Tamper detection technologies: it takes a thief to catch a thief

Conversation with the expert

This fictional conversation is based on an interview with Gabriele Coppolino, tamper detection expert at Telsy.

telsy.com

Thank you for taking the time to discuss anti-tampering technologies with us. It is indeed a topic that is catching the attention both of experts in security as well as of the general public.
The first question that comes to mind is almost obvious: what does anti-tampering mean? Let's start with a simple and intuitive definition at first, so that also non-experts can wrap their head around the issue.

I would start with a terminological note first. At least to describe what we do here, it is better to talk about tamper detection, rather than of anti-tampering. Tampering occurs when an enemy, an attacker, gets hold of our device and opens it to see what's inside, how it is composed, to stimulate certain intermediate inputs to see how the output looks like, and to manipulate it with the aim of exfiltrating sensitive information. In a nutshell, we do not prevent the attacker from tampering with the device, so it would be partially improper to talk about anti-tampering, but we rather implement mechanisms of tamper detection, which – as the name says – allow us to detect attempts to physically attack our machine and to initiate appropriate defensive actions.
We will see some examples of attacks later, but let me first take you a step further. A key question to understand the concept behind tamper detection is: why is it important? It would be foolish of us to think that no one would be able to open our machine. Hence, we need to set ourselves in the mindset that our device may be vulnerable; ask ourselves how it could be vulnerable and how to construct a machine that already embeds physical security in its design. In the face of this fact, we need to implement methods that, when an attacker gets hold of the device, in the first place, detect inappropriate usage, and, immediately after, activate defense actions to protect our data. For example, if the top part of the machine is removed, we want the sensitive content, such as the cryptographic material, to be instantly erased, preventing attackers to access the specific cryptographic algorithm, the secret keys, or the memorized data.
To understand what tamper detection is and how it works, we always want to keep in mind that security is, in the first place, a mindset. When we design a machine, we do it by making sure that it is ready to respond to a wide variety of attacks: security is part of the design practice.

Indeed, security requires a particular mindset. Security experts need to see the world differently: it seems like a security expert cannot design or use something without wondering about its possible security vulnerabilities. Intuitively, we could say that, in the case of tamper attacks, these possible vulnerabilities and mechanisms resemble that of a full-fledged burglary: an enemy aims at "breaking in" our machines and at stealing what is inside. But what kind of information can be gained by physically breaking in the machine? What would an attacker do and why?

The metaphor of the thief is quite appropriate in this scenario. We can think of tamper attacks as attempts to get into someone's house, tear it apart if needed, and walk out with the jewelry. Out of metaphor, the physical machine (or device) is the house, and the encryption keys and the memorized encrypted data are the jewelry. As you can imagine, getting access to the encryption key gives access to the content of the information it is used to encrypt. Specific levels and systems of protection need to be designed to face the not-so-unlikely scenario in which someone gets hold of our device. When someone tries to force access to your house, you want there to be defense methods in place that both keep the thief out and secure your goods.sure that it is ready to respond to a wide variety of attacks: security is part of the design practice.
Tamper detection methods allow us to identify the attack when it takes place, and to predispose defense actions, such as the erasure of sensitive information.
Tamper detection methods are implemented so that, in a figurative sense, the machine itself is able to detect misuse and to act consequently, without further human instructions.

Exactly like when you are not at home and someone tries to break in, you want your alarm system to set off independently and to activate some kind of immediate response, say, for example, to dial your mobile phone number and alert you. Tamper detection methods behave more or less in the same way: when they detect unexpected behaviors of the machine, they set off the alarm, hence activating the relevant defensive measures.
Let's now turn to how tamper-thieves operate. As I have mentioned before, the main goal of an attacker is to access the information stored in the device or transmitted by means of it. One way to gain access to sensitive data, called modification, is to change the path on which information flows in the device by directly connecting the "red" and "black channels bypassing the encryption algorithm and accessing the non-encrypted information directly. This kind of attack is useful when the device is put back in the hands of the unaware user after it has been manipulated; in this case, information would flow non-encrypted through the device, and thieves could simply access them. In other cases, modification attacks may not be useful or feasible and attackers may try to extract the encryption keys to decipher sensitive information. Given that extracting the encryption key directly by trying all possible combinations, the so-called brute force attacks, would take an unreasonable amount of time and a definitely effort consuming venture, savvy attackers make use of side-channels to retrieve information that would allow the to reconstruct the key. A side-channel attack is a security exploit that involves collecting information about what a computing device does when it's performing an operation, such as a cryptographic operation, and using that information to, for example, extract the key.

To do this, the attackers take advantage of various factors, such as time elapsed, power consumption, system clock glitches, emitted radiations, focused laser beams, to extract information about the internal activity of the components, which in turn may lead to the extraction of the secret keys. There are, of course, various kinds of side-channel attacks, which exploit particular design aspects of the algorithm and/or of the device. One type of action that an attacker could take to exploit side-channels is fault injection: the attacker modifies input to an intermediate function to obtain outputs that leak information about the encryption key. Other kinds of attacks exploit physical properties of the implementation to extract information about the secret keys. One classical example is the analysis of power consumption: in electronic devices, the instantaneous power consumption is dependent of the operations performed by the device. Hence, profiling power consumption of an under-operation system would, for example, allow a thief to get information on a specific variable (i.e. an encryption key) used in a specific operation. In such attacks, it would be like you could figure out where the jewelry is stored in the house, and eventually steal it, just by looking at the intensity, and so on, of the electricity that flows in and out of the building. Of course, this is easier said than done, but it is nonetheless possible. Another attack that we can mention here is reverse engineering: when not only the encryption key is secret, but also the algorithm, attackers need to figure out the latter as well. In the case of breaking into a house, it would be like when not only the key but also the structure of the lock is initially unknown, and the latter has to be established only by looking at how the lock is composed by the outside. With encryption algorithms, the process of – broadly speaking - reconstructing the source code from the machine language is what we call reverse engineering. Within this scenario, tamper detection can, for example, prevent attackers from laying hands on the physical storage medium that contains the program.

Correct us if we are wrong but, as far as we can understand, we are mostly talking about attacking – and hence protecting – the hardware components, the physical machine. Can you give us some practical example of what aspects of a machine can be exploited by attackers and, consequently, which are the most common detection methods?

Yes, we are mostly talking about attacks to the hardware components of the machine. But, as I was telling you earlier, by looking at how the machine works, as for example by looking at power consumption, we can eventually extract important information about the processes being executed and, eventually, the encryption algorithm and the encryption keys.
Try to imagine this scenario. Someone gets hold of a machine and, most probably, the first thing he or she would do is to remove the cover and take the machine apart. As when someone wants to break into a house, the first thing to do is to find a way in. What we can do is to apply detection methods on the machine that precisely aim at detecting when and where someone is trying to break in. For example, the physical case surrounding the machine is the first line on which to apply tamper detection methods. Hence, we can apply pressure switches under the front mask, which is the first part of the external case that will be disassembled. These switches remain pressed until the case is closed. When you take the mask off, the pressure on the switches is released. The processor is sensitive to the switches status and an alarm is activated when there is an unexpected change of state. However, we should not forget that if the attacker does not plan to put the device back in the hands of the unsuspecting user, he or she might as well use other, more invasive, methods to open the device, such as for example by drilling holes in it. Consequently, different kinds of sensors may be implemented to, in the case of drilling, detect vibrations.

Other widely used mechanisms that follow the same conceptual principles are temperature and light sensors. Imagine that the thief actually passes the first line of defense – the front door – and gets into the house. You want to have other detection methods scattered around the house, and especially in the close vicinity of sensitive components. Temperature and light sensors can be thought of as traps installed in the house and cooperate, together with other detection methods, in creating layers of security. The operative principle of these sensors is fairly straightforward: when the sensors detect levels of light or temperature that go beyond given thresholds, the machine knows that a tamper attack is in motion. If someone illicitly opens the machine, the light sensors will be exposed to a different and unexpected light source, triggering the alarm. A similar mechanism underlies temperature detection; in general, we are interested in detecting if the temperature raises or decreases quickly above or beyond given thresholds. For example, a sharp drop of the temperature may freeze the physical memories that contain the information and “trap” the electrical charges enabling thieves to read them even if the processor is shut down. In this and similar cases, temperature sensors, which have to be battery-backed so that they are constantly active, detect that temperature variations and set of the internal process that leads to the activation of the alarm. Another kind of attack, which is becoming more and more common, involves the modulation of the power supply voltage so to take the machine to operate outside of its standard conditions. This may lead to malfunctioning of the machine that, sometimes, reveal information about the secret key or about the plain text. To inhibit this kind of attack and to prevent data exfiltration, defenders can design circuits that already embed protection systems: for example, if voltage drops below a given threshold, the protection mechanism can be programmed to shut down the entire downstream circuit.
This is clearly just a brief overview of tamper attacks and of subsequent detection methods. However, as you can imagine, the list could go on as each detection methods targets a specific modality in which tampering may take place.

This makes us think at your job in these terms: to some extent, it is like you have to impersonate a thief – at least as a thought experiment - to try to imagine how a malicious actor would actually try to "break in". It seems rather fair to suppose that, at least in the case of tamper detection, only a thief can catch a thief. Gaining the correct perspective to foresee possible threats and burglary seems almost like a mind game between the defender and the offender. The main question now is: how do you choose the appropriate methods?

To do this job properly and to create and implement effective detection methods, you definitely need to put yourself in the perspective of the offender and to imagine all possible ways in which your machine might be or become vulnerable. Even those attacks that may look like magical tricks are actually based on physical or procedural processes that take place in the machine, and as such they can be foreseen and the correct defensive measure can be implemented. At this level, we cannot prevent malicious actors to get hold of the machine, but we can make sure that, even if they do, no sensitive information will be accessible.
The answer to how we choose the appropriate methods depends on the type of machine we are dealing with and, especially, on the level of security we need to assure; that is, on what we know or can imagine of the technological capabilities and resources of a potential attacker. The level of protection needs to be also tailored to the sensitivity of the information.

Very much in the same way in which you would not use the same kind of lock to protect the door of our house and the door of the National Bank vault, we would not implement the same methods independently from what a device is used for. Out of metaphor, not all devices require the same level of security and hence it would not be worth to implement the same tamper detection methods to all kinds of machines. First of all, detection methods can be expensive. Second of all, they have costs in terms of both usability and design. On the one side, the implementation of detection methods complicates the machine and can make it susceptible to malfunctioning. In the second place, tamper detection methods cannot be placed anywhere in the machine without restrictions. For example, temperature sensors are inexpensive but they cannot be installed close to components that produce a lot of heath. In the third place, the implementation of tamper detection methods can have consequences on storage and usability modes. For instance, if you implement temperature sensors with a threshold for low temperature, you cannot store the machine in an environment where the temperature is likely to fall below that threshold. And so on. Therefore, what and where tamper detection methods are implemented depends on a range of considerations. In very general terms, the more sensitive is the information the device protects, the more refined and interlocked the tamper detection measures are, making it almost impossible – or indeed very hard – to get unauthorized access to the information.
However, there are methods, like tamper evidence, that can be applied more broadly. Tamper evidence is a sort of subfamily of tamper detection methods, that simply allows us to know in a very direct, mostly visual way, whether someone has, for example, opened a machine. An example may be that of numerated seals issued by specialized firms, or glues applied on the screws that are hard to replicate by the attackers. Just to point out the widespread use of such methods, the phone you are carrying has one as well: if you remove the battery of your mobile phone, you would lose warranty on the device, as an intact battery is a tamper evidence. Tamper evidence methods do not prevent data exfiltration, but if we know that an information has been compromised, then we can change the keys (or take other appropriate action) and make the attack worthless. Besides these additional methods and their various possible fields of applications, the richness of tamper detection methods is that, when combined to protect a machine, it is very difficult to neutralize all of them simultaneously.

In more practical terms, an attacker may even be able to keep to switches on the front mask pressed when disassembling it, and hence to bypass the first layer of protection, but he or she will nonetheless have to find ways – which are not at all straightforward – to overcome other various mechanisms, scattered in the device, that might set off the alarm. In tamper detection, unity is strength: the physical protection of the device is more effective when the various tamper detection methods and responses work synergically one with the other at different levels of security. If, on the one side, each tamper detection method is specific for a defined type of attack, it is the interlocking of the various mechanisms that allows us to secure the device and the information it contains. The metaphor of protecting your house from a burglary may help us again to clarify this concept: installing an alarm that tells you if someone opened the door – or a window – is not enough to protect the jewelry. The thief may in fact be skilled enough to pass this line of protection without even setting off the alarm. Hence, you want there to be various levels of protections, such as various kinds of “traps” in the house, that, even if designed to protect again one specific issue, only when combined and taken together, they substantially increase the level of security of the device.

Thank you so much. You definitely gave us a comprehensive overview of possible attacks and detection methods. What is indeed interesting is the combination of concepts that are implicated in tamper detection, from the need of a security mindset to the choice of the right layer of protection. One last question I would want to ask you is: how do you see the future of tamper detection?

As for now, tamper attacks are extremely expensive and require highly specified technologies. For example, examining the chip at the electronic microscope to understand if a bit is set to 0 or 1 and to understand if the single cell is powered or not, is not something you do easily, but rather a process that requires expertise, advanced technologies, and quite a large budget. We are talking about at least of hundreds of thousands of dollars for a single attack, depending on the level of sophistication of the machine and hence the classification of the information it protects. So, for now, tamper detection methods interest just some very specific sectors. But things are already changing. Today’s landscape in tamper detection methods is in line with security issues and with the technological advancement of our times. But, if we think of a tomorrow not very far from now, like in 10 or 15 years, the situation may be very different. Solutions that are now designed and available just for elite-types of machines, may become the standards for those products that we will then find in department stores. In 15 years from now, tamper attacks may become easier to launch and overall less expensive, forcing the consumer market to change accordingly. In many ways, working on tamper detection now is a pioneering endeavor that is posing the basis to prevent threats that, today, have direct impact only on specific types of machines, but that are likely to become the everyday threats of tomorrow. It does not just take a thief to catch a thief, but it takes a thief with a security mindset to catch a thief that is constantly evolving and refining its methods.