

2 Basic Concepts and Terminology

Revision 4.1, August 2025.

2.1 Objectives

- Establish a common vocabulary for security since “language shapes thought”.
- Understand the pillars of security: Confidentiality, Integrity, and Availability (C-I-A).
- Understand the kind of threats that your system faces.
- Understand how attacks to software security can come through vulnerabilities in the code.

2.2 An unsolved problem

As you begin studying software security, you can rest assured that there is no danger of the problem being completely eliminated or your expertise made moot. Security vulnerabilities are, if anything, on the increase. And years after some of the earliest most fundamental flaws were first discovered, they continue to recur and cause serious harm with no end in sight.

The number of incidents registered in the National Vulnerability Database¹ (NVD) maintained by the US National Institute of Standards and Technology (NIST) provides a very rough measure. The figure below clearly demonstrates that security problems are ongoing with no evidence of becoming a solved problem anytime soon. Common Vulnerabilities and Exposures (CVE) numbers are assigned to issues reported to the NVD (for example, the well-known Heartbleed² bug in OpenSSL that compromised the security of computers and related equipment worldwide is identified as CVE-2014-0160³). There is no well-established statistic that captures the overall trend in software security, and CVE counts are not entirely accurate as such; nonetheless, the NVD may be the best publicly available data.

¹ <https://nvd.nist.gov>

² <http://heartbleed.com/>

³ <https://nvd.nist.gov/vuln/detail/CVE-2014-0160>




Figure 1: CVE (Common Vulnerability and Exposures) counts from 2009 – 2022

2.3 The C-I-A Pillars

The three pillars of information security are Confidentiality, Integrity, and Availability (C-I-A), expressing the ideals we seek to protect and that attackers seek to undermine. These are the fundamental expectations for any information system: good security means that data will not be disclosed inappropriately (confidentiality), data and code will not be tampered or destroyed (integrity), and proper access will be promptly given at all times (availability).

These principles are so basic that it is worth defining and giving examples of each of the three pillars.

2.3.1 Confidentiality

Confidentiality refers to protecting against the disclosure of information to unauthorized parties. Potentially, even slight information leakages that provide any clues about the protected data represent a compromise of this principle. For example, revealing just the first letter or the number of letters in a name, combined with other information or guesses represents a partial disclosure.

Examples of attacks on confidentiality include:

- Reading stored data without authorization.
- Performing statistical analysis on restricted data.
- Reading messages (communication) without authorization.

- Detecting if communication traffic is present or not.

2.3.2 Integrity

Integrity refers to protecting information from being created, modified, or destroyed by unauthorized parties. Potentially, even the plausible threat that data may have been tampered with (and no means of knowing whether it was or not) could be considered an exploit of Integrity.

Examples of attacks on integrity include:

- Modifying, destroying, or creating new data files.
- Changing program code with the goal of causing errors, failures, crashes, or unexpected behavior.
- Modifying, deleting, replicating, change of order, or creating extra messages.

2.3.3 Availability

Availability refers to ensuring that authorized parties are able to access the system and its resources when needed. Successful compromise of availability is often called a Denial-of-Service (DoS) attack. While destruction of a system or its data certainly does make it unavailable, this is normally considered loss of integrity; the term “availability” is generally concerned specifically with temporary impediments to access.

Examples of attacks on availability include:

- Swamping an internet-connected service with requests, impeding its ability to serve clients.
- Removing applications, components, or files, rendering the system completely or partially inoperable.
- Crashing a server, preventing it from serving clients.
- Overloading some part of a system, including the hardware or software service.

2.4 Attack Paths

Understanding how to achieve security inherently begins with a look at how it potentially could be compromised. Consider by analogy an example of a bank vault containing a lot of money that must be protected from bad guys.

There are many attack paths available to a bank robber: enter through the front door during business hours, break into the employee-only rear entrance, dig a tunnel, and so forth. The sum of all possible starting points is called the **attack surface**, so that is where building defenses begins. The money is an **asset** — the object of protection — that in software systems is often data, code, equipment or other digital resource.

In this bank robbery example, suppose the vault was given the easily guessed combination of 1-2-3 (that's the specific vulnerability), and so the robber poses as a customer, walks up to the vault and in a few seconds gains access while nobody is looking (insufficient access control, surveillance, and guarding are additional contributing weaknesses). So, the **attack path** now consists of the sequence of front door entry during business hours, walking up to the vault, and trying to guess the combination. The sum of these steps completes the exploit.

Once in the vault, in terms of the C-I-A fundamentals, the robber has many options:

- Count the money to learn the bank's funds. (confidentiality)
- Replace the money with fake bills. (integrity)
- Hide the money somewhere inside the bank. (availability)

Obviously, real bank robbers are focused on taking money for themselves rather than the C-I-A model, which is intended for information systems, but the same principles can be applied.

2.5 Security Terminology

The field of information security has developed specific terminology for the underlying fundamental principles. As with software in general, we reuse existing English words adding specific meanings, so it is easy to pick up the vocabulary, but it is worth taking the time to explicitly define the specialized meanings clearly.

Unfortunately, the specialized terminology and meanings of these words is not uniformly applied so you may hear some variations and there is no authoritative right and wrong. Some of the common alternative terms will be noted here parenthetically when the meaning is similar. Generally speaking, these differences in language usually are harmless and easy to adapt to, so it is not something to try to fix. In this book, we use what we believe to be the most common and logical words and meanings from our diverse experiences and make every effort to use these terms consistently throughout the text.

"Software bugs are errors, mistakes, or oversights in programs that result in unexpected and typically undesirable behavior." — The Art of Software Security Assessment

"All software has bugs." The origin of this statement seems lost in time, but after many decades of software development, it remains truer than ever and hardly controversial. Given a set of bugs (also called defects or **weaknesses**) in a piece of software, some subset of those bugs will inevitably be useful to an **attacker** (bad guy, malefactor) to cause harm: these bugs are thus termed **vulnerabilities**.

“A vulnerability is a defect or weakness in system security procedures, design, implementation, or internal controls that can be exercised and result in a security breach or violation of security policy.” - Gary McGraw, Software Security

Reducing vulnerabilities is the central goal of this book. Eliminating a vulnerability, or somehow making it less likely to be found or exploited, is called **mitigation**. Since the complete elimination of all bugs in a complex piece of software usually is not practical, we try to build in multiple layers of defense so that a vulnerability is more difficult to find or has less impact.

The various kinds of hypothetical harm an attacker might achieve on a system are called **threats**. These threats are inherent — if the system holds a secret, there will always be a threat of it being revealed — and cannot be eliminated. Threats are identified with respect to protected resources called **assets**, things of value. In a computer system, these things are usually data, though they might also include the computing resource itself or a physical device controlled by the computer. The reason that it is important to recognize threats is that they can be an extremely useful lens for determining where important vulnerabilities may exist.

“The process of attacking a vulnerability in a program is called exploiting.” — The Art of Software Security Assessment

When a real attacker uses a vulnerability to actually cause harm, that attack is called an **exploit**. Getting past protective defenses is called an **intrusion**.

2.5.1 Threats

Secure software is achieved by anticipating the many attack scenarios and protecting against hypothetical attacks that might violate confidentiality, integrity, or availability. Here is a technical definition of an information security threat:

“A potential cause of an incident that may result in harm of systems and organization.” — ISO 27005

Put simply, threats are the answers to the question: what could go wrong? It is important to remember that threats are hypothetical and cannot be eliminated unless you are needlessly doing something risky and can dispense with it entirely. If your system holds an important secret, then the threat of disclosure or destruction of that secret is inevitably a threat to be dealt with.

Security improvements begin with evaluating potential threats to your system. Since these threats will be diverse and many, you need to prioritize them to focus your efforts appropriately. From this strategic holistic perspective, you begin the process of identifying places where threats can be

reduced, finding and then fixing the actual vulnerabilities. It is critical to add here that the vulnerabilities can be anywhere along the attack path of a possible exploit, so the location of the vulnerability can be distant and seemingly unrelated to the actual asset the attack may compromise.

We will discuss threats in more depth in Chapters 2.2 and 2.3, covering the topic of threat modeling.

2.6 Security Risk

The threats to large information systems are great. Such threats arise based on the extent to which such a system processes sensitive data, has a large attack surface, and has a high level of implementation complexity. Given the difficulty of securing such a system, it is essential to take a strategic approach focusing our protection efforts on the parts of the system with highest risk. Since time and resources available for security work are always limited, the work inevitably runs up against practice constraints — this is why all software is said to have vulnerabilities. Ideally, security effort continues working in priority order until further effort enters the realm of diminishing returns.

Risk management is an established practice that arose in the financial sector (insurance, investment, and gaming) to identify sources of risk and attempts to quantify them for the purpose of predicting potential losses and implementing cost effective mitigations. Put simply, risk management strategy accepts that risk is unavoidable and attempts to “put it in a box” in the sense of putting a cap on the worst case, spreading the loss over time, investing in preventative efforts, and monitoring actual losses so that management can accurately understand their actual risk stance.

Financial risk management allows quantitative assessments of risk by converting impacts into dollars; however, this is more challenging in the information security space for several reasons.

- Financial assets are based on money, which is easily measured in precise amounts.
- Money is fungible (one dollar is as valuable as another) but information is not.
- Money lost can be paid back but information once lost can be irretrievable.
- Insured risks have been accurately measured by actuaries and are well understood, but software security risks are newer and change rapidly.
- Financial institutions are carefully regulated, but software is mostly unregulated, and companies often conceal details of security incidents, so little public data is available.

- The financial impact of security incidents is difficult to measure, and there are additional risks such as losses to reputation or customer confidence that are difficult to quantify.

Ideally, you would like to have concrete criteria for these probabilities. These criteria might be related to measurable loss of service availability or cost of data loss. For example, electronic healthcare record (EHR) providers can accurately quantify the cost per record based on the black-market value of such records and the liability insurance cost to cover the provider.

However, risk can also be more subjective, when the assignment of risk is based on the judgment of the analyst. In our EHR example, another risk is loss of reputation, which might have a resulting loss of sales or cause the creation of burdensome new regulations.

When it is difficult to assign accurate numerical values, one common approach is to use “T-shirt sizes” to assign probabilities and gauge impacts as Small, Medium, Large, or X-Large. The use of such categories makes it clear that these categories represent an approximation of the risk.

2.7 Risk = Impact × Likelihood

The impact x likelihood formula is the foundational way to quantify risk (though modern financial enterprises use more complex formulas). To understand this formula, consider a couple of examples from the financial sector. Auto insurance rates are calculated based on historical data of collision statistics that detail the frequency that these collisions occur and the distribution of settlement costs. Interest rates are calculated in a similar way: a bank may make small loans without much analysis, knowing that a certain percentage of those loans will not be repaid, but the loss is relatively limited. For large loans, the bank will study the borrower, ask for collateral, and generally work harder to either lower the risk of non-payment or have some other recourse since so much money is at stake. In both cases, the lower impact events may be more likely, and that balances out against making the higher impact risks less likely to occur.

Now, let’s consider risk assessment for software security. Impact can be assessed against the three C-I-A pillars:

- Confidentiality: how sensitive is the data and how much of it might be disclosed?
- Integrity: how valuable is the data, how much could be tampered with, and how likely would that be detected?
- Availability: how long and to what degree would the system be unavailable?

There is a wide range of how quantitative and how accurately these can be measured. For example, if the private key to a Bitcoin wallet is compromised, then that is worth the present dollar value of the cryptocurrency it protects. On the other hand, if a private family photo is publicly revealed, it could be upsetting, lead to bad publicity, but no amount of money or effort can undo the damage done (so this is entirely subjective).

Evaluating the likelihood of an exploit is usually more complicated to determine. It is impossible to fully anticipate attackers to know their motivations, amount of inside information they have, and their level of expertise and capability. In addition, there may be a large amount of luck involved in whether or not a vulnerability is discovered and then exploited maximally. With open source software, you can assume that the attacker has the same knowledge that you do. However, it is risky to assume that you have much of an advantage if your software is based on proprietary source code. This amounts to **Security by Obscurity**, which is a dangerous strategy.

The one exception to the difficulty of assessing likelihood is when detection is virtually certain. CVE-2018-7602⁴ is a good example of this: a critical remote code execution vulnerability in the popular website system Drupal. Since, by design, it is trivial to scan a website and determine that it runs on Drupal and learn what version is installed, all vulnerable sites were well known. Once the fix was deployed publicly, attackers easily learned about the vulnerability in detail and how to exploit it. Simply by scanning the installed version they could automate attacks on any unpatched systems. So, the likelihood of such an attack was high – and, in fact, was quite common. Obviously, such high-impact high-likelihood risks are of the utmost priority.

The multiplicative risk formula is a useful lens through which to consider the value of various defenses. The risk values allow us to compare and rank different risks to help establish the priority in addressing them. For example, well designed password-based login systems⁵ should restrict the number and frequency of login attempts as a defense against brute-force password guessing. Since the risk of a lucky guess (and weak passwords make this all easier) cannot be eliminated, making it harder to guess does significantly reduce the likelihood, therefore reducing the value of $\text{Impact} \times \text{Likelihood}$. It is a little more inconvenient for a good user who has mistyped their own password to wait a little longer between attempts, but the risk reduction is arguably worth it.

⁴ <https://nvd.nist.gov/vuln/detail/cve-2018-7602>

⁵ Note that, in general, we discourage the use of pure password systems (one factor authentication) and recommend the adding of a second factor to the authentication process.

Reducing high risks on the impact side with specific mitigations is another valuable approach, focusing on the other multiplier in the equation. Any enterprise that maintains a large customer database needs to defend against the prospect of a major disclosure of all that sensitive data (as has happened famously to companies like Target, Equifax, and many more). A standard database configuration will include an administrative function of dumping the full database given the proper authorization, representing a huge risk. Possible mitigations might include only allowing download of an encrypted copy of the data (requiring a key to which only administrators would have access), or restricting the rate of data export such that it would take a long time to get all the data (yet fast enough to be useful for routine backup purposes). In these ways, the impact of a compromise may be reduced, lowering the overall risk.

A counterpoint about the limits of applying the standard risk formula to software security is worth mentioning. Quantifying both impact and likelihood in a meaningful way can be quite difficult. By the commutative principle of multiplication, this risk formula implies that in theory the case where the likelihood is extremely low, and the impact is very high is equal to the case where the likelihood is high and the impact is very low. Clearly these are very different risks, since an example of the former might be the world's chances of nuclear destruction and the latter an obvious bug in an obscure software application. The point here is that software security risk assessment is ultimately subjective and should not be reduced to a mindless calculation that could miss seeing the forest for the trees.

2.8 Cost of Insufficient Security

In theory, it is hard to imagine anyone seriously arguing that software security is frivolous and not important, yet somehow in practice diligent security professionals find themselves forever needing to justify the need for this important work. So, it is worth elaborating why security is important, which amounts to pointing out the significant costs of insufficient security. Part of the problem is the psychological effect. In practice, good software security is difficult to appreciate because it is so hard to notice; good security results in the absence of bad things happening. Yet, when the inevitable exploit does happen, then and only then is the spotlight on security in a bad light.

Software security is critical to protect our digital systems especially as these become more ubiquitously part of so many growing aspects of modern life, business, government, and culture at a time when everything is increasingly interconnected on the Internet (where attacks are always possible from literally anywhere at any time).

Attacks are expensive and may directly or indirectly affect assets including

money, reputation, data, communications, services, infrastructure, intellectual property. Examples of significant losses due to poor software security regularly appear in the mainstream media, and the potential for future damage is even larger and growing.

The explosive growth of digital systems, exemplified but by no means limited to giant platforms such as Google, Amazon, Apple, and Microsoft, is only possible because of the great value digital information has transformed the economy and most aspects of modern life. The value of the data they process greatly exceeds the cost of the physical infrastructure (such as data centers and networks) supporting them, and that high value represents high risk both in terms of massive complexity as well as attractiveness to malicious attackers. Experience shows time after time that if anything we need to practice more software security to higher standards in order to keep up with growing threats.

2.9 Summary

- Software security remains an unsolved problem; even old well-known vulnerabilities persist.
- Security terminology expresses subtle meanings, and it is important to use terms correctly.
- The pillars of information security are confidentiality, integrity, and availability.
- Attackers exploit vulnerabilities in the code to successfully exploit them.
- Every system (and all assets) faces many threats.
- Software security is important because the cost of insufficient security can be devastating.

2.10 Exercises

1. List your own creative examples of the three C-I-A pillars of information security.
2. Study the design of a well-known open source technology or component and enumerate some of its attack paths. Suggested subjects: OpenSSH, Sendmail, OpenVPN, VirtualBox.
3. Pick a few issues from an open bug tracking system (or make up some hypothetical bugs yourself) and prioritize them for security risk. Suggestions: Chrome bugs⁶, Firefox bugs⁷, and Linux kernel bugs⁸.

⁶ <https://bugs.chromium.org/p/chromium/issues/list>

⁷ <https://bugzilla.mozilla.org/>

⁸ <https://bugzilla.kernel.org/>

4. Choose a well-publicized software security incident (feel free to invent undisclosed details) and enumerate all of the resultant costs that resulted from fixing and recovering from the exploit.
5. Write a short essay on why computer security problems persist or even grow, despite much study and great efforts fixing them.