Trusted Computing and DRM

Nezer Zaidenberg, Pekka Neittaanmäki, Michael Kiperberg, Amit Resh

Department of Information Technology, University of Jyväskylä, Jyväskylä, Finland

1 Introduction

Trusted Computing is a special branch of computer security. One branch of computer security involves protection of systems against external attacks. In that branch we include all methods that are used by system owners against external attackers, for example Firewalls, IDS, IPS etc. In all those cases the system owner installs software that uses its own means to determine if a remote user is malicious and terminates the attack. (Such means can be very simple such as detecting signatures of attacks or very complex such as machine learning and detecting anomalies in the usage pattern of the remote user).

Another branch of attacks requires protection by the system owner against internal users.

Such attacks include prevention of users to read each other's data, use more than their allotted share of resources etc. To some extent anti-virus/anti-spam software is also included here. All password protection and used management software are included in this branch.

The third branch, *Trusted Computing*, involves the verification of a remote host that the user machine will behave in a certain predictable way, i.e. protection against the current owner of the machine. The most common example for this kind of requirement is distribution of digital media. Digital media is distributed in some conditional access mode (rented, pay per view, sold for personal use, etc.). Obtaining digital media usually does not entitle the user to unlimited rights. The user usually may not redistribute or edit the digital media and may not even be allowed to consume it himself after a certain date. (Media rentals, pay per view) However, as the user is consuming media on his private machine. How can the media provider assure himself that a malicious user does not tamper with the machine so that contents are not replicated? The problem of security against the owner of the machine is the problem region of Trusted Computing. In trusted computing as opposed to other branches of security the "attacker" is not limited to some attack surface that was exposed to him but can also use a soldering iron to tap into busses, replace chips and other system parts etc.

Trusted computing also includes other protection tools against the current owner (or possessor of the machine if not the legal owner). For example protection of

sensitive data or disk encryption solutions for laptops and mobile phones that can potentially be stolen.

Trusted computing can also be used on the cloud to ensure that the host does not inspect a cloud server and the software running on the server is not stolen. Latest trusted computing technology involves means to ensure commands are sane and are not malicious, for example in computers on cars and avionics. In this chapter we will review DRM and Trusted computing solutions from multiple sources.

2 Ethics – Trusted or Treacherous computing

Users don't like trusted computing.

First and foremost, the concept of conditional access leads to numerous digital rights debates. For example, if I legally purchased contents, shouldn't I be allowed to make backups of said contents? Especially as no media vendors are currently proposing to offer free (or even cheap) replacements of corrupted media contents! However, if we allow media to be replicated then how can we disallow illegal copies? What is stopping the user from redistributing copies or "backups"? How can we distinguish legal use of copies (backups) and illegal copies of the same content?

Secondary, as many trusted computing devices requires the user to actively install something on his machine (a TPM chip, EFI firmware, etc.). And the said hardware component does not contribute to the end-user system features at all (if anything trusted computing only limits the user). Why would the user willingly spend money and install some piece of hardware in his computer that only serves to limit what she can and cannot do?

All these reasons have lead Richard Stallman to call trusted computing treacherous computing and numerous hackers to try attacks on TPM chips and trusted platforms.

As of writing this chapter there is no clear cut winner in this technology battle. On the one hand there is still no massive install base for trusted computing solutions and on the other hand the trusted computing group is still alive and still releasing new trusted platform modules and specs.

3 The Trusted Processing Module by TCG

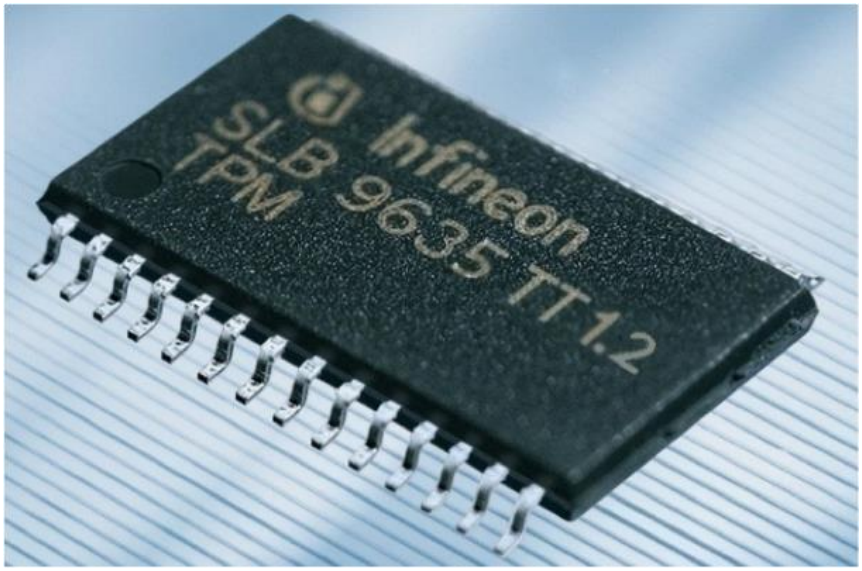

Figure 1: Trusted platform module

The trusted platform module, as demonstrated in Fig. 1 is a separate computer chip that is added to the computer motherboard and is frequently connected to the CPU using the LPC (low pin count) bus. The TPM is a cryptographic co-processor that is able to perform several cryptographic functions as well as generate and store keys.

The TPM can also verify the hardware and software that the system runs on and attest for the system's sanity.

When a remote host is querying the system for sanity it can use the TPM to verify that the software that it runs on was not tampered. TPM supports two attestation methods: Remote Attestation and Direct Anonymous Attestation.

3.1 Remote Attestation

The attestation solution proposed by the TCG (TPM specification v1.1) requires a trusted third-party, namely a *privacy certificate authority* (privacy CA). Each TPM has an RSA key pair called an Endorsement Key (EK), embedded inside the TPM (that the user cannot access – at least not easily.)

In order to attest itself, the TPM generates a new RSA key pair (Attestation Identity Key or AIK).

Remote Attestation is a typical trusted 3rd party process. Assuming Alice wants to attest Bob and Bob wants to be recognized by Alice but neither wants to reveal his private identification code to each other, remote attestation is suggested. It requires a trusted party that both can trust.

Bob attests himself by signing the public AIK using the EK, to the trusted 3rd party. The trusted 3rd party then verifies Bob by using Bob's public EK. Of course the CA may and should blacklist TPMs if it receives too many requests using the same key simultaneously. Alice can later verify with the CA Bob has is indeed attested.

3.2 Direct Anonymous Attestation

The Direct Anonymous Attestation (hereby DAA) protocol was only added to the TPM standard in version 1.2. DAA is based on three entities and two steps. The entities are the TPM platform, the DAA issuer and the DAA verifier. The issuer is charged to verify the TPM platform during the Join step and to issue DAA credentials to the platform. The platform uses the DAA credentials with the verifier during the Sign step. The verifier can verify the credentials without attempting to violate the platform's privacy (zero knowledge proof [5, 6]). The protocol also supports a blacklisting capability, so that verifiers can identify attestations from TPMs that have been compromised.

DAA allows differing levels of privacy. Using DAA Interactions is always anonymous, but the user/verifier may negotiate as to whether the verifier is able to link transactions (with the same user but not a specific user). Verifying transactions would allow persistent data to be saved over sessions and would allow profiling and tracing multiple logins.

4 Intel TXT instructions

Intel TXT technology defines unique extensions for the CPU instruction set to allow trusted execution [2]. Using intel TXT, one can attest the hardware, OS and software currently running and ensure a stable (as opposed to tampered) system state. Intel TXT uses the TPM for measurements and cryptographic functions to attest to a 3rd party and ensure that system software or the OS that is currently running is indeed trustworthy and non-tampered with.

The PCR registers on the TPM contain measurements and SHA-1 hashes of various system stages and code and by checking and verifying these measurements the system can be trusted to boot a non-tampered software.

5 AMD/ARM Trustzone

AMD/ARM Trustzone is the ARM/AMD implementation of trusted computing. It is roughly corresponding to Intel TXT. The Trustzone implementation is used by both AMD and ARM. Trustzone allows signed secure OSs to be loaded, for example, by using AMD/ARM SVM.

6 Other architectures for "Trusted computing"

These architectures provide means to prevent replication of data and thus introduce trust on various systems. We focused mainly on Video content delivery in this chapter. Different systems for preventing homebrew and pirated software on game consoles (which is another form of trusted computing) are covered in Chapter 3.

6.1 HDMI and HDCP and its predecessors

The Video industry has always been interested in mixed goals:

1. It searched for ways to deliver high quality video to the user's home. Generating a new revenue stream from videos that no longer appeared in cinemas (Video/DVD rentals).
2. It searched for ways to prevent the user from obtaining permanent access to the video equipment she rented by making illegal copies.

To some extent the battle was a lost cause to begin with because the user could always point a standard camera to the screen and just record using the camera (or create low quality copies using older, already broken technology). However, the industry was interested in preventing the user from making high quality copies (for example, digital quality copies in the case of HDMI).

This approach led to several technologies whose purpose was to circumvent the user's ability to create illegal copies

6.2 Macrovision

Old VHS video devices had a macrovision device that prevented direct creation of copies of VHS media by connecting two video devices to each other.

The Macrovision devices modified the output stream in a way that was unnoticeable to users but prevented VHS devices to create VHS cassettes copies by daisy chaining devices.

6.3 CSS and DeCSS and improvements

CSS or Content Scrambling System is an encryption system that is used on all major DVDs.

CSS was 40 bit encryption system.

The use of CSS was supposed to make it impossible to copy video content directly from DVD to a video. This was done as the encryption keys were kept in unreadable (by data DVD players) location.

CSS also allowed for DVD regions, Macrovision etc.

DVD CSS was broken at 1999, about 3 years after it was introduced with the introduction of DeCSS software. An inherent bug was used to reduce the keys from 40bit to only 16bit long and most players were able to break this encryption in less than 1 minute by brute force.

Two of DeCSS authors remain unknown even today. The 3rd was a Norwegian teenager: Jon Lech Johansen. Mr. Johansen was brought to trial and acquitted by the Norwegian court. The prosecution appealed and Mr. Johansen was acquitted for the second time.

When DVD was superseded by Blu-Ray and HD-DVD CSS was replaced with the AACS (Advanced access contents system, which was broken using leaked keys).

6.4 HDMI and HDCP

HDMI or High Definition Media Interface is a high quality media interface allowing high quality media transfer to monitors and screens. HDMI raised the problem of creating exact or near exact high quality replicas of video content.

To avoid copying the contents, HDCP will encrypt the content travelling between two end points of HDMI and will only provide contents to devices with trusted keys. These keys can later be revoked if they are stolen.

By 2010 the master key for HDCP had been leaked, rendering all revocation list useless.

It is possible that the revocation key was used too many times and provided sufficient data that made breaking the key easier.

7 Other uses for trusted computing

Several attacks on the user can be done after the attacker has obtained full or even partial control on the end user device. For example the attacker may be interested in the contents of the user hard drive after obtaining control of the user laptop (for example, by stealing it)

Such trusted computing content protection methods involve the usage of protected (encrypted) storage were the keys are saved on the TPM.

7.1 Microsoft Bitlocker and similar products

Microsoft Bitlocker is a full disk encryption solution that can be used on computers (especially laptops) to ensure that the disk contents are unreadable to an attacker, even if the computer/laptop was stolen. The complete disk is encrypted and the key to decipher the disk content is unreadable and saved on the TPM.

7.2 Protection on Mobile phone data

Mobile phones contain private data that can be exposed if the phone is lost or stolen.

Numerous technologies have been generated by various sources from using TPM and encryption on the device to a more biometric approach.

Examples include apple usage of fingerprint reading devices on the iPhone device that are required to unlock a stolen phone.

Other technologies include a kill code that is used to wipe the device and prevent it from connecting to the network ever again.

8 Attacks on trusted computing

8.1 Reset Attacks on the TPM chip

The TPM is often connected to the LPC (low pin count) bus. A legacy slows bus that exist on virtually all PCs. Attacks on this exist for over 10 years. One of the first cases of attacks on this bus occured on the first XBox.[9] By connecting

and eavesdropping to the LPC bug several hackers have been able to intercept and reset the TPM[8].

8.2 Attacks on the implementation of TPM

In 2010 [3] and later in 2012 [4] Chris Tarnovsky demonstrated physical attacks against TPM chips by Infineon and ST microelectronic. Tarnovsky attacked the TPMs by eliminating parts of the TPM chips thermal casing and attacking (i.e. connecting external devices) to the chips itself.

Tarnovsky demonstrated that ST and Infineon chips are made of older processors chips from their past. He demonstrated that by physically attacking the chips itself he could expose and modify content on the TPM chip itself.

Tarnovsky's methods require a special lab, chemicals and equipment which may not be in every hacker's reach. But it is definitely not beyond the reach of professional attackers and hackers.

8.3 Other attacks on trusted computing

One of the features of the Intel CPU is SMM or software maintenance mode. SMM is used to allow updating CPU code (microcode). SMM code is executed at higher permissions then user, kernel or hypervisor code on the Intel platform. Therefore, SMM is considered to run at permission ring -2 (if Ring 3 is userspace code, ring 0 is kernel code and ring -1 is hypervisor)

Rafal Wojtczuk and Joanna Rutkowska have demonstrated breaking TXT limitations using SMM [1]. These attacks may have been Intel's main reason for devising the SGX extension.

These attacks are possible because TXT protection blocks execution and permission in rings 3 (user space), 0 (kernel) and -1 (hypervisor) but TXT memory defense is still vulnerable to attacks on Ring -2 using SMM permission level which does not require any special permissions and can be used even after the OS has been attested by the TPM.

9 Beyond Trust – SGX

SGX or Software Guard Extension is an innovative technology from Intel that will be implemented in future chips. SGX provides a solution to the trusted computing problem on Intel platforms [2]. SGX technology allows creating an execution container for each process in which the process memory is contained. This

approach is similar to the approach taken by Qubes OS development to create separation using hypervisor code between applications so different applications are running on different virtual OSs [6] and by Trusted computing software such as TrulyProtect, which keeps secrets in the hypervisor layer [7]. At the time of writing this chapter SGX is not available with any Intel CPU on the market (thus there are no known attacks on SGX).

Bibliography

[1] Rafal Wojtczuk and Joanna Rutkowska. Blackhat DC 2009 “Attacking Intel® Trusted Execution Technology”

[2] Intel Trusted execution Technology - whitepaper hardware based technology for advanced server protection http://www.intel.com/content/www/us/en/trusted-execution-technology/trusted-execution-technology-security-paper.html

[3] Chris Tarnovsky Defcon 2012 “DEF CON 20 - Attacking TPM Part 2 - Chris Tarnovsky”

[4] Chris Tarnovsky. Blackhat DC 2010 “hacking the smartcard chip”

[5] Quisquater, Jean-Jacques; Guillou, Louis C.; Berson, Thomas A. (1990). "How to Explain Zero-Knowledge Protocols to Your Children". Advances in Cryptology – CRYPTO '89: Proceedings 435: 628–631.

[6] Blum, Manuel; Feldman, Paul; Micali, Silvio (1988). "Non-Interactive Zero-Knowledge and Its Applications". Proceedings of the twentieth annual ACM symposium on Theory of computing (STOC 1988): 103–112

[7] Nezer Zaidenberg ECIW 2013 “TrulyProtect 2.0 and attacks on TrulyProtect 1.0”

[8] TPM Reset Attack Evan Sparks http://www.cs.dartmouth.edu/~pkilab/sparks/

[9] Michael Stiel “17 mistakes microsoft made with the xbox security systems”