

InCyT Training Module

Information Security for Employees

Unit 2 Week 3

Unit name: Introduction to Cryptography

<i>Learning Activities</i>	☒ Webinar ☐ Exercises ☐ Workshop ☐ Presentation ☐ Other
<i>Learning Responsible</i>	Claudio Fornaro
<i>Learning Activity Goals</i>	1. Cryptography 2. Cryptoanalysis 3. One time pad algorithm 4. Symmetric algorithms 5. Asymmetric algorithms 6. Hash functions 7. Signatures & public key infrastructures 8. Cryptographic Tokens 9. Challenge-response mechanism
<i>Skills to be Gained</i>	• Basic knowledge about cryptography algorithms and infrastructures Technical Skill 2 - TS2
<i>Duration of Learning activity</i>	1 week
<i>Learning requirements</i>	What is needed? • No previous specific knowledge is required

Brief information about the module

Cryptography pertains to methodologies for hiding information from disclosure, both during a communication and for archival purposes. Its objectives are achieved with the design and implementation of protocols that prevent an unauthorized third party from retrieving the information. In a secure communication, the adversary third party cannot access the information transmitted, in the case of archiving the third party cannot trace the content of what has been archived.

In Cryptography, the third party is always a malicious entity that aims to retrieve information that it must not have access to, whether for privacy, trade secret or national security purposes. The data confidentiality and integrity, the authentication of the parties involved and the non-repudiation of what the data source has transmitted or stored are the fundamental principles of modern cryptography.

After an introductory part, where the fundamental concepts are described, we present a progression of basic cryptographic techniques, starting from the *One Time Pad Algorithm* and then building up for considering the two cryptographic families of algorithms: Symmetric Algorithms and Asymmetric (or Public Key) Algorithms with a comparison of them. The role of the Hash Functions is also discussed to introduce Digital Signatures. A description of the key concept of Public Key Infrastructures is presented in depth with some details on the standard mechanisms and formats. To conclude, some information about Cryptographic Tokens and the Challenge-Response Mechanism is provided.

Content of the Learning Material

Introduction

Cryptography is the mechanism by which the content of a message is hidden from other entities. In other words encryption deals with encrypting messages, i.e. taking a plaintext, using an algorithm and generally a key, making the content of the message unintelligible.

Cryptanalysis studies how to decrypt the message without knowing the key. In an attack, cryptanalysis consists in trying to trace the content of the message without having the key. Cryptographers and cryptanalysts are actually the same people: who writes the cryptographic algorithms also analyzes them to find their weaknesses or validate their robustness. As always, when dealing with security, the idea of being able to deal with protection issues without knowing how attack techniques work is not very effective.

Information in messages is generally redundant, in a formal sense, messages generally use a much higher number of bits than that really necessary to carry the information. **The redundancy of information facilitates cryptanalysis:** redundancy in the information means that the messages can show a structure that makes it easier to trace the content of the message, as not all messages are equally likely. If we think of a text represented with normal ASCII characters in byte we know that the combinations of bits that can represent meaningful messages are relatively few compared to all possible combinations. The true messages that can be represented with a certain number of bytes are very few. This makes it possible to identify valid messages by discarding those that do not carry true information. All this vastly helps the activity of the cryptanalyst, so the purpose of the encryption algorithm is also to hide this redundancy to better hiding the information.

Example: Encoding *a sequence* of the four colors BLUE, GREEN, RED and YELLOW.

These can be coded:

- with 2 bits (00, 01, 10, 11)
- as 4 different uppercase letters
- encrypting the color names

In the first case, the minimum number of entities (bits) is used, by substituting each color name with the corresponding 2 bits will conceal the original sequence, but not the order as there is a correspondence between COLOR and the 2 bits. In the second case, as each letter uses 8 bits in the ASCII code, 6 bits are redundant for every color → 75% redundancy and we achieve the same grade of hiding as the 2 bits. We still cannot associate colors and values.

Consider the following very simple cipher, the so-called *Caesar cipher*, which does nothing but scroll the letters in the alphabet by a certain number of positions. For example, if we slide them by 3, all the 'A's become 'D's, all the 'B's become 'E's, and so on. This is a very simple cipher; actually

attributed to Julius Caesar, but if the color names are encrypted with this method, the repeating pattern of their encoded letters are easily identified (the ASCII code for R followed by the ASCII code of E followed by the ASCII code of D are easy to recognize as an entity). So the sequence can be identified, and if the words are also known in this case just counting the number of letters is enough to identify the original words.

This simple example allows us to understand that cryptanalysis is not an analysis that is independent of the content of the message: a completely randomized message would be much more difficult to identify and analyze than a message that has a recognizable structure.

Some types of cryptographic attacks and algorithms rely on knowing part of the text. In some cases, knowing part of the text makes it possible to trace the key and then access the rest of the text. This is very significant when it comes to network communications, because in network communications the protocols at different levels usually have a certain amount of information that is in any case recognizable, e.g. in an IP packet we know that the header contains the sender address and the recipient address.

An encryption algorithm hides also redundant information, so applying a compression algorithm to an encrypted message does not bring much advantage, because it would reduce compression rate. If we want to use compression and encryption, we must first use compression and then encryption.

The One Time Pad Algorithm

The One Time Pad (OTP) is an algorithm, so to speak, “perfect”. It is a very simple encryption algorithm, very simple to implement and has a characteristic that makes it particularly interesting: it is proven that without knowing the key it is impossible to go back to the original text. Now we look at it and then we discuss why, as we have such a good algorithm, we have to find others.

The EXCLUSIVE OR $\oplus$ (XOR) is a logical function where the result is True when the two operands are different, it operates on logical values but translated to binary digits, usually True=1 and False=0, its *truth table* is:

$$0 \oplus 0 = 0$$

$$0 \oplus 1 = 1$$

$$1 \oplus 0 = 1$$

$$1 \oplus 1 = 0$$

From this table it is clear that if $X \oplus Y \rightarrow Z$, then $Z \oplus Y \rightarrow X$ and $Z \oplus X \rightarrow Y$. The EXCLUSIVE OR is a

very simple function and also easy to implement as an electronic circuit.

The OTP mechanism is used to transfer a message from two parties. In the scientific literature it is common practice to identify the actors with personal names, they are usually called Alice (A) and Bob (B). So we will describe how Alice can securely transmit an n -bytes message (*PlainText*) to Bob.

The One Time Pad algorithm requires a setup phase:

- A random sequence of exactly n bytes is generated, this is the *One Time Pad* (OTP) or simply the *Key*;
- The OTP is securely shared between the two parts in some way (given in advance, sent in an already secured channel, etc.).

When both parties possess the key, they can use it to transmit messages. In case Alice wants to send a message to Bob:

- Alice calculates a *bit by bit* EXCLUSIVE OR of the *PlainText* with the *Key*, producing a *CypherText*:
Encryption: $\text{PlainText} \oplus \text{Key} \rightarrow \text{CypherText}$
- The resulting *CypherText* is transmitted by Alice to Bob over a possibly insecure channel
- Bob calculates the *bit by bit* EXCLUSIVE OR of the *CypherText* and the *Key*, so recovering the original *PlainText*:
Decryption: $\text{CypherText} \oplus \text{Key} \rightarrow \text{PlainText}$

Example

Let us suppose both Alice and Bob have already agreed upon a random OTP (*Key*).

Alice wants to send “HELLO” to Bob (in ASCII Code, spaces are here introduced for ease of reading):

	H	E	L	L	O
<i>PlainText:</i>	01001000	01000101	01001100	01001100	01001111
			$\oplus$		
<i>Key:</i>	10101001	01110010	10110010	10000110	11000110
			=		
<i>CypherText:</i>	11100001	00110111	11111110	11001010	10001001

Bob receives the *CypherText* and applies the *Key*:

CypherText: 11100001 00110111 11111110 11001010 10001001

$\oplus$

Key: 10101001 01110010 10110010 10000110 11000110

=

PlainText: 01001000 01000101 01001100 01001100 01001111

H E L L O

The *PlainText* has been recovered.

As the *Key* is a random sequence, there is no way to say if a certain bit of the sequence is '0' or '1', as they are equally probable. Without the *Key*, given a *CypherText* bit, we have no chance of trying to guess what the original bit might be. This is true for single bits, for the whole message, for any fragment of our message.

In case a part of the *Key* is known, it is possible to recover the corresponding part in the *PlainText*, but without the rest of the *Key* there is no lack of information on the other bits of the *PlainText*. Therefore, the cryptanalysis mentioned before is not effective. In this case, to discover the remaining part of the original message, the remaining part of the *Key* is needed. All the information conveyed by the *Key* is needed in order to recover all the information of the text.

One Time Pad Problems

The strength of the OTP mechanism relies on the length of the *Key* being exactly equal to the length of the text and on using it just one time. The main drawback is that the transmission of the OTP has the same problems as the message.

In some cases this problem is easily solved, for example a diplomatic briefcase with a ribbon containing a very long key could be provided to an embassy and subsequently the messages that will be exchanged with that embassy can be encrypted by consuming the bits of the key gradually. These are very peculiar uses, but quite effective. It is clear that if we consider network communication, the real uses of the One Time Pad are very limited. There is a need for something more flexible, something that allows the exchange of a key that can be used for a very long time.

Let us consider each of the problems.

Offline transmission in advance

This is a general problem with keys: keys have to be agreed in some way by the parties. This need to agree on keys is the most important problem of using cryptography. The problem with cryptography is not so much algorithms and cryptanalysis, as well as key management: that is being sure that all and only the right subjects have access to the keys, then exchange them in advance, and be sure that the keys exchanged are the correct ones.

One time use

The One Time Pad is used only once, as the name suggests, this is an important aspect and limitation. If we tried to use the same key on two messages, we would weaken the security of the method.

Generating random numbers in a deterministic system is very difficult

General purposes computers use *pseudorandom* number generation where each number is computed from the previous one with some computation, this produces a sequence of numbers that repeat after a very long period. These are not true random numbers as a computer is a deterministic machine. In cryptography, however, this feature is deleterious. If it is possible to derive the next random number from a previous one, the protections offered by cryptographic algorithms or protocols become much less secure. Writing code that implements cryptographic algorithms is not trivial: there is a great deal of detail to watch out for, all the flaws that have been identified over the years by code developers that only the code developers in this specific area have learned to manage. This means that it is always a bad idea to develop and/or implement in-house cryptographic algorithms instead of using well-tested, widely used, extensively validated libraries.

Algorithms other than OTP

The other algorithms use fixed length keys. If we use an n bit key, we have 2^n different keys, if the value of n is big enough and you just apply a brute force attack (that is trying all the possible keys), it would require so many years to find the right one so that “probably” the content of the message is no more of interest. For example, with one of the fastest supercomputers available, checking a 256 bit key would require 3.67×10^{55} years. Clearly, using keys of sufficient length has some additional computational cost even for those who legitimately encrypt and decrypt, but the impact is limited.

There are Symmetric and Asymmetric algorithms.

Symmetric Algorithms

The most typical and old family of cryptographic algorithms, the ones we typically think of when we think about encrypting messages, are the ***symmetric algorithms***, called after the symmetry of the encryption and decryption operations. These algorithms use a single key that must be known to both the sender and the recipient; it is used to encrypt a message and then to decrypt it.

A simple symmetric schema could be the same as the OTP, as OTP is a symmetric algorithm, with the big difference that the *Key* has a fixed length (e.g. 256 bits) and it is repeated in some way, possibly after some complex change each time:

Encryption: *encryption_function*(PlainText, Key) → CypherText

Decryption: *decryption_function*(CypherText, Key) → PlainText

More in general, even if we use the same *Key* the function used to decrypt is not necessarily the same used to encrypt. For example, if we consider Caesar's cipher, it is clear that the function that decipheres the message slides the alphabet to the left and not to the right, but decryption uses the same key: number 3.

Some important symmetric key algorithms

DES (Data Encryption Standard)

This historical algorithm, now outdated, uses a 56-bit key that in the late 1970s (when the algorithm was developed) was absolutely adequate for many uses, because a brute force attack on a 56-bit key to get the message with the computing resources available at the time and in the contexts in which that algorithm was used, was impractical. 40 years later the DES algorithm would not last even a few minutes of a home computer. This is a very important aspect: the length of the key is chosen based on the *actual* available computing capabilities. A key that is considered robust enough now may no longer be that strong in 20 years. In general we tend to abound with the length of the key in order to ensure that there is some margin, both because computing capabilities increase, but also because cryptanalysis as a discipline is making progress and may be able to find weaknesses in the algorithm. This is quite an interesting aspect. Contrary to what one might think, normally a cryptographic algorithm does not fail suddenly, i.e. the day a flaw is discovered in an algorithm it instantly turns from secure to completely insecure and anyone can find the plaintext in a moment. What happens is that researchers analyze the algorithm and maybe found small weaknesses that can ease attacking it in certain cases. Then someone possibly finds some algorithm optimization to solve the problem. As time passes the algorithm is gradually weakened due to unsolvable problems, until it is no longer adequate for certain uses or at all. There is a progressive erosion of the defenses of cryptographic algorithms, and this progressive weakening is a more difficult problem to manage, because obviously we

cannot know what progress there will be in the future from the point of view of cryptanalysis.

AES (Advanced Encryption Standard, also known by its original name Rijndael)

AES was chosen based not only on robustness, but also on the efficiency of implementation in typical architectures. No secret algorithm was sought because a secret algorithm is not considered more robust than a non-secret one: on the contrary, the algorithm was intended to be analyzed by as many valuable cryptanalysts as possible to ensure that it was intrinsically robust. An algorithm that is safe but too slow, too expensive in terms of computing resources, would be not usable. This constraint was set to be as the most difficult for competitors to overcome, compared to that of robustness.

Asymmetric (or Public Key) Algorithms

Asymmetric algorithms do not use just one key, but two: one key is used to encrypt and the other to decrypt. Therefore, in the exchange of messages, sender and receiver do not use the same key. These keys are generated and managed in pairs: what a key (any of the pair) encrypts can be decrypted only by the other key of the pair.

The reason to use asymmetric algorithms derive from the problems symmetric cryptography has:

1. Confidential communications between 2 entities requires that a unique key be shared between them only, if the entities are n then $n \cdot (n-1)/2$ different keys are required to allow for confidential communication between each pair of the parties
2. Both sender and recipient know their shared key, so *non-repudiation* (repudiation is denying of being the author of a certain message) is not possible

The first problem is not about the number of keys that each subjects has to manage, the problem is *exchanging* them. Each subject must agree on a key with each of the other $n-1$ subjects, which is actually a complex operation; this makes symmetric cryptography impractical for many uses.

The second problem, the *non-repudiation*, lies in the fact that we often want to make sure that a certain message was actually sent by a certain subject. It is clear that if a subject receives a message from some other subject and the two subjects are the only ones that know the key used to encrypt it, each subject is sure that the other subject is the author of the message. Vice versa, since they only know the key, each subject knows that only the other subject can read the message. The problem appears when there is a need to prove this to a third party, for some kind of litigation, as this mechanism does not work: each party is able to create a message, encrypt it, and then say that the other party is the author. The *non-repudiation* problem is not solvable by using symmetric cryptography.

The asymmetric schema is similar to the symmetric one, but the keys used are the two that are

generated in pair:

Encryption: $encryption_function(PlainText, Key1) \rightarrow CypherText$

Decryption: $decryption_function(CypherText, Key2) \rightarrow PlainText$

The names **Key1** and **Key2** have been used instead of *EncryptionKey* and *DecryptionKey* because the two keys are mathematically interchangeable. The subject that generates the keys generally keeps one private/secret (called the *private key*) and makes the other public (called the *public key*), hence the name of the algorithm class that is also called the *public key algorithm*. More on the term *public* will be discussed later.

The function used to encrypt is not necessarily the same as the one used to decrypt. The important thing is that together they represent the cryptographic algorithm, one part relies on one key and the other part relies on the other key.

The two keys have such mathematical properties that discovering one key from the other is quite complex and requires an enormous amount of time. In particular, even knowing the public key, the *PlainText* and the *CypherText*, reconstructing the private key would require an unreasonable amount of time.

Asymmetric Algorithm Uses

Suppose (details will be given later) that the *public key* is publicly available and its owner is the sole to know the corresponding private key. Being the public key available to anyone, anyone can use it to encrypt a message, but only the owner of the corresponding private key will be able to decrypt the message, this provides **confidentiality** of the message transmission.

So we have moved from a situation where one subject needed n different keys to confidentially communicate with n other subjects to a situation where just one key (the receiver's public key) is needed for everyone to send a confidential message to a subject.

In order to allow n subjects to communicate to each other, each subject must have a pair of keys, so the total number of keys is $2 \cdot n$. The interesting point is not so much that the keys have decreased in number; but that being public keys these can be communicated in a much more simple way. If someone else gets to know those public keys, it is not a problem at all. Therefore, these keys can be really published, made available to everyone and everyone can use them without weakening the protection of the messages.

If the subject is the only owner of a private (secret) key, whose public counterpart is publicly available, anyone can verify (decrypt) messages encrypted with the private key of the subject. Then if it is possible to use the public key of a subject to decrypt a message, it means that only the subject that has the corresponding private key could have encrypted the message, this provides

for the **non-repudiation** of the message.

A problem is still to be solved: ensuring that a public key is really associated to the legitimate subject. Here a *super partes* certification method must be provided, this is provided by a *Public Key Infrastructure* (PKI), discussed later.

As a side note, asymmetric algorithms can also be used as a *challenge-response algorithm*. In order for subject A to *authenticate* (prove that a subject is genuinely who claims to be) subject B, A sends (in the clear) a random number to B; then B encrypts it with its private key and sends it back to A. If A is able to decrypt the message with the public key of B, A is sure that the other party is B, as it is the only one that can encrypt something the public key of B can decrypt.

Advantages and Disadvantages

The total number of keys is significantly decreased, which decreases the complexity of guaranteeing the correct management of all the keys (and only the owner knows the secret key). Since public keys are made publicly available, there is no need for specific agreements between different subjects. This use of public keys is, for example, very important in e-commerce, because e-commerce cannot require that there is an a priori agreement of the seller with customers.

To obtain this link between the two keys and, at the same time, guarantee that one cannot be computed from the other, and that what is encrypted with one is decrypted with the other, make asymmetric algorithms much more complex in terms of mathematical problems than symmetric algorithms. These mathematical problems essentially require slower algorithms than symmetric algorithms and the use of longer keys, at least with the currently used algorithms. Therefore, if on one hand there is a greater flexibility, on the other they are more onerous in terms of computing resources.

Prominent algorithm: RSA (from the initials of Ronald Rivest, Adi Shamir, and Leonard Adleman)

This algorithm is based on the complexity of prime factorization: it is easy to multiply two primes and obtain the result, but it is complex/intractable to find the two primes from the result. Prime factorization is a difficult problem. We also know that it is not easy to find very high prime numbers, as there is no rule that allows finding out prime numbers other than to try to factor them and discover if they are prime. This means that working on very large numbers, finding out which are the two prime numbers from which a given number is derived is a problem that requires significant computing resources. On the other hand, legitimate operations can be done with limited computing resources.

Prominent algorithm: Elliptic Cryptography (ECC - Elliptic Curve Cryptography)

The “difficult problem” in this case is the computation of the discrete logarithm (in a finite field). A

“finite field” means that integers values are in a limited set (very rough definition).

The ECC problem is more complex than prime factorization, so you get the same security with shorter keys. ECC algorithms are faster than RSA.

Comparison of Symmetric and Asymmetric Algorithms

Asymmetric algorithms exploit mathematical problems other than those of symmetric cryptography. This means that the lengths of the keys are not comparable to those of symmetric cryptography. The problem with symmetric encryption is just that the key is long enough to prevent trying all possible numbers that could match a key, so with a 128-bit key there are 2^{128} (that is $3.4 \cdot 10^{38}$) possible keys. Today’s symmetrical algorithms typically use 128-, 256- and sometimes 512-bit keys.

If we consider RSA, for example, which exploits prime numbers, it is clear that not all keys are possible, because they are related to the presence of prime numbers. Not all numbers are prime, so not all numbers are possible keys. Very simplifying, to have the same algorithm robustness (therefore the same difficulty from the computational point of view) of a 128-bit symmetric algorithm, the RSA algorithm needs a key of the order of 1024 bits. Today’s asymmetrical algorithms typically use 1024, 2048 and 4096 bit keys.

Hash Functions

Hash functions calculate a short fixed length summary of a text of arbitrary length; this summary is called a *digest*. A hash function must be simple and fast to calculate.

Security properties

Hash functions are “one- way functions” in the sense that it is very complex and slow from the digest to recover a *possible* text that generates it, but this text is hardly useful. This is a consequence of the fact that the original text digest is (and should be) larger than the digest, so virtually infinite sequences of bytes can result in a specific digest. Even if a text is found to have the given digest, it is extremely improbable it is the original text or even makes sense. Building two meaningful messages with the same digest or building a meaningful messages out of a given digest is extremely difficult, any modification of a message, even a single bit, will give a completely different digest (about 50% of its bits are different).

Integrity check

Given a message M, its digest D is computed by means of a hash function H: $D = H(M)$

A secure transmission of digest D enables to ensure that message M is correct: the receiver of message M calculates the digest on it and compares it with the one (D) received; if they are the identical, the message has not been modified. It is essential that the digest is transmitted in a

secure way, otherwise there is no protection at all.

This mechanism is also used as a verification of a downloaded software instead of other less powerful (but faster) algorithms known as CRCs (Cyclic Redundancy Checks), common in network protocols.

Some Important Hash Algorithms

The following are two of the most known hash algorithms.

MD5 (Message Digest #5, 1991)

MD5 requires a low computational power (compared to the majority of the other hash algorithms) and produces a 128-bit hash. Over time it has been gradually weakened to be completely unreliable for security purposes (at present it is possible to find two distinct messages with the same MD5 hash – which is called a *collision* – in seconds), anyway it is still useful in non-cryptographic applications, for example to discover transmission errors instead of the more common (and weak, but faster) CRC-32.

SHA-3 (Secure Hash Algorithm #3, 2015)

SHA are actually a family of secure hash algorithms, SHA-3 is the last version and the result of a competition among cryptographers. The algorithm can be tuned balancing security with speed; it produces hash values of 224, 256, 384 and 512 bits. It is 2-3 times slower than MD5 but no weakness have been found so far.

Signatures (Digital Signatures)

Digital signatures, called electronic signature in the legislation of some countries, take a message, calculate a hash, and encrypt it (just the hash) with the secret key of an asymmetric algorithm. This hash signed with the secret key is called a *signature*. A digital signature then does not encrypt the original message, just allow to guarantee that the corresponding text has not been modified, because otherwise the digest would change. It also guarantees that the origin can be verified (*non repudiation*), as it was actually encrypted with the private key of an asymmetric algorithm. This means that anyone can verify that the text is intact and has actually been generated by a certain subject.

Similar to digital signatures are **Message Authentication Codes (MACs)**. They also aim to ensure the integrity and authenticity of a message. The digest is computed hashing the text to be protected and a *symmetric algorithm* key (there is no assurance of non-repudiation, so it is suitable only when this characteristic is not needed). One of the most known algorithm is HMAC: Keyed-hash MAC.

Public Key Infrastructures

The confidentiality, integrity and authenticity requirements of asymmetric encryption and digital signature operations require establishing a relationship of trust between the parties involved. As often the parties had no contact before, a trusted (by them) third party must provide the needed guarantees, in other words this third party must assure the authenticity of the owner of a public key.

A *Certificate Authority* (CA) is a part of a hierarchical trust model that guarantees for the user's identity. It must be recognized as “trusted” by all entities involved in the communication.

In order to assure about the authenticity of a subject, a CA signs a *digital certificate* that links the subject to its public key (it is conceptually similar to an Identity document of a person). The signature applied by the CA is a digital signature, so it is computed by using CA's private key. Of course, the digital certificate must be validated by using CA's public key.

A *Public Key Infrastructure* (PKI) is a hierarchical structure of CAs, where every CA has its public key signed by the previous level CA (which uses its private key for this purpose), the process goes up to the *Root Certificate Authority* that self-sign a certificate holding its public key by using its private key.

A hierarchical model allows building hierarchies of any depth. For example, within a Country we could have the municipal, ..., regional certification authorities, up to a national root CA.

The root CA certificate must be made public in many locations and ways in order to make it impossible to substitute it with a forged one. As an example, the Operating Systems themselves are shipped with a set of root and trusted intermediate CAs. Figure 1 shows an excerpt of the Root Certificates list from a MS Windows 10 operating system.




Figure 1. Windows 10 Certificate Manager

Components of a PKI

The main actor is the **Certification Authority (CA)**, it manages the cryptographic keys and certificates of the PKI. It is the component that holds the real keys and that then performs the cryptographic operations. Obviously, in this sense it is also the most critical element of the certification authority.

The **Registration Authority (RA)** manages certificate requests and provides issued certificates. It deals with the distribution of cryptographic tokens and software (it is the front-end of the CA). It is like a counter, for example, to which a citizen can go to request a digital certificate and from which it can be obtained. It can be a physical counter or a website, the type is associated to different levels of security. In addition an RA generally can also assist the user in generating the key pair and managing these certificates securely, for example by providing ad hoc software or hardware tools that support the applicant in generating signatures, encrypting documents, verifying signatures, and decrypting information.

A **Certificate Server** is directory where user certificates are published and in some way made searchable. The directory follows the X.500 standard and is queried by using the LDAP protocol.

The Digital Certificate

The digital certificate structure and format is described by the ITU-T X.509 standard (it is a standard of the International Telecommunication Union numbered X.509).

An X.509 digital certificate must contain at least the following fields:

DN (distinguish name) user: Gaius Julius Caesar

CN (common name): Julius Caesar

O (organization): Government

OU (organization unity): Senate

C (country): Roman Empire

This is the classic pattern of a certificate that meets the X.500 standard, of which X.509 is a member of the family. A certificate can have other attributes (e.g. email). There may be an IP address or a URL that corresponds to an X.500 Directory path (e.g. LDAP – Lightweight Directory Access Protocol):

IT / Roman Empire / Government / Julius Caesar / Gaius Julius Caesar /

An X.509 digital certificate must also contain at least the following fields:

DN Issuer: name of the Authority in X.500 format (the certification authority also has its own distinguished name, which appears in the certificate. Obviously the format with which it is represented is the same as the distinguished name of the subject to which the certificate applies.

Serial Number: serial number of the certificate held by the Authority (each certificate is uniquely identified by the serial number assigned to it by the authority. On the other hand, the pair (*serial number & issuer distinguish name*) uniquely identifies a certificate throughout the plethora of certificates that can be issued worldwide as a unique pair).

NotBefore: activation date

NotAfter: expiration date

These are two information strictly related to each other. A certificate therefore has a date of beginning of life and an expiration date, beyond which it is no longer valid. The reason derives from the desire to protect the keys from excessive use.

Public Key: the public key of the certified subject (the certificate contains personal data and the public key)

Public Key Algorithm: the algorithm the public key is suitable for

Signature Algorithm: the algorithm used by the Authority to sign the Certificate

Signature: the signature affixed to the Certificate by the Certificate Authority (therefore the

element that guarantees the authenticity of the certificate itself)

KeyUsage: the purposes the key can be used for

Generation and Revocation of Certificates

The RFC 2314 standard (from PKCS # 10) defines the format of the *certificateRequest* message and contains the public key and data of the applicant (both given to an RA). The request message requires the so-called *proof of possession* of the private key: the *certificateRequest* message must be signed with the private key of the applicant. The authority, by verifying this signature substantially verifies that the applicant actually possesses this key (otherwise the applicant would not have been able to produce this signature).

On the date corresponding to the "notAfter" attribute, the Certificate is considered "expired". The expiration is a mechanism that also serves to ensure that the same asymmetric keys associated with the digital certificate cannot be used for an excessive period of time, which could expose them to a series of cryptanalytic attacks.

The expiration can refer to:

- **the certificate** (the association subject-key)
- **the associated keys**, in relation to the use for which they are intended (a key can be used for certain goals only and not for others)

An expired certificate cannot be used by its owner, since its use is in practice prohibited. In reality, nothing prevents the holder of the certificate from using it for other purposes but all those who communicate with him will somehow detect that this certificate has expired and will not accept it. The certificate is of course available to verify and decrypt messages issued before expiration. In some cases, it is possible to "renew" a certificate, which is usually not automatic.

Sometimes a certificate can be made invalid prior to expiration through *revocation* (this happens before the "notAfter" attribute value. The reason could be, for example, the key associated with the certificate has been compromised, the subject is no longer associated to the company or changed unit, the CA that issued the certificate (or one of its higher level CAs) has been compromised or other reasons. A revoked certificate cannot be used anymore, except when it is just kept "on hold". This is a particular form of revocation in which the certificate is temporarily inserted in the CRL because of specific reasons (such as "possible key compromise", "request of revocation by a person other than the owner", "temporary non-use of the certificate" e.g. for vacation)", etc. and its validity could be restored later. In case of "reactivation" from suspension, a new entry in the CRL is inserted with a special reason code: *removeFromCRL*, as no entry can ever be deleted from the CRL.

To inform the PKI users that a certificate has lost validity, its serial number is listed in a *Certificate Revocation List* (CRL). This list is issued periodically by the issuer Certification Authority and propagated (actively or passively) to all the users, so that all the subjects who have such a certificate can no longer use it. The CRL complies with ITU-T X.509 standard - rfc3280 - rfc4325 which means that its format is standardized. It is signed periodically by the CA and includes a field called the “nextUpdate” which contains the date and time in which the new Certificate Revocation List will be available. If this date belongs to the past it means that our CRL is “expired”, if the date belongs to the future we know when we will have to update our list of revoked certificates. Each revoked certificate includes the reason of revocation. The CRL is provided as a service via the LDAP protocol.

As CRLs can become quite large, especially if the certification authority has hundreds of thousands or even millions of users, it can be partitioned and the CA just provides pieces of the Certificate Revocation List, for example referring to a certain period of time. Therefore, when it is necessary to verify a certificate, the citizen downloads the only piece of CRL of interest. Instead of downloading a CRL or a part of it, a service implementing the *Online Certificate Status Protocol* (OCSP - RFC 2560) can be used to find if a certain certificate number is in the CRL. The responses are signed by the CA in order to be trusted

The Chain of Trust

To verify a certificate, a subject must verify all the chain of the intermediate CAs starting from a trusted one, possibly the root CA. First the certificate to check is verified against the issuer CA, then the CA certificate is verified against its parent CA, and so on up to an intermediate trusted (by the user) CA or the root CA itself.

KeyUsage

Actually, keys can be used for different purposes and a certificate specify the allowed uses, among them we have:

- **nonRepudiation:** the key can be used for producing a signature with legal value which can then be opposed to third parties as a handwritten signature
- **dataEncipherment:** the key can be used to encrypt data
- **keyAgreement:** the key can be used to securely exchange keys

TimeStamping

A PKI can also be used to time stamping a document according to the Time Stamp Protocol (RFC 3161) to prove its existence in a point in the past. A CA will sign the document and includes a temporal mark. The signature also ensure that the document has not been altered after affixing the timestamp. A Time Stamp can extend the validity of an electronic signature beyond the lifetime of the digital certificate associated with the keys used, for example by requesting the affixing of a new time stamp on the contract every year that passes. The affixing will be done

every year using cryptographic techniques with a strength adequate for the period in which the signature itself is applied.

Standard Mechanisms and Formats

The following are the two standard mechanism and formats commonly used.

PKCS #7 - Public key cryptography standard #7

PKCS #7 describes the format of messages/files signed, encrypted (with symmetric and asymmetric encryption) and verifiable by MAC (digested data)

S/MIME - Secure Multipurpose Mail Extension

Secure Multipurpose Mail Extension (RFC 3851 - SMIMEv3) is an extension of the MIME standard (RFC 2045, 2046 and 2047 as an extension of RFC 822) and describes the encryption of mail messages. PKCS #7 entities are substantially translated into new MIME entities and therefore describe new containers in which we can find portions of signed or encrypted e-mail messages or signed or encrypted attachments.

Exchange and Storage Formats

The **PKCS #12 (Personal Information Exchange Syntax Standard #12)** defines a format for the exchange and secure storage (on file) of cryptographic keys and digital certificates (trusted certificate chain). Storage is safe because it is based on encryption of both the keys and the certificates, and of the entire chain of certificates a subject chooses to trust. It is the format used by browsers, e-mail clients and operating systems when a user exports a signature and encryption key and a certificate, for example for a backup copy or to copy it to another device.

A PKCS #12 encrypts the data with a symmetrical algorithm whose key is derived from a password provided by the user. It guarantees the integrity of the content through HMAC, or through a digital signature.

The benefits of using PKCS #12 are:

- there are **no implementation costs** as it is a format available in many libraries for many devices (from computers to cellphones)
- they are **already supported by most operating systems** and software, this means that it is immediately available
- it **uses “strong” symmetric algorithms** (the encryption that is applied is a strong, very resistant encryption, so they are keys that resist for example brute force attacks)
- it is **portable on different devices**, it is possible to move a file to other devices or copy it to all the devices we use to have the keys always available

There are also some disadvantages:

- **it can be deleted, copied, transmitted**, this is an important problem because it can be stolen, erased, copied to take it away or just made unavailable
- **the private key must be used directly by the signing/decryption software**, this means that the software uses decrypts this file and accesses the key directly, so there is a moment when the key is in the clear on the system, unprotected, although maybe it is just the volatile memory which is the least easily accessible area; however it is available in clear text on the system in some way and this is a very strong limitation of this format
- **it possible to perform a brute force attacks**, for example, in the event that someone steals the PKCS #12 file, they can try all possible passwords in a separate computer system

Cryptographic Tokens

Cryptographic tokens are hardware devices that natively implement encryption algorithms. They do not allow direct access to cryptographic material from outside, they are the only enabled to access and use the key. Moreover, these commands can generally be executed only after the user, and the system on behalf of the user, authenticate to the token, for example by entering a PIN. Tokens can have a small Operating System through which the user requests key generation, encryption and data signature operations. The secure destruction of the token often means not only the cancellation, but also the inability to access the memory areas where the key was written, so to avoid any form of attack that we could not conceive today. The operating system and token hardware can be certified according to specific security standards.

The Challenge-Response Mechanism

A challenge/response mechanism allows two parties to authenticate without exposing useful information on the communication channel. The idea is that one party A “challenges” the other party B to operate on a random number (called *nonce* because it is a number used just one time) with its own key, for example with a One-Way function H . Then B returns $H(\text{nonce})$ to A. If B returns what A expects, A is sure about the identity of B. In the simplest case, B encrypts the nonce with a secret key known to A. A, which has the same key, in turn applies the function to the nonce using the same key, and verifies that the number given to it by B matches the one it calculated.

This is an authentication based on a password that has not been sent over the channel, so it is not possible to trace the key. What makes the mechanism effective is the fact that the random number is different for each session. Since the information read on the channel is different for each session, it is not useful for subsequent authentication. An advantage is that no synchronization is needed (differently from the One-Time-Password mechanism, where it is important to know the starting point of the Pad to use). As it is uneasy to be used by humans, it is

typically used in machine-to-machine authentication. A challenge-response mechanism is found in the network authentication mechanisms of ADSL routers to the provider authentication server.

Questions:

In separate file.

Attachments:

Powerpoint in a separate file.

References:

Every topic has its own page on Wikipedia, sometimes already too thorough for the purposes of this course; books are too specialized and are useful only after a very deep training on the specific subject.

InCyT Training Module

Information Security for Employees

Unit 2 Week 4

Unit name: Malware and Most Common Attacks

<i>Learning Activities</i>	☒ Webinar ☐ Exercises ☐ Workshop ☐ Presentation ☐ Other
<i>Learning Responsible</i>	Claudio Fornaro
<i>Learning Activity Goals</i>	10. Introduction to the various types of malware 11. Knowledge of the most common types of attacks
<i>Skills to be Gained</i>	Basic knowledge about what is malware and the most common attacks to computer systems Soft Skill 2 - SS2
<i>Duration of Learning activity</i>	1 week
<i>Learning requirements</i>	<i>What is needed?</i> No previous specific knowledge is required

Brief information about the module

When it comes to malware, the name already implies something malicious. This term generally refers to software designed intentionally to cause damage of various kinds, in particular (but not limited to): blocking of services and communications, interception of transmissions, theft of private information (personal, corporate or government) and of personal identities, unauthorized access to information or systems, deprivation of users of access to information.

After an introductory part, the fundamental concepts of malware are outlined, along with the description of the main types, the risks that they entail and how these can be mitigated.

The second part of the module intends to offer an overview of the most common attacks that a system can suffer, this list being by no means complete as new types of attacks are discovered every day. These attacks are attempts to gain access to a computer system or disrupt its services. The description starts from the most common and easiest attacks (Phishing) to more complex ones: Password Attacks, Buffer Overflow, SQL Injection, Cross-site Scripting (XSS), Cross-Site Request Forgery (CSRF), File Related Attacks, Denial of Service (DoS) and Distributed Denial of Service (DDoS), Man in The Middle (MITM), ARP Spoofing or ARP Poisoning, DNS Spoofing, Drive-by Attacks,

Credential Reuse, Web Session Hijacking, TCP Session Hijacking, IP Spoofing, Zero-day Exploits, and the consequences of Leak of Code.

Content of the Learning Material

In computer science, "hacker" and "cracker" are two terms that are used to indicate people with a great knowledge of computer systems. The main difference between the two lies in the attitude with which they act. A **hacker** (the term was born in the second half of the 1900s at the MIT) acts by exploring systems in order to better understand how they work and trying to discover new flaws in order to solve them and therefore increase the security of a system. Some actions can only be demonstrative or for ethical reasons, but never for profit. A **cracker**, on the other hand, uses his abilities illegally trying to break into systems to profit by stealing data or in a destructive way.

Hackers are constantly trying to acquire new knowledge of computer systems in order to better understand them, or improve their functions, performance and security. Some of them also engage in demonstrative and ethical actions, but never for profit. Hacker actions are performed almost always without any permission, so the Law punishes them (with great differences among the countries), for this reason hackers usually hide behind nicknames. Sometimes, wise security or IT companies instead of prosecuting hackers, often hire them to take advantage of their high skills and creativity (possibly after paying part of their debt to Justice). They are then employed to test

applications, security systems and programs both in production and already on the market in order to block attacks, and to prevent problems that could arise in the future from system bugs.

Real hackers have a reputation as experts, and are highly respected in their community. Competitions are organized with the aim of "piercing" a system or program in the shortest possible time. These competitions are called *hackathons* and are often sponsored by the giants of the computer industry.

The term "cracker" is attributed to Richard Stallman (a well-known figure of the free software movement, creator of the GNU project and member of the hacker community since the 1970s). He wanted to avoid the abuse of the term "hacker", differentiating who acted without any willing to harm and who, on the contrary, acts in order to have profit or just give damage. A cracker is a cyber-criminal who tries to attack computer systems in an unauthorized and illegal way, with malice, trying to take fraudulent advantage of them by exploiting vulnerabilities not yet known and resolved. Crackers are often involved in industrial espionage, data theft, and blocking of networks and websites, vandalism, etc.

The terms hacker and cracker have subtle bounds that divide them. For example, an intrusion can be considered an illegal act for some people and not for others, and the Law reflects these opinions; but this document is not about ethic or legislative aspects. Instead, we must consider that today, unfortunately, the two terms are considered synonymous.

In this document, we will *unhappily* use the more common and misused term hacker for both categories, being sure that the wise reader will be quite able to distinguish between a *hacker* and a *cracker*.

"Hackers" can be roughly classified into several categories, the most widely used are the following.

Black hat hackers use their knowledge for malicious or criminal purposes, being careful not to be traced.

White hat hackers test systems in order to identify possible vulnerabilities and eliminate them to prevent unwanted access. They often are security experts paid for this purpose.

Gray hat hackers are similar to white hats, but they often conduct research on system vulnerabilities without authorization and once these vulnerabilities are discovered they use them to force vendors to fix the problem, usually for free. Some vendors reward them with money, others prosecute them.

Suicide hackers are usually less competent than black/white hats, but they act for ideological reasons and have no interest in hiding their identity.

State-sponsored hackers are paid by a Government to launch cyber attacks against another Country; they have access to large resources as provided by the State as well as protection from prosecution for the offenses committed. They act to subtract confidential information, plans or projects, or to create generalized disorders and damage the population, the Army, the Government management, hospitals, plants for the production of energy, water and power supply, etc. These actions are considered real war actions.

Script kiddies do not have or have not yet acquired great skills, they usually make use of scripts and tools created by other hackers. All hackers before they acquire great skills are script kiddies.

Cyber terrorists are individuals or groups with hostile intentions that generally aim to create chaos in a Country; they are similar to state-sponsored hackers but do not have (yet) a Government to support them, or the support is not evident.

Malware

Malware is the acronym of “**malicious software**”, a generic term that identifies any malicious program harmful to a system. Malware is the software hackers create to pursue their goals: making money, stole data, disrupt activities, etc.

One of the most common malware is the so called a **Trojan horse**¹ (or simply **Trojan**), a program that seems useful software and often act as such, but carries a hidden function that is executed to do along with the normal activated when the application is started. These are often shareware programs or cracked programs that are distribute modified in order to carry the Trojan horse. In other cases, it is just malware with a misleading name that makes it seem a useful software. Trojans are often spread by emails or by downloading programs through the Web and usually do not spread by themselves to other computers. Malicious code examples are *keyloggers* (programs that steals passwords and send them to the hacker), *backdoors* (programs that work as a terminal server and allows the remote hacker to access the computer in some ways bypassing the normal authentication procedures), *cryptominers* (programs that mines cryptocurrency), *bots* (short for robots, used for example for a distributed denial of service attack).

Rootkits are software installed by malware that try to hide the presence of a malware itself from detection; they achieve this by modifying the operating systems.

The malware that infect other software can be a virus or a worm. These get their name from the way they spread rather than from any specific types of behavior. A computer **Virus** is software similar to Trojan horse but can produce copies of itself by infecting (introducing itself into) other software, including the operating system files. The virus is spread to other systems when infected software is copied and run on these systems.

¹ The term is derived from the Ancient Greek story of the Trojan horse used to invade the city of Troy by stealth

Similar to a virus is a **Worm**, this is a stand-alone software that operates on a computer and infects other computers through the network, but without infecting files.

A **Ransomware** is a type of malware that restricts access to the device it infects, requiring a ransom to be paid to remove the restriction. The main types of ransomware are 1) *Screen-locking ransomware* that blocks screens with an intimidating accusation of illegal behavior and asking the victims to pay a fee; 2) *Encryption-based ransomware*, where the malicious software some type of files (pictures, documents) or all files on an infected machine. When the full encryption is done, a pop-up usually informs the user that (all) files have been encrypted and that a fee must be payed to recover them, usually in cryptocurrency. Examples of encryption-based ransomware are CryptoLocker and WannaCry. A variant known as a **Wiper** looks like a ransomware but just destroys the data, even if a payment is often required.

Greyware are a different category of malware: these programs can worsen the performance of computers and may cause security risks, behave in an annoying or undesirable manner, but usually are less dangerous (which is a relative opinion) than the previous categories of malware. Spyware (e.g. keyloggers, programs that record all the keys pressed on the keyboard), adware (advertisements windows on the screen or in the browser, fraudulent dialers, etc.), Potentially Unwanted Programs (PUPs, applications that the user installed just because of a mistake in a windows of agreement during another program installation) are typical form of greyware. They can be stand-alone programs or plugins for other programs, for example browsers. They sometime limits themselves to change the settings of other programs, for example in order to bring the user to a specific ad website or show an ad page each time the browser is started or occasionally.

Antiviruses

As antivirus engines enhance their power to detect malware, a combination of various techniques are used to try to avoid detection and analysis. More recent malware is able to adopt sophisticated countermeasures such as:

- analyzing and detecting by fingerprinting the environment to adopt specific hiding methods;
- trying to obfuscate themselves in order to confuse antivirus detection based on malware signature;
- adopting a timing-based evasion in order to fool detection of anomalous behavior, the malware is normally quiescent and activates itself in very specific periods (for example during the boot process) or following certain actions taken by the user.

A very hard to discover technique for evasion is *Fileless malware* or *Advanced Volatile Threats*. This kind of malware runs in memory and does not require a file to operate. It uses existing system tools to perform its activity. The absence of any file on the file system makes them quite

difficult to detect and analyze, except, partially, for those monitoring tools that check the resource usage and behavior of the system. The increase of this type of attacks is very high.

Risks

Vulnerable software can be exploited to gain privileges or bypass defenses, until the problem has been patched; this is the main reason that software should be up-to-date, in particular the Operating System. Firewalls and Intrusion Detection Systems monitor traffic patterns on the LAN.

The rule of the least privilege says that users and programs should have no more privileges than they require, otherwise malware can take advantage of this.

Mitigation

Antivirus software block and possibly remove some types of malware, they can provide real-time protection against the installation of malware software on a computer and a full scan of files and configuration (e.g. Windows registry) looking for potential threats.

Personal firewalls analyze any network connection and can spot unusual activity.

Sandboxing is running an application isolated from the rest of the system, some application can run a specific application while in more complex systems a virtual machine could be more appropriate.

Website vulnerability scans is performed by some antivirus software against a list of well-known dangerous sites and prevent the connection to them.

Network Segregation dividing a network in parts and enable only that traffic between them that is known to be legitimate, this prevents malicious software to spread across the network. Modern Software Defined Networking techniques are able to implement realize this segregation.

“Air gap” isolation or “parallel network” brings Network Segregation to the limit by completely isolate a portion of a network and provide enhanced controls over the entry and exit of software and data from the outside world. Even this solution is not the definitive one, as software and data must enter the isolated system, for example for patching. A well-known example of breaching an isolated network is the famous **Stuxnet** malware that was introduced to the target environment via USB drives “casually” forgotten nearby, picked up by the plant staff and inserted into the plant computers to see the contents. Stuxnet has three modules: a worm that controls the attack; a link file that automatically executes the worm when inserted in a Windows computer; and a rootkit component responsible for hiding all malicious files and processes, to prevent detection. Other ways of crossing air gaps analyze electromagnetic, thermal and acoustic emissions.

Most Common Types of Attacks

This part describes some of the most common attacks to a computer system. It is not and cannot be complete as the variety of attacks is extremely huge, each one has many variants and many attacks and variants are created and discovered (and solved) every day. This means that the aim of this part is not to provide a complete classification, but to give an idea of the high level of knowledge required to both perpetrate and counter attacks.

Phishing

Phishing attacks are very common and easy to perform. They consist of sending e-mails that appear to come from a reliable sender, such as a colleague or a bank, with the aim of tricking the victim into doing a certain type of operation.

The attacker tries to persuade the victim to open an attachment or provides a link that leads to a page very similar to the original one, for example a bank home page, but created by the attacker, then to ask the victim to enter their credentials. To optimize this type of attack there are a whole series of techniques that are part of the so-called social engineering, since, in this case, the vulnerability of the computer system is to be identified in the user himself, who is deceived and unintentionally provides his credentials.

To defend against these attacks, a careful inspection of the link is needed (it is near impossible to distinguish between lowercase L 'l', an uppercase I 'I', and the digit one '1', or to notice the difference between .it and .It in a URL). Moreover, one should be aware that serious companies never put links in their email that bring to a login page.

Password Attacks

To get passwords from a victim a hacker can try social engineering, as seen in the previous paragraph, or launch a password attack. Passwords are usually stored as hashes, and hashes are very hard to invert. There are essentially two cases: the encrypted password is available or is unknown and this type of attack can be performed in different ways:

- **Brute force attack:** in this case the program will try all the possible combinations of characters out of a set;
- **Attack with dictionary:** the program will test a sequence of commonly used passwords (dictionary);
- **Attack with rainbow files:** a rainbow table is a precomputed table listing the output of hash functions of password. Use of a key derivation that employs a salt makes this attack infeasible.

System managers should test the password files of the users of their systems on a regular basis with programs capable of performing password attacks, both by brute force and by dictionary.

To protect against these attacks, a system should check the quality of the password before allowing the user to set it, requiring the use of a minimum number of characters (e.g. 10) from different sets (upper and lower case letters, digits, punctuation characters). If the password entered is incorrect for some subsequent attempts, it is appropriate to adopt an account lockout or delay policy on the next attempt, thus preventing the attacker from quickly trying many combinations.

Buffer Overflow

Buffer overflow (or overrun) occurs when you send more bytes to a buffer than it can hold, these bytes overflow the buffer intended to hold them and change other bytes or, depending on the Operating System and file type, even inject executable machine code. While modern Operating Systems may prevent running an executable file with mixed data and code, this is not true for the simpler Operating Systems often found in embedded systems, where in many cases there is no Operating System at all.

SQL Injection

A very common attack is SQL Injection. The SQL language is used in relational databases; hence, these types of attacks are aimed at information theft. In essence, due to the interpreted nature of the SQL query, malicious SQL code is sent to a web page that is known to use an SQL query to gather the data to present to the user. The malicious SQL code modifies (sometimes just adding other syntactic parts) the predefined SQL query that results in a modified query that exposes other data than those the original query was written for.

While preventing this type of attacks is not complex (for example preventing the access to other tables of the database besides the ones needed to perform the original query), SQL injection is still one of the most common attacks.

Cross-site Scripting (XSS)

In a cross-site scripting (XSS) malicious executable scripts are injected into the code of a trusted web application (plugin) or website page, with the result that it is downloaded by other people (viewers). When the server sends back the compromised web page, the user's web browser considers it been delivered from a trusted source, and then the injected code is not identified as being malicious. The active content is usually a script (programming code that is executed by the browser) that can allow an attacker to get access to sensitive page content, session cookies, or other information the browser keeps on behalf of the user. With XSS attacks, it is possible to steal

sensitive data of all users who access infected sites. For example, stealing the SESSIONID of a website session (which is located in the user browser cookies that in turn are used to avoid that each connection to the same website require authentication beyond the first one), could allow an impostor to bypass the login phase and illicitly operate on behalf of the user.

Cross-Site Request Forgery (CSRF)

A CSRF is an involuntary http request (wanted by the attacker) to a website performed by a user already authenticated to that website; the website at this point will execute the request without further verification, allowing the attacker to carry out his attack. This type of attack allows modifying user data or performing unwanted transactions, such as unwanted purchases, if it were for example an e-commerce site. The attacker must have very good knowledge of the site in question to carry out a CSRF attack.

An example of a CSRF attack is performed by changing an URL. For example the source code of the web page is analyzed by the malware to find the “Change password” link and verify that the http GET method is used, then a new link with a new password (chosen by the attacker) can be formed and sent to the victim, for example via an email attachment. If the link is clicked and it happens that the victim is already authenticated to the website, a hidden start of password changing procedure is activated, effectively giving the account to the hacker. Good password changing methods ask to insert the old password too (and the GET method is not used). Moreover, the use of a CAPTCHA method (a test to verify whether the user is a human or a computer) is quite effective.

File Related Attacks

There are many types of attacks on files. The most known vulnerabilities are:

- **Insecure Direct Object References:** occurs when an application requests a resource from a server calling it by its name, but allowing the user to modify the latter, in order to request another resource.
- **Local File Inclusion (LFI):** exploits the vulnerability of the GET and POST functions of the PHP language to insert malicious files into the code.
- **Path Traversal:** allows a malicious user to access directories with limited access and execute commands. Often, they are used following an LFI to get deep into the system and obtain the necessary privileges.

To protect against these types of attacks, it is important to prevent file uploads to servers, disable remote code execution, restrict PHP access to the file system, and update software periodically.

Denial of Service (DoS) and Distributed Denial of Service (DDoS)

Maybe the most know network attacks, DoS attacks basically consist in sending many packets to

block the possibility to receive other packets, thus preventing access to the machine for regular users, blocking the services offered.

The difference between a DoS and a DDoS attack lies in the fact that in the first case a single host is used to carry out the attack, while in the second the attack is carried out with many hosts (called zombies) at the same time. These hosts have been previously infected by a malware and at the command of the attacker start the connection to the victim machine.

Among the DoS attacks, we can remember:

- **Syn-Flood:** consists of sending a large amount of connection requests. The machine that receives a request (a SYN packet) from a client will respond to the client by sending a message, allocating some resources for this connection, and waiting for the acknowledge packet (that will never arrive). Therefore, the connection request remains active until a timeout occurs. If all the possible connections have been started, the service provided by the machine will no longer be available. The timeout occurs within seconds but the huge rate of requests keeps the machine clogged.
- **Smurf:** this attack is a combination of IP spoofing and ICMP requests. IP spoofing refers to the sending of an IP packet in which the sender's IP is forged. By continuously sending *ICMP echo requests* (the message sent by the ping command, used to verify if a machine is reachable) to broadcast IP addresses after having substituted the sender IP address with victim's IP address, the victim machine will receive a response for every request from any host active on the network, thus blocking the connection (flood). To prevent these attacks from occurring, it is possible to disable in the routers the forwarding of broadcast IP addresses.
- **Ping of death:** consists in sending a fragmented packet with a total IP size greater than the maximum allowable to a victim machine, once the machine reassembles the packet it can incur buffer overflows.

It is common for a DDoS attacks to use a **Botnet** that is thousands or millions of computers infected with malware, which at the command of the attacker, execute a Syn-Flood attack in unison on a victim server.

Man in The Middle (MITM)

With MITM we mean a whole typology of attacks that aim to intercept, analyze and retransmit a communication between two parties who believe they are communicating directly to each other.

This type of attack is usually carried out using IP spoofing: network sniffers eavesdrop a connection and new packets created by changing the IP address of the sender and recipient so that the sender sends data to the attacker who in turn it will re-send them to the recipient and vice-versa.

There are many ways to defend against these types of attack; the one most effective and used

today is adopting *end-to-end encryption*, where the end subjects perform the encryption/decryption and all the intermediate systems can only see encrypted data.

ARP Spoofing or ARP Poisoning

The Address Resolution Protocol (ARP) enables network communications to reach a specific device on the local network. It is used to map IP addresses to Media Access Control (MAC) addresses (used to identify the hosts on a LAN). A host must use the ARP protocol to know the MAC address of the router/gateway in order to connect to the Internet. Each host maintain an *ARP cache* to map already requested and known IP/MAC pair. If the host does not know the MAC address for a certain local IP address, it sends out a broadcast ARP request packet, asking all the other hosts on the local network who has the IP address. The host with that IP address will answer, thus providing its MAC address.

An ARP spoofing/poisoning attack can be classified as a local network MITM attack. The attacker must have access to the local network and scans it to determine the IP addresses of the router and another host. The attacker sends forged ARP responses that make the router and the host to believe that the attacker's MAC address is the one of the host and the router respectively. This fools both router and workstation to connect to the attacker's machine, instead of to each other. As the two devices update their ARP cache entries, host and router communicate with the attacker instead of directly with each other.

DNS Spoofing

The Domain Name System (DNS) is the naming system used to convert Internet names (more properly FQDN - Fully Qualified Distinguished Names, e.g. www.server.net) to the corresponding IP address. The system is hierarchical and decentralized with some top-level servers. Each time a system attempts to reach a network resource, a connection is made to a DNS server in order to convert the name of the resource to an IP address. The resource records usually contains administrative information and other data used for example by mail servers.

A hacker will want to alter DNS records in order to send traffic to a fake or "spoofed" website instead of the legit one. The spoofed machine could then do MITM attacks, spoofing, etc. To prevent DNS spoofing, DNS servers must be kept up-to-date.

Drive-by Attacks

In a drive-by attack, malicious code is embedded into an insecure website. When the infected page is visited, the script is automatically executed on the victim computer and compromises it, there is no need to do anything on the site or enter any information. In this case, the use of an updated recent and decent antivirus and/or personal firewall is the only protection as they can

detect from their database if a site is unsafe before a user visits it.

Credential Reuse

The need to have credentials for several services brings users to reuse some of their passwords. Once an attacker has a list of usernames and passwords of compromised websites or services (for example obtained from the Dark Web), it is possible that the same user could use the same user/passwords pair on other systems. This will greatly reduce the number of password to check with a dictionary attack. In order to contrast this type of attack, the same countermeasures adopted to counter password attacks are effective. When many complex passwords are needed, some people find password managers useful.

Web Session Hijacking

A Session Hijacking attack consists of the exploitation of session tokens, used by web browser as session control mechanisms. The web server needs to relate different connections from the same user, as for each web page a TCP connection is created for each element (text, images, etc.). The most used method is based on a token that the Web Server sends to the client browser after a successful client authentication. A session token is just an identifier, a long unique string, and can be used to avoid the need to perform an authentication procedure for each connection. This attack is performed by stealing or predicting a valid session token to gain unauthorized access to the Web Server.

TCP Session Hijacking

TCP connections must guarantee to deliver packets in the right sequential order, in order to achieve this it uses acknowledgement packets and sequence numbers. These acknowledgement packets are just numbers that are incremented in some way after the 3-way initial handshake. The goal of the TCP session hijacker is to create a state where the client and server are unable to exchange data; enabling him/her to forge acceptable packets for both ends, which mimic the real packets. Thus, the attacker is able to gain control of the session. If the sequence can be predicted, an attacker can send forged packets that are recognized as valid from the victim host.

IP Spoofing

An IP spoofing is usually performed with a DoS attack: a huge number of connections blocks the legitimate server and, at the same time, a fake server sets its IP address with the number of the legitimate server. Therefore, the user is driven to the fake server that sends messages with an IP address indicating that the message is coming from a trusted host.

Blind Hijacking: In cases where source routing is disabled, the session hijacker can also use blind

hijacking where he injects his malicious data into intercepted communications in the TCP session. It is called blind because he cannot see the response; though the hijacker can send the data or commands, he is basically guessing the responses of the client and server.

Zero-day Exploit

Strictly speaking, zero-day exploit is not a specific attack as cyber-criminals learn of a vulnerability that has just been discovered in certain software and operating systems and then target organizations who are using that software before a fix becomes available.

Leak of Code

The source code of applications can sometimes contain sensitive information and give valuable insights for a cyber-attack, in addition to the intrinsic economic value of the code itself.

Questions:

In separate file.

Attachments:

In separate file.

References:

Every topic has its own page on Wikipedia, sometimes already too thorough for the purposes of this course; books are too specialized and are useful only after a very deep training on the specific subject.

1. Malware, or malicious software
URL: <https://www.malwarebytes.com/malware>
2. Cybersecurity for everyone.
<https://www.avast.com/c-malware>
3. What are Cyber Security Threats?
URL: <https://www.imperva.com/learn/application-security/cyber-security-threats/>
4. 21 Top Cybersecurity Threats and How Threat Intelligence Can Help
URL: <https://www.exabeam.com/information-security/cyber-security-threat/>