

OS11: Security

Including parts of Chapter 11 and Section 9.6.3 of [Hai19]

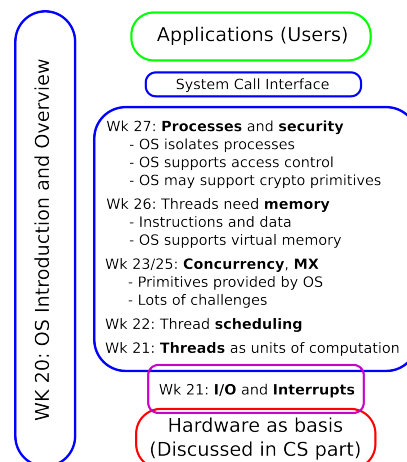
Jens Lechtenbörger

Computer Structures and Operating Systems 2019

1 Introduction

1.1 OS Plan

- OS Motivation (Wk 20)
- OS Introduction (Wk 20)
- Interrupts and I/O (Wk 21)
- Threads (Wk 21)
- Thread Scheduling (Wk 22)
- Mutual Exclusion (MX) (Wk 23)
- MX in Java (Wk 23)
- MX Challenges (Wk 25)
- Virtual Memory I (Wk 26)
- Virtual Memory II (Wk 26)
- Processes (Wk 27)
- Security (Wk 27)
- Wrap-up (Wk 28)



Contents

1	Introduction	1
2	Cryptography	5
3	Message Integrity	13
4	In-Class Meeting	16
5	Conclusions	20

1.2 Today's Core Questions

- How can I ensure that my downloaded software has not been manipulated?
- What is e-mail self-defense?

1.3 Learning Objectives

- Explain confidentiality and integrity as security goals
 - Discuss differences between end-to-end and hop-by-hop goals
- Explain use of hash values and digital signatures for integrity protection and discuss their differences

1.4 Retrieval Practice

- Security — So far
 - Hardware building blocks
 - * Kernel mode vs user mode: Restrict instruction set
 - Protect kernel data structures
 - Enable access control via system call API
 - * Timer interrupts
 - Transfer control periodically back to OS
 - Process as major OS abstraction
 - * Virtual address spaces
 - Isolate processes from each other
 - * Access rights

1.5 Information Security

- **Safety**: Protection against unintended/natural/random events
 - (Not focus here; requires proper management, involves training, redundancy, and insurances)
- **Security**: Protection against deliberate attacks/threats
 - Protection of **security goals** for objects and services against **attackers**

1.5.1 Security Goals

- Classical security goals: **CIA triad**
 - **Confidentiality**
 - * Only intended recipient can access information
 - * Typically guaranteed by encryption mechanisms
 - (Or, e.g., with envelopes and protecting laws)
 - **Integrity**

- * Detection of unauthorized modification
- * Typically guaranteed by cryptographic checksumming mechanisms
 - (Or, e.g., with signatures and/or seals)
- **Availability**
 - * Information and functionality available when requested
 - * Supported by redundancy
- Further goals
 - * Accountability, authenticity, anonymity, (non-) deniability, ...

1.5.2 Relativity

- Security is **relative**
 - You need to **define your goals** and risks for **specific pieces** of information, e.g.:
 - * How much confidentiality for course slides vs course exam?
 - * Apparently, it's easy to keep the slides "secure"
 - Harder for the exam
 - Also: Who is the **attacker** with what **resources**?
 - * Appropriate security mechanism, typically with **risk acceptance**
- Security via **design process** and **management**
 - BSI (Germany) and ISO standards
 - * IT-Grundschutz
 - Topic in its own right

1.5.3 Attacker Models

- Sample classifications of attackers
 - Strategy
 - * Targeted (specialized, looks for "weakest link")
 - E.g., espionage, blackmailing
 - * Opportunistic (standardized, looks for "weakest target")
 - E.g., phishing, extortion, bot/zombie creation (DDoS, spam, bitcoin mining, proxy)
 - Financial resources
 - Compute capacity
 - Time
 - Knowledge (insider and position?)

1.6 Design Principles for Secure Systems

- Selected principles based on [SS75]
 - Fail-safe defaults (whitelisting): If no explicit permission, then deny
 - Least privilege (need to know): Subject has only those privileges that are necessary for given task
 - Economy of mechanism: Security mechanisms should be as simple as possible
 - Complete mediation: All accesses need to be checked
 - Open design: Security should not depend on secrecy; instead open reviewing
 - Separation of privilege: Permission not based on single condition
 - Psychological acceptability: Security should not hinder usage
 - And more, see [SS75] or [Hai19]

1.7 End-to-End Security

- Security goals may have varying **scope**
 - Hop-by-hop
 - End-to-end
- Integrity and confidentiality are **end-to-end goals**
 - Beware: That's not generally understood!
 - * (See next slide...)
 - Consider hop-by-hop confidentiality
 - * Alice wants to send confidential message M to Bob via one hop, Eve
 - Alice encrypts M for Eve, sends encrypted M to Eve
 - Eve decrypts M, encrypts M for Bob, sends encrypted M to Bob
 - * Security gain or loss? (Compared to what?)
 - Hop-by-hop integrity similarly

Suppose that you want to send some e-mail to a friend, where the e-mail's contents are a private matter. In this case, the security goal confidentiality needs to be protected. Quite likely, you want confidentiality as an end-to-end goal meaning that only the communication endpoints, namely you and your friend, can read the message, independently of the number of hops or intermediary machines (such as Internet backbone routers) that forward the message from you to your friend.

If you send the e-mail as usual, sender and recipient need a password to access their accounts and e-mails at their providers' servers. Thus, some protection is offered for e-mails at their destinations. However, obviously also the providers' administrators and everybody else with access to their infrastructures (such as intelligence agencies violating human rights and other criminals) have access to the e-mails. Thus, those parties can access your draft folder as well as the recipient's inbox to access messages, violating confidentiality.

Besides, in the case of e-mail it is not clear whether e-mails forwarded between providers are encrypted or not. In response to the Snowden revelations there is a major shift towards encryption in transit; however, this type of encryption is not guaranteed. Thus, your e-mail

might also traverse the Internet in plaintext, and on its way it typically passes a couple of computers owned by parties that are unknown to you and that might copy or change your e-mails. Actually, when e-mails cross country borders it's almost certain that intelligence agencies copy the messages, again violating confidentiality. Obviously, this type of confidentiality violation can be prevented if providers encrypt their message exchanges, which would guarantee confidentiality on a hop-by-hop basis.

Clearly, encryption on a hop-by-hop basis is better than no protection, while you need to take protection into your own hands if you are interested in end-to-end goals.

1.7.1 (Counter-) Example: De-Mail

- De-Mail is a German approach **defining** legally binding, “secure” e-mail

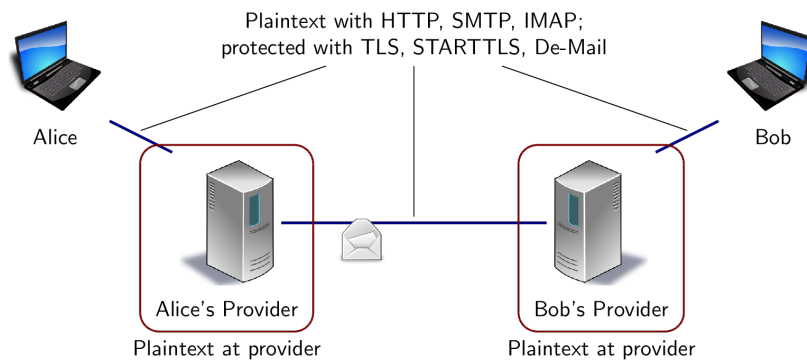


Figure 1: Hop-to-hop security of e-mail

- General picture
 - * **Strong** (hop-by-hop) security for each of the three blue links
 - * **Plaintext** at both providers (and **broken** approach towards integrity, see [Lec11])
 - End-to-end encryption allowed
 - Digital signatures used in special cases

De-Mail serves as example for hop-by-hop security and as counter-example for end-to-end security. Key characteristics are shown on this slide. While De-Mail may be attractive for legal reasons when it allows to replace paper with digital communication, I don't see much value for individuals.

The broken aspect of integrity protection mentioned here is that the technical specification for De-Mail includes a step “Metadaten setzen und Integrität sichern” which adds a simple hash value that is later checked in a step called “Integritätssicherung prüfen”. As part of a JITT assignment you should convince yourself that such a hash value provides no integrity protection against attackers.

2 Cryptography

2.1 Key Notions

- Cryptography = Art of “secret writing”
- Set of mathematical functions
 - Cryptographic **hash** functions

- Classes of **encryption** algorithms
 - * **Symmetric, secret key**: en- and decryption use the same shared secret key
 - * **Asymmetric, public key**: participants own **pairs** of secret (decryption, signature creation) and public (encryption, signature verification) keys
 - * **Hybrid**: asymmetric initialization, symmetric encryption
- Basis for various security mechanisms
- Performance
 - Hashing > Symmetric Enc. > Asymmetric Enc.
 - * (One can hash more data per second than one can encrypt)
 - * (One can encrypt more data per second symmetrically than asymmetrically)

2.1.1 Basic Assumptions

- Fundamental Tenet of Cryptography from [KPS02]
 - “If lots of smart people have failed to solve a problem, then it probably won’t be solved (soon).”
 - The problem to solve here: Break specific crypto algorithm
 - * If that did not happen for a long time, probably the algorithm is strong
 - * (Lots of crypto algorithms come without security proof)
- Kerckhoffs’ Principle (1883)
 - Security of crypto systems should not depend upon secrecy of en- and decryption functions (but on secrecy of the used keys)
 - “Open Design” principle from [SS75]
 - * Not respected in national security/military/intelligence settings in Germany
 - From Enigma through Libelle (approved for “Streng geheim”; developed by BSI, not published)
 - Opposite: Security through obscurity

2.1.2 Names

- Alice and Bob; Charlie, Carol, Dave, ...
 - Communicate frequently
 - Value their privacy
 - Have limited trust in third parties
 - Appeared to be subversive individuals in the past
 - * Growing understanding in general public
 - And, of course, politically correct names instead of “A” and “B”
- Eve, Mallory, Trudy
 - Eavesdropper, malicious attacker, intruder

2.1.3 Notation

- M, C: Message and ciphertext (encrypted messages)
- K: Key (random bits, maybe with certain structure)
- E, D: En- and decryption functions
- K_{AB} : Secret key shared between Alice and Bob
- K_{A-} : Alice's private key
- K_{A+} : Alice's public key
- $K(M)$: Message M encrypted with key K (if function E is clear from context)
- $[M]_K$: Message M signed with key K

2.2 GnuPG

- GNU Privacy Guard
 - Free software for (e-mail) encryption and digital signatures
 - Implementation of OpenPGP standard
 - * Secure e-mail based on hybrid cryptography
 - In addition, lots of cryptographic algorithms via command line
 - * `gpg --version ... gpg (GnuPG) 2.1.13 ...` Öff. Schlüssel: RSA, ELG, DSA, ECDH, ECDSA, EDDSA Verschlü.: IDEA, 3DES, CAST5, BLOWFISH, AES, AES192, AES256, TWOFISH, CAMELLIA128, CAMELLIA192, CAMELLIA256 Hash: SHA1, RIPEMD160, SHA256, SHA384, SHA512, SHA224
 - Start by creating key pair: `gpg --gen-key`

2.2.1 E-Mail Self-Defense

- My suggestion: Try out OpenPGP
 - Create key pair, upload public key to server, send/receive encrypted (possibly signed) e-mails
- More specifically, follow Email Self-Defense
 - GnuPG and Thunderbird with Enigmail plugin
 - Of course, other implementations exist
 - * The choice is yours
 - Note: That guide contains instructions concerning the **e-mail robot Edward**, which can reply to your encrypted (and signed) test e-mails

2.3 (Cryptographic) Hash Functions

- Hash function (or message digest): One way function H
 - Input: Message M (bit string of arbitrary length)
 - Output: Hash value $H(M)$ (bit string of **fixed** length)
 - Collision: Different messages mapped to same hash value
- Cryptographic hash value \approx digital fingerprint
 - Collision **resistant** (different hash values for different messages)
 - Weak collision resistance of hash function H
 - * Given message M it is computationally infeasible to generate M' such that $H(M) = H(M')$
 - (Computationally infeasible means that attackers should not be able to create collisions due to resource limitations)
 - Strong collision resistance of hash function H
 - * Computationally infeasible to generate M and M' such that $H(M) = H(M')$

I suppose that you remember hash functions for fast searching. Recall that hash collisions are to be expected.

With cryptographic hash functions, collisions are a Bad Thing since hash values are supposed to serve as digital fingerprints. Ideally, each message (or document or piece of data or code) should have its own, unique fingerprint. When a message is changed, also its fingerprint should change. However, if a hash collision occurs and two messages produce the same cryptographic hash value, the fingerprint becomes unusable to distinguish them.

On the slide you see two versions of collision resistance. Please take a moment to convince yourself that the strong version implies the weak version.

2.3.1 On Collision Resistance

- Later: Hash values are essence of digital signatures
 - Consider contract between Alice and Mallory
 - * “Mallory buys Alice’s used car for 20,000€”
 - Contract’s text is message M
 - Digital signatures of Alice and Mallory created from $H(M)$
 - Suppose H not weakly collision resistant
 - * Mallory may be able to create M' with price of 1€ such that $H(M) = H(M')$
 - * As $H(M) = H(M')$ there is no proof who signed what contract
- Birthdays, collisions, and probability
 - Hash people to their birthdays (day and month, without year)
 - (a) Weak collision resistance: Anyone sharing *your* birthday?
 - (b) Strong collision resistance: Any pair sharing *any* birthday?
 - * Birthday paradox

The importance of weak collision resistance is best understood in the context of digital signatures, which are used to create legally binding digital contracts proving who signed what. Without going into details of digital signatures right now, it is sufficient to know that the contract's text is a message M and that digital signatures on M are created from the cryptographic hash value $H(M)$.

Suppose Alice and Mallory agree that Mallory buys Alice's used car for 20,000€. Both digitally sign the contract's message M . However, Mallory changes his mind and does not want to buy the car any longer.

If hash function H is not weakly collision resistant, Mallory may be able to create a second contract M' which includes the price of 1€ for Alice's car such that $H(M) = H(M')$. In this situation, as digital signatures are derived from hash values, the digital signatures of Alice and Mallory created for M are also valid for M' . Thus, Alice has no proof that Mallory signed M in the first place.

So: If a message M is given, nobody should be able to create a second message M' with the same hash value under weak collision resistance.

For strong collision resistance, nobody should be able to create any collision at all, even if those collisions only occur for messages that look like gibberish without practical value.

A different angle on collision resistance is provided by the following birthday analogy. Consider the hash function mapping each person to his or her month and day of birth. Essentially, there are 366 different hash values (including February 29), and a collision occurs when two people share the same birthday.

Suppose you are in class. When you wonder whether some of your fellow students shares *your* birthday, you consider weak collision resistance. In contrast, when you ask whether *any pair* of students shares the same birthday, you consider strong collision resistance.

For simplicity, ignore leap years and consider just 365 different birthdays, all with the same probability. I'm confident that for a class of 30 students you can compute the probabilities of (a) somebody sharing your birthday as well as (b) any pair sharing a common birthday. If you do the math for the first time, you may be surprised by the high probability in case (b), which is known as the birthday paradox (whose essence is the fact that the number of pairs grows quadratically, about which you can read more at Wikipedia). As the probability of case (b) is larger than that of case (a), it is harder to defend against case (b). Thus, hash functions targeting strong collision resistance must be "stronger" than those offering weak collision resistance.

2.3.2 Hash Applications

- Avoidance of plain text passwords
- Integrity tests
- Digital signatures

2.3.3 Hash Standards

- MD4, MD5, SHA-1: Broken
- SHA-1, SHA-2: Designed by NSA
 - Bruce Schneier, 2004: "Algorithms from the NSA are considered a sort of alien technology: They come from a superior race with no explanations"
 - Cryptographic hashing is extremely difficult, quote from 2006
 - * "Joux says that we do not understand what we are doing and that we do not really know what we want; there is agreement from all the panelists."
 - 2017: SHA-1 **SHattered** (deprecated by NIST in 2011)
 - * <https://shattered.io/>

- SHA-3 (Keccak)
 - Winner of public competition from 2007 to 2012
 - Standard released in 2015

2.3.4 Sample Message and Fingerprints

Hi Bob,

let's get started tomorrow!

Best wishes

Alice

- Sample hash values with GnuPG
 - `gpg --print-md SHA1 alice.txt`
 - * `alice.txt: 6FC1 F66C 598B D776 BA37 1A5C 2605 06CB 4CF9 0B89`
 - `gpg --print-md SHA256 alice.txt`
 - * `alice.txt: 84E500CB 388EE799 05F50557 43C5481B 08B0BF17 1A2AE843 F4A197AD 2BA68D2E`
- (Besides, specialized hashing tools exist, e.g., `sha256sum`)

2.4 JiTT Assignment

Submit a solution for the following task in [Learnweb](#).

2.4.1 Ticket to Exam

Discuss the correctness of the following statements.

- If messages of unbounded size are hashed cryptographically, an infinite amount of hash collisions is guaranteed.
- If a hash function is weakly collision resistant, it is computationally infeasible to compute hash collisions.
- If a hash function is strongly collision resistant, it is computationally infeasible to compute hash collisions.
- If I download a piece of software along with its hash value produced by a weakly collision resistant function and if the downloaded software has that precise hash value, I can be pretty sure that I obtained the “correct” software (without accidental or malicious changes).

2.5 Symmetric Encryption

- Sender and recipient share secret key, K_{AB}
- **Encryption** of **plaintext** message M with K_{AB} into **ciphertext** C
 - $C = E(K_{AB}, M) = K_{AB}(M)$
 - * Bits of M and K_{AB} are mixed and shuffled using reversible functions (e.g., XOR, bit shift)
 - * Simplest, yet provably secure case: **One-time pad** with XOR of **random** bit string and M
- **Decryption** with same key K_{AB}
 - $M = D(K_{AB}, E(K_{AB}, M))$
 - Notice: Need to exchange secret key ahead of time
- Typical symmetric algorithms: **AES**, **3DES**

2.6 Intuition of Asymmetric Encryption

- Participants own **key pairs**
 - Private key, e.g., K_{B-} : **secret**
 - Public key, e.g., K_{B+} : **public** / **published**
 - En- and decryption based on “hard” mathematical problems
- Think of key pair as **safe/vault** with numeric **key pad**
 - Open safe = public key
 - * Everybody can deposit messages and lock the safe
 - Opening combination = private key
 - * Only the owner can open the safe and retrieve messages

While symmetric encryption with shared keys, in particular the one-time pad, may seem intuitively clear, asymmetric cryptography requires some thought. Every participant needs a key pair, which consists of a private key and a public key. As the names suggest, a private key needs to be kept secret and must only be accessible by its owner, whereas the public key can be published, e.g., on web servers or special key servers.

This slide offers an analogy of public key cryptography with physical safes, which might help to convey essential ideas: The public key of Alice is used by others to encrypt messages to her, while she uses her private key to decrypt them. Similarly, she might offer opened safes in the real world, into which messages can be placed and which can be locked by everyone. Only Alice is able to open the safe using its opening combination to retrieve and read contained messages. Thus, the opening combination corresponds to her private key.

In class, you will experience another analogy, where participants and their key pairs are identified by colors. To that end, please bring along pens with different colors, e.g., text markers.

A noteworthy challenge of asymmetric cryptography, which is mentioned on the next slide, is the reliable distribution of public keys: How does Bob know that he really obtained Alice's public key and not one created by Mallory and distributed in her name? Or in the above analogy: How does he make sure that he places his messages into Alice's safe and not into one owned by Mallory to which Mallory attached the name tag “Alice”? Answers to this question are provided under the term “public key infrastructure”, and they frequently rely on the idea that Bob needs to verify a fingerprint of Alice's public key through an out-of-band communication channel. This highly relevant, fascinating, and challenging topic is beyond the scope of this presentation, though.

2.7 Asymmetric Encryption

- Participants own **key pairs**
 - Private key, e.g., K_{B-} : **secret**
 - Public key, e.g., K_{B+} : **public / published**
- **Encryption** of message for Bob with Bob's **public key**
 - $C = E(K_{B+}, M) = K_{B+}(M)$
 - Notice: No secret key exchange necessary
- **Decryption** with Bob's **secret key**
 - $D(K_{B-}, K_{B+}(M)) = K_{B-}(C) = M$
 - Notice: Only Bob can do this
- Challenge: Reliable distribution of public keys
 - Solution: Certificates in **Public Key Infrastructure, PKI**

2.7.1 Sample Asymmetric Algorithms

- **Diffie-Hellman Key Exchange** (1976)
 - Used, e.g., in IPsec, SSL/TLS, Tor, OTR
 - * **RFC 7568**, June 2015: SSLv3 **MUST NOT** be used
- **RSA** (Rivest, Shamir, Adleman 1978; Turing award 2002)
 - Most famous, PGP, GnuPG
- **ElGamal** (1984)
 - Based on Diffie-Hellman
 - GnuPG and newer PGP variants
- **Elliptic curves**
 - Newest class, shorter keys
 - GnuPG

2.7.2 GnuPG: Hybrid Encryption

- Create asymmetric key pair
 - `gpg --gen-key`
 - Various options/alternatives
- Encryption for Bob
 - `gpg -e -a -r Bob file`
 - * Creates `file.asc`; more precisely:
 - * Creates random secret key K_{AB}

- * Symmetric encryption of file with K_{AB}
 - Specific algorithm obtained from Bob's public key
- * Asymmetric encryption of K_{AB} with K_{B+}
 - Beware! No naïve encryption, but, e.g., **PKCS #1**
- * Result: $K_{B+}(K_{AB}) + K_{AB}(\text{file})$
 - (“+” between ciphertexts denotes string concatenation)

3 Message Integrity

3.1 Situation and Goal

- Alice sends message M to Bob
 - (Parts of) Network controlled by unknown parties (Eve and Mallory)
- Goals of integrity
 - Bob is sure that **M came from Alice**
 - * Notice: Need authentication!
 - Bob can **detect modifications to M**
- Non-goals: Alice cannot be sure
 - that no third party receives M
 - that Bob receives M
 - that Bob receives M in unchanged form

3.2 General Idea

- Alice sends message along with its fingerprint
 - Recall: A hash value is not good enough
 - Instead: Use some ingredient that is **unknown to the attacker**
- Bob receives message and fingerprint and verifies whether both match
 - If message changed by Mallory, he cannot produce a matching fingerprint
- Typical techniques
 - Message authentication codes
 - * E.g., Alice and Bob share secret K_{AB} , concatenate that to message before hashing
 - Digital signatures (next slides)

3.3 Digital Signatures

- Based on asymmetric cryptography
 - En- and decryption **reversed**
- Basic idea
 - **Signature** created by encryption with **private** key: $K_{A^-}(M)$
 - * Only Alice can create this!
 - **Verification** via decryption with **public** key: $D(K_{A^+}, K_{A^-}(M))$
 - * Everyone can do this as public key is public!
- Practice: Encrypt hash value of M, e.g., $K_{A^-}(\text{SHA-3}(M))$
 - Recall
 - * Performance
 - * Hash collisions

3.3.1 Some Details of Digital Signatures (1/2)

- **Signing** of M by Alice with private key K_{A^-}
 - Signature $S = K_{A^-}(h(M))$
 - * Only Alice can do this
 - Transmit signed message $[M]_{K_{A^-}} = M + S = \text{message} + \text{signature}$
 - * (“+” is concatenation)




Figure 2: Creation of digital signature

3.3.2 Some Details of Digital Signatures (2/2)

- **[M]** received by Bob
- **Verification** whether [M] sent by Alice and unchanged along the way




Figure 3: Verification of digital signature

- Split [M]: $[M] = M' + S'$
- Hash M' : $H = h(M')$
- Decrypt S' : $H' = K_{A+}(S')$
 - * Bob needs public key of Alice to do this
 - * Everyone can do this
- Verify $H = H'$

3.3.3 GnuPG: Digital Signatures

- `gpg --sign -a -b alice.txt`
 - Creates `alice.txt.asc`
- `gpg --verify alice.txt.asc`
 - Expects to be verified content as `alice.txt`
 - Verifies signature
 - Frequently used to verify integrity of downloads

3.4 Electronic Signatures

- “Signatures” of varying legal impact in IT environments
 - Different types, e.g., simple (e.g., sender’s name in e-mail), advanced (digital signature as discussed above), qualified
 - Qualified electronic signatures may replace paper based signatures (e.g., dismissal, invoice)
 - * Subset of advanced electronic signatures
 - * Based on qualified certificates (with qualified electronic signature, issued by accredited organization; law prescribes rules concerning infrastructure and processes)
 - * Created on secure signature-creation devices (nPA may store qualified certificate; additional reader necessary)

3.5 JiTT Assignments

Answer the following questions in Learnweb.

3.5.1 Asymmetric Cryptography

1. Mark the correct statements.
 - If the hash function used to create a digital signature produces a hash collision for messages M and M' , recipients will not know whether M was signed or M' or both.
 - Alice encrypts messages to Bob with her public key.
 - Alice encrypts messages to Bob with her private key.
 - Alice encrypts messages to Bob with Bob's public key.
 - Alice encrypts messages to Bob with Bob's private key.
 - Alice needs her public key to sign messages.
 - Alice needs her private key to sign messages.
 - Alice needs Bob's public key to sign messages addressed to him.
 - Alice needs Bob's private key to sign messages addressed to him.
 - Bob needs Alice's public key to verify her signatures.
 - Bob needs Alice's private key to verify her signatures.
 - I generated a key pair with GnuPG and sent an encrypted e-mail.

3.5.2 Preparation and Feedback

2. To prepare the in-class meeting, please bring along some colored pens, e.g., text markers.
3. What did you find difficult or confusing about the **contents** of the presentation? Please be as specific as possible. For example, you could describe your current understanding (which might allow us to identify misunderstandings), ask questions that allow us to help you, or suggest improvements (maybe on [GitLab](#)). You may submit individual questions as response to this task or ask questions in our Riot room and the [Learnweb forum](#). Most questions turn out to be of general interest; please do not hesitate to ask and answer in forum and Riot room. If you created additional original content that might help others (e.g., a new exercise, an experiment, explanations concerning relationships with different courses, ...), please share.

4 In-Class Meeting

4.1 Asymmetric Cryptography with Colors

- Based on [Security Protocol Game](#) by Len Hamey
- [Instructions](#)
 - Form groups of 3
 - * Every group member with a different color
 - Mark message with your color to sign
 - Place message into colored envelope to encrypt

4.2 (In-) Security News

- July 2018, installation of rogue CA certificates by malware
 - (In German)

4.3 Basic OS Security Services

4.3.1 Service Overview (1/2)

- Rights management, **authorization**
 - Discussed already: [Access rights](#)
 - * What is Bob allowed to do?
- **Logging**
 - Who did what when?
 - (Not considered in following)
- Basic **cryptographic** services
 - Offering selection of above techniques: a/symmetric techniques, hashing

4.3.2 Service Overview (2/2)

- **Identification/Authentication**
 - Identification: Claim of identity
 - * I'm Bob ...
 - Authentication: Proof of identity
 - * My password is "p@ssw0rd"
 - (Bad idea, easily broken!)
- **Integrity protection**

4.3.3 Authentication

- Proof of identity
 - Something the individual knows
 - * Password, PIN, answers to security questions
 - Something the individual possesses
 - * Smartcard, iTAN
 - Something the individual is
 - * Static biometrics, e.g., fingerprint, iris scan
 - Something the individual does
 - * Dynamic biometrics, e.g., voice or typing pattern
- Necessary prerequisite to enforce access rights
 - Who is allowed to perform what operation on what resource?

4.3.4 Two-Factor Authentication

- Combinations of above categories
 - Physical banking
 - * Bank card (possession) plus PIN (knowledge)
 - Online banking
 - * Password for login (knowledge) plus mTAN or iTAN (possession)
 - Beware: Must keep factors separate
 - * Do not record PIN on card
 - * Do not perform online banking on device that receives mTAN

4.4 Passwords

4.4.1 Authentication with Passwords

- **Randomness** of passwords
 - Measured via entropy
 - Ca. 4 bits per character in natural language
 - * (Contrast this with 8 bits, which are necessary for typical representation)
 - Thus, 32 bits per 8 character password
 - * Instead of 64 bits if characters were randomly distributed
 - * 2^{32} is about a factor of 10^9 smaller than 2^{64}
- Consequently, security of memorable passwords is **limited**

4.4.2 Password Advice

- **Write down** strong and diverse passwords
 - Strong: Use entire keyboard alphabet, no real words, > 10 characters
 - Where to store?
 - * Martin Hellman: “Write down your password. Your wallet is a lot more secure than your computer.”
 - * Attackers?

4.4.3 Password Checking

- **Never** store passwords in **plain text**
- Typical setup
 - System stores (salted – next slide) **hash value** of password
 - * GNU/Linux: `/etc/shadow`
 - * Windows: Security Accounts Manager (SAM) or Active Directory (AD)
- Verification

- System computes hash value for entered password string
- Compare computed to stored value
- (Alternative, e.g., Kerberos)
 - Use hash value of password as symmetric encryption key

4.4.4 Sample Password Attacks

- Social engineering
- Guessing
- Observe or record
 - Shoulder surfing, (web)cam, keylogger, network sniffing, extract from RAM, deduce from cache, etc.
 - [Friend or Foe?: Your Wearable Devices Reveal Your Personal PIN](#)

4.4.5 Brute Force and Salting

- Brute force = Try “all” combinations
 - Various heuristics, implemented in tools
 - * Try common patterns first
 - Databases of stolen passwords (e.g., “123456”)
 - Variations of dictionary words (e.g., “p@ssw0rd” above)
 - * [Rainbow tables](#): precomputed hash values
 - * Counter measure: **Salting**
 - Concatenate password with random number (=salt), then compute hash
 - Store random number and hash

4.5 Key Security Best Practices

- Consult others
- Adopt a holistic risk-management perspective
- Deploy firewalls and make sure they are correctly configured
- Deploy anti-virus software
- Keep all your software up to date
- Deploy an IDS
- Assume all network communications are vulnerable
- ... (see Sec. 11.8 in [Hai19])

5 Conclusions

5.1 Summary

- Security is complex, requires design and management
- Cryptography provides foundation for lots of security mechanisms
 - Don't implement cryptographic protocols yourselves!
 - Use proven tools, e.g., GnuPG
- Asymmetric crypto with key pairs
 - Public key for encryption and signature verification
 - Private key for decryption and signature creation
- Hash functions and digital signatures for integrity

Bibliography

- [Hai19] Max Hailperin. *Operating Systems and Middleware – Supporting Controlled Interaction*. revised edition 1.3.1, 2019. URL: <https://gustavus.edu/mcs/max/os-book/>.
- [KPS02] Charlie Kaufman, Radia Perlman, and Mike Speciner. *Network Security: Private Communication in a Public World, Second Edition*. Second. Upper Saddle River, NJ, USA: Prentice Hall Press, 2002. ISBN: 9780137155880.
- [Lec11] Jens Lechtenbörger. “Zur Sicherheit von De-Mail”. In: *Datenschutz und Datensicherheit* 35.4 (2011), pp. 268–269.
- [SS75] Jerome H. Saltzer and Michael D. Schroeder. “The protection of information in computer systems”. In: *Proceedings of the IEEE* 63.9 (1975), pp. 1278–1308. DOI: 10.1109/PROC.1975.9939. URL: <http://web.mit.edu/Saltzer/www/publications/protection/>.

License Information

This document is part of an Open Educational Resource (OER) course on Operating Systems. Source code and source files are available on GitLab under free licenses.

Except where otherwise noted, this work, “OS11: Security”, is © 2017, 2018, 2019 by Jens Lechtenbörger, published under the Creative Commons license CC BY-SA 4.0.

No warranties are given. The license may not give you all of the permissions necessary for your intended use.

In particular, trademark rights are *not* licensed under this license. Thus, rights concerning third party logos (e.g., on the title slide) and other (trade-) marks (e.g., “Creative Commons” itself) remain with their respective holders.