Computer Security Fall 2023
Network Security Homework

Srdjan Matic, Juan Caballero
IMDEA Software Institute

Abner José García Ferrufino

Problem 1: Initial Sequence Number Guessing *(14 points)*

Which of the following ISN selection methods used by a server are secure? If you believe that a method is not secure, explain the reason.

a. The server uses as ISN a constant value selected at random at boot time. *(2 points)*

Using an Initial Sequence Number (ISN) as a constant value selected at random at boot time is not secure. The purpose of a randomized ISN is to enhance security and prevent potential TCP sequence prediction attacks. However, if the ISN is defined at boot it means that the value will be defined only once every time the server boots up, thus sending the same ISN to every session, making it possible to obtain the future value from an existing session.

b. The server generates its ISNs by using the 32 least significant bits of a known secure hashing (i.e., SHA256) calculated over a constant value selected at random at boot time. *(2 points)*

Generating the ISN using the least 32 significant bits of a known secure hashing is not secure. This is mainly because of two reasons. The first reason is like the one described in point a) where you define the ISN at boot time only once, therefore making it possible to obtain the future value of the ISN until the server reboots. The second reason why this is not secure is because of the trimming of the result of the secure hashing algorithm, this action makes it vulnerable to collision attacks; an attacker could be able to define a value x' such that $LSB32(SHA256(x')) = LSB32(SHA256(x))$, where x is the random number generated by the server at boot time, not to mention that the amount of possible values for this mechanism is 2^{32}, which makes it relatively easy to a computer to obtain the valid ISN generated by the server.

c. The server chooses a random ISN for each connection. *(2 points)*

Choosing a random ISN for each connection is the most secure ISN selection method only if the random number generator does not show a pattern on the results, such as the ones described in the Zalewski's analysis of ISN.

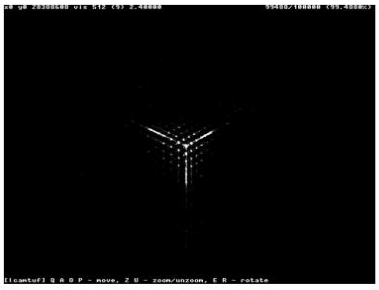

Test date: April 2001
OS: Windows 98 SE
Attack feasibility: 100%

Test date: April 2001
OS: MacOS 9
Attack feasibility: 89%

Test date: October 2002
OS: MacOS X
Attack feasibility: 0%

d. The server generates its ISNs by performing a XOR among the following four variables: source IP, source port, destination IP, destination port. *(2 points)*

Generating an ISN by performing a XOR among the source IP, source port, destination IP and destination port is not secure. The function to generate this ISN will be as follows:

$$ISN = sIP \oplus sp \oplus dIP \oplus dp$$

Where $sIP = Source\ IP; sp = Source\ Port; dIP = Destination\ IP;\ dp = Destinatino\ Port$

The vulnerability of this method is that relies on the scenario where the source port is equal to the destination port, such that:

$$ISN = sIP \oplus sp \oplus dIP \oplus dp$$

$$ISN = sIP \oplus dIP \oplus sp \oplus dp$$

$$ISN = sIP \oplus dIP \oplus 0$$

$$ISN = sIP \oplus dIP$$

Therefore, if we manage to clone the IP address of a user (or if we use the same proxy or network with the same NAT) we could scan the ports to obtain a valid ISN for an attack only matching the source port with the destination port.

e. The server generates its ISNs by using the 32 least significant bits of a known secure hashing (i.e., SHA256) calculated over the concatenation of the following four variables: source IP, source port, destination IP, destination port. *(2 points)*

The generation of the ISN using the 32 least significant bits of a known secure hashing calculated from the source IP, source port, destination IP and destination port is not secure. This is because of the same reasons described in point b), the 32 least significant bits value of the hashing function is vulnerable to collision attacks, and it makes it relatively easy to an attacker to find a valid ISN without knowing the IP and port of the victim.

f. The server generates its ISNs by performing a XOR among the following five variables: source IP, source port, destination IP, destination port, current time. *(2 points)*

Generating an ISN by performing an XOR between the source IP, source port, destination IP, destination port and current time is secure only if the time used is verbose enough. Assuming that the current time value is based on the connection of each client, if we use a timestamp that counts to milliseconds, it will be hard for an attacker to generate a valid ISN to attack a user, regardless of using the same value for the source and destination ports, as discussed in point d).

g. The server generates its ISNs by performing a XOR among client ISN and a random value selected at boot. *(2 points)*

Generating an ISN based on each client ISN and a random value selected at boot is not secure. First, any client could be able to generate a valid ISN for each connection and complete it with the random number generated at the server. Moreover, as discussed in points a), if the ISN is defined at boot it means that the value will be defined only once every time the server boots up, thus sending the same ISN to every session, making it possible to obtain the future value from an existing session.

Problem 2: SYN Flooding *(18 points)*

SYN Cookies and Client Puzzles are two techniques developed for defending against SYN flooding attacks.

1. Describe the idea behind SYN cookies. *(4 points)*

The idea behind a SYN cookie is to store the information of the state after establishing a connection between the server and the client to avoid SYN flooding attacks. The SYN cookie is generated using a timestamp, a maximum segment size and a cipher from a secure hashing algorithm which takes as input the timestamp, maximum segment size, source IP, source port, destination IP and destination Port; once the SYN cookie is generated, it returns the value in the SYN-ACK to the client, which returns an ACK that the server validates, if correct, stores the state of the connection.

2. Describe the idea behind client puzzles. *(3 points).*

The idea behind client puzzles is to force the client to perform some computational work before continuing with the connection between a server and a client over a network, this technique is defined with the objective to slow down attackers that performs SYN flooding attacks. The puzzle to solve must be relatively hard to compute for the client but easy to verify from the server's side.

3. Flooding of a target can be performed in multiple ways: (a) packets can be sent from the IP address of the attacker, (b) the attacker can spoof the source of each packet, or (c) the attacker might leverage a large network of compromised machines (i.e., a botnet) to send packets. Explain in which of those three scenarios SYN cookies and client puzzles are an effective solution against flooding attacks. *(3 points)*

SYN cookies are effective against scenarios a) and b). this is because the ISN is generated based on both the source IP and the source port of the client, so every time an attacker wants to request a new connection, the server checks if its there a valid SYN cookie for the client, to minimize the workload of the server for each connection. For scenario c), since there are multiple machines (therefore, multiple IP addresses), SYN cookies are not so effective defending against this kind of attack, because it will generate one SYN cookie for each client that is trying to connect, thus increasing the workload of the server.

On the other hand, client puzzles are effective against scenarios and c) because each of the compromised machines must be able to solve a problem/puzzle, which will require a considerable amount of computational capacity, thus, experimenting delays of the connection, mitigating the risk of potential denials of service attacks.

4. Can UDP be affected by a similar attack? *(2 points)*

Although UDP cannot be affected by a SYN flooding, since it does not relies on acknowledgement of messages and sequence number, it can be affected by other types of denial of service attacks, such as the Distributed Denial of Service (DDoS) attack or UDP floodings by spoofing the IP address machines in a network.

5. What would be the problem if the attacker manages to guess the value of the SYN cookie or the solution to a client puzzle? *(2 points)* Can you think of any network scenario where this would happen? *(2 points)*

If a SYN cookie or the solution to a client puzzle is guessed by an attacker, then a connection could be established between the server and the attacker, thus compromising the availability, integrity or confidentiality of the information stored and transmitted by the server.

This could happen if in a botnet (trying as many values as possible for the SYN cookie or the client puzzle) one or more nodes are able to establish a connection with the server, hence providing them with a channel to perform the attack.

6. Are SYN cookies and client puzzles useful also in presence of *link flooding*? *(2 points)*

They are not useful for link flooding arracks, since this attack relies on cutting off the connection by flooding the network links, not the server themselves.

Problem 3: Scanning *(10 points)*

A security researcher has identified the IP address of a command-and-control (C&C) server and wants to gather information about the botnet controlled through this server. To this end, the researcher performs a scan of the services running on the server. After sending several probes, the researcher observes a "RST" packet being returned from port 194.

1. Which protocol typically runs on port 194? *(2 points)*

Internet Relay Chat (IRC) Protocol, which allows communication in the form of text between multiple machines.

2. Which transport layer protocol did the research use in his probe packet to trigger this response? *(2 points)*

TCP, since the RST is a flag used on this protocol.

3. What can the researcher infer about the port status? Is the port "open", "closed" or "filtered" (i.e. , is the server is behind a firewall)? *(6 points)*

It can be inferred that port 194 is closed.

Problem 4: Denial-of-Service Amplification *(18 points)*

1. Describe one *network layer* amplification attack that we discussed in the class. *(4 points)*

Reflector amplification attacks work by sending traffic to a reflector using the IP address of a victim as the source IP. A reflector is a machine used by the attacker to reflect the traffic to a victim.

The attacker must select requests with relatively small size that can produce a bigger response. This is usually measured with the BAF (Bandwidth Amplification Factor) and the PAF (Packet Amplification Factor).

A Smurf Attack is one type of reflector attack, imagine an attacker who has access to multiple compromised machines, the attacker could send multiple ICMP requests to a victim (example, ping the server and request to return verbose information of itself) to return information massively to every node.

2. Describe two reflector attacks with amplification that do not use the DNS or the NTP protocol. Detail (*i*) all the involved parties, (*ii*) the type of requests and responses sent by each party, and (*iii*) the packet amplification and bandwidth amplification factors achieved by the attacker. *(14 points)*
(*Hint*: There are several attacks in the NDSS'14 paper by Rossow, but I want you to provide the details of the requests involved.)

a. NetBios: the attack was performed by using a name lookup, for which a receiving Windows system responds with its current network and host name configuration. The average BAF was 3.8 and the average PAF was 1.0.
b. Quake 3: Quake 3 is a video game that connects to servers around the world, each player creates a session to the server before playing the game. The attack to these servers was performed by asking a server for its status, a 15-byte-wide request. The reply was significantly larger because it included the detailed server configuration and a list of current players. The average BAF was 63.9 and the PAF was 1.01.

Problem 5: HTTPS *(12 points)*

A security researcher executes a piece of malware inside a Virtual Machine (VM) to study its network communication. The malware uses HTTPS to communicate with a remote command-and-control (C&C) server.

1. Since the C&C communication uses HTTPS, can the researcher recover the domain name and the IP address of the C&C server by monitoring the traffic that the malware generates? Justify your answer. *(4 points)*

Yes, it can recover the domain name and IP address. The domain name usually comes in the certificate of the C&C server (IP address can be inserted there, but it's not so common), for the IP address it just need to perform a ping to the server to get the address.

2. Explain how the researcher could set up a Man-in-The-Middle HTTPS proxy to decrypt and log the communication between the malware running in the VM and the C&C server. *(4 points)*

Considering the diagram discussed in class:

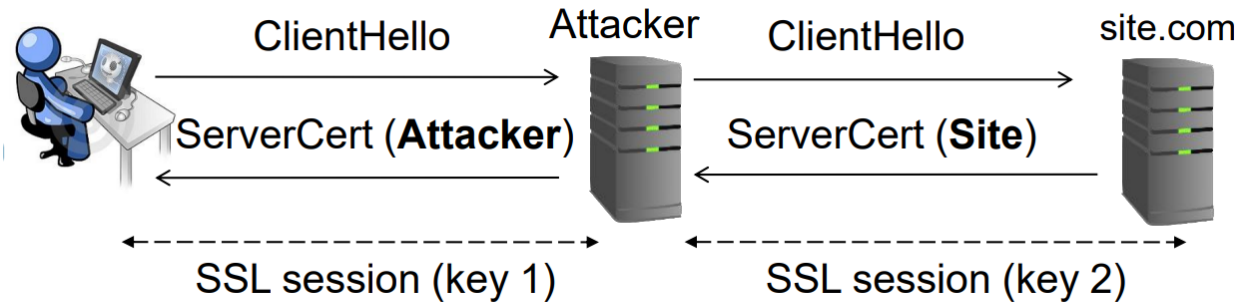

The researcher could act as a proxy for the VM with an SSL session (Key 1), while requesting the C&C server with a different SSL session (Key 2). Hence, the researcher will use the certificates and keys of both sessions to decrypt the traffic coming from the VM and the C&C server to log them in the researcher's machine.

3. Explain how the malware developer could protect against such MiTM interception. *(4 points)*

A developer could protect against this MiTM by validating the digital certificate of the server. This is because the researcher's certificate will be either self-signed, use a certificate installed locally or getting it from a rogue trusted CA.

Problem 6: Certificates *(28 points)*

Check the HTTPS certificate for `https://software.imdea.org` and answer the following questions. For each question, you need to provide a textual answer and a screenshot where the answer is visible.

1. How many certificates are there in the certificate chain? *(3 points)*

There are 3 certificates for ISRG Root X1, R3 and software.imdea.org.

2. What is the name of the Certification Authority that produced the leaf certificate? *(3 points)*

The Certification Authority that produced the leaf certificate (software.imdea.org) is R3

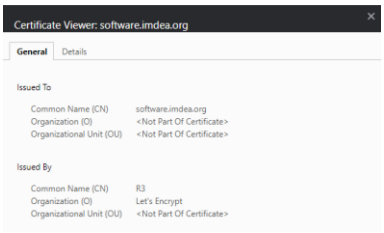

3. What is the organization name of the Certification Authority that produced the root certificate in the trusted store? *(3 points)*

It is ISRG Root X1

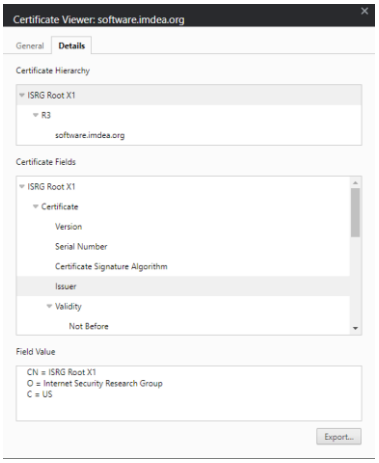

4. What is the validity period of the intermediate certificate? *(3 points)*

Issued on: 9/4/20, 2:00:00 AM GMT+2

Expires on: 9/15/25, 6:00:00 PM GMT+2

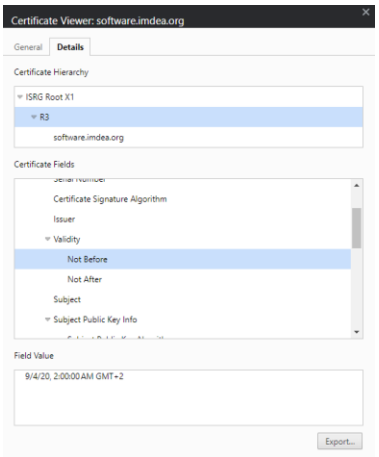

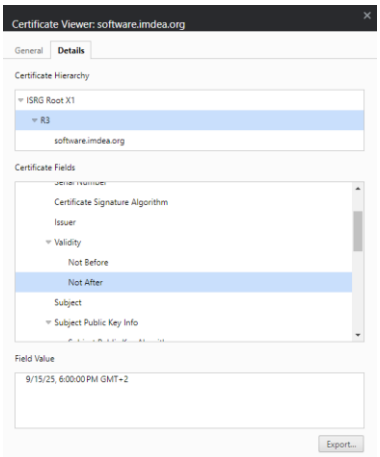

5. How many domains (*i.e., fully qualified domain names*) is the leaf certificate valid for? List the domains alphabetically *(4 points)*

- 2022.software.imdea.org
- hugo.software.imdea.org
- imdeasoftware.com
- imdeasoftware.es
- imdeasoftware.eu
- imdeasoftware.net
- imdeasoftware.org
- software.imdea.org
- www.imdeasoftware.com
- www.imdeasoftware.es
- www.imdeasoftware.eu
- www.imdeasoftware.net
- www.imdeasoftware.org
- www.software.imdea.org

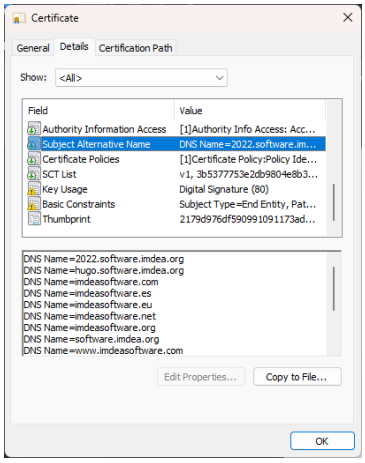

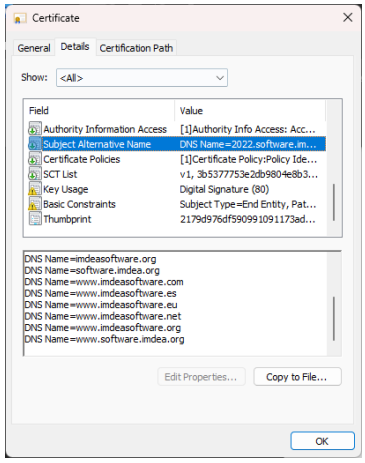

6. What hash algorithm is used by the intermediate certificate? *(3 points)*

SHA-256

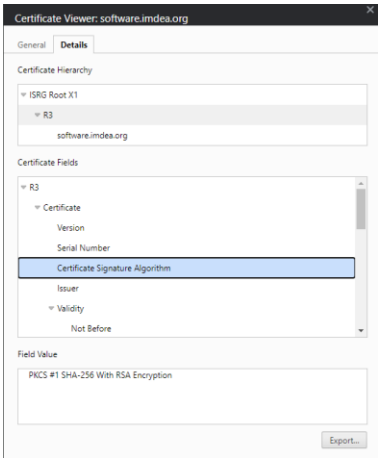

7. For the leaf certificate, what are the URLs of the OCSP server and the certificate revocation list (CRL)? *(3 points)*

The leaf certificate shows the following values for the Authority Information Access:

- OCSP Responder: URI: http://r3.o.lencr.org
- CA Issuers: URI: http://r3.i.lencr.org/

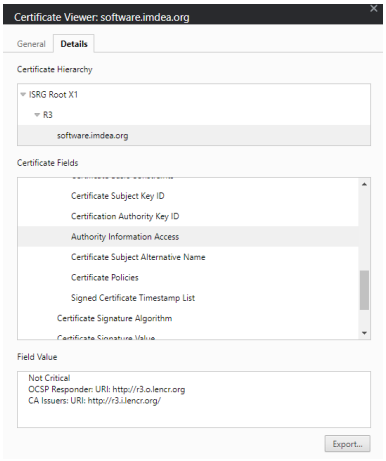

Neither of those URI gets the CRL file, however, the CA Issuer URI does contain the .DER file of R3, wich also has the CRL Distribution Point: URL=http://x1.c.lencr.org/.

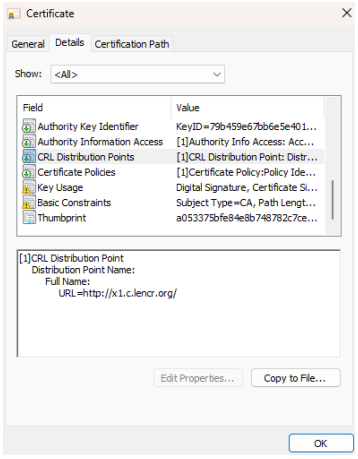

8. Download the certificate revocation list from the above URL. How many revoked certificates are there in the file? Note that how to view the contents of the CRL file varies depending on your OS and could require you to install a tool. Please provide the command used and the human-readable content of the CRL file in the answer. *(6 points)*

Windows already comes with a built-in GUI to visualize .CER, .DER and .CRL files.

The content of the downloaded CRL file is:

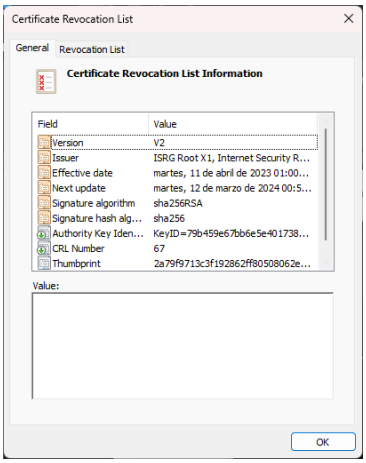

However, the revocation list appears empty.