Computer Security 1, COMPGA01, 2010-11 **ANSWERS PAPER, CONFIDENTIAL**

Answer ALL questions. 2.5 hours.

Answers should be short, straightforward and clear. 100 marks total.

Marks for each part of each question are indicated in square brackets

Calculators are NOT permitted

THIS IS THE ANSWERS PAPER, NOT TO BE PRINTED, CONFIDENTIAL

Question 1. Questions about network security.

1. A workstation with an IP address of 10.0.0.1 carries out a ping 10.0.0.2 command. Shown below (in order) is the network traffic on the LAN (Local Area Network). The commands between square brackets are written in an informal way making them easier to understand for the reader.

step 1 10.0.0.1 sends [ARP who-has? 10.0.0.2] to the LAN broadcast address.

step 2 10.0.0.2 replies [ARP is-at 0A:BC:00:00:00:1F] to 0A:BC:00:00:00:1A

step 3 10.0.0.1 sends ping packet to 10.0.0.2

Answer the following five sub-questions:

(a) What does the acronym MAC stand for? What is the main purpose of it? This is in the context of this exercise (not in cryptology and not in access control theory).

(b) How many bits does it have in wired networks (Ethernet)?

(c) Security-wise, is it better to have this MAC value implemented in hardware or as a software feature?

(d) What is the MAC destination address in the ping packet in the example above?

(e) What is the MAC address of the workstation initiating the ping request?

[5 marks]

Answer on the next page:

Answer:

(a) MAC stands for Media Access Control. It is a physical address which should be unique for each network card in one Local Area Network (LAN).

(b) As on examples seen here, it is 48 bits in Ethernet. Typically the first 24 bits are fixed for one manufacturer.

(c) It is better to have a MAC address fixed in hardware and nobody can change it. The manufacturer makes sure that these addresses never repeat. Fixed MAC addresses can be logged and allow greater trace-ability of communications.

(d) 0A:BC:00:00:00:1F

(e) 0A:BC:00:00:00:1A.

2. An attacker notices that a particular network request “Request A” from a client to a software application running on a server is always 30 bytes long (including UDP headers, IP headers and Ethernet headers). In response to this request, the server always provides the reply “Reply B”, which is 600 bytes long (also including UDP+IP+Ethernet headers). The attacker decides to spoof the IP address of the client and use this as the source IP address when sending many requests of type “Request A” to the server. We assume that all communication between the client and the server is based on UDP (User Datagram Protocol). If the attacker has 1 Mbits/s of outgoing bandwidth, assuming he can use all of it to mount an attack against a victim, and neglecting any additional overheads, how much of the victim’s bandwidth will be used by the attack?

[3 marks]

Answer: Because 600 is 20 times bigger than 30, ignoring any additional overheads, we expect about 20 times more, which is about 20 M bits per second.

3. What is a firewall? State at least 3 and at most 5 key fundamental principles of computer security which apply to firewalls and elsewhere. For each of these principles explain using one example, or in one way, how firewalls may reflect, play a certain role in, or help to implement, the application of this principle.

[4 marks]

Answer on the next page:

Answer: A firewall is a system which securely interconnects networks with different security requirements (other definitions may also be accepted). They are meant to make the propagation of attacks harder and to isolate networks from threats and attacks, while allowing legitimate well understood patterns in network activity. The key principles are:

(a) **Least Privilege** or Limitation Principle. Each user or module should only have minimal rights necessary to carry out their tasks. Firewalls separate different networks and greatly limit access to MAC and IP addresses and other ressources.

(b) **White-list Principle or Default Deny**, also quite similar: Fail-Safe Defaults. Firewalls, depending on how sophisticated they are, are here to rather disallow access just to any network address and network resource from anywhere, expect for those which are specifically authorized by the firewall. In firewalls it seems easier and more natural to write rules to describe what is allowed rather than try to disallow all possible attacks.

(c) **Defense in Depth** Principle. Just one layer of defense may fail because it was disabled, badly configured, badly understood, hacked, circumvented etc.. For example even if a firewall blocks all entering FTP connections (destination port =21) it is still important to deactivate all FTP servers which are maybe installed by default on certain machines.

(d) **The (False?) Principle of the Weakest Link**. The weakest link can fail and get around the security we put elsewhere. For example, it may be useless to have a very sophisticated firewall to inspect HTTP traffic for viruses and exploits, if it forgets also to inspect email traffic.

(e) **Simplicity** Principle. Firewall, as in the notion of Reference Monitor, needs to be simple enough in order to build correctly, documented and understood correctly, and configured correctly.

(f) **Acceptability** Principle. A firewall should be acceptable and accepted by users, system administrators and business people who pay for them. For this they need to strike a balance between security and functionality and other business benefits.

4. For each of the following three words: X = Software, Hardware or Stateful explain a) what X means as an adjective in a context of firewalls. b) What is one thing which is good and positive about it, mostly from the security perspective. c) What is one thing which is problematic or dangerous about it, again from the security perspective. Overall we expect 9 answers.

[4 marks]

Answer:

(a) **Software** a) implemented in software on a workstation by installing a firewall software b) easy to upgrade to a more secure version c) inherits all standard attacks on the workstation and its OS

(b) **Hardware** a) implemented in a hardware "metallic box" b) very hard to hack, certain features are impossible to alter c) hidden inside it could have a bug we cannot fix or be controlled by a rogue manufacturer

(c) **Stateful** a) aware of connections, packets may be rejected because are not part of an existing connection b) can inspect packets further and be very smart to distinguish between legitimate traffic and attacks c) system administrators will trust the firewall and forget USB sticks 3G phones and other channels, or be mystified by all the fancy possibilities and use the default setting well known by the attackers.

5. Explain the term DMZ in the context of network security.

[2 marks]

Answer: DMZ stands for a De-Militarised Zone. Typically it is a zone which is not inside the internal (secure and protected) network but a separate logical and physical zone. It hosts machines which will be communicating with the Internet and insecure outside networks: proxies, email public FTP and HTTP servers, etc. These machines are exposed to attacks, it is a sort of “electronic frontier” of an organisation.

[Total 18 marks]

Question 2. Questions about computer security, data security, cryptography, storage and management of confidential data, passwords and keys.

Prof. Sandwich has email client software installed on a local PC, which always connects to his company's mail servers situated in a DMZ through and over the classical SMTP protocol. Before Prof. Sandwich can use his PC, he enters his password, which is unknown and assumed to follow a probability distribution with high entropy.

1. List between 3 and 4 essential requirements or recommendations for the secure storage of this password on the local hard drive.

   [3 marks]

   Answer: The OS should use a file which is not easy to capture, it should not even be readable to an administrator. The password should be stored hashed with a one-way function. This with salting (preferred) or a mechanism based on unique user IDs. And also making sure that the file cannot be copied to another computer, for example hash also same unique data different for each computer.

2. Assume that the entropy of his password is as high as 80 bits. Is this enough? Assuming that passwords are NOT reused in multiple systems, what else do we need to look at? Given several different probability distributions with entropy being equal to 80 bits, can some of these distributions be much weaker and how?

   [3 marks]

   Answer: Some distributions can contain some very frequent passwords, some will not. This is measured by Min-entropy. It is defined as the logarithm in basis 2 of the probability of the most frequent passwords. This will matter in a scenario where many passwords are generated from the same distribution (multiple users following similar password policies).

3. Explain how dropping the privileges can be used to protect an email application against a rogue email containing an exploit.

   [3 marks]

   Answer: For example, when we start an email client, it may open certain files and then drop privileges to open these files. In Unix for example a program can permanently drop root privileges by using setuid(). Dropping the root privileges will make that the exploit can do less harm. A modern program would be made with several modules which have very small privileges each.

4. Is SMTP traffic encrypted and are emails sent by Prof. Sandwich encrypted and authenticated?

   [2 marks]

   Answer: SMTP is very insecure, SMTP traffic and commands are not encrypted, the user is sometimes authenticated with a password but emails are still sent in cleartext and are not authenticated.

5. Can a hacker somewhere on the wider Internet (outside Prof. Sandwich's company) use an ARP poisoning attack to intercept his SMTP traffic?

   [2 marks]

   Answer: No. The ARP poisoning attack can work inside a company's LAN but not from the outside.

6. If we choose some sort of a secure SMTP extension (could be based either on SSL, SSH, Kerberos or other protocol) what is a necessary **trusted setup** condition which should not be forgotten at the moment of installing or configuring the email system?

   [2 marks]

   Answer: Many answers will be accepted and many solutions exist, they all require some cryptography. He needs to have some **verification keys** allowing the SMTP server to be somewhat authenticated, either explicitly, or implicitly as part of TLS or authenticated Diffie-Hellman session, or as part of Kerberos-style authentication etc. Both symmetric and public-key solutions exist, in the latter the minimum trust requirement is to have one authentic public key. If he doesn't, anybody on the local network can pretend to be an SMTP server, capture his password and intercept outgoing email traffic, and maybe even insert spam.

Question 2 continued.

Prof. Sandwich sometimes uses PGP, GNUPG as a plug-in for his email software.

7. Explain what is the meaning of Hybrid Encryption.

[2 marks]

Answer: Hybrid Encryption is composed of a symmetric encryption system to encrypt long messages (this is called Data Encapsulation Module or Mechanism, DEM) and only uses (expensive) public key cryptography to convey a short and ephemeral secret key used for symmetric encryption in this session and generated at random. The latter mechanism is called Key Encapsulation Module, KEM. Both layers of encryption should be probabilistic. See slide 126 in part 6 of our course slides on Moodle.

8. Why, when sending an encrypted email message with PGP, do we want to add ourselves to the recipient list?

[2 marks]

Answer: When sending an encrypted email message we want to add ourselves to the recipient list, otherwise we are unable to read (decrypt) an old message stored in our own Outgoing mailbox.

Question 2 continued.

One day Prof. Sandwich lost his private key but he still has the corresponding public key. His hard drive crashed due to a mechanical fault, and he cannot recover all the data, but he still has the hard drive and it was neither stolen nor otherwise compromised. He decided not revoke this key because he is too busy and he is confident that it was never compromised.

9. Provide some advice about how to really destroy the data on a hard drive so that he can be confident the nobody will recover his private key. Provide one argument to the effect that destruction is in this case arguably better or more secure than revocation.

[2 marks]

Answer: First, data are still there even if with errors, they might be recovered by a specialist. If possible, data should be overwritten many times with random patterns. Then we can use an an alternating magnetic field and/or microwaves, to try to erase the data recorded on the magnetic medium. In addition we can use things such as a mechanical high-pressure press, very high temperatures and corrosive chemicals. Furthermore it is good to keep what remains rather than throw it away, which is an additional risk.

We can argument or consider that the destruction of the private key is better than revocation because not everybody checks the revocation lists.

10. On the next day Prof. Sandwich was able to restore his computer, his emails and his settings to work back again from his backups, but the private key was never a part of his backups. Can he still send encrypted emails? Can anyone decrypt any of the encrypted emails sent to him?

[2 marks]

Answer: He can still send encrypted emails to anyone for whom he can obtain a trusted or certified key. Neither him nor anyone else can decrypt any of the encrypted emails sent to him.

11. What does he need to do now, to be able to recover all the benefits of having a secure email system? What assumptions and maybe related hardware requirements would we like to have here or/and what could go wrong with a process of secure key **generation**? We omit problems related to secure key storage and trust/certificates.

[2 marks]

Answer: He needs to generate a new private PGP key, making sure he has some good sources of randomness, and that this key will not be leaked due to some malware on his computer.

12. How can this private key be securely stored on a hard drive (not using any special hardware or smart cards)? What security remains if his computer is now stolen or hard drive partition cloned? Can you propose one extra enhancement?

[3 marks]

Answer: The private key should be stored encrypted, to be able to decrypt it at any moment. Hashing would not work because a hash cannot be decrypted back to get the initial value. Typically a password-based encryption system is used, where the key remains secure if the computer is stolen, provided the password has very high entropy and was not captured or stored somewhere. Furthermore we can use TrueCrypt or similar whole drive encryption software, which will add a whole second layer of security, now two passwords are needed to get the private key.

[Total 28 marks]

Question 3. Explain in detail the following notions, definitions, technical terms or acronyms. Between 2 and 10 lines of explanation are expected for each sub-question.

1. Memory protection, segmentation and paging.

   [4 marks]

   Answer: The key principle of Memory Protection is that one process should not access memory of other processes. Even less so for Kernel memory. It is achieved by segmentation and paging and made possible by Intel CPUs starting from 386. Paging means that when a program accesses a page of memory, it is an interval of a virtual not real memory space. There is a one-to one mapping between these small intervals (pages, for example 4 K bytes) and pages in the actual RAM, some being swapped onto the hard drive, in a way transparent to the process. In additional to improve the memory protection, segments and increasingly so also individual pages can be marked as impossible to write, or impossible to execute etc.

2. A Trusted Path Mechanism.

   [2 marks]

   Answer: It is some sort of mechanism which provides some level of confidence that the user is communicating with the right entity, it is an attempt to obtain an "unspoofable and incorruptible" channel. what the user intended to communicate with, ensuring that attackers can't The I/O protection built in computers allows to implement the "trusted path" as follows, this is one example of a trusted path system. If we press Ctrl+Alt+Del, the computer through interruptions and the OS (Windows NT or later, etc), is designed in such a way that only the WinLogon process, a trusted system process, can receive notification of this keystroke combination.

3. DEP and NX bits. Why or how it became an 'NX bit' to be set and not an 'X bit' to be disabled?

   [2 marks]

   Answer: Since i386, W/R permissions exist at the page table entry level, 4 K pages typically, but only very recent computers are able to prevent the execution (NX bit, Never eXecute) from a given individual memory page, and this type of protection was for years implemented in form of for example software patches. Now it is a hardware mechanism and existed since ever for i64 CPUs, but then Pentium 4 Prescott and AMD adopted this for general-public computers around 2004. In order to preventing the execution with the NX bit, and for the compatibility reasons, by default X is enabled, and it must be explicitly disabled: the NX bit must be set to prevent execution from this page. This is called Data Execution Prevention (DEP) and works only if the OS supports it (using 64-bit page tables with PAE, roughly speaking Windows XP SP2 and later and Linux kernel since release 2.6.8.).

4. Explain step by step a typical code injection attack with buffer overflow and stack smashing. The student may draw a pictures with text heap and stack, and explain what happens at different stages.

   [10 marks]

   Answer: We want to inject and run arbitrary code through standard input channels of the program. We provide to the program binary data longer than the buffer allocated from the stack by the current C sub-routine. Pictures which illustrate this are found in lecture notes, slides part 10. If the C program uses strcpy, the copy will continue as long as there is no byte at 0. This will overwrite the return address on the stack with an address of the code injected into that stack, part of the same crafted long binary input. When the current C sub-routine exists, it will jump to the injected code written by the attacker.

5. Explain how the concept of $W \oplus X$ can be used to prevent these buffer overflow attacks.

[2 marks]

Answer: The executable part of the program virtual memory space (a.k.a. text) should be marked as X, not W, there is no need to ever write it except when the program is loaded to memory which is done by the OS. The all the data pages of this space (stack and heap) should be marked as W, not X. Never to be executed. Interestingly there is no need here for pages with both W and X active.

6. Explain how ASLR and Random Canaries can be used to prevent these buffer overflow attacks.

[4 marks]

Answer: ASLR means Address Space Layout Randomisation, it will randomize the order and the layout of elements on the stack, at the runtime. It is not a perfect solution, typically just makes attacks work with a smaller probability, Canaries, or Random Canaries, which for example can be implemented as a compiler extension, or as a manual care taken by the programmer creating additional variables himself, are used mostly to detect the buffer overflows and other intrusions by checking if the canary value at a given location in memory has not been modified by the attacker.

[Total 24 marks]

Question 4. Questions about Biba model.

1. Explain the term MAC in the context of access control theory (not in cryptography and not in networking).

   [2 marks]

   Answer: Mandatory Access Control (MAC): a system-wide policy with mandatory rules always obeyed, to restrict access (possibly for example denying users full control over the access to resources they created).

2. Given a lattice and a totally ordered set, write a definition of the Bell-LaPadula or Biba product lattice.

   [3 marks]

   Answer: Let H be a set of classifications with a total ordering $\leq_H$, and let Cat be a set of categories/classes, and let $C = P(Cat)$ be a set of compartments, being arbitrary subsets of Cat. The Bell LaPadula or Biba product lattice, or just a product lattice, is defined as a poset $H \times C, \leq$ where $\leq$ is defined by:

   $$(l_1, c_1) \leq (l_2, c_2) \Longleftrightarrow l_1 \leq_H l_2 \quad \wedge \quad c_1 \subseteq c_2$$

3. In Biba model, explain what is the semantics or our understanding of what it means for a file to be at a high integrity level in some lattice.

   [2 marks]

   Answer: The data is accurate, reliable and not contaminated (nothing about its secrecy).

4. Explain in terms of information flow, the mandatory requirements in Biba model. Give a correct interpretation of the two main rules regarding reading and writing.

[3 marks]

Answer: In Biba model the information flow goes down, and only through pairs related in the product lattice. No read-down really means can only read up if the subject's level is dominated by the level of the object in the product lattice. No write up really means can only write down if the object's level is dominated by the level of the subject writing it.

5. Consider the set of integrity levels $L = \{low, admin, kernel\}$, where $kernel > admin > low$. Furthermore consider the set of categories Cat containing

$$Cat = \{HRandAdmin(H), MarketTraders(T), ITEngineers(E)\}.$$

Compute the exact number of security classes in the product lattice.

[3 marks]

Answer: $3 \cdot 2^3 = 24$

6. Compute the bottom element $\perp$ in the product lattice.

[2 marks]

Answer: $\perp = (low, \{\})$

7. Name and compute the following bound in the product lattice:

$$(admin, \{E\}) \vee (kernel, \{H\}).$$

[2 marks]

Answer: LUB, Least Upper Bound, $= (kernel, \{H, E\})$

8. Name and compute the following bound in the product lattice:

$$(low, \{E\}) \wedge (kernel, \{T, E\}).$$

[2 marks]

Answer: GLB, Greatest Lower Bound, $= (low, \{E\})$

9. List all the security classes that a subject with classification $(admin, \{E\})$ can write.

[3 marks]

Answer: Can write at levels below himself and related in the product lattice $(admin, \{E\})$; $(admin, \{\})$; $(low, \{E\})$; $(low, \{\})$

10. A security software package with anti-virus functionality was installed by a person with clearance $(kernel, \{H, T, E\})$ and all its key components have level $(kernel, \{H, T, E\})$. Explain how the Low-Water-Mark Policy for Subjects could allow a sub-process spawned, potentially running at a different integrity level than the father process, to access any file in the system and check them for viruses.

[4 marks]

Answer: The security software and all its key components have level $(kernel, \{H, T, E\})$, however when it runs a virus detection sub-process, it could run at the level $\perp = (low, \{\})$ and access any file. Alternatively, with the Low-Water-Mark Policy for Subjects each time some random file is accessed, the scanning process/subject will be downgraded to the GLB of the two, at the end it will probably end up running also at $\perp = (low, \{\})$.

11. Explain how the Biba's optional Controlled Invocation (Ring Invocation) policy could potentially be used to prevent a trader who logs at level $(admin, \{T\})$ from injecting an exploit into his trading console program running as $(kernel, \{T\})$ which would allow him to override the limit of 5 millions dollars per day which is the maximum amount of money he is allowed to spend in one day on purchasing some trading options on an external automated market exchange system.

[4 marks]

Answer: Imagine that the trader runs a software console at level $(admin, \{T\})$. The Controlled Invocation (Ring Invocation) policy says he is allowed to use a trading console above his level at $(kernel, \{T\})$, which is indeed related and above $(admin, \{T\})$ in the lattice. Here it is expected that **IF one is allowed** to install this trading console at this high level of $(kernel, \{T\})$, it should be immune to exploits through careful design and mandatory built-in intrusion detection countermeasures.

[Total 30 marks]

[Total For The Whole Exam 100 marks]

THIS IS THE ANSWERS PAPER, NOT TO BE PRINTED, CONFIDENTIAL

END OF PAPER