting Speedy Results Safely

In cooperation with guest writer: Alexandru Mircea Rotaru

In fields like information security, where every second matters, using anything to save time can mean the difference between success and bankruptcy. The Agile Framework promises greater speed for projects by eliminating redundancies wherever possible. However, many believe that information security is one of the redundancies the organization can do away with. This chapter will show you why this is not the case, and how you can implement Secure Agile in your organization.

What is the Agile Framework?

The Agile Framework is a means of distributing tasks and delegating roles in a way that gets the job done as efficiently as possible. Much of how its principles stem from the Agile Manifesto from 2001[87]. Agile spreads the work over multiple cycles, called sprints, that progressively add components that come together to form the finished product; in many ways an Agile project is in itself a sprint-like entity, as they both share many of the same key components.

A sprint begins with definitions: goals (definition of done), prioritizing tasks (must have vs. should have vs. could have), the timeline (the duration of the sprint), resource allocation (in particular, which team goes where), task allocation, and many others that are specific to the organization size or any project particularities (such as working with toxic chemicals needing dedicated personnel and constant monitoring). Many of these will have been broken down as a backlog, and a planning session at the beginning of the sprint will determine which items of the backlog will get done by the

[87] See https://agilemanifesto.org/

end of the sprint. Consensus drives the Agile Framework, so the team must gather to decide how exactly to achieve the goals they set out to do for the sprint, and who does what, in order to have everyone on board with what needs to happen.

Once everyone is on board, the team moves into the execution phase. Though the work in and of itself is different every time and comes with its own challenges, a daily meeting, called a scrum, must always occur. The scrum serves to have everyone in the scrum team report on progress from the previous day, decide who will tackle which item will that day, and tackle any disruptions to business as usual.

At the end of the sprint period, the team will gather once more for a Review and Reflect stage. The Review involves presenting progress to the end-customer for approval. The team also needs to Reflect, and figure out how to do better in next sprints based on the one that just ended. The cycle then repeats from definitions, until the project's definition of done gets fulfilled.

Looking at the whole Agile process in general, there are few security concerns; they mostly revolve around keeping the sensitive data within team conversations and the Review safe. Should the employees work remotely, additional concerns regarding unauthorized access to the company system arise. Most information security complications arise from Agile's prerequisites to functioning properly and from scaling, discussed below.

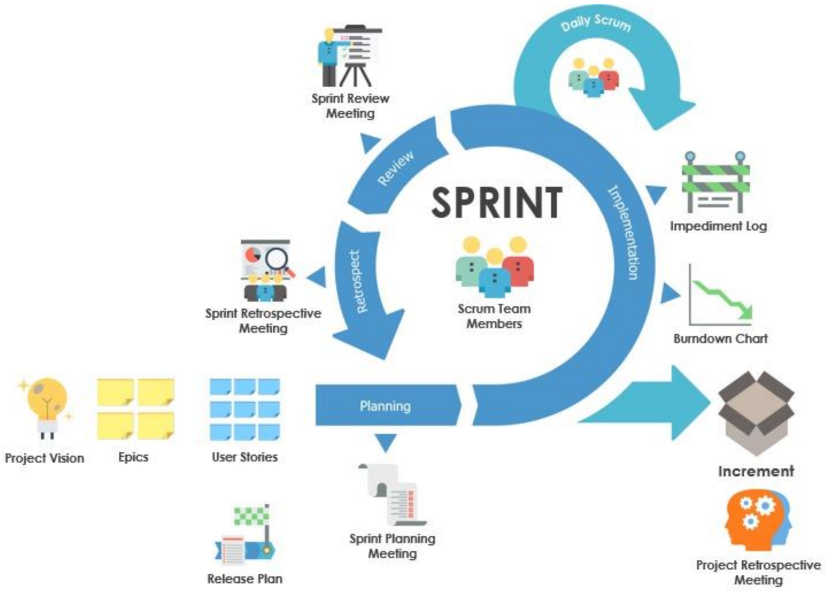

Figure 29[88]. General Structure of a Sprint.

Agile Looks Simple until You have to Scale It

The phrase “two is a couple, three is a crowd” perfectly illustrates how quickly scaling Agile can turn into a manager’s worst nightmare. Most organizations employ multiple teams, which adds a layer of complexity. While the project will mostly run as a series of parallel sprints, with all teams meeting for the Definition and Review and Reflect stages, managing the whole thing is a whole mess in and of itself.

There are multiple frameworks for Agile at scale, such as SAFE and LeSS that can help manage all concurrent sprints to ensure everything runs smoothly. The biggest issues in scaling Agile are communication across sprint teams and dependencies. To ensure effective communication, everyone should be able to maintain interactions with everyone else involved in the project, which is why Agile project should have no more than Dunbar’s Number, or ~250 people. Dunbar’s Number is the number of connections any person can maintain at any given time.

Dependencies occur when one sprint team needs to complete a task in order for another to begin theirs. All Agile projects should eliminate every possible dependency from their project. However, sometimes, they need to exist; in those situations different sprint teams need to coordinate their actions, which slows down progress.

From an information security perspective, scaling a project creates more pathways for attackers to breach your organization’s cybersecurity - with more people to target and more information being passed around. Further, there may be differences in approaches for bits and pieces of the IT solution that are being developed in different sprints (especially parallel sprints). Thus, scaling Agile also requires scaling security with it, which can be a burden if the team is too large.

That being said, the ideal scenario would employ minimizing the number of teams working on a given project; that way, scaling becomes easier to handle. To keep the other teams from idling, your organization can run multiple independent projects at the same time. Granted, removing too many employees slows down the project, which is why you need to find the ideal balance.

[88] Source: https://www.scrum.org/resources/what-is-scrum

Nothing in Life is Free. So, What am I Trading for Speed?

The main thing you will be trading for speed is scope - features that may not be critical to the project's success, but which your organization may find useful. For instance, a car cannot function without an engine or breaks, but it can work without an in-built entertainment system. Agile is an end-justifies-the-means (within moral and legal limits) kind of framework, where the only thing that matters is if the outcome of the project is able to address the issues at hand and have the features as written in the project definition of done.

Agile also needs consensus and trust to work. This, in turn, requires adjusting the organization culture to embrace speed, trust, and communication. This kind of investment is not unlike electrifying a heavily-used rail corridor to reduce the rail company's carbon footprint as part of their commitment to be environmentally-friendly. Therefore, Agile cannot function without having a strong HR Department; a culture of Diversity, Equity, and Inclusion; and a commitment to fostering communication and allowing dissent.

One thing you should never trade for speed is security. Every second that passes represents a chance for an attacker to find a vulnerability or a loophole that compromises your systems; once that happens, everything your company is legally obligated to protect becomes vulnerable, and the sprint grinds to a halt until the incident gets closed. So, either way, losing a little speed by implementing information security protocols is worth it, compared to losing days and millions in fines and lawsuits if you don't implement them. Also, the cost of implementing the information security protocols scales with the sensitivity of a project that gets compromised during a cybersecurity attack; therefore, if you tailor the information security protocols to the project's sensitivity and specifics, you should have little trouble getting buy-in from both upper management and the scrum teams.

Trust Your People Enough to Tell You when the Ship has a Hole, so You Don’t Sink with It

The big boon for information security that comes from Agile is its reliance on free and open communication: that way, people will be able to voice their concerns about something and usually help mitigate a vulnerability before it even becomes a problem.

That being said, there’s no room for “ego” in “Agile.” Your teams need to be able to openly criticize anything wrong with the project, otherwise you could be leaving vulnerabilities undiscovered; this is why consensus must drive a project: everyone brings something to the table through their unique perspective, meaning that any team member might miss something that another might find obvious.

Begin with the End in Mind: Getting Something Done

The whole purpose of using Agile is solving a problem: that’s the entire reason you have a project to begin with. Any problem has a multitude of solutions, some better than others, but your project needs to have at least one by the end - which can then be improved as part of a new project. This leaves room to play with scope, as some things may be done too soon, while others may get delayed for any number of reasons - the most time-consuming of which is when integrating the various modules that each sprint team produces doesn’t go according to plan.

That being said, even though you can vary scope to your organization’s desire to stay on schedule, there are some things you should never compromise on - chief of which being information security. One way in which you can make sure your organization doesn’t turn information security into something beyond the project’s scope is to introduce it in the definition of done.

Just Writing "Keep the Project Secure from Cyber Threats" in the Definition of Done Sounds Unusually Simple. What's the Catch?

Keeping a project secure is a layered approach. You need to have every step of your sprint secure, then the sprint as a whole secure, and finally the project as a whole secure. Securing Agile is similar in many ways to a window: the windowpane itself represents the individual steps of every sprint - one hole and you need to replace it; the window locks represent the sprints as a whole - if one is compromised, the window cannot be closed; and the window hinges represent the project as a whole - if they fail, your window is useless.

Setting-up from the beginning, based on the architecture of the IT solution and the IT landscape of the organization, certain baseline requirements for specific components can speed the analysis and implementation phases for each sprint. Further, it helps in establishing a coherence in approach and in the level of security throughout the newly developed IT solutions.

In addition, for each sprint and each component of the IT solution, additional analysis from a security perspective (e.g. secure software design, network and integration security) should be had in mind. In practice, this can be achieved through review by security staff of the stories for each sprint and, ideally, participation in the sprint meetings. Of course, certain aspects can be automatized, such as the code review part. Similar analysis should be performed after partial/full integration of all components of the IT solution and when the preparation for deployment in production takes place. In case of shortage of resources, certain milestones can be set for further apart for security review to take place.

This Information Security Infrastructure is Massive. How can I keep it Going at the Same Speed as Everyone Else?

As mentioned before, Agile begins to slow down at scale. So, one way to manage the information security infrastructure is to tailor it to every project within the first sprint. At the same time, you would need a company-wide information security team monitoring every project, to maintain the larger picture. It's like the window and hinges analogy from before: you need a solid window that locks to keep you safe, as well as sturdy hinges so it doesn't just fall off.

Agile allows you to model every project as an independent mini-organization. Therefore, you can simply use the information security infrastructure already presented in previous chapters and apply it to a project as if it were an organization - and remove anything that doesn't apply. The key difference here is that you need to integrate their work in the scrum team, which also runs the risk of creating dependencies. So, again, you need to reach consensus with management, the scrum team, and the safety team before moving forward in any way.

I need Speed to Protect my Data. Can I use Agile to do that?

Not exactly. Though Agile does provide incredible speed, protecting data is more expert-based, and the tasks are unpredictable. The Lean Framework fits the situation much better. Lean is basically a ticketing system that many companies are already using: write down every task on a ticket, prioritize the tickets received, and then execute the tickets from most to least important, minimizing work in progress.

So, Agile Can't Make Me Breakfast?

Sadly, when it comes to shiny new methodologies or technologies, people rush to implement them without a thought - which is incredibly problematic. The same goes for Agile: it's only effective in the settings it was designed to operate - which are mostly projects that require some degree of innovation, that aim to resolve concrete, specific issues.

There are other methodologies out there that can be better suited for your situation. For example, restaurants and tech support frequently use Lean to manage their operations, simply because of how unpredictable the timing of their issues is, and because they usually operate with smaller-scale projects.

Traditional methodologies work incredibly well for replicating already-existing results. This usually works in construction and mass-production, which made up most of any nation's economic output until Agile came around. Traditional methodologies also work with anything that has incredibly stringent requirements - which is usually the case in things like sculpting statues of historical figures.

So, before you dig into Agile, make sure to ask yourself: is this the kind of project that best fits the Agile framework? If the answer is yes, then, by all means, go ahead! If you're not sure, hopefully, this section can shed some light as to when to use which methodology.

Conclusion

To sum up, the Agile framework breaks down your organization's work into its atoms while also maintaining a larger picture of a project's end goal, which is also what your information security team needs to do to keep it secure. With this framework, treating every project as its own independent mini-company will help you reuse already existing information to your advantage - in true Agile style. Finally, you can simplify the work your information security and scrum teams do by minimizing dependencies.

14. Health Sector Specific Aspects: Handling Electronic Health Data – The Main Security and Privacy Guidelines

In last few decades, medical entities, such as hospitals, clinics, and individual doctor offices (either public or private) have been collecting, processing, and transferring more and more data about patients. Furthermore, given that patients usually visit multiple medical entities for the same medical issue or for multiple medical issues throughout their lifetime, the issue of data security and privacy for the transfer of data arises.

Firstly, the manner in which medical entities send data to patients and vice-versa has evolved from paper to electronic formats. Secondly, in certain cases, the data is transferred within the same medical entity or towards another medical entity. For example, the second opinion market that has been growing in recent years would benefit from a swifter manner of centralizing all data about a patient in one location and obtaining data from a single source for review. Each of these data flows has specific privacy and security aspects to take into consideration.

This chapter analyzes these types of data flows and provides insight into the main privacy and security principles to have in mind when designing such data flows and also how to approach existing legacy systems.

Consequently, the aim of this chapter is to outline the main security and privacy aspects to take into account for medical entities within their IT perimeter and when sending data to third parties (other medical entities, patients, specific doctors).

This chapter does not cover transfer of data to or from authorities or otherwise provided by law, such as the electronic patient records held by relevant authorities, clinical trials data handling or prescription handling. Further, the chapter does not cover the data collected, stored and otherwise handled by medical devices (which is detailed in the next chapter).

14.1 Know your Data Flows

As in all cases involving security, IT governance or data protection, the first step in setting-up proper internal infrastructure in line also with security and privacy requirements is to know where the data is stored and how it travels throughout your organization and towards third parties.

This may be difficult to maintain in case of large organizations. In order to have this data flow map updated, look first at the organizational culture towards security and privacy, and take a risk-based approach in which the main business processes are mapped to the IT resources used to support them. Of course, shadow IT risks are also of concern in this case and measures of mitigating them should be applied before this mapping takes place and also afterwards.

For completeness of the analysis, the locations where scanned copies of documents containing patient data need to be identified in order to properly secure them.

The main steps for the preparatory work before implementation are shown in the figured below:

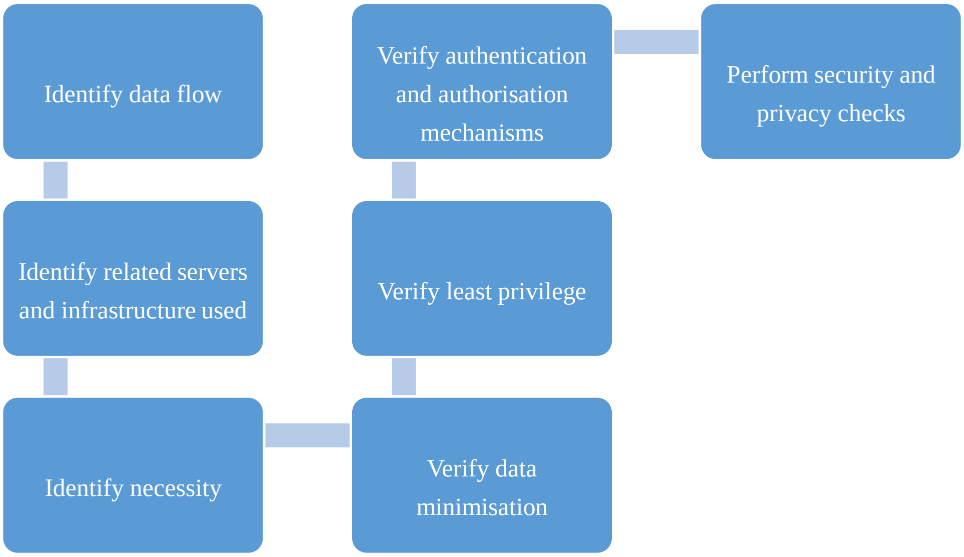
Identify data flow
Identify related servers and infrastructure used
Identify necessity
Verify data minimisation
Verify least privilege
Verify authentication and authorisation mechanisms
Perform security and privacy checks

Figure 30: Pre-design phase – identify data flows

Once these data flows are identified, an assessment of the business owner has to be made with respect to the necessity of these flows.

Further, the interactions with third parties outside of the organization (patients, other medical organizations) should be identified in order to further analyze the manner in which the data is sent and accessed from a security and privacy perspective.

In addition, the internal data flows between various departments are relevant in order to verify the "need-to-know" requirements for accessing data and in order to ensure accuracy of identifiers used for patients and of medical tests performed.

Whenever third parties are used for data storing, maintenance of IT systems holding data or data handling on behalf of the healthcare organization (defined as data processor under data protection legislation):

- Proper contract has to be in place concerning the data to which they have access, including from a data protection legislation perspective;
- Data minimization and need-to-know principles have to be implemented;
- Retention periods and deletion triggers are established and implemented;
- Ensure that any transfer of data outside the EU (to a country without an adequate level of protection of personal data) complies with data protection legislation;
- The third-party warrants for and undertakes to hold proper security measures, audit or monitor the implementation of such security measures;
- The third-party undertakes to replicate all of the above requirements with any of its sub-contractors;

Whenever data is sent to third parties holding or analyzing it on their own behalf, defined as data controllers under data protection legislation (e.g. for other healthcare services offered to the patient, for clinical trials):

- Adequacy of basis for transferring should be analyzed and confirmed;
- Transfer only the data needed under the identified transferring basis, thus implementing the data minimization principle;

- Ensure adequacy of transferring data outside the EU (to a country without an adequate level of protection of personal data), if this transfer occurs;
- Ensure proper information of individuals whose personal data is being transferred (e.g. patients) about the data transfer.

14.2 Communication between Systems within the same Organization

Healthcare organizations can have multiple departments that need to be interconnected in certain situations – e.g. a patient is transferred from one department to another.

This transfer of data can be done in multiple manners and it definitely depends on the legacy systems in place within the healthcare organization. Some organizations use share drives, some use paper, some use email and some have a centralized IT system.

Whereas we can build security on top of each of these to some extent, the preferred solution considering efficiency, time and costs on the long run and limitation of organizational risks (e.g. sending data to the wrong recipient or losing data) is the centralized IT system (or interconnection of existing IT systems within each department).

For this centralized IT system, the access management principles in the following sections are essential, which also entail proper management of user creation and deletion.

Aside from the data access, the main point to consider is the identification of a patient within all departments alike. This can be done by using a unique identifier for that patient, either unique within the hospital or uniquely issued by state authorities.

Especially for paper using organizations, the proposed solution, aside from enhanced security of the data, can ensure also proper accuracy of data, as data will not need to be re-typed in a different IT system and proper identification of medical tests and patients can be performed much easier.

In terms of access rights, of course, each department has to have access only to its patients and only to the data needed for their purpose. For example, the cardiology department may need access to the entire medical record of a patient transferred from the internal medicine department, as it needs to know all medical issues to avoid prescription of medicine that can negatively affect the patient's health. However, the laboratory that needs to process a test may need access only to a limited amount of data in order to identify the patient for which the medical test is performed.

These type of intricate access rights may seem complex to design, but, generally, from the data flow mentioned above the need for data access should be clear. If additional situations appear, a specific one-off transfer permission with choosing data to be transferred may be implemented.

If clinical trials are conducted within the healthcare organization, separate IT systems, interconnected or not with the healthcare organization ones should be considered, as this represents a distinct data processing activity from the general healthcare service provider.

14.3 Communication with Patients

There are many ways in which communication with patients can be implemented. The method we are proposing in this chapter involves a designated space for communicating.

A designated space can be created for each patient or for each medical test taken by a patient. The model chosen depends on the needs of the healthcare organization.

Account per medical test: For organizations that offer patients medical tests, the approach with one user account per medical test may be easier to implement and more secure. It is easier to implement because patients may come at various times for various medical tests and the user accounts can be generated at each interaction separately and automatically, with username and passwords that can be communicated to the patient when the medical test is taken. Having a single user account for all medical test in this case, where there is no constant relation between the patient and the healthcare organization, may result in authentication issues, such as forgetting passwords and requesting password resets over the phone, potential unauthorized access to the account, higher degree of risk of unauthorized access to the account from having all health data (which is sensitive data under the data protection legislation) in a centralized location easily accessible to users outside of the healthcare organization.

For this situation, an ID of the user (e.g. patient ID) and an ID of the medical test can be provided to the patient when the medical test is taken.

Account per patient or patient visit: For organizations that have a constant relation with a patient (e.g. long terms monitoring of the patient due to a disease) and need for the patient to have at hand all of his/her medical records, it may be easier, from an organizational perspective, to have all the patient data (and test results) in a single user account.

Another hybrid option is to have, for instance, an account per patient visit. Thus, is a patient is committed to the hospital for a week, all of the relevant data pertaining to this period is stored in a separate account. Then, if the patient comes with a different medical issue two years later, a distinct account is created for this patient visit.

For this situation, a patient user ID can be setup when the account is first created and the password can be chosen then by the patient.

Further, in both use cases, the SCA (strong customer authentication) mechanism should be had in mind, with at least two out of the following three information being requested for authentication: knowledge (something only the patient knows – CNP, date of the medical test, ID of the medical test), possession (something only the patient possesses – a telephone number) and inherence (something the patient is – biometrics).

The above distinction is more of a business use case distinction. From a technical perspective, a platform that has these two options configured can be created, with automation of the first use case (test-based user accounts) and manual creation of the second (or partial automation, if the situations to which it is applicable can be easily identified).

In either case, access to the content of this designated patient space should be limited to the patient. In this respect, best practice is for the individuals in contact center and reception not to have access to the authentication credentials of the patient or to the content of the designated patient space.

Further, access of patient family members to this designated space (in either cases) should be done only after careful analysis of legal requirements on confidentiality and exemptions in this respect. The option of family members should be embedded in the initial architecture of the designated space application, with proper segregation of duties in place (e.g. proposal of account access granting by one person and approval by another person).

The proposed solution ensures a higher degree of security and privacy (especially in terms of individuals having access to data) than giving medical test results via email (to the email provided by the patient) or to the email given by the patient in the contact center of the healthcare organization. Having such types of communication methods in place increases the risk of data being sent to other individuals than the patients.

Nevertheless, if further contact is needed with the patient - e.g. notifications about upcoming checkups, repeating tests, such information can be sent to the contact details provided by the patient (e.g. email, telephone number), with minimal health data being included in this communication, to limit risks of interception or wrong recipient due to inaccuracy of contact details.

Irrespective of the legacy systems in place within the healthcare organization, the adapting of existing IT systems and development of new IT systems as per the above requirements represents improved quality of patient care and security of data held by the healthcare organization.

The below figure shows the main principles to consider when designing and developing the designed patient space.

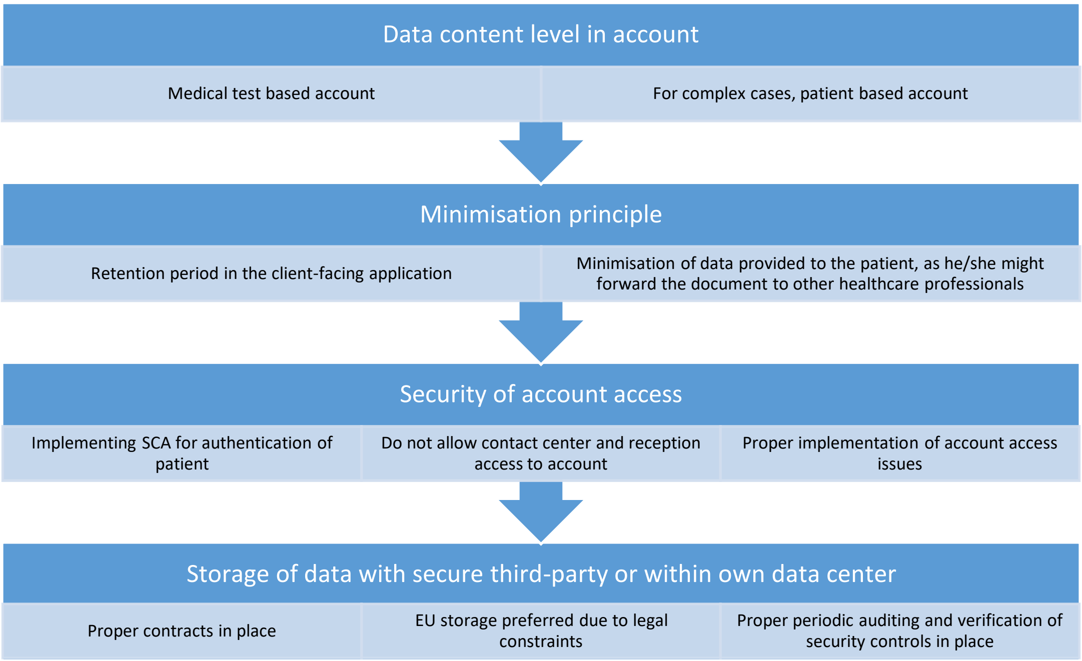
Data content level in account
Medical test based account
For complex cases, patient based account
Minimisation principle
Retention period in the client-facing application
Minimisation of data provided to the patient, as he/she might forward the document to other healthcare professionals
Security of account access
Implementing SCA for authentication of patient
Do not allow contact center and reception access to account
Proper implementation of account access issues
Storage of data with secure third-party or within own data center
Proper contracts in place
EU storage preferred due to legal constraints
Proper periodic auditing and verification of security controls in place

Figure 31: Main principles for Patient-access data platform

14.4 Communication with other Healthcare Organizations

The specific solutions proposed in the previous section with respect to communicating data to patients can be applied to communication to other healthcare organizations.

Given the need to observe the confidentiality obligation concerning patient data, there are limited situations in which data can be shared directly with other healthcare organizations. Aside from urgent situations when the vital interest of the patient is at stake, in other situations, the consent of the patient for this transfer is needed.

Therefore, the use of an online platform is essential, with proper access management in place. This can be based on the specific patients, with an account being created per third-party healthcare organization and for each patient separately. Alternatively, the access can be granted for the specific third-party healthcare organization and for a specific medical test result.

Regardless of the manner in which the granularity of data is established, before granting third-party healthcare organizations access to data, proper identification and authentication mechanisms have to be set to ensure that the individual receiving the authentication credentials is in fact the representative of the third-party healthcare provider. This can be performed by face-to-face means or, remotely, by specific authentication questions and double checks through various means – e.g. telephone, emails, official registries.

For transfer from third-party healthcare organizations, the same principles can be applied. If the third-party entity cannot provide a proper and secure manner of transfer, this can be created by the receiving healthcare organization.

When receiving health data concerning patients, confirmation from the sending entity about the consent (or fulfilment of other legal requirements) should be obtained prior to receiving the data together with confirmation that only the data needed by the receiving entity is being sent.

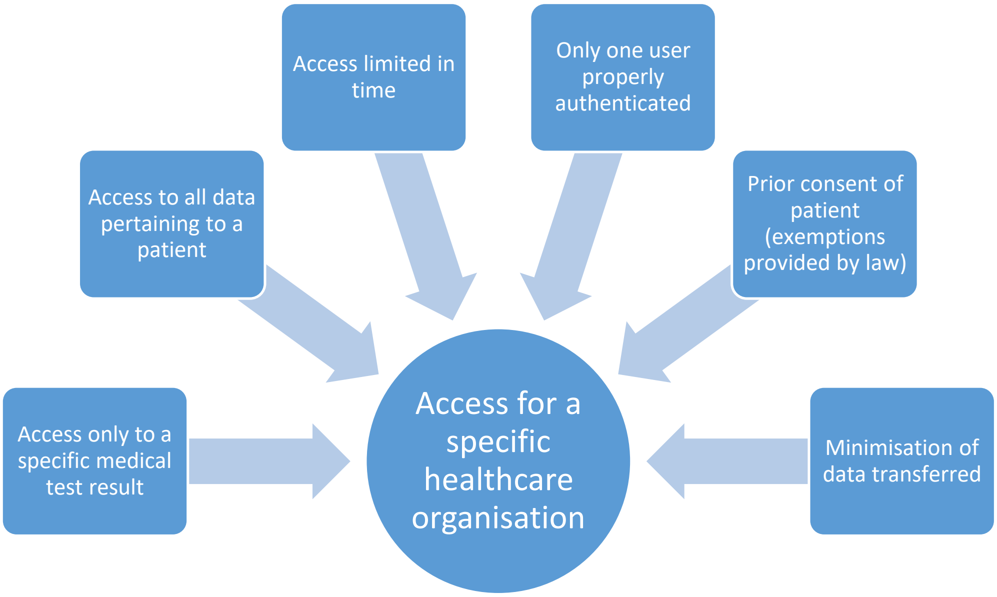
Access limited in time
Only one user properly authenticated
Access to all data pertaining to a patient
Prior consent of patient (exemptions provided by law)
Access only to a specific medical test result
Access for a specific healthcare organisation
Minimisation of data transferred

Figure 32: Requirements for granting access to third-party healthcare organization

In case of continuous transfer of data between healthcare organizations, APIs or dedicated user accounts for the other healthcare organization can be set in place or, at least a SFTP (e.g. for large files that need to be transferred) in order to transfer data securely. This should be accompanied by internal policies and procedures on the situations in which data can be transferred, which data can be transferred and the time limitation for the transfer.

14.5 Healthcare Professional Accessing Data from outside the Organization's Network

There has been an increasing use of BYOD in the past decade and this tendency has reflected also in the medical field. Even if email and share drives may be among the preferred options for viewing patient data, remote access to specific healthcare applications used by the healthcare organization is also being considered.

The practical recommendations outlined in this section refer to the relevant NIST publication on BYOD and the specifics of the healthcare sector.[89]

The first step relates to using only devices that are approved from a policy perspective, as they fulfil all the requirements in terms of security.

The second step refers to having the relevant Enterprise Mobility Management (EMM) software installed on the devices. This may be performed afterwards.

Using an EMM can help in a correlated and integrated approach towards BYOD and remote connection to resources of the healthcare organization. In this manner, proper BYOD policy can be implemented, with EMM in place and lack of possibility for the users to download the documents outside of the professional area.

Proper encryption or remote-wipe solutions for these devices can be implemented, as they add layers of protection for the confidentiality of data. The remote-wipe solution is generally included in EMM solutions.

Further, EMM solutions generally have a policy checking in place that can analyze the mobile device and flag certain vulnerabilities. The policy can be configured such that the EMM solution does not permit access to the IT systems of the healthcare organization if the conditions in the policy are not met –e.g. non-rooted devices, no mobile apps ranked as potentially having malware are installed, display protection is on.

[89] NIST, SP 1800-22 (Draft), "Mobile Device Security: Bring Your Own Device (BYOD)".

14.6 General network structuring and access management principles

General principles about network structuring and user management should be observed, including the following main directions:

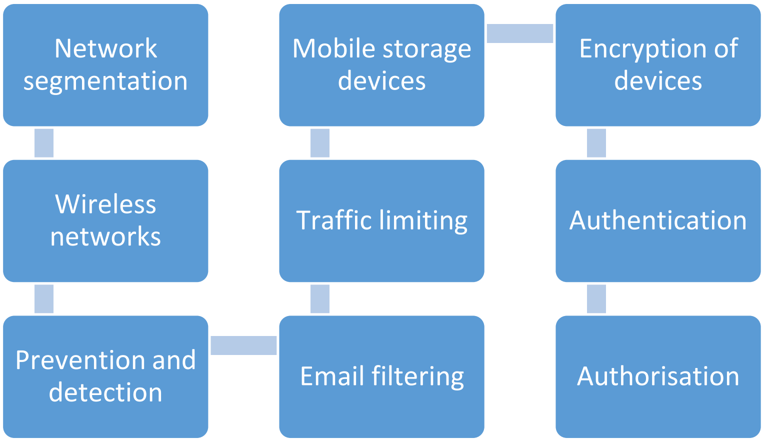

Figure 33: General entry-level security principles for networking structuring and access management

Network segmentation: Network segmentation is essential in keeping successful attacks from affecting multiple areas within the healthcare organization. The sections can be created based on the business areas within the healthcare organization and on the various data classification labels. For example, the medical devices will be kept in a separate network from the computers used by staff for communicating with the patient application server. Further, in case of clinical studies being conducted in the healthcare organization, the IT systems and data used for these should be held in a separate network.

In addition, VPN should be used for remote access to data/IT systems of the healthcare organization.

Wireless networks: Wireless networks are frequently used as entry-points for attackers and their design should be correlated to the network segmentation part. In addition, specific designated and separate wireless networks should be created for internally used devices (e.g. medical devices, computers), for the staff (with a hidden SSID) and for visitors (a general guest network).

Prevention and detection: Prevention and detection mechanisms should be set in place at specific locations within the network structure (e.g. specific nodes, close to firewalls or at the host level in certain cases). Whereas IDS is aimed at detection of potential attacks, IPS can be used for prevention as well. In this case, there are many off-the-shelf customizable options. Further, a monitoring tool aimed at analyzing at least the network traffic and server communication should be considered.

Email filtering: Emails are classic entry points for social engineering attacks and, on the long run, proper email filtering may be worth investing in.

Traffic limiting: In view of helping further on the phishing side, aside from awareness training, one may contemplate limiting traffic towards certain types of websites from being accessed through work devices (e.g. computers, tablets). These may use a significant amount of bandwidth (e.g. social media, video streaming) and may also be entry points for social engineering attacks.

Mobile storage devices: Even if this is more on the host side, a relevant aspect from a security perspective is the prohibition of using USBs on the organization devices. This ensures that the staff uses the designated channels for sending documents and limits exposure of data through mobile storage.

Encryption of devices: The organization devices used within the healthcare activity should include proper encryption so that, in case they are lost, the data they contain cannot be (easily) retrieved by third parties.

Authentication: For authentication in IT systems holding patient data, multi-factor authentication should be used in order to prevent access by attackers in case the passwords or account are compromised. Generally, this can be doubled by allowing authentication with a user account only on one device at a time. In case the staff has to access multiple IT systems at a time, SSO (single sign-on) may be contemplated for the remaining IT systems, after the first authentication takes place.

Authorization: The staff should have proper access to data based on their job description. This can be set when a new employee (or an independent co-operator) starts work in the healthcare organization and the process for granting authorization can be coordinated with the HR onboarding process. This assumes that, within the healthcare organization there is a central (or distributed) record of all employees and independently contracted staff for all departments.

It is important to keep the authorization level updated throughout the employment lifecycle. Any change in position or department has to result in a verification of the authorization needed for the new position. The rights of any leavers should immediately be removed. Any temporary staff should have temporary authorizations set. Thus, for this employment lifecycle, the authorization updating can be correlated with the HR process in order to ensure accuracy of the data access/change rights granted to the employee.

This authorization level is to be checked by each IT system/storage server accessed by staff members.

14.7 Availability of Data and Availability of IT Systems - Working towards Resilience

An essential part of the protection of health data is its availability, as this may prove critical in most situations of interaction with patients. Further, the availability of IT systems is equally essential.

In order to achieve this, the first step is the identification of the time interval after which back-ups to be created at a different and separate location than the original servers.

The main purpose of such back-ups is for easy restoration into production if the original IT systems are compromised or fail. Thus, back-ups can be created online at a different location or offline (in case of data that can be easily restored – e.g. xml files, small databases). In case cloud solutions are used for storage of IT systems, the back-up can be integrated into the cloud service solution purchased.

Regardless of the back-up option chosen, regular testing should be made to ensure that the back-ups can be restored to production easily and swiftly.

14.8 Conclusions

Whereas there are many areas to have in mind in order to ensure security, privacy and business continuity within a healthcare organization, this chapter aims at providing a starting point for the first steps in this respect, with minimal resources and costs.

The main focus in a healthcare organization is on availability of data, ease of secure communication within the organization and outside the organization (to patients or to other entities). In addition, given the mobility of staff members within the building of the healthcare organization and outside it, remote access to patient data has become more and more useful. The chapter includes basic bring your own device principles that can be easily implemented and scaled within the healthcare organization.

There are various manners in which the data (e.g. test results, medical records) can be transferred to patients or to other medical organizations. This chapter proposes specific solutions in this respect, outlining the practical implications and advantages thereof from a security (especially availability) and privacy perspective.

We further outline, for either solution chosen, main networking, user management and business continuity principles to be observed, with practical examples of implementation, tailored to a healthcare organization and having in mind the usage of less resources for implementation.

It is also important that the above principles are replicated by all the entities involved in the supply chain of the healthcare organization in order to ensure, on the one hand, compliance with applicable legal requirements.

15. Health Sector Specific Aspects: Medical Devices Security–How to Protect Them

Medical devices are fixed-function devices, designed to perform a specialized task. The lifecycle for medical devices may be as long as 10, or even 20 years. These are optimized devices to minimize processing cycles and memory usage, so they lack extra processing resources. In this context, the traditional information assurance approach Confidentiality-Integrity-Availability (the CIA triad) needs to be enriched with Safety-Reliability-Availability of the processes, devices and connected systems to strengthen the overall security posture. Nevertheless, we must include any safety functions and assess the consequences of malfunction to people, equipment and environment.

In our days, medical devices become a tightly integrated systems with complex data flows not only between devices, but also between devices and hospital IT systems. We should be aware that these devices are not only more vulnerable and difficult to protect, but also that the security compromise of any of these could result in patient harm or impact on care delivery – as we mentioned before, in addition to the traditional security concerns around data confidentiality, integrity, and availability. Considering the medical device lifecycle, the situation becomes even worst because we have to deal with different "generations" of devices and to build appropriate security measures around them in order to harden overall security.

According to NIST SP 1800-1, "All healthcare organizations need to fully understand the potential risk posed to their information systems, the bottom-line implications of those risks, and the lengths that attackers will go to exploit them… Assessing risks and making decisions about how to mitigate them should be continuous to account for the dynamic nature of business processes and technologies, the threat landscape, and the data itself…. We recommend that organizations implement a continuous risk management process as a starting point for adopting this or other approaches that will increase the security of the Electronic Health Records (EHR). It is important for management to perform regular periodic risk review, as determined by the needs of the business."

In the circumstances presented above, our security approach to protect our healthcare organization should have two directions:

- A strategic direction through a risk-based approach towards addressing threats.
- A tactical direction to harden the operational environment through Defense in Depth.

Threat modelling plays a vital part of the Security Development Lifecycle (SDL) process because it helps in identification of system vulnerabilities and threats and helps in establishing appropriate mitigation techniques. Threat modelling methodology involves optimization of Network/Application/Internet security through identifying objectives, threats, and defining countermeasures to mitigate the effects of the threat. Thus, threat modelling can be used in medical devices to optimize mitigations through identification of threats and vulnerabilities to a specific device from an organization supply chain that can harm the patient.

The context of our work is depicted by the following security metamodel in the figure below:

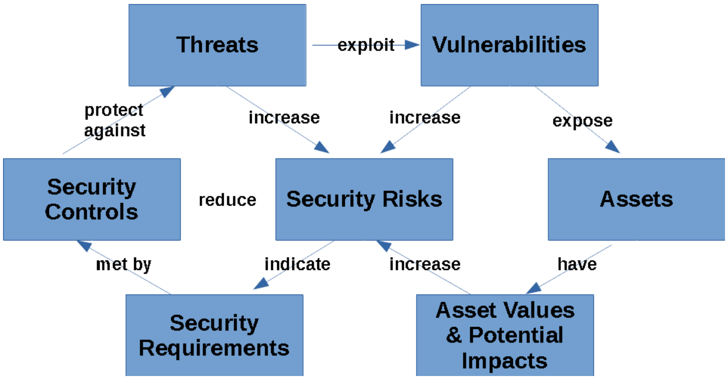

Figure 34: Security Risks Metamodel

We should be aware that the sources of the incidents are definitely worth an investigation in order to understand the most likely human based attacking vector that should be prevented or mitigated. According to any threat report, current employees are still the highest risk to cause new security incidents; however, but former employees and hackers are common threats as well. Addressing these threats with priority will improve the organization's risk posture.

15.1 Risk-based Approach Towards Addressing Threats

We need to keep permanently in our minds the fact that a compromised medical device may also serve as access point for the hospital networks with the purpose of stealing confidential data. Software incorporated in connected medical devices such as defibrillators, cardiac pacemakers, and network-connected X-ray machines is vulnerable to cybersecurity threats - some exploits could affect integrity of health data, availability of patient care, or even the manner in which the medical device operates.

Also, we need to be aware that there are different generations of technology we have to protect in a hospital, all of them doing great functional jobs, but also lack important security features and this might transform them in entry points for potential attackers.

The main objective is to design, develop and implement medical technical solutions that are secure throughout the whole life cycle without compromising patient safety. Therefore, we need to establish a framework to help us adjusting the hospital's risk posture. The framework we propose is presented in the above figure and has the following main pillars:

Design Control – through a Cybersecurity Risk Assessment we identify the various information assets that could be affected by a cyber-attack (such as hardware, systems, and patient data), and then we identify the various risks that could affect those assets. Either for a new technical solution or to redesign the security of the existing one, we need to define the requirements and applicable standards and then use them to design/redesign our technical solution, implement and test it. Before moving our project to the operational environment, we need to perform a penetration test in order to assure that our technical solution has no vulnerabilities based on current penetration testing best practices.

Operations Control – through Identity Access Management and Logging and Monitoring we keep track of the user activity. It is essential to keep our technical solution up to date (through Vulnerability and Patch Management). In case of incidents, the ability to respond and fix them is very important (through Incident Response process).

A particular care should be taken on Decommissioning. When we either need to clean-up the residual information from a business process or we need to decommission the equipment, we need to ensure that the data is wiped properly using specialized, HIPAA (Health Insurance Portability and Accountability Act) certified, tools. The same applies for equipment disposal – in this case, considering the embedded risks, we would advise you to pass through the overall process (assess the risks, redesign the solution, implement it, test it, and the decommission the equipment, wipe the data and dispose it).

The selection of appropriate security controls at various life cycle stages of a medical device depends on:

- Type of the medical device (e.g., device that contains software/firmware, device that contains programmable logic, software that is a medical device, mobile medical app, device that is considered part of an interoperable system, legacy device);
- Device classification;
- Intended use of the device;
- Operating characteristics of the device;
- Sensitivity levels of contained data;

Therefore, appropriate Inventory Management is an important process you should have in place and a very useful source of information for the security team. This should be reflected also in the internal procurement procedures to ensure that this information is clearly obtained for each medical device purchased by the organization and is properly stored in an accessible form. Further, the specific EU legislation in place or pending entrance into force (such as Regulation (EU) 2017/745) for medical device production and maintenance has to be considered in terms of information and support received from medical device producers.

Regardless how well we will perform the risk management in our healthcare organization, there will always be residual risks we will need to address. The ongoing assessment of security threats, is needed to ensure that security controls and countermeasures match the potential risks. Therefore, Cybersecurity Management Plan should be in place in order to provide guidance, effective planning, and proper risk management for the overall organization.

When we speak about cybersecurity risk assessment process, we refer to the process of evaluation of the risks to the system posed by specific threats and vulnerabilities. This process consists of four tightly linked activities:

- Analyze the capability of existing security controls to prevent an occurrence of the adverse event, detect an occurrence, and contain the impact of an occurrence.
- Estimate the likelihood (high, medium, low) of an event given the nature of the vulnerability, the capability of existing threats, and the strength of current controls.
- Analyze the potential damage an event would inflict to the system, its data, and the healthcare organization's goals. Rate the potential impact as high, medium, or low.
- Derive the risk rating from the combination of likelihood and impact.

Another important aspect in the security management of our healthcare organization is the communication with the employees. All the findings, threats, recommendations, guidance you consider useful should be assembled in a security whitepaper and made available to all employees. This is not a replacement for awareness training.

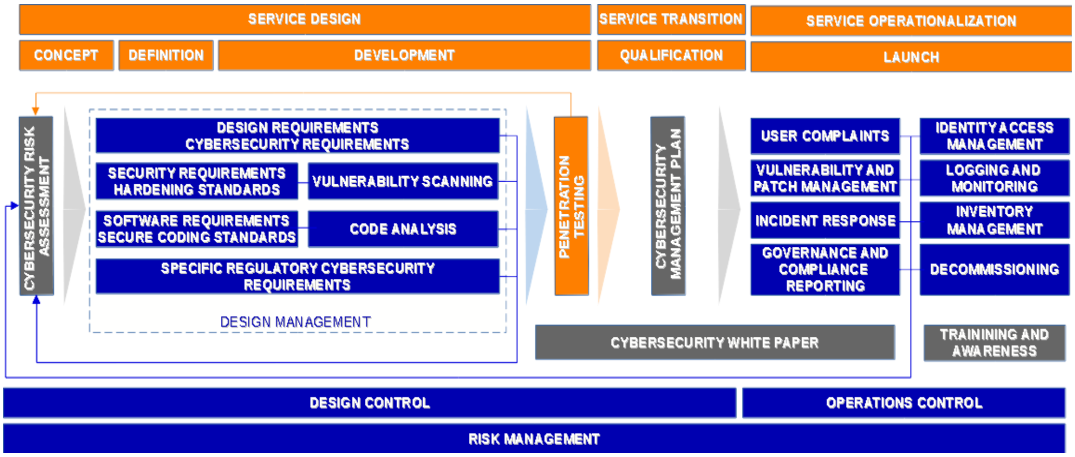
SERVICE DESIGN
SERVICE TRANSITION
SERVICE OPERATIONALIZATION
CONCEPT
DEFINITION
DEVELOPMENT
QUALIFICATION
LAUNCH
CYBERSECURITY RISK ASSESSMENT
DESIGN REQUIREMENTS CYBERSECURITY REQUIREMENTS
SECURITY REQUIREMENTS HARDENING STANDARDS
VULNERABILITY SCANNING
SOFTWARE REQUIREMENTS SECURE CODING STANDARDS
CODE ANALYSIS
SPECIFIC REGULATORY CYBERSECURITY REQUIREMENTS
DESIGN MANAGEMENT
PENETRATION TESTING
CYBERSECURITY MANAGEMENT PLAN
USER COMPLAINTS
VULNERABILITY AND PATCH MANAGEMENT
INCIDENT RESPONSE
GOVERNANCE AND COMPLIANCE REPORTING
IDENTITY ACCESS MANAGEMENT
LOGGING AND MONITORING
INVENTORY MANAGEMENT
DECOMMISSIONING
CYBERSECURITY WHITE PAPER
TRAININING AND AWARENESS
DESIGN CONTROL
OPERATIONS CONTROL
RISK MANAGEMENT

Figure 35: Risk-based Approach towards Addressing Threats

15.2 Hardening the Operational Environment through Defense in Depth

One of the most effective ways to ensure CIA is to take a defense-in-depth or layered approach when addressing privacy and security issues. A tiered approach avoids a single point of failure and supports layered controls in case one of the controls is compromised or does not operate as intended. Considering the heterogeneous environment we have to protect, Defense in Depth is the information assurance concept to apply in order to build appropriate layered security mechanisms able to keep threat actors as far as possible from our critical assets.

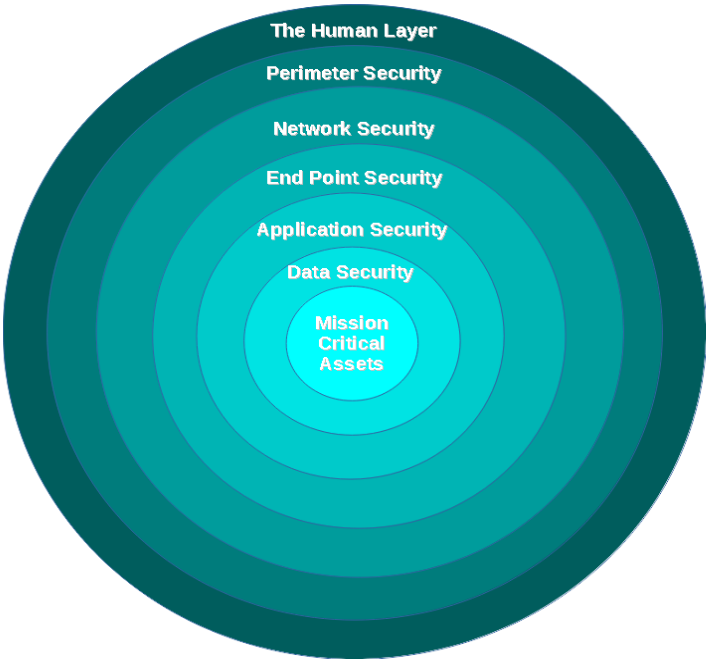

Figure 36: Security Areas to Protect Mission Critical Assets

Human Layer – Humans are considered the weakest link in any cybersecurity posture. Only through user training and awareness we can establish a security conscious culture within organization and deliver security knowledge to employees.

How can the risk be managed?

- *Produce a user security policy*: Develop a user security policy, as part of the organization’s security policy with the aim to have adequate authentication and authorization mechanisms in place, and by considering the need-to-know and data minimization principles. Security procedures for all systems should be produced with consideration to different roles and processes.

- *Establish a staff induction process*: New users, including contractors and third-party users, should be made aware of their personal responsibility to comply with the organization’s security policies as part of the induction process, with practical examples tailored to their role in the organization.

- *Maintain user awareness of the security risks faced by the organization*: All users should receive regular refresher training on the security risks to the organization together with specific messages on particular threats at a given time (e.g. a new ransomware targeting hospitals) through a delivery means appropriate for such important messages.

- *Monitor the effectiveness of security training:* Establish mechanisms to test the effectiveness and value of the security training provided to all users. This will allow training improvements and the opportunity to clarify any possible misunderstandings.

- *Promote an incident reporting culture*: The organization should enable a security culture that empowers staff to voice their concerns about poor security practices and security incidents to senior managers, either anonymously or not. This is also in line with upcoming legislation on whistleblowers.

- *Establish a formal disciplinary process*: All staff should be made aware that any abuse of the organization’s security policies will result in disciplinary action being taken against them. All sanctions detailed in policy should be enforceable at a practical level.

- *Support the formal assessment of security skills*: Staff in security roles should be encouraged to develop and formally validate their security skills through enrolment on a recognized certification scheme. The same approach should be taken for IT-related staff in order to ensure proper implementation of security requirements.

- *Perimeter Security* – the point in which we have control of our network, technology, and data. It is a defense system around our network designed to stop malicious attacks from entering.

How can the risk be managed?

There are many technologies and mechanisms available to you to help secure your network perimeter:

- Firewalls.
- Intrusion Prevention System (IPS).
- Intrusion Detection System (IDS).
- Unified Threat Management (UTM).
- Messaging Security (Antivirus, Antimalware).
- Data Loss Prevention (DLP).
- Secure De-Militarized Zone (DMZ).
- Virtual Private Network (VPN) solution for remote users' access.

Network Security - the techniques and tools to protect your network data from malicious threats and save your organization from losses. Your techniques require you to know how to protect, detect, respond, and predict a broad range of attacks.

How can the risk be managed?

Key techniques and tools include:

- Access control: To improve your network security by restricting user access and resources to just the sections of the network that clearly relate to. Role based access control upon the level of need-to-know of each defined role. Proper management of new arrivals and departures/changes in user role are an essential part of proper implementation of access control, with a specific internal process generally being designed in this respect.
- Antimalware and antivirus software: To detect malicious programs and stop them from spreading.

- Anomaly detection: Implement network anomaly detection engines (ADE) to evaluate your network, recognize anomalies and respond to them.
- Application security: Implement additional security measures for critical applications to your network security.
- Data Loss Prevention (DLP): Prevent personnel and other users from abusing and potentially compromising valuable data.
- Endpoint security: Additional layer of defense between organizational networks and remote devices.
- Intrusion prevention systems: IPD/IDS protect from known attack vectors so threats can be recognized easier.
- Network segmentation: Split the network in similar parts to improve security and give appropriate access to the appropriate traffic while controlling the traffic from other users or sources.
- Web security: Prevent web-based threats such as malicious scripts, or adware programs that leverage browsers as access points to penetrate your network.
- Two-factor authentication: Adds an additional layer of security to the authentication process by making it harder for attackers to gain access.
- Virtual Private Network (VPN): An encrypted connection over the Internet from a device to a network to ensure that sensitive data is safely transmitted.
- Data at rest encryption: To prevent unauthorized access to your stored data.
- Encrypted backups: An extra security measure to protect your data.

End-Point Security - refers to securing end-user devices like desktops, laptops, and mobile devices. It includes data security, network security, advanced threat prevention, forensics, endpoint detection and response (EDR), and remote access VPN solutions.

How can the risk be managed?

There are many technologies available to you to help secure your endpoints:

- Content Security: Anti-virus & anti-malware software.
- Patch Management.
- DLP.
- Endpoint Detection and Response (EDR): EDR “watches” the endpoint, looking for security problems or anomalous behavior that might indicate an attack or compromise.
- Virtual Private Network (VPN).
- Local/host firewall.
- End Point security policy enforcement.

Application Security - measures taken to improve the security of an application often by finding, fixing and preventing security vulnerabilities.

How can the risk be managed?

The following technologies and techniques are available to help you secure your applications:

- Static Application Security Testing (SAST): scans the application source code files, accurately identifies the root cause and helps remediate the underlying security flaws.
- Dynamic Application Security Testing (DAST): simulates controlled attacks on a running application or service to identify exploitable vulnerabilities in a running environment.
- Server Patch Management.
- Web Application Firewall (WAF): helps protect your web applications or APIs against common web exploits and bots that may affect availability, compromise security, or consume excessive resources. A WAF is deployed to protect a specific web application or set of web applications.

Data Security - the practice of protecting digital information from unauthorized access, corruption, or theft throughout its entire lifecycle. It refers to data security controls required to protect the storage and transfer of data.

How can the risk be managed?

The following technologies and techniques are available to help you secure your data:

- Identity & Access Management.
- DLP.
- Data Integrity Monitoring.
- Data Wiping Tools.
- PKI (Public key infrastructure).
- Data at rest (DAR)/ Data in use (DIU)/ Data in motion (DIM) Protection.
- Data/Drive Encryption.
- Data Integrity Monitoring.

Mission Critical Assets – devices, applications, databases and data that are crucial for your organization – without them your organization cannot operate. The following links https://www.enisa.europa.eu/publications/methodologies-for-the-identification-of-ciis/at_download/fullReport and https://www.enisa.europa.eu/publications/security-and-resilience-in-ehealth-infrastructures-and-services/at_download/fullReport describe ENISA methodology and recommendations for eHealth Critical Asset Protection.

How can the risk be managed?

The following steps should be fulfilled in order to identify and protect your critical assets:

- Perform asset management
- Perform risk assessment
- Identify critical assets
- Implement appropriate measures to protect your critical assets

The following mindmap depicts the Layered Security to Protect Mission Critical Assets presented above.

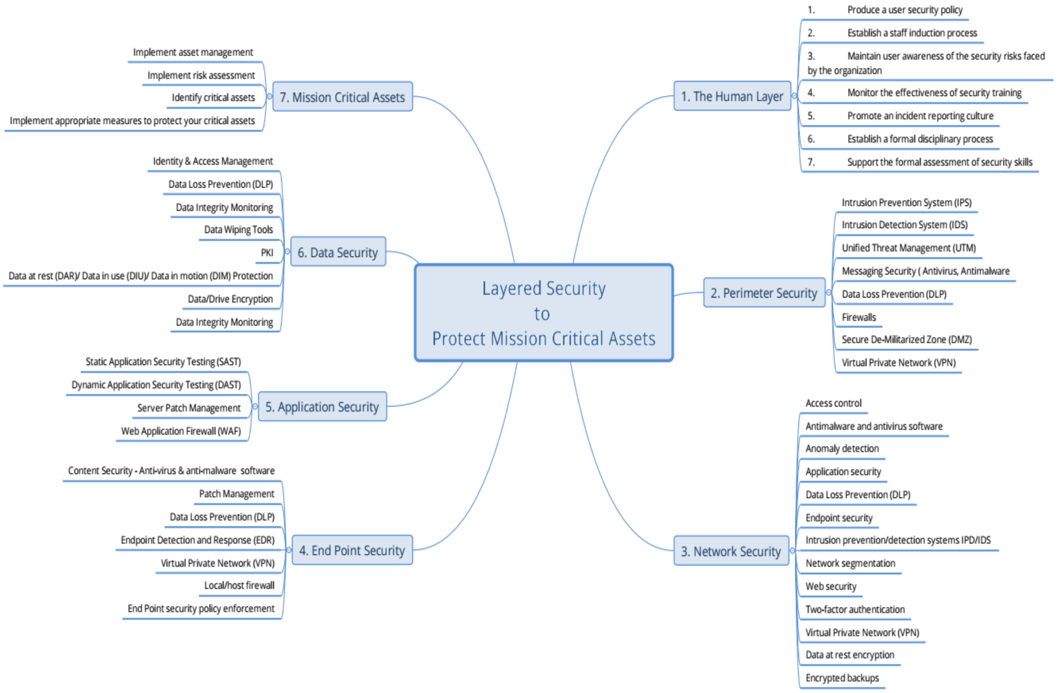
Layered Security
to
Protect Mission Critical Assets
1. The Human Layer
1. Produce a user security policy
2. Establish a staff induction process
3. Maintain user awareness of the security risks faced by the organization
4. Monitor the effectiveness of security training
5. Promote an incident reporting culture
6. Establish a formal disciplinary process
7. Support the formal assessment of security skills
2. Perimeter Security
Intrusion Prevention System (IPS)
Intrusion Detection System (IDS)
Unified Threat Management (UTM)
Messaging Security (Antivirus, Antimalware
Data Loss Prevention (DLP)
Firewalls
Secure De-Militarized Zone (DMZ)
Virtual Private Network (VPN)
3. Network Security
Access control
Antimalware and antivirus software
Anomaly detection
Application security
Data Loss Prevention (DLP)
Endpoint security
Intrusion prevention/detection systems IPD/IDS
Network segmentation
Web security
Two-factor authentication
Virtual Private Network (VPN)
Data at rest encryption
Encrypted backups
4. End Point Security
Content Security - Anti-virus & anti-malware software
Patch Management
Data Loss Prevention (DLP)
Endpoint Detection and Response (EDR)
Virtual Private Network (VPN)
Local/host firewall
End Point security policy enforcement
5. Application Security
Static Application Security Testing (SAST)
Dynamic Application Security Testing (DAST)
Server Patch Management
Web Application Firewall (WAF)
6. Data Security
Identity & Access Management
Data Loss Prevention (DLP)
Data Integrity Monitoring
Data Wiping Tools
PKI
Data at rest (DAR)/ Data in use (DIU)/ Data in motion (DIM) Protection
Data/Drive Encryption
Data Integrity Monitoring
7. Mission Critical Assets
Implement asset management
Implement risk assessment
Identify critical assets
Implement appropriate measures to protect your critical assets

Figure 37: Layered Security to Protect Mission Critical Assets

15.3 Conclusions

Both perspectives presented in this chapter, Design Control and Operations Control (through Defense in Depth) have something in common – the need of risk assessment before taking any action. Considering that we deal with limited resources, it underlines even stronger that risk assessment, vulnerability analysis and incident response are the minimum required three main pillars you should have in place in order to protect effectively your organization.

16. Remote access: Endpoint Protection

A quick google search will reveal that endpoint security “is an approach to the protection of computer networks that are remotely bridged to client devices”[90]. To put it simply, it is the practice of ensuring that devices connected to your network are protected against malicious activity, which typically translates to deployment of malware and/or exploitation of exiting vulnerabilities.

16.1 Guiding Principles

The business landscape has evolved significantly over the last few years, determined, of course, by the rapid development of new technologies. It is not uncommon in today’s business to have people using laptops, smartphones or tablets – either from the office premises or from home or while on the road. Recently, with the wide adoption of IoT devices, the landscape has become even more complex, as these “gadgets” are included into the scenery – if not connected to your network, then at least connected to a device that is connected to your network (just to give a simple example, a smartwatch connected to the company provided phone of tablet). And this are just the new additions. While making use of these “feats of technology”, we still have the traditional office network, with servers, network equipment, and end user computers. Not to mention perhaps old, deprecated devices that might still linger on the network due to some important tasks they might be performing.

As one can observe, the scene is complex and protecting it entails a fair amount of effort, creativity and, most probably, a decent budget (to put it mildly). While we

[90] https://en.wikipedia.org/wiki/Endpoint_security

cannot cover effort and specific designs, nor can we cover budgets, we will try to list some of the most used and known endpoint protection solutions out there.

Given the complex nature of the environment and the shift towards mobility, it is required that we introduce several concepts at this point:

- Defense in depth
- Zero Trust Model

16.2 Defense in Depth

“Defense in depth” principle refers, in a nutshell, to the practice of ensuring the security of a network/asset through the deployment of multiple independent controls, the logic behind it being that if one control fails, the next one will continue to protect the asset(s). It is a wide known principle and typically should be considered in all security aspects.

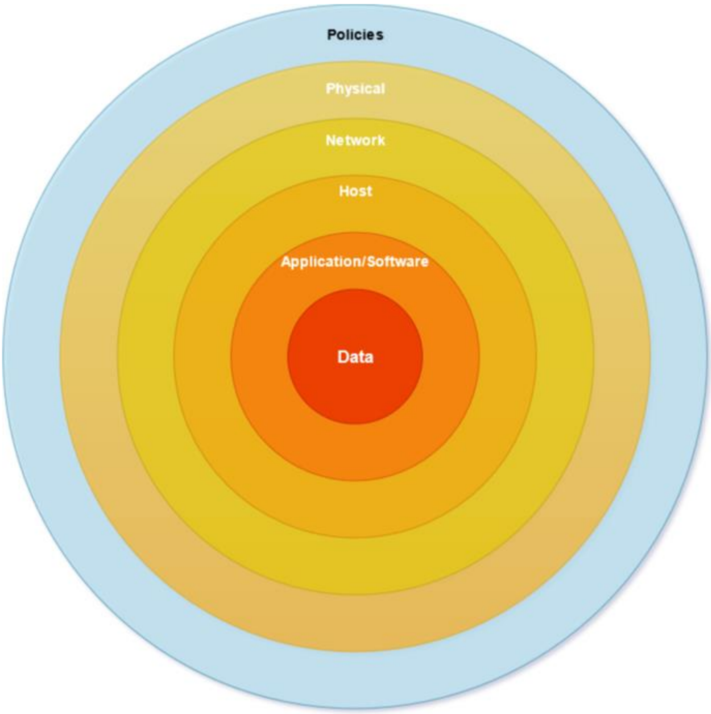

Figure 38: Defense in Depth principle

16.3 Zero Trust Model

"Zero Trust is a security concept centered on the belief that organizations should not automatically trust anything *inside* or *outside* its perimeters and instead must verify anything and everything trying to connect to its systems before granting access."[91] The concept is rather a new one and differs significantly from the traditional approach studied for decades in which you had a network and whatever was outside was malicious while the inside was to be protected. At a high level, of course. You would still have the various internal network areas with varying degrees of security requirements, but overall, that was the idea. You would have the outside zone, the DMZ and the inside network. The Zero Trust Model entails that you trust nothing and verify everything, including internal or own resources. That's the mantra. And it is rapidly gaining popularity due to the complex nature of the present-day business landscape we've just described in the chapter in terms of technology deployed.

Vendors and organizations have their own proposals and architectures for zero trust implementations. Some of the common aspects to all implementations, and that should be considered are:

- Authenticating the user
- Authenticating the host
- Definition and enforcement of access control policies for subjects, objects and data
- Enterprise Public Key Infrastructure – for managing digital certificates issues by the organization to subjects, objects, services and applications which typically are used to authenticate them and grant access to resources
- Identity Management system – creating and managing user accounts.
- Diagnostics and mitigation systems – gathering information about the assets and applying updates to configuration and software components[92]

[91] https://www.csoonline.com/article/3247848/what-is-zero-trust-a-model-for-more-effective-security.html

[92] NIST Special Publication 800-207 – Zero Trust Architecture - https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf

One thing to note here is that regardless if you go for the traditional models or the zero-trust model, the defense in depth principle should still be considered and applied.

16.4 Endpoint Protection Strategies

Now that we have covered some of the main principles and protection strategies (and overall differences between them), we'll take a more closed look at the actual endpoint protection strategies available. These can be used independently, in any number of combinations and, why not, in whole, depending on budget.

Active Directory Policies

Although you might not find these in many of the lists, we consider that starting with the basics is the way to go. Any discussion regarding NextGen AV, SIEMs, DLP and others is useless if we don't get the basics right. Especially if mobile computing is a part of your day-to-day business life. Active Directory group policies are a good starting point to ensure basic security such as user access and rights management (considering you do not employ a more complex access management solution). Besides this, there are some good basic protection strategies that every business should make use of. Just to name a few of the most well-known ones:

- Password policy (complexity, history)
- Failed logins allowed
- Local Admin account disable
- Audit policies enabled (account management, object access, logon/logoff, system, etc.)
- Restrict access to local resources such as registry editor, configurations, etc.
- Enforce multi factor authentication – at least for admin accounts
- Lockdown service accounts or disable unused ones
- Disable SMB, if possible – or at least SMBv1

- Restrict USB/IO port access

For a more comprehensive security stance, using security benchmarks is a good starting point for defining your AD security policies. Security benchmarks are created by independent organizations, well known in the industry. These represent recommendations or best practices of how different operating systems or software should be configured. Some of the recommendations made in the benchmarks may be conflicting with your needs in terms of functionality. However, consulting them is a good way to ensure a starting point in securing the environment. You can, of course, tailor the recommended profiles to your needs. And, as always, document the deviations for future reference.

Multi Factor Authentication

Multi Factor Authentication (MFA) has already been mentioned in the section above. With that in mind, we believe this should have its own dedicated section. With the wide adoption of cloud technologies and remote work, MFA should be considered for a company-wide implementation, rather than ensuring such practices only for admin accounts.

Gone are the days in which the workforce is needed to come into the office to have access to resources. Nowadays, most of them are just a click away from accessing internal resources, sometimes even confidential data that resides in the company cloud instance. Attacks targeting the cloud environment, such as brute force attacks or password spray attacks are a regular thing. This setup, combined with perhaps not so secure user passwords may lead to unauthorized access to resources (such as email or collaboration spaces such as Sharepoint online) or, even worse, to malware attacks and compromise.

The principle for multi factor authentication is simple: the users must pass multiple authentication steps in order to gain access to resources. It is important to note that two passwords are not considered a form of multi factor authentication. For the purpose of completeness, let’s review some forms of authentication and what they entail:

- One factor: something you know – typically a password
- Two-factor: something you know + something you have – a password and a token, access card, etc.

- Three-factor: something you know + something you have + something you are – a password + a token/access card + biometric verification (retina scan, fingerprint, palm scan, face recognition)

Probably the most widespread MFA form is the two-factor authentication. This is usually employed through the use of a password in combination with some kind of token that will generate an additional access code to use in conjunction with your password (usually requires input or validation of a code after the password has been provided). Tokens may be hardware tokens or software tokens. In recent years, software tokens have been used more and more due to their practicality – they may be installed on your smartphone, without the need to carry around a hardware token. Authentication software such as Google or Microsoft Authenticator are widespread in usage. For company-wide implementation, covering all users, these are some of the solutions that may be considered for implementation.

Antivirus

Next point on the list is, of course, the antivirus. Not the home use or "free" type, definitely. Organizations of all sizes should consider deploying enterprise level antivirus solutions. They offer central management capabilities and definitely will ease administration tasks. Antivirus solutions are software designed to detect and remove malware such as viruses, trojans, worms, rootkits, keyloggers etc. Traditional antivirus solutions typically employ signatures as detection mechanisms. Although these solutions may protect against a wide array of malware, they are not bullet-proof. They may be less effective against polymorphic or armored viruses which are designed to encrypt parts of themselves, create differing copies of themselves or, as in the case of the latter, deploying various methods to avoid detection.

For this purpose, in recent years the shift has been towards NextGen AV solutions which are rather signature-less. "Numerous approaches to address these new forms of threats have appeared, including behavioral detection, artificial intelligence, machine learning, and cloud-based file detonation."[93] The afore-mentioned are some of the newer techniques employed to detect malicious software and, potentially, even zero-day attacks.

[93] https://en.wikipedia.org/wiki/Antivirus_software

To note that with the growing complexity and number of threats, new security solutions have been developed in the endpoint security space. It is worth going through each of them to define them and understand their use.

EPP – Endpoint Protection Platform

Endpoint Protection Platforms have the role of preventing endpoints from threats such as malware, zero-days, etc. EPPs are typically the antivirus solutions we've just covered.

EDR – Endpoint Detection and Response

EDR solutions come into play when an incident has occurred. While EPPs are rather more…passive if you will, EDR tools are designed to help analysts respond to incidents. They will help analysts during response activities to identify IoCs (indicators of compromise), provide real time alerts, etc. They also have the ability to trigger automated responses such as host isolation or reimaging or manual responses.

To note that modern EPP solutions nowadays typically contain an EDR solution as well.

XDR – Extended Detection and Response

XDR is a term that was coined rather recently and defines complex platforms that have the ability to ingest and correlate log data from various sources such as servers, network equipment, and various security solutions including EEPs and EDRs and endpoints. The goal is to provide a unified view of the security stance of your environment. "Gartner defines XDR as *"a SaaS-based, vendor-specific, security threat detection and incident response tool that natively integrates multiple security products into a cohesive security operations system that unifies all licensed components."* Improved protection, detection capabilities, productivity, and lower ownership costs are the primary advantages of XDR."[94]

[94] https://en.wikipedia.org/wiki/Extended_detection_and_response

DLP – Data Loss Prevention

Data loss prevention solutions are another option to protect your endpoints. This time it is not to protect against an infection, but rather to prevent potential data breaches or data exfiltration attempts. DLP solutions have several characteristics

- Ability to label data based on predefined classification levels
- Data discovery
- Ability to scan outgoing traffic (network, email)
- Block potential data exfiltration attempts based on defined policies

Enterprise level DLP solutions are typically deployed on hosts in the form of an agent that monitors endpoint traffic.

Encryption

Another important aspect of endpoint security is encryption – specifically disk encryption. Usually employed through the use of native applications such as Bitlocker on Windows systems or dedicated software specifically built for this purpose by various vendors. To note that disk encryption will not protect against malware or other software-based attacks – it rather mitigates the risk of device loss by blocking access to the data stored on it.

Mobile Device Management

Mobile Device Management solutions, MDM for short, should not miss from any organization's arsenal if they make use of significant numbers of mobile devices (smartphones, tablets or laptops). These are usually special-purpose built software created for such tasks. Some of the most encountered MDM characteristics are as follows:

- Ensure a configuration standard for all devices
- Managing updates for devices, applications, and applied policies
- Monitoring and tracking of equipment
- Remote troubleshooting and administration in a consistent manner

Vulnerability Management & Patch Management

Last, but not least, a sound vulnerability management and patch management program is a must have for all organization. Vulnerability management is the continuous work of identification, classification and remediation of known software vulnerabilities. The program should cover all assets under your management, regardless whether they are located in a data center or devices being used by your mobile workforce.

To note that some vulnerability management solutions nowadays also have the ability to scan your endpoints not only for known vulnerabilities, but also for configuration flaws. This may be achieved by using benchmarks and special built scanning profiles based on those benchmarks to identify deviations from the standard (remember, we have mentioned the use of benchmarks is the beginning, when talking about security configurations).

Vulnerability management and patch management programs, at the very least, should consider:

- Including all assets in scanning
- Define regular vulnerability scans and reporting on findings
- Regular application of security patches and software updates
- Applying security patches or updates based on the risk level to your organization. While all vulnerability scanners will provide a rating for the vulnerabilities they have found, an assessment that takes into account the environment and network configuration is needed to ensure that the most relevant updates are applied first. For example, an internal assessment might determine that a vulnerability that was marked as “high” by the scanners is actually of medium priority, based on the specific environment, while a “medium” marked vulnerability might take precedence as it has been determined by your internal assessment that it is of critical importance

16.5 Conclusions

There is a variety of solutions out there that will help protect endpoints in an organization. We've just covered a few of the most important ones. However, we want to underline that there is no bulletproof solution and none of the ones we have discussed, nor many others, will ensure security by itself. As we've already noted, defense in depth is an absolute guiding principle in all security related aspects and deploying as many of these as possible and covering as many layers as possible will ensure a more robust security stance.

17. GIS and its Information Security Applicability

In cooperation with guest writer: Alexandru Mircea Rotaru

17.1 What is GIS?

Geographic Information Systems (GIS) represent data as points on a map. Most commonly, it helps monitor how incidents cluster and correlate, thus informing its users about where they need to invest effort and resources next. This way, your organization ensures every dime and every second used is accounted for at all times, keeping all wastes at a minimum.

17.2 A Brief History of GIS

Though the first modern GIS only appeared in 1963[95], its principles date back to Victorian England. A doctor by the name of John Snow (sadly, no iron throne appears in this story) sought to find the source of a cholera outbreak in London's suburb of Soho. He mapped out the cases he found, and most converged at a nearby water pump. Snow then used these findings to show how cholera transmits through water. As a result, when the authorities removed the handle from said water pump, the nearby cholera cases sharply decreased[96].

[95] https://www.esri.com/en-us/what-is-gis/history-of-gis

[96] https://www.ph.ucla.edu/epi/snow/snowcricketarticle.html

17.3 GIS Today

With time, more and more fields adopted GIS, for a wider range of uses. Public health still uses GIS to monitor various vector-borne diseases, such as West Nile Virus and Malaria. Part of the Public Health Accreditation Board requirements for any Health Department in the United States to receive accreditation involves using GIS in daily operations[97]; this includes over 200 state, local, and tribal health departments in the US spanning 46 states and the District of Columbia[98]. Some nations also use GIS to map COVID-19 cases for contact tracing.

Beyond public health, GIS can be used in land surveying and mapping, in natural resource exploitation, in logistics and transportation, in housing and urban development, and in everything else that needs to identify geographic areas where data clusters or correlates.

That's all Fine and Dandy, but what does GIS have to do with Information Security?

There are two ways in which GIS and information security management can rely on each-other to work better. On the one hand, you can use GIS to monitor and prevent threats by identifying which geographic areas pose the greatest vulnerabilities for your organization. Once you have collected the data, you need to represent it somehow, and if a map makes the most sense, you need to use GIS. You can then use this information to intervene before anything goes wrong.

GIS will also help you correlate data, and incidents in your area as a whole. That way, your organization will know that, should they impact factor A, it could also have a collateral impact on factors B and C. For instance, if you see that in communities with high rates of smoking there are low rates of high-school graduation, implementing programs to keep children in school would likely reduce the rate of smoking in the neighborhood. Likewise, if you see that in communities with high alcohol consumption rates there are low anti-depressant consumption rates, implementing campaigns to reduce the amount of alcohol consumed would likely increase anti-depressant usage.

[97] https://phaboard.org/wp-content/uploads/PHABSM_WEB_LR1-1.pdf, p. 44

[98] https://phaboard.org/who-is-accredited/

In short, the various factors GIS monitors do not happen in a vacuum, and any collateral effects of any interventions (or lack thereof) are the organization's responsibility. Using GIS to identify reasonable correlations would at least help you identify a solid amount of these collateral effects, so that you can intervene.

Therefore, GIS can be used to provide a greater overview of past incidents affecting the confidentiality, integrity, or availability of data, taking into account incident data available through incident reports, logs, and circumstantial information. This is useful to identify root causes and remedies given the entire set of circumstances surrounding an incident. Furthermore, previous incidents and their geographic locations can be correlated with other information, such as identified vulnerabilities and threat modelling, to prevent similar incidents from happening in the future. This is particularly useful for certain types of organizations, such as agriculture companies using IoT and petroleum extraction companies using sensors for quality of services and security, as well as for business continuity and disaster recovery planning.

On the other hand, GIS still collects and processes data – often personal data – no matter what you use it for, so you need to keep that data safe. Most countries have devised some sort of data privacy laws, such as GDPR in the European Union, and HIPAA and FERPA in the United States. Every information system risks getting breached on the daily, so, when working with personal data, you need to make sure it is safe wherever it may be used, including in GIS.

For the European Union, these are the most essential principles you need to take into account:

- **Data minimization:** The organization should limit the data collected or processed through GIS only to that personal data needed for the purpose of the GIS use. For example, if it is relevant to know the household contact details or behavior data for the purpose of the organization's activity, these should be collected. Otherwise, if such data is not needed (i.e. the activity of the organization can be fulfilled without this data), it should not be collected.

- **Need-to-know principle:** The access to the GIS data should be limited only to the individuals that need access to perform their job description, at the lowest level of access that would allow them to fulfill their functions within the organization. This applies to both employees of the organization and the employees of any third parties working on the GIS data.

- **Purpose limitation and basis for processing:** Data collected for one purpose (such as traffic monitoring) generally cannot be used for other purposes (such as marketing, creating heat maps of households). In addition, the organization has to identify if there is a basis for data processing before it starts any actual processing of the data; these include fulfillment of contract with individuals, legitimate interest, and public interest.

- **Limitations in data disclosure:** The organization should have limited third-party involvement, only when needed and within the scope and purpose of the data processing. When transferring to other data controllers, the organization should analyze if there is a basis for such data processing.

- **Proper transparency of data processing:** The organization must inform any individuals whose data gets processed, per the transparency requirements, about how their personal data is being used. Furthermore, when consent is needed, it must be properly obtained prior to any data collection.

17.4 How can you Avoid a Collision with Something you don't Notice?

Per Murphy's law, risks never go away. This is why many organizations today prefer to be reactive when addressing threats and breaches. Though, wouldn't it make more sense to know where threats can come from so that they don't catch you off-guard? That's exactly where GIS comes in.

Still, organizations may not see why they would need to invest time, energy, and funds into getting an operational GIS system in order. Going with the collision example, your eyes can help you see obstacles and incoming traffic if you're driving a car, and some may say that's enough. However, what if you are flying a plane at half the speed of sound? Cyber security incidents are more like an airplane flying through a flock of Canada geese than a road accident, as they can occur very quickly, and you are likely going down and have to prepare for disaster by the time you notice them.

So, how can anyone expect pilots flying at half the speed of sound to notice the flock of geese in the way using nothing but their eyes? They don't. Instead, the airline industry adapted, with radars, forecasts, and air traffic control making sure planes don't crash into each-other, fly into dangerous updrafts, or collide with any local birds. Also, in case something does go wrong, airplane manufacturers design engines that withstand the impact of flying through birds as a last resort.

So, why shouldn't an organization invest in something similar to prevent cyber security incidents and other risks? This is where some sort of monitoring and prevention system comes into play, particularly when operating with personal or confidential data. If the threats are related to roads, geographic areas, or underground resources, a GIS system will serve that role.

17.5 Data Security is Key

Naturally, every monitoring and prevention system comes with drawbacks of its own. Otherwise, we would all be using these perfect systems, and information security as a field would not even exist. Sadly, we do not live in that perfect world, and your organization still needs to look out for threats whenever using information systems, including GIS.

In many cases the data that gets analyzed through GIS is personal data and/or confidential, mainly through its identifiers - such as names, addresses, phone numbers, social security numbers (or equivalents), and bank account numbers. By identifier, we mean any data that can reasonably lead to inferring an individual's identity; this definition includes personal data. Particularly when doing analytics using GIS to see which geographic areas or points of interest on a map need the most attention from the company, you need to secure the data in case of a cyber security incident. This is in line with the data minimization principle, which entails collecting, storing and processing only the data needed for the data processing purpose.

The easiest way to do so is to remove all identifiers before analysis. In the United States, studies in the sciences and social sciences require that the final results be stripped of any and all identifiers, so it may very well be part of the local regulations to do that. However, if that is not possible, you need to secure the data and the identifiers.

17.6 Increase Security the More Confidential the Data

With GIS, you can layer data, and select who gets to see what. That way, if your organization's information systems get breached, you can contain the attackers to the level they penetrated. With GIS, you can have the lowest level of security be the heatmap with no identifiers, and then provide higher level access to more confidential data as needed. That way, if the breach only reaches the "dots on a map with exact address labels" without knowing anything about what exactly each dot represents or the details of any sub-category said dot may be in, you can contain the breach at that level. However, every second matters, and if the breach gets to the highest level of security, it's time to call the company's management, lawyers, and incident handling team.

17.7 Nobody Believes Organizations that Publish Data that's been Tampered with

One reason your organization's GIS system may be breached is to steal personal information. However, particularly in research, public policy, and public health, using data that's been tampered with is the fastest way to kill your organization. Tampering with data can happen for multiple reasons, which mostly boil down to your competition trying to undermine your organization's credibility, a cover-up of how bad the situation truly is (or how good it is, to have access to more funding that would otherwise be inaccessible), or an attempt to increase or decrease the monitored issue's level of urgency for the stakeholders by eliminating existing correlations or fabricating ones that don't actually exist.

This is why you need to set clear protocols as to saving multiple backups of data and working with the data, to ensure that, when someone does tamper with the data, your organization will be able to pick up from where it left off. You also need to make sure your organization is able to identify which sets of data have been tampered with - a quick cross-reference with the backups and/or logs should do.

Also, if you find that your data has been tampered with and there is no unauthorized access to data from outside your organization, you and/or the organization managers might want to have a long conversation with the staff. In the best situation, a quick refresher on how to manipulate data without altering it should do. If you have a malicious employee, you need to make unpopular decisions that limit liability. If nobody knows how the data got tampered with, you might want to upgrade your organization's locks, security cameras, logging and monitoring, as prevention is key in cyber security.

17.8 The Data Correlation Process is a (Hacker's) Dream

Even though data correlation is one of GIS' biggest assets, it's also one of its greatest vulnerabilities. The entities that breach your informational systems aim to get as much out of it as possible, and correlations provide multiple sets of data that are also somehow connected. Your organization's data is at the greatest risk when you do the correlation process, so you need to pay extra attention to what is going on during those times. Also, make sure that the data sets you correlate are then separated as soon as you finish the analysis, thus limiting the amount of time they stay correlated.

17.9 Don't Forget the Obvious – Confidentiality

Imagine you're presenting your findings about anti-depressant use in your area. You made sure your data is secure, and that nobody breached or tampered with it. Only, you left the data as dots, and one of them is on your address. Next thing you know, you have a PR scandal and everyone calls your sanity into question because you forgot the obvious: when presenting GIS data, make sure to eliminate all identifiers – especially the addresses. Anonymisation techniques should be applied to the data. However, this requires a balance between generalizing the data in order to reduce the possibility of individuals being identified and keeping the data useful for the data processing purpose. GIS data provided to the public should, generally, be anonymized completely; many organizations today use heatmaps instead of the actual dots to do just that.

17.10 Insight without Action is Worthless

The whole point of identifying threats and weaknesses in your organization's system is for your organization to do something about it. GIS will help inform your decisions and the consequences they will likely have, but it will be all for nothing if your organization doesn't use them. So, before investing in a shiny new GIS, make sure you know what you are using it for and that you have a process in place for integrating data from GIS in its usefulness in the incident prevention process, the business continuity analysis, and the risk management process. As a rule of thumb, if you begin with the end in mind, your work will yield far better results.

17.11 So, how do I incorporate GIS in my Organization?

First, you need to ask yourself: what would I use GIS for? If you work with maps, natural resources, transportation tracking, the weather, and the like, GIS will prove ideal for you. Furthermore, for some of these use cases, GIS can assist with implementing legal requirements such as those under the NIS Directive. Then, you need to implement protocols for using the GIS, for handling the data without altering it, for security levels, and for what to do when a cyber security incident happens. Also, people from different backgrounds using GIS can each develop their own set of abbreviations, which would turn using GIS into a nightmare akin to translating dead languages. To avoid these kinds of situations, a list of all organization-wide standard abbreviations for the GIS would help make sense of what is going on.

Finally, you need to have a GIS technician on your team. You can either hire someone from outside or get one of the present employees certified; it's a very intuitive system that shouldn't take much time to master, and many institutions in Europe, the United States, Canada, Australia, New Zealand, and other parts of the world offer such certifications and/or coursework for credit.

18. Discussing Information Security with the Board of Directors

Given the growing importance of data and IT systems for businesses and the growing number of vulnerabilities, cyber threats, data beaches and cyber security regulatory landscape, information security has become in recent years a topic that should be addressed at the level of board of directors. This is relevant for both SMEs and large corporations.

According to a recent study, respondents asked what puts the highest pressure on cybersecurity management responded competition for budget (46%) and severity of attacks (49%). Thus, the management of security risks, including prioritization, budget and management buy-in weight almost as much as the ever-changing security threat landscape.[99]

18.1 What Content to Present to the Board of Directors?

Usually, investments numbers for all initiatives are considered in relation with revenue numbers or cost savings/avoidance that support the company strategy. If we regard the IT risk related initiatives as just hygiene, they will not be taken as a compelling call for action, but at best as an urgent fix (such as wearing a mask or washing the hands, these days) and at worst as not necessary or not now.

The usual reasons for IT security investments are related to potential negative impacts from fines, and this was even more so after GDPR legislation enforcement. Reputation hits after data leaks or successful hack are easier to imagine, since there

[99] 451 Research, commissioned by Kaspersky, "Cybersecurity Through the CISO's Eyes: Perspectives on a Role"

were (in)famous incidents that are globally known, due to their big impacts. The loss of productivity is also considered with the systems unavailability. With exercises such as ORSA (Operational Risk Self-Assessment), these types of impact can be quantized in money so the conversation becomes more intuitive and the relation with the organization's objectives clearer. All these need strategic decisions and investments only the board of directors can make and decide accordingly.

Prevention is likely less costly than remediation. Thus, the board of directors has to analyze opportunity of a cost avoidance due to the implementation of certain security controls or items in the security program.

Generally, loss of clients' trust is more damaging to revenues than the budget needed for the security investment required.

In B2B commercial relationship, security is actually considered both as risk evaluation and as professional assessment of existing measures. In a competitive bid, there are RFP check lists that an organization passes or not, thus qualifying for new business proposals.

So, how do we show all that in a convincing presentation to your executive board?

That depends on the organization's objectives, on defined (or not defined) risk appetite and the board members experience with security concepts understanding.

Considering these, you need to choose the "money only" argument and / or the risk tolerance reason and / or the market security best practices. [100]The arguments are there, the cause is worthy but the best suited plead is the one that fits your audience best.

Thus, the main pain-points encountered relate to differences in terms and presentation angles used by the board and by the information security team (e.g. CISO). In this respect, the following main environmental aspects have to be taken into account:

Specifics of the organization's business field:

The approach to be taken in discussions with the board depends on the industry in which the organization acts, as organizations in certain industries may hold extensive amounts of data and/or complex IT systems or infrastructures, whereas

[100] https://www.gartner.com/smarterwithgartner/5-security-questions-board-will-definitely-ask/ , last accessed on 15 April 2021.

other industries may hold less data and/or have less complex IT systems or infrastructures.

Maturity level of the organization:

Further, this approach depends also on the maturity level of the organization in terms of security and information technology. For each level of maturity, the CISO has to take into account the existing internal processes, the maturity level of other areas within the organization, the organization culture in terms of technology and security and adapt proposed improvements accordingly.[101] Not having the maturity level in mind may lead to inadequate proposals, as the building blocks necessary to implement the new proposal do not exist yet in the organization. For instance, if the CISO wishes to implement a SIEM solution, but some of the IT systems within the organization do not record proper logs, the CISO should first focus on proper logging of relevant aspects for each IT system.

Team contribution:

As information security does not exist in isolation from the other areas within the organization, security proposals should have in mind the impact in such other areas. Aside from the business objective, which is essential, some of the most relevant areas with which the CISO has to coordinate for information security projects are the operational side of information technology and the data protection officer. This ensures correlation with other initiatives within the organization and support from the teams that will be involved in the implementation of the project.

[101] ENISA, "NIS Implementation Report", 2020.

18.2 How to Discuss with the Board of Directors? Approach Tips and Tricks

Continuous updating on projects:

Create a list of current and finished projects since the last meeting and explain how they have positively impacted the company. This ensures that management understands the relevance of information security for the business objectives and, also, the progress that have been and that remains to be made in this respect. Of course, non-technical, summary presentations should be used. Thus, the emphasis should not be on the budget amount for the security projects, but, rather, the status within the security plan based on the agreed security strategy. It is useful for management to understand from a quantification perspective, how the organization is more secure than before. The quantification may include types of vulnerabilities closed together with their business impact, number of incident alerts generated for specific vulnerability exploitations.

Non-technical language:

The topics to be presented before the board of directors should not be at a granular technical level, but, rather, at a high-level, outlining the business impact (and regulatory, if the case). Thus, the presentation should be aligned with the business goals and reflect business perspectives so it will reach the business people we have as audience during our presentation, with emphasis on information security risks with high impact and probability from a technical perspective and from a business perspective.

The consequences related to risks should be presented in a measurable manner. For certain consequences, the calculations may prove to be difficult.

For example, for a penetration testing report, the presentation should not outline the technical vulnerabilities identified and the technical consequence mentioned in the report together with the probability mentioned in the report. In terms of probability, the report is originally created in the context of the penetration testing and may need to be adjusted having in mind the methodology used by the organization to calculate probability and having in mind the entire IT infrastructure and security controls in place. Furthermore, a general consequence of data leakage may not be sufficiently clear for non-technical individuals, such as the members of the board of directors. In this case, additional context on the criticality of the systems affected and the consequences based on the data affected should be detailed as quantifiable as possible.

Focusing on risk overview:

This can be achieved with emphasis on specific, measurable impacts on the business or on the organization. As per a survey conducted by ISACA, only 21% of senior management is briefed on risk topics at every senior management meeting.[102] Thus, the technical analysis and specific technical pain points are useful. However, in terms of presentations before the board of directors, emphasis has to be placed on impact and probability of a risk, with such risk being described in plain language and not technical jargon. The risk-based approach highlights that information security is not a one-off exercise and not just a question of compliance with legal/best practice requirements, but an exercise of risk assessment and risk mitigation in a continuously changing environment.

Preventive vs. incident costs:

This depends significantly on the type of incident. For certain types of incidents estimations may be made – e.g. recovery from ransomware where no data exfiltration occurred. But for most incidents, aside from the vulnerability fixing and disaster recovery steps (which relate to actions taken by the company itself), other costs/damages/fines that may be incurred by the organization can prove difficult to calculate. Nevertheless, this comparison can be used in certain specific situations. In terms of comparison of costs for implementing controls, one might try to make a comparison between the costs for control implementation and the costs for investigating and remediating an incident.

The information on incident costs may be obtained from previous incidents, bug fixing activities, vulnerability fixing activities and market statistics on incident fixing costs. On the financial side, indirect damages, such as loss of productivity due to downtime, loss of sales can also be taken into account for the relevant calculations.

It is difficult to establish value of data in an IT system, but criticality for business may point in the right direction. Insurance companies themselves have just started in the last decade to explore more comprehensive models for calculating cyber-risk in view of calibrating their insurance products.

Benchmarking against other companies

This type of analysis (especially in the same sector) is often requested by management. However, this is difficult and tricky in practice. In terms of governance

[102] ENISA, "Survey Strong tech governance drives improved business outcome", 2017, https://www.isaca.org/why-isaca/about-us/newsroom/press-releases/2017/survey-strong-tech-governance-drives-improved-business-outcomes , last accessed on 4 April 2021.

and risk management, such information is not publicly available. Even if, for instance, the budget amount for information security would be mentioned in the financial statements of the competitors, this does not detail the internal steps taken, nor the risk assessments performed by the organization. In terms of incident handling, public information is scarce, without proper information on costs or mitigation actions taken. The only relevant information that may provide guidance on certain trends is included in various data breach costs or security posture surveys performed in the market. Nevertheless, these can provide only limited information about the trends in terms of information security, with the internal strategy and approach to be decided by the organization based on its specifics.

For controls related to incident identification steps and early response in case of incidents, quantification of the breach costs reduction may be useful to be included in the business case.

Statistics about past events:

Management generally asks about number of data breaches for a particular IT system or type of vulnerability as a manner of calculating the probability of a risk occurring. This may be included as insight into presentation of risk, however, it should not be relied on as a metric in this respect. In the same manner, management inquires about financial loss, fines or damages incurred in the past by the organization (or other entities in the market) for a specific type of vulnerability exploitation or security incident. This may be factored into the decision about risk response, however, it is not necessarily relevant, as these are not indications about future situations, which may differ based on consequences of the incident, evolving threat landscape and evolving guidance and requirements in terms of level of security to be implemented by organizations.

Hands-on testing:

Testing incident response and business continuity/disaster recovery processes together with relevant stakeholders and the board of directors can also help with further understanding of information security risks and potential consequences of the occurrence of such risks. Given that the testing involves use cases from real-life scenarios together with role-playing and detailing of business impact, it is a good approach towards more information security awareness among the board members.

Certifications

These are useful in terms of client's perspective on the organization and in order to ensure standardization and state-of-the-art process in place within the organization. One approach that may be easy for management to relate to are the steps to be implemented in order to achieve a security certification, which generally helps also with increasing the maturity level of the organization. Therefore, information security may also be used as a selling point for the organization in terms of branding, showing the effort the organization makes to ensure security and privacy of its customer data. Certification has to be viewed in correlation with the maturity level of the organization in terms of information security.

Further, in case of service companies or highly regulated sectors, certifications (especially ISO ones) are starting to become the norm, with large demand from relevant stakeholders and authorities in this respect.

Risk appetite calculation:

The risk appetite established by management should be based on proper knowledge of the threat landscape and the security posture of the organization by reference to risks identified and evaluated. This ensures that proper risk appetite is chosen by management and this is the key in setting the tone in information security within the organization. The role of the board is to set the direction of the company and make the decisions in terms of business development and operation, while taking into account all compliance, legal, security risks and without going into details about the day-to-day operational activity.

Cooperation in cybersecurity:

Management has to be aware of the need for cooperation in ensuring information security. Information security is not achieved with effort solely from the IT security department, but also from other departments from IT operations, IT development, legal, compliance, risk management, data protection, business owner, audit etc.[103] For this reason, for each initiative presented before the board of directors, proper emphasis has to be placed on the internal skills and departments that need to be involved in the initiative. This includes also reliance on and dependencies with vendors, in terms of request for proposals needed, additional services that should be added to existing agreements, special SLAs for information security and business continuity, proper liability and undertakings clauses etc. Emphasis should be placed on the resilience that is provided to the organization through the information security steps taken.[104]

Provide insight into the need for orchestration of people, processes and technology in order to ensure information security, while underlining the role of the board of directors in this respect.[105]

Preparing the meeting with the board of directors

Before having the presentation before the board, it is useful that all matters be discussed with the relevant stakeholders, so that all are aware of the risks and the proper internal security assessment process has been completed. Having all stakeholders involved in the risk assessment and risk addressing already knowledgeable can bring clarity to the discussions with the board of directors and relevant (and measurable) input for the CISO to prepare his/her presentation.[106] This also ensures that each stakeholder understands their role in the information security program, especially those that are responsible for taking appropriate steps as per the relevant RACI matrix.

[103] Khalid Kark, Caroline Brown, Jason Lewris, “Bridging the boardroom’s technology gap”, Deloitte University Press, June 29, 2017

[104] World Economic Forum, “Principles for Board Governance of Cyber Risk”, Insight Report, March 2021.

[105] RSA, “Security and Risk: How to talk digital risk with the board”, https://www.rsa.com/content/dam/en/analyst-report/gartner-how-to-talk-digital-risk-to-the-board.pdf , last accessed on 15 April 2021.

[106] Tony Kontzer, “C-suite cybersecurity awareness may be the key to taking a bite out of breaches”, RSA Conference, July 19, 2018.

The responsibility of each stakeholder (e.g. business owner) about information security control implementation ensures that security is included in the project budget, such as budget for initial and periodical penetration testing and for periodical vulnerability assessment. This leads to swift implementation of appropriate controls and consistency in budgeting methodology between business changes and information security controls. The alternative would involve approving a specific budget for the security department and may lead to situations where the budget does not match the actual needs of the business. Further, it useful to have in place a cyber-risk management process that interacts with the operational processes.

Further, a correlation with the business objectives and alignment with these should be implemented.[107] One approach in this direction is to link top initiatives to top business risks.

There is regulatory landscape on security (in certain domains such as banking, insurance, energy, entities falling under the NIS Directive), and there are also implications of GDPR in terms of security steps to be taken. Nevertheless, in most cases, the legislation provides the aim of the security steps to be taken by organization, but leaves the specific operational approaches to be decided by each organization, based on its structure, activity and IT landscape.

A balanced score card can be used to outline the financial aspects of security (e.g. supplier management, efficiency in internal security management with task allocation mechanisms, use of security to grow the business and reach business objectives), customer (availability of service, security of data, confidence and trust in the services offered by the organization), operational (proper IT solution s and automation, proper change management process) and human factor (proper awareness of risks).

Thus, when analyzing the risk impact/probability and the risk acceptance, compliance with legal provisions, potential fines/damages, business impact and trust of clients should be considered. In this respect, the business strategy can use the security controls in place.

The risk addressing proposals and any improvement in the information security program should outline three different approaches:

- Minimum: This option presents the bear minimum requirements that can be implemented efficiently within the given budget. It often includes only

[107] Isaca, "Reporting Cybersecurity Risk to the Board of Directors", 2020.

critical risks or legal requirements with more impact in terms of risk mitigation.

- Moderate: This is often the option chosen for implementation. It goes beyond the bear minimum requirements with additional controls. However, it is based on the level of maturity of the organization and a balance between the potential benefits it can bring and the investment it requires.

- Bullet proof: a solution that includes state-of-the-art implementation of all requirements and best practices. Often this solution entails huge costs, which have to be analyzed by reference to the benefits it brings to the organization (either profit or reduction of potential losses/costs).

A what-if scenario-based on each of the three approaches is highly desirable to reveal the organization's risk posture including possible financial losses.

The request for board decision(s) should be clearly articulated and should emphasize security team's recommendation(s). This request should touch briefly the following structure – organizational change required, program/project/project change required, training/awareness required, investment required, and nevertheless what are the costs required for implementation.

Nevertheless, the controls to be set in place (technical, organizational or other types) should be presented in a measurable manner, outlining the risk reduction level, any cost or FTE reduction and any dependencies on other departments or assistance needed from other departments in the organization.

Frequently asked questions by the board of directors

The following questions are often asked by the board of directors when faced with information security decisions.[108] We have included a short recommendation for approaching them. Nevertheless, the specific factors concerning the organization's business sector, security maturity level and structuring have to be taken into account in order to best address such questions and provide relevant information for a qualified decision.

[108] RSA, "Security and Risk: How to talk digital risk with the board", https://www.rsa.com/content/dam/en/analyst-report/gartner-how-to-talk-digital-risk-to-the-board.pdf , last accessed on 15 April 2021.

- Can this be automated?

Explain what can be automated and what cannot. Further, detail what are the advantages of automating certain steps in security or in other business process to reduce cyber-risks, but also highlight the areas where automation cannot bring added value by reference to its costs and the areas in which automation can increase exponentially operation risk if not implemented properly.

- How should we approach this risk?

Explain the business impact of each risk and the risk response options ranging from the responses that best suites business needs to the ones that best suit security/risk mitigation needs.

- What are our options?

Generally, the board of directors would like to understand the limits of what can be implemented: the minimum controls to be implemented for partial mitigation of the risk, the moderate option that includes best cost-benefit balance and the bullet proof option that eliminate all or most of the risk, but may be very costly and time consuming to implement.

- We thought this was already taken care of. Why is this a recurring point in our meetings?

Explaining that there is a constant change in the IT and security landscape, with new threats, vulnerabilities and IT systems. Further, for certain risks, the mitigation process may take multiple steps to implement all relevant controls that can be implemented.

- Are we 100% secure against all type of threats?

The threat landscape changes daily. The best approach is to prioritize the items with the highest risk rating first and take into account the organization’s risk appetite.

- Are we secure against incidents similar to the ones incurred by our competitor?

One cannot speculate about the root cause of such incidents. Most of the information about the incident, causes and consequences is not made public. Only official information stated by the company or authorities can be mentioned and taken into account.

- Do we know our risks? How can we have better insight into our risks?

This is the moment to emphasize any shortcomings in the risk assessment process and to request improvements. Detail the role of each stakeholder involved in the risk assessment process and the concrete steps going forward.

Further, highlight that the risk assessment process has to be done periodically, given the changing threat landscape and the changing landscape within the organization itself (e.g. new IT systems, new products, new processes).

- How did this happen?

Explain the facts, the root cause and how these could not be prevented based on the controls in place and the risk assessment/risk addressing mechanisms in place.

- Why are we spending so much?

Explain the progress in information security projects/programs and their impact on/correlation with business objectives. This helps the board of directors to understand the need for continuous investment in information security, given the continuous changing variables of the respective risk assessment process.

18.3 Conclusions

Have continuous, focused, discussions with the board of directors and to ensure emphasis on the risk-based approach for prioritization of steps to be implemented from an information security perspective.

The discussions can be supported by key indicators on the implementation process, in terms of correlation with best practices and with potential costs, damages and losses.

The continuous feedback loop ensures on the one hand that the board of directors is aware about information security risks when taking decisions and, on the other hand, provides further alignment between business objectives and information security goals/steps, while including all relevant stakeholders in the decision-making process.

19. Conclusions

This book includes recommendations for setting-up frameworks for Penetration testing, Red Teaming and Blue Team, with practical examples for each of them and specific use cases in certain domains (e.g. healthcare).

The first chapters in this book describe a risk-based view of the Penetration testing and of the Red Teaming strategy and practices; the difference between the two ad their principles and methods are explained.

When an organization is considering including penetration testing and Red Teaming in its existing processes, this has to be correlated with the risk management process. This is essential firstly in order to analyze and address the IT systems with the highest risk (and, thus, impact on the organization's business) and, secondly, to manage cost-benefit properly after having a holistic risk overview for a specific IT system or process.

This angle of offensive security is closely correlated with the Blue Team angle. The Blue Team has to prepare from a technical perspective (relevant technical tools) and from an organizational perspective (e.g. internal processes for identification of threats and incident handling, proper setting-up of SOC, supply chain management).

The legal and data protection aspects concerning offensive security have to be analyzed and implemented before starting any activity. They refer to collection of data, manner of processing the data and how the data can be transferred to third parties (vendors, authorities, community partners).

All of these actions have to be presented before the management of the organization and the management has to be actively involved in security decisions. Have continuous, focused, discussions with the board of directors to ensure emphasis on the risk-based approach for prioritization of steps to be implemented from an information security perspective.

The aspects mentioned in this book ensure that a proper level of security (including offensive security) is in place within the organization, firstly from a governance perspective and, secondly, from an implementation and monitoring perspective. Continuous monitoring of internal procedures and implemented controls, together with continuous improvements to the process/controls from lessons learnt are essential in any organization, regardless of the maturity level of the organization from a security perspective.