


Author
Michael Preisach
1155264

Submission
**Institute for Networks
and Security**

Supervisor
Univ.-Prof. DI Dr. **René
Mayrhofer**

September 2021

Project Digidow: System Integrity and Attestation for the Biometric Sensor



Statutory Declaration

I hereby declare that the thesis submitted is my own unaided work, that I have not used other than the sources indicated, and that all direct and indirect sources are acknowledged as references.

This printed thesis is identical with the electronic version submitted.

Linz, September 2021

Abstract

What is it all about? Why is that interesting? What is new in this thesis? Where is the solution directing to?

Kurzfassung

Das am Institut für Netzwerke und Sicherheit entwickelte Projekt *Digital Shadow* benötigt in vielen Bereichen ein prüfbares Vertrauen um eine Erkennung von Nutzern anhand ihrer biometrischen Daten zu erkennen und Berechtigungen zuzuteilen. Das Vertrauen soll dem Nutzer die Möglichkeit geben, die Korrektheit des Systems schnell und einfach zu prüfen, bevor er/sie dieses System biometrische Daten zur Verfügung stellt. Diese Masterarbeit beschäftigt sich nun mit den existierenden Werkzeugen, die ein solches Vertrauen schaffen können. Das implementierte System kombiniert diese Werkzeuge, um damit sensible Daten von Nutzern aufzunehmen und im Netzwerk von Digital Shadow zu identifizieren. Es soll dabei sicher gestellt sein, dass eine fälschliche Verwendung der sensiblen Nutzerdaten ausgeschlossen wird. Anhand dieses Systems werden die Eigenschaften einer vertrauenswürdigen Umgebung für Software diskutiert und notwendige Rahmenbedingungen erläutert.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Trust	2
1.3	Project Digidow	2
1.4	Our Contribution: Deriving Trust from the Biometric Sensor	5
1.5	Organization	6
2	Background	8
2.1	Trusted Platform Module (TPM)	8
2.1.1	Using the TPM	10
2.1.2	The Hardware	10
2.1.3	TPM Key Hierarchies	11
2.1.4	Endorsement Key	11
2.2	Trusted Boot	12
2.2.1	Platform Configuration Register	12
2.2.2	Static Root of Trust for Measurement	14
2.2.3	Platform Handover to OS	14
2.2.4	Secure Boot	15
2.2.5	Intel TXT	16
2.3	Integrity Measurement Architecture	16
2.3.1	Integrity Log	17
2.3.2	IMA Appraisal	18
2.3.3	IMA Policies	18
2.3.4	Other System Integrity Approaches	19
2.4	Direct Anonymous Attestation	19
2.4.1	Mathematical Foundations	20
2.4.2	DAA Protocol on LRSW Assumption	22
2.4.3	Standardization of DAA	26
3	Concept	27
3.1	Definition of the Biometric Sensor	27
3.2	Attack Vectors and Threat Model	28
3.3	Prototype Concept	30
3.3.1	Integrity and Trust up to the Kernel	30

3.3.2	Integrity and Trust on OS Level	33
3.3.3	Proving Trust with DAA	33
4	Implementation	37
4.1	Hardware Setup	38
4.2	Select the Operating System	39
4.3	Trusted Boot	40
4.3.1	Install and use Trusted Boot	41
4.3.2	Detailed description of the scripts	42
4.4	Integrity Measurement Architecture	45
4.5	Dynamic System Analysis	48
4.6	Prove TPM2 Certificate Chain	49
4.7	Using the DAA Protocol	52
4.7.1	Provision Hosts of Test Setup	52
4.7.2	Installing Xaptum ECDAA Library	53
4.7.3	Installing ECDAA network wrapper	53
4.8	DAA Demo Application	54
4.8.1	DAA Network Protocol	55
5	Testing	58
5.1	Trusted Boot	58
5.2	IMA	59
5.3	Processing and Sending Biometric Data	60
5.4	Limitations	64
6	Conclusion and Outlook	67
6.1	Future Work	67
6.1.1	Closing the chain of trust between TPM manufacturer and DAA issuer	67
6.2	Outlook	68

This work has been carried out within the scope of Digidow, the Christian Doppler Laboratory for Private Digital Authentication in the Physical World, funded by the Christian Doppler Forschungsgesellschaft, 3 Banken IT GmbH, Kepler Universitätsklinikum GmbH, NXP Semiconductors Austria GmbH, and Österreichische Staatsdruckerei GmbH.

1 Introduction

1.1 Motivation

We all live in a world full of digital systems. They appear as PCs, notebooks, cellular phones or embedded devices. Especially the footprint of embedded computers became so small that they can be used in almost all electrical devices. These embedded systems form the so called *smart* devices.

All these new devices made life a lot easier in the past decades. Many of them automate services to the public like managing the bank account, public transportation or health services. The list of digital service is endless and will still grow in the future.

The downside of all these digital services is that using these services generate a lot of data. Besides the intended exchange of information, many of the services try to extract metadata as well. Metadata answers some of the following questions. Which IP address is sending or receiving? What kind of device is that? Is the software of the device up to date? Was this device here in the past? What else did the owner on the device? This set of questions is not complete.

Aggregating metadata is not required to fulfill the function of the requested service. However, aggregating and reselling the metadata brings the provider more margin on the product and hence more profit. Consequently, the market for metadata is growing and yet only partly regulated. Since metadata aggregation is one downside of using smart services, providers try to downplay or to hide these aggregation features where possible. Often a proprietary layer is used either on the client or the server side to hide those functions. The result is a piece of software which is provided as binary and the user cannot prove what this software is exactly doing besides the visible front end features.

There are of course other purposes for delivering software in a closed source manner. Firmware of hardware vendors is usually not disclosed. Instead, vendors provide an API where an *Operating System* (OS) can connect to. Some companies deliver complete closed source devices with internet connection. In such cases, the feature of closed source is to protect the intellectual property of those companies. Any user of these closed source products must use them as black box and needs to *trust* the vendor that it is working correctly.

There is, however, a special need for users to keep sensitive data secret. Especially when providing confidential data like passwords or biometric data, a certain level of trust is required. This means that the user assumes that the provided sensitive data is handled properly for only the designated usage. One may argue that a password can easily be changed when revealed to the public. Unfortunately, this does not apply to a fingerprint since a human usually has only ten of them during lifetime.

1.2 Trust

When using a system with an authentication method, trust plays a key role. For black box systems this trust is cast to the vendor of the system or device. There is however no mathematical proof that the device is indeed executing the software as intended from the vendor.

This thesis will therefore use the term *trust* as a cryptographic chain of proofs, that a system is behaving in an intended way, a so called *Chain of Trust*. By providing a Chain of Trust, a user can ask the vendor for a certification of its devices and consequently comprehend the state of the system at hand. The Chain of Trust will be separated into two parts, namely the creation of trust on a certain system, and the transfer of trust over the network for verification purposes.

1.3 Project Digidow

The Institute for Networks and Security is heavily using the cryptographic form of trust in the project *Digital Shadow* (Digidow). Digidow introduces an electronic authentication




Figure 1.1: Overview of the Digidow network with its interactions

system, which aims to minimize any generation of metadata on system and network level and hence maximizes the level of privacy for their users. The project furthermore aims to specify a scalable solution for nationwide or even worldwide applications including provable trust and integrity to the user.

The picture in Figure 1.1 provides an overview of the Digidow network. The nodes using the Digidow network and their interactions are not fully specified at the time of this writing. Therefore the processes may be adapted when necessary. DigiDow introduces five main parties which are involved in a common authentication process.

- *Individual*: The human user who wants to be identified via the Digidow network.
- *Personal Identity Agent (PIA)*: The PIA is the digital shadow of an individual. This individual is also the owner of the PIA and should be in control of sensitive data and software on it.
- *Issuing Authority*: This party acts as an authority for the individual's attributes. These attributes show an aspect of the individual's identity. After identifying the

individual via the Digidow network, these attributes may be used to allow or deny a certain action.

- *Verifier*: This is the party that verifies the whole authentication process and may finally trigger the desired action. It is usually strongly connected with the sensor which starts the identification process.
- *Sensor*: For authentication, an individual has to be uniquely identified. Therefore, the sensor records biometric data from the individual and passes it to the PIA via the Digidow network. Sensors are not limited to sensing biometric data. However, we focus in this thesis on developing a prototype of a biometric sensor (BS).

When an individual wants to be identified by Digidow, he will eventually step in front of a sensor. This defines the beginning of a Digidow transaction. The procedure is as follows:

1. The sensor will start recording a unique digital representation, triggered by the directly connected verifier or by hardware detection.
2. The sensor eventually finds one or a small group of eligible PIAs, where a secure communication channel is established.
3. After receiving the digital representation of the sensor, the PIA identifies the individual if possible.
4. When identification was possible, the PIA eventually sends a proof of identification and a claim of the requested attributes to the verifier.
5. The verifier will check the cryptographic proofs of the claim and the sensor data. If successful, the verifier will grant the desired action for the asking individual.

The above illustration is an early draft of the whole setup and is under constant development. Latest developments the whole system will be published on the Digidow project page¹.

¹<https://digidow.eu>

1.4 Our Contribution: Deriving Trust from the Biometric Sensor

The Digidow network is designed to preserve privacy and to build trust for any user. A key feature is to show the user that all involved parts of the system are working as intended. So we design a prototype based on the common x86 architecture and use the cryptographic features of *Trusted Platform Modules* (TPM). A TPM is a passive crypto coprocessor available on many modern PC platforms which has an independent storage for crypto variables and provides functions to support above mentioned features.

We define a solution for installing and booting a Linux kernel with TPM-backed integrity measurements in place. We use an attached camera as example for a biometric sensor hardware to create the dataset to continue with the authentication process. This dataset will be combined with the integrity measurements of the system and a signature from the TPM and finally sent to the PIA for verification and further computation.

By building a system with an integrated TPM, the system should be able to provide the following properties:

- *Sensor Monitoring.* The system should be able to monitor the hardware sensor (fingerprint sensor, camera, etc.) itself.
- *System Monitoring.* It should be possible to track the state of the system. Especially every modification of the system at hardware level should be detected.
- *Freshness of Sensor Data.* To prevent replay attacks, the system should prove that the provided biometric data is captured live.
- *Integrity of Sensor Data.* As it is possible for an adversary to modify the provided data during the identification process, integrity should guarantee that the data is unmodified until identification is done.
- *Confidentiality of Sensor Data.* It should not be possible to eavesdrop any sensitive data out of the system. Furthermore almost all kinds of metadata (e.g. information about the system or network information) should not be published.
- *Anonymity.* Given a message from a BS, an adversary should not be able to detect which BS created it.

- *Unforgeability.* Only honest BS should be able to be part of the Digidow network. Corrupt systems should not be able to send valid messages.

The thesis focuses on a working setup as basis for future research. Since the Digidow protocols are not yet finalized, some assumptions are defined for this work and the prototype implementation:

- Any network discovery is omitted. BS and PIA are assumed to be reachable directly via TCP/IP.
- We look into a protocol which proves trustworthiness from BS to PIA. Any further proofs necessary for a Digidow Verifier are also not focused in this thesis.
- The sensitive datasets will be transmitted in cleartext between BS and PIA. It is considered easy to provide an additional layer of encryption for transportation. However this should be considered in the Digidow network protocol design. This thesis focuses only on the trust part between BS and PIA.
- Any built system is considered secure on a hardware level. Any threats which are attacking the system without changing any running software on the system may remain undetected. This includes USB wire tapping or debug interfaces within the system revealing sensitive information.

1.5 Organization

In the next chapter, we will introduce and discuss existing contributions in the targeted scientific area. This includes especially the theoretical foundations of the network protocol which is part of our contribution. Together with that, we will introduce our theoretical solution for the previously stated problems in chapter 3. chapter 4 introduces then a working implementation with all necessary parts for provisioning the environment and the used hosts accordingly. Finally we will present the results and limitations in chapter 6 and give an overview of future work.

- What exists in the field?
- Strongswan Attestation –

- Intel TXT
- Trusted Execution Environment (TEE)
- nanovm (nanovms.com)

2 Background

In this chapter we describe four main concepts which will be combined in the concept of this thesis. The TPM standard is used to introduce trust into the used host platforms. *Trusted Boot* and the *Integrity Measurement Architecture* (IMA) are two approaches to extend trust from the TPM over the UEFI / BIOS up to the OS. The generated trust should then be provable by an external party—in our case the PIA—by using the protocol of *Direct Anonymous Attestation* (DAA).

2.1 Trusted Platform Module (TPM)

The *Trusted Platform Module* (TPM) is a small coprocessor that introduces a variety of cryptographic features to the platform. This module is part of a standard developed by the Trusted Computing Group (TCG), which released the current revision 2.0 in 2014[13].

The hardware itself is strongly defined by the standard and comes in the following flavors:

- *Dedicated device*. The TPM chip is mounted on a small board with a connector. The user can plug it into a compatible compute platform. This gives most control to the end user since it is easy to disable trusted computing or to switch to another TPM.
- *Mounted device*. The dedicated chip is directly mounted on the target mainboard. Therefore, removing or changing the TPM is impossible. All recent Intel and AMD platforms supporting TPM2.0 are able to manage a TPM within the BIOS, even as mounted device.
- *Firmware TPM (fTPM)*. This variant was introduced with the TPM2.0 Revision. Instead of using a dedicated Coprocessor for the TPM features, this variant lives as

firmware extension within Intel's Management Engine or AMD's Platform Security Processor. Both Intel and AMD provide this extension for their platforms for several years now. When activating this feature on BIOS level, the user gets the same behavior as when using a mounted device.

- *TPM Simulator.* For testing reasons, it is possible to install a TPM simulator. It provides basically every feature of a TPM but cannot be used outside the OS. Features like Trusted Boot or in hardware persisted keys are not available.

Dedicated and mounted devices are small microcontrollers that run the TPM features in software giving the manufacturer the possibility to update their TPMs in the field. fTPMs will be updated with the platform updates of the CPU manufacturers.

The combination of well constrained hardware and features, an interface for updates and well defined software interfaces make TPMs trustworthy and reliable. When looking up the term *TPM* in the Common Vulnerabilities and Exposures database, it returns 23 entries[8]. Eight of them were filed before the new standard has been released. Another seven entries refer to vulnerabilities in custom TPM implementations. Six entries refer to the interaction between the TPM and the operating system, especially the TPM library and the shutdown/boot process. The last two entries describe vulnerabilities in dedicated TPM chips, which are mentioned in further detail:

- *CVE-2017-15361:* TPMs from Infineon used a weak algorithm for finding primes during the RSA key generation process. This weakness made brute force attacks against keys of up to 2048 bits length feasible. According to Nemec et al.[18], 1024 bit keys required in the worst case scenario 3 CPU months and 2048 bit keys needed 100 CPU years when using one core of an Intel Xeon E5-2650 v3 CPU. Infineon was able to fix that vulnerability per firmware update for all affected TPMs.
- *CVE-2019-16863:* This vulnerability is also known as "TPM fail" [17] and shows how to get an elliptic curve private key via timing and lattice attacks. The authors found TPMs from STMicroelectronics vulnerable, as well as Intel's fTPM implementation. Infineon TPM show also some non-expected behaviour, but this could not be used for data exfiltration. STM provided an update like Infineon did for their TPMs. Intel's fTPM required a platform firmware update to solve the issue.

2.1.1 Using the TPM

On top of the cryptographic hardware, the TCG provides several software interfaces for application developers:

- *System API (SAPI)*. The SAPI is a basic API where the developer has to handle the resources within the application. However, this API provides the full set of features.
- *Enhanced System API (ESAPI)*. While still providing a complete feature set, the ESAPI makes some resources transparent to the application like session handling. Consequently, this API layer is built on top of the SAPI.
- *Feature API (FAPI)*. This API layer is again built on top of the ESAPI. It provides a simple to use API but the feature set is also reduced to common use cases. Although the interface was formally published from the beginning, an implementation is available only since end of 2019.

The reference implementation of these APIs is published on Github[7] and is still under development. The repositories are maintained by members of TCG. At the point of writing stable interfaces are available for C and C++, but other languages like Rust, Java, C# will be served in the future. The repository additionally provides the tpm2-tools toolset which provides the FAPI features to the command line. Unfortunately, the command line parameters changed several times during the major releases of tpm2-tools[20].

2.1.2 The Hardware

With the previous mentioned software layers the TCG achieved independence of the underlying hardware. Hence, this design made the different flavors of TPMs possible.

With the TPM2.0 standard, TCG defined a highly constrained hardware with a small feature set. It is a passive device with some volatile and non-volatile memory, which provides hardware acceleration for a small number of crypto algorithms. The standard allows to add some extra functionality to the device. However, the TPMs used in this project provide just the minimal set of algorithms and also the minimal amount of memory.

Since TCG published its documents, several IT security teams investigated the concept and implementations of TPMs.




Figure 2.1: The manufacturer certifies every TPM it produces

2.1.3 TPM Key Hierarchies

A TPM comes with four different key hierarchies. These hierarchies fulfill different tasks and are used in different use cases on the whole platform. Will Arthur et al.[2] provide a more detailed description on how the hierarchies work together.

- *Platform Hierarchy:* This hierarchy is managed by the platform manufacturer. The firmware of the platform is interacting with this hierarchy during the boot process.
- *Storage Hierarchy:* The storage of a platform is controlled by either an IT department or the end user and so is the storage hierarchy of the TPM. It offers non-privacy related features to the platform although the user may disable the TPM for her own use.
- *Endorsement Hierarchy:* This is the privacy-related hierarchy which will also provide required functionality to this project. It is controlled by the user of the platform and provides the keys for attestation and group membership.
- *NULL Hierarchy:* The NULL hierarchy is the only non-persistent hierarchy when rebooting the platform. It provides many features of the other hierarchies for testing purposes.

Each of the persistent hierarchies represent its own tree of keys, beginning with a root key. Since TPM2.0 was published, these root keys are not hard coded anymore and can be changed if necessary. The process of key generation described below is similar to all three persistent hierarchies.

2.1.4 Endorsement Key

The *Endorsement Key* (EK) is the root key for the corresponding hierarchy. Figure 2.1 illustrates the certificate chain of building a new EK. Every TPM has, instead of the full

EK, a unique key seed to derive root keys from. The EK can then be generated with a *key derivation function* (KDF) which allows to add additional entropy to that of the TPM. According to Arthur et al.[i]n chapter 15, this additional entropy is not used and the parameter field is filled with zeroes. As a result, although it is possible to generate an arbitrary number of EKs, the TPM generates always the same by default. Only this default EK comes with a corresponding certificate which is signed by the TPM manufacturer by using its own root *Certificate Authority* (CA). Since the platform supports generating and using multiple root keys at a time, it is also possible to encrypt temporarily unused keys and store it on an external storage, e.g. on the platform disk. Consequently it is quite easy to have different EKs at once to address privacy features also between different functions of the endorsement hierarchy.

2.2 Trusted Boot

A boot process of modern platforms consists of several steps until the OS is taking over the platform. During these early steps, the hardware components of the platform are initialized and some self tests are performed. This is controlled by either the BIOS (for legacy platforms) or the UEFI firmware. There exists no source of trust and hence no check for integrity or intended execution in this common boot procedure.

2.2.1 Platform Configuration Register

The *Trusted Computing Group* (TCG) introduced their first standard for a new Trusted Computing Module (TPM) in 2004. As part in this standard, TCG defined a procedure where every step in the early boot process is measured and saved in a *Platform Configuration Register* (PCR). In this context, *Measuring* means a simple cryptographic extension function:

$$\text{new_PCR} = \text{hash}(\text{old_PCR} || \text{data}). \quad (2.1)$$

The function "||" represents a concatenation of two binary strings and the hash function is either SHA1 or SHA256. In recent TPM-platforms, both hashing algorithms can be

performed for each measurement. Consequently, both hash results are available for further computations.

The formula shows that a new PCR value holds the information of the preceeding value as well. This *hash chain* enables the user to add an arbitrary number of hash computations. One can clearly see that the resulting hash will also change when the order of computations changes. Therefore, the BIOS/UEFI has to provide a deterministic way to compute the hash chain if there is more than one operation necessary. The procedure of measurements is available since the first public standard TPM1.2. For TPM2.0, the process was only extended with the support with the newer SHA256 algorithm.

A PCR is now useful for a sequence of measurements with similar purpose. When, for example, a new bootloader is installed on the main disk, the user wants to detect this with a separate PCR value. The measured firmware blobs may still be the same. So the TPM standard defines 24 PCRs for the PC platform, each with a special role and slightly different feature set. The purpose of every PCR is well defined in Section 2.3.3 of the *TCG PC Client Platform Firmware Profile*[12] and shown in table 2.1. Especially those PCRs involved in the boot process must only be reset according to a platform reset. During booting and running the system these registers can only be *extended* with new measurements.

Table 2.1: Usage of PCRs during an UEFI trusted boot process

PCR	Explanation
0	SRTM, BIOS, host platform extensions, embedded option ROMs and PI drivers
1	Host platform configuration
2	UEFI driver and application code
3	UEFI driver and application configuration and data
4	UEFI Boot Manager code and boot attempts
5	Boot Manager code configuration and data and GPT / partition table
6	Host platform manufacturer specific
7	Secure Boot Policy
8–15	Defined for use by the static OS
16	Debug
17–23	Application

When TCG introduced Trusted Boot in 2004, UEFI was not yet available for the ordinary PC platform. Consequently, TCG standardized the roles of every PCR only for the BIOS platform. Later, when UEFI became popular, the PCR descriptions got adopted for the new platform.

2.2.2 Static Root of Trust for Measurement

The standard defines which part of the platform or firmware has to perform the measurement. Since the TPM itself is a purely passive element, executing instructions provided by the CPU, the BIOS/UEFI firmware has to initiate the measurement beginning with the binary representation of the firmware itself. This procedure is described in the TCG standard and the platform user has to *trust* the manufacturer for expected behavior. It is called the *Static Root of Trust for Measurement* (SRTM) and is defined in section 2.2 of the TCG PC Client Platform Firmware Profile[12]. As the mainboard manufacturer do not publish their firmware code, one may have to reverse engineer the firmware to prove correct implementation of the SRTM.

The SRTM is a small immutable piece of the firmware which is executed by default after the platform was reset. It is the first piece of software that is executed on the platform and measures itself into PCR 0. It must measure all platform initialization code like embedded drivers, host platform firmware, etc. as they are provided as part of the mainboard. If these measurements cannot be performed, the chain of trust is broken and consequently the platform cannot be trusted. When PCR 0 is zeroed or filled with the hashed representation of a string of zeroes, the SRTM did not act as expected. This indicates a broken chain of trust and should only appear when using the TPM simulator.

2.2.3 Platform Handover to OS

The BIOS or UEFI performs the next measurements according to table 2.1 until PCRs 1–7 are written accordingly. Before any further measurements are done, the control of the platform is handed over to the kernel of either a bootloader or the OS when booting without any bootloaders. In any case, these binaries are stored in the *Master Boot Record* (MBR) or provided as EFI blob in the EFI boot partition. It is noteworthy that the bootloader itself and its configuration payload is measured in PCR 4 and 5 before the handover is done. This guarantees that the chain of trust keeps intact when the bootloader/OS takes control.

The bootloader has to continue the chain of trust by measuring the kernel and the corresponding command line parameters into the next PCRs. The support and the way of how the measurements are done is not standardized. GRUB, for example, measures all

executed GRUB commands, the kernel command line and the module command line into PCR 8, whereas any file read by GRUB will be measured into PCR 9[9].

The whole process from initialization over measuring all software parts until the OS is started, is called *Trusted Boot*. The user can check the resulting values in the written PCR registers against known values. These values can either be precomputed or just the result of a previous boot. If all values match the expectations, the chain of trust exists between the SRTM and the kernel.

2.2.4 Secure Boot

Secure boot is another technology to prevent malware from being executed before the OS kernel is loaded. Microsoft describe on their Documentation for Windows and UEFI what requirements are needed and how the secure boot process looks like[16]. It is part of the UEFI specification and uses, similar to trusted boot, checksums of firmware, option roms and the boot loader. These checksums are checked against a signature database, which is held within the platform's NVRAM. The signatures are created with the platform key (PK) which is by default owned and managed by Microsoft. Although it is possible to install a new own PK and sign relevant software with it, you can only boot software signed from Microsoft by default when secure boot is enabled.

Shim is the gatekeeper for OSes not maintained by Microsoft. The binary is signed with the official PK and uses itself a self signed CA to sign further executables. A detailed description how shim works on Ubuntu is shown on their corresponding Wiki page[10]. Only this workflow enables secure boot when using Linux OSes.

Secure boot uses instead of the TPM conventional NVRAM to store private parts of its signing keys. Although there exists a signature chain for any software involved in Microsoft's boot process, using alternative OSes breaks the signature chain. When using an own PK, you lose the benefit of having externally created and signed hash values for checking during booting. Secure and trusted boot can, however, exist side by side on one system. The benefit of using it seems to be very limited when not using a Microsoft OS.

2.2.5 Intel TXT

Intel developed a solution to build a trusted environment on a hypervisor which they call *Trusted Execution Technology* (TXT). It requires an enabled TPM on the hypervisor as well as an activated trusted boot workflow. **ToDo!**

2.3 Integrity Measurement Architecture

The *Integrity Measurement Architecture* (IMA) is a Linux kernel extension to extend the chain of trust to the running application. IMA is officially supported by RedHat and Ubuntu and there exists documentation to enable IMA on Gentoo as well. Other OS providers may not use a kernel with the required compile flags and/or do not provide userland software required to manage IMA. The IMA project page describes the required kernel features for full support in their documentation[21].

The process of keeping track of system integrity becomes far more complex on the OS level compared to the boot process. First, there are far more file system resources involved in running a system. Even a minimal setup of a common Linux Distribution like Ubuntu or RedHat will load several hundred files until the kernel has completed its boot process. Second, all these files will be loaded in parallel to make effective use of the available CPU resources. It is clear that parallelism introduces non-determinism to the order of executing processes and, of course, the corresponding system log files. Hence when using PCRs, this non-determinism results in different values, as stated in subsection 2.2.1. The system, however, might still be in a trustworthy state.

Finally, the user might know some additional data to the current value in the PCR register. Since the value itself does not tell anything to the user, a measurement log must be written for every operation on this PCR index.

IMA comes with three property variables which set the behaviour of the architecture:

- `ima_template` sets the format of the produced log.
- `ima_appraise` changes the behaviour when a file is under investigation.
- `ima_policy` finally defines which resources should be analyzed.




Figure 2.2: Overview of generating an entry in the integrity log

These settings will be discussed in more detail in the following.

2.3.1 Integrity Log

IMA uses the *integrity log* to keep track of any changes of local filesystem resources. This is a virtual file that holds every measurement that leads to a change on the IMA PCR. When IMA is active on the system, the integrity log can be found in `/sys/kernel/security/ima/ascii_runtime_measurements`.

Before a file is accessed by the kernel, IMA creates an integrity log entry as it is shown in Figure 2.2. Depending on the settings for IMA, a SHA1 or SHA256 hash is created for the file content. The resulting *filedata hash* will be concatenated with the corresponding metadata. This concatenation will again be hashed into the so called *template hash* which is independently of the previous algorithm a SHA1 checksum. Finally, the template hash is the single value of the whole computation that will be extended into the PCR. The integrity log holds at the end the filedata hash, the metadata and the template hash as well as the PCR index and the logfile format. Unfortunately, recomputing the hash chain was not possible in the demonstration setup. We discuss that problem in chapter 6.

IMA knows three different file formats, where two of them can be used in recent applications. The only difference between these formats lies in the used and logged metadata:

- `ima-ng` uses, besides the filedata hash, also the filedata hash length, the pathname length and the pathname to create the template hash.
- `ima-sig` uses the same sources as `ima-ng`. When available, it also writes signatures of files into the log and includes them for calculating the template hash.

The older template `ima` uses only SHA1 and is fully replaceable with the `ima-ng` template. Therefore, it should not be used for newer applications.

The first entry in every measurement file is called `boot_aggregate`. It is the trust link between trusted boot and IMA representing a cumulative hash of the PCR values 0 – 7. Consequently, the PCR result of trusted boot is also embedded in the measurement log and the corresponding hash chain of PCR 10.

2.3.2 IMA Appraisal

IMA comes with four different runtime modes. These modes set the behaviour especially when there exists no additional information about the file in question.

- `off`: IMA is completely shut down. The integrity log just holds the entry of the boot aggregate.
- `log`: Integrity measurements are done for all relevant resources and the integrity log is filled accordingly.
- `fix`: In addition to writing the log file, the filedata hashes are also written as extended file attribute into the file system. This is required for the last mode to work.
- `enforce`: Only files with a valid hash value are allowed to be read. Accessing a static resource without a hash or an invalid hash will be blocked by the kernel.

2.3.3 IMA Policies

The IMA policies define which resources are targeted with IMA. There exist three template policies which can be used concurrently:

- `tcb`: All files owned by root will be measured when accessed for read.
- `appraise_tcb`: All executables which are run, all files mapped in memory for execution, all loaded kernel modules and all files opened for read by root will be measured by IMA.
- `secure_boot`: All loaded modules, firmwares, executed kernels and IMA policies are checked. Therefore, these resources need to have a provable signature to pass the check. The corresponding public key must be provided by the system manufacturer within the provided firmware or as Machine Owner Key in shim.

In addition to these templates, the system owner can define custom policies. Some example policies can be found in the Gentoo Wiki[11]. It is, for example, useful to exclude constantly changing log files from being measured to reduce useless entries in the measurement log.

2.3.4 Other System Integrity Approaches

Schear et al. [22] developed a full featured trusted computing environment for cloud computing. They show in their paper how a TPM of a hypervisor can be virtualized and used by the guest operating system. This includes trusted bootstrapping, integrity monitoring, virtualization, compatibility with existing tools for fleet management and scalability. The concept of a well known virtual environment does, however, not apply to our use case. Instead, the system should be self contained as good as possible and it should be possible to get information about the system via anonymous attestation.

2.4 Direct Anonymous Attestation

Direct Anonymous Attestation (DAA) is a cryptographic scheme which makes use of the functions provided by the TPM. DAA implements the concept of group signatures, where multiple secret keys can create a corresponding signature. These signatures can be verified with a single public key when private keys are member of the same group.

The scientific community is researching on TPM-backed DAA since the first standard of TPM went public in 2004. Since then many different approaches of DAA were discussed. According to Camenisch et al. in [4] and [3], almost all schemes were proven insecure, since many of them had bugs in the protocol or allowed trivial public/secret key pairs. This also includes the implementation of DAA in the TPM1.2 standard.

This section describes the concept by Camenisch et al. [4] including the cryptographic elements used for DAA. Unlike the description in the original paper, we describe the practical approach, which will be used in the following concept.

2.4.1 Mathematical Foundations

The following definitions form the mathematical building blocks for DAA. It is noteworthy that these definitions work with RSA encryption as well as with *Elliptic Curve Cryptography* (ECC).

Discrete Logarithm Problem

Given a cyclic group $G = \langle g \rangle$ of order n , the discrete logarithm of $y \in G$ to the base g is the smallest positive integer α satisfying $g^\alpha = y$ if this x exists. For sufficiently large n and properly chosen G and g , it is infeasible to compute the reverse $\alpha = \log_g y$. This problem is known as *Discrete Logarithm Problem* and is the basis for the following cryptographic algorithms.

Signature Proof of Knowledge (SPK)

A SPK is a signature of a message which proves that the creator of this signature is in possession of a certain secret. The secret itself is never revealed to any other party. Thus, this algorithm is a *Zero Knowledge Proof of Knowledge* (ZPK).

Camenisch and Stadler [6] introduced the algorithm based on the Schnorr Signature Scheme. It only assumes a collision resistant hash function $\mathcal{H} : \{0, 1\}^* \rightarrow \{0, 1\}^k$ for signature creation. For instance,

$$SPK\{(\alpha) : y = g^\alpha\}(m)$$

denotes a proof of knowledge of the secret α , which is embedded in the signature of message m . The one-way protocol consists of three procedures:

1. *Setup*. Let m be a message to be signed, α be a secret and $y := g^\alpha$ be the corresponding public representation.
2. *Sign*. Choose a random number r and create the signature tuple (c, s) as

$$c := \mathcal{H}(m || y || g || g^r) \quad \text{and} \quad s := r - c\alpha \pmod{n}.$$

3. *Verify.* The verifier knows the values of y and g , as they are usually public. The message m comes with the signature values c and s . He computes the value

$$c' := \mathcal{H}(m || y || g || g^s y^c) \quad \text{and verifies, that } c' = c.$$

The verification holds since

$$g^s y^c = g^r g^{-c\alpha} g^{c\alpha} = g^r.$$

This scheme is extensible to prove knowledge of an arbitrary number of secrets as well as more complex relations between secret and public values.

Bilinear Maps

Bilinear Maps define a special property for mathematical groups which form the basis for verifying the signatures in DAA. Consider three mathematical groups $\mathbb{G}_1, \mathbb{G}_2$, with their corresponding base points g_1, g_2 , and $\mathbb{G}_T$. Let $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ that satisfies three properties [4, 5]:

- *Bilinearity.* For all $P \in \mathbb{G}_1, Q \in \mathbb{G}_2$, for all $a, b \in \mathbb{Z} : e(P^a, Q^b) = e(P, Q)^{ab}$.
- *Non-degeneracy.* For all generators $g_1 \in \mathbb{G}_1, g_2 \in \mathbb{G}_2 : e(g_1, g_2)$ generates $\mathbb{G}_T$.
- *Efficiency.* There exists an efficient algorithm that outputs the bilinear group $(q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2)$ and an efficient algorithm for computing e .

Camenisch-Lysyanskaya Signature Scheme

The Camenisch-Lysyanskaya (CL) Signature Scheme [5] is based on the LRSW assumption and allows efficient proofs for signature possession and is the basis for the DAA scheme discussed below. It is based on a bilinear group $(q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2)$ that is available to all steps in the protocol.

- *Setup.* Choose $x \leftarrow \mathbb{Z}_q$ and $y \leftarrow \mathbb{Z}_q$ at random. Set the secret key $sk \leftarrow (x, y)$ and the public key $pk \leftarrow (g_1^x, g_2^y) = (X, Y)$.

- *Sign.* Given a message m and the secret key sk , choose a at random and output the signature $\sigma \leftarrow (a, a^y, a^{x+ym}) = (a, b, c)$.
- *Verify.* Given message m , signature σ and public key pk , verify that $a \neq 1_{G_1}$, $e(a, Y) = e(b, g_2)$ and $e(a, X) \cdot e(b, X)^m = e(c, g_2)$.

Camenisch et al. stated in section 4.2 of their paper [4] that one has to verify the equation against $e(g_1, b)$ and $e(g_1, c)$ which is not correct.

2.4.2 DAA Protocol on LRSW Assumption

DAA is a group signature protocol, which aims with a supporting TPM to reveal no additional information about the signing host besides content and validity of the signed message m . According to Camenisch et al. [4], the DAA protocol consists of three parties:

- *Issuer $\mathcal{I}$.* The issuer maintains a group and has evidence of hosts that are members in this group.
- *Host $\mathcal{H}$.* The host creates a platform with the corresponding TPM $\mathcal{M}$. Membership of groups are maintained by the TPM. Only the key owner (TPM, passive) and the message author (Host, active) form a full group member.
- *Verifier $\mathcal{V}$.* A verifier can check, whether a host with its TPM is in a group or not. Besides the group membership, no additional information is provided.

A certificate authority $\mathcal{F}_{ca}$ is providing a certificate for the issuer itself. The basename bsn is some clear text string, whereas nym represent the encrypted basename $bsn^{g^{sk}}$. $\mathcal{L}$ is the list of registered group members which is maintained by $\mathcal{I}$. The paper of Camenisch et al. [4] introduces further variables that are necessary for their proof of correctness. These extensions were omitted in the following to understand the protocol more easily.

- *Setup.* During setup, $\mathcal{I}$ is generating the issuer secret key isk and the corresponding issuer public key ipk . The public key is published and assumed to be known to everyone.

1. On input $SETUP, \mathcal{I}$

- generates $x, y \leftarrow \mathbb{Z}_q$ and sets $isk = (x, y)$ and $ipk \leftarrow (g_2^x, g_2^y) = (X, Y)$. Initialize $\mathcal{L} \leftarrow \emptyset$,
 - generates a proof $\pi \xleftarrow{\$} SPK\{(x, y) : X = g_2^x \wedge Y = g_2^y\}$ that the key pair is well formed,
 - registers the public key (X, Y, π) at $\mathcal{F}_{ca}$ and stores the secret key, and
 - outputs SETUPDONE.
- *Join.* When a platform, consisting of host $\mathcal{H}_j$ and TPM $\mathcal{M}_i$, wants to become a member of the issuer's group, it joins the group by authenticating to the issuer $\mathcal{I}$.
 1. On input JOIN, host $\mathcal{H}_j$ sends the message JOIN to $\mathcal{I}$.
 2. Upon receiving JOIN from $\mathcal{H}_j$, $\mathcal{I}$ chooses a fresh nonce $n \leftarrow \{0, 1\}^\tau$ and sends it back to $\mathcal{H}_j$.
 3. Upon receiving n from $\mathcal{I}$, $\mathcal{H}_j$ forwards n to $\mathcal{M}_i$.
 4. $\mathcal{M}_i$ generates the secret key:
 - Check that no completed key record exists. Otherwise, it is already a member of that group.
 - Choose $gsk \xleftarrow{\$} \mathbb{Z}_q$ and store the key as $(gsk, \perp)$.
 - Set $Q \leftarrow g_1^{gsk}$ and compute $\pi_1 \xleftarrow{\$} SPK\{(gsk) : Q = g_1^{gsk}\}(n)$.
 - Return (Q, π_1) to $\mathcal{H}_j$.
 5. $\mathcal{H}_j$ forwards JOINPROCEED(Q, π_1) to $\mathcal{I}$.
 6. Upon input JOINPROCEED(Q, π_1), $\mathcal{I}$ creates the CL credential:
 - Verify that π_1 is correct.
 - Add $\mathcal{M}_i$ to $\mathcal{L}$.
 - Choose $r \xleftarrow{\$} \mathbb{Z}_q$ and compute $a \leftarrow g_1^r, b \leftarrow a^y, c \leftarrow a^x \cdot Q^{rxy}, d \leftarrow Q^{ry}$.
 - Create the prove $\pi_2 \xleftarrow{\$} SPK\{(t) : b = g_1^t \wedge d = Q^t\}$.
 - Send APPEND(a, b, c, d, π_2) to $\mathcal{H}_j$

7. Upon receiving $\text{APPEND}(a, b, c, d, \pi_2), \mathcal{H}_j$
 - verifies that $a \neq 1, e(a, Y) = e(b, g_2)$ and $e(c, g_2) = e(a \cdot d, X)$, and
 - forwards (b, d, π_2) to $\mathcal{M}_i$.
 8. $\mathcal{M}_i$ receives (b, d, π_2) and verifies π_2 . The join is completed after the record is extended to $(gsk, (b, d))$. $\mathcal{M}_i$ returns **JOINED** to $\mathcal{H}_j$.
 9. $\mathcal{H}_j$ stores (a, b, c, d) and outputs **JOINED**.
- *Sign.* After joining the group, a host $\mathcal{H}_j$ and TPM $\mathcal{M}_i$ can sign a message m with respect to basename bsn .
 1. Upon input $\text{SIGN}(m, bsn), \mathcal{H}_j$ re-randomizes the CL credential:
 - Retrieve the join record (a, b, c, d) and choose $r \xleftarrow{\$} \mathbb{Z}_q$.
Set $(a', b', c', d') \leftarrow (a^r, b^r, c^r, d^r)$.
 - Send (m, bsn, r) to $\mathcal{M}_i$ and store (a', b', c', d') .
 2. Upon receiving $(m, bsn, r), \mathcal{M}_i$
 - checks, that a complete join record $(gsk, (b, d))$ exists, and
 - stores (m, bsn, r) .
 3. $\mathcal{M}_i$ completes the signature after it gets permission to do so.
 - Retrieve group record $(gsk, (b, d))$ and message record (m, bsn, r) .
 - Compute $b' \leftarrow b^r, d' \leftarrow d^r$.
 - If $bsn = \perp$ set $nym \leftarrow \perp$ and compute
 $\pi \xleftarrow{\$} \text{SPK}\{(gsk) : d' = b'^{gsk}\}(m, bsn)$.
 - If $bsn \neq \perp$ set $nym \leftarrow H_1(bsn)^{gsk}$ and compute
 $\pi \xleftarrow{\$} \text{SPK}\{(gsk) : nym = H_1(bsn)^{gsk} \wedge d' = b'^{gsk}\}(m, bsn)$.
 - Send (π, nym) to $\mathcal{H}_j$.
 4. $\mathcal{H}_j$ assembles the signature $\sigma \leftarrow (a', b', c', d', \pi, nym)$ and outputs **SIGNATURE** (σ) .

- *Verify*. Given a signed message, everyone can check, whether the signature with respect to bsn is valid and the signer is member of this group. Furthermore, a revocation list RL holds the private keys of corrupted TPMs, whose signatures are no longer accepted.

Upon input $VERIFY(m, bsn, \sigma), \mathcal{V}$

- parses $\sigma \leftarrow (a, b, c, d, \pi, nym)$,
 - verifies π with respect to (m, bsn) and nym if $bsn \neq \perp$,
 - checks that $a \neq 1, b \neq 1, e(a, Y) = e(b, g_2)$ and $e(c, g_2) = e(a \cdot d, X)$,
 - checks that for every $gsk_i \in RL : b^{gsk_i} \neq d$,
 - sets $f \leftarrow 1$ if all test pass, otherwise $f \leftarrow 0$, and
 - outputs $VERIFIED(f)$.
- *Link*. After proving validity of the signature, the verifier can test, whether two different messages with the same basenamespace $bsn \neq \perp$ are generated from the same TPM.

On input $LINK(\sigma, m, \sigma', m', bsn), \mathcal{V}$ verifies the signatures and compares the pseudonyms contained in σ, σ' :

- Check that $bsn \neq \perp$ and that both signatures σ, σ' are valid.
- Parse the signatures $\sigma \leftarrow (a, b, c, d, \pi, nym), \sigma' \leftarrow (a', b', c', d', \pi', nym')$.
- If $nym = nym'$, set $f \leftarrow 1$, otherwise $f \leftarrow 0$.
- Output $LINK(f)$.

Camenisch et al. [4] extend the general group concept scheme with their concept. The feature of linking messages together requires further security features within the DAA scheme, which the authors also prove in their paper along with the other properties of the scheme:

- *Non-frameability*: No one can create signatures that the platform never signed, but that link to messages signed from that platform.

- *Correctness of link*: Two signatures will link when the honest platform signs it with the same basename.
- *Symmetry of Link*: It does not matter in which order the linked signatures will be proven. The link algorithm will always output the same result.

2.4.3 Standardization of DAA

The *Fast IDentity Online* Alliance (FIDO) is an organization which standardizes online authentication algorithms. When the first generation of TPMs were available, the consortium defined a standard for Direct Anonymous Attestation with Elliptic Curve cryptography (ECDAA). When the newer standard, TPM 2.0, was published, FIDO decided to update their algorithm to be compatible with recent developments. This standard is still in development; a draft version from February 2018 is published on the FIDO website[1] This standard implements a close variant of the previously described concept.

3 Concept

In this chapter we define the constraints for the *Biometric Sensor* (BS) as well as a generic attempt for a prototype. The constraints include a discussion about the attack vectors to the BS. We explain which requirements can and will be addressed and how sensitive data is processed in the BS.

3.1 Definition of the Biometric Sensor

The BS itself is defined as edge device within the Digidow network. According to the schema shown in ??, the BS will be placed in a public area (e.g. a checkpoint in an airport or as access control system at a building) to interact directly with the Digidow users. There, the BS acts as interface to the Digidow network. By providing a biometric property, the user should be able to authenticate itself and the network may then trigger the desired action, like granting access or logging presence. Depending on the biometric property, the sensor may not be active all the time, but activated when an authentication process is started.

The following enumeration shows the steps of the BS for identifying the interacting person.

1. *Listen*: Either the sensor hardware itself (e.g. a detection in a fingerprint sensor) or another electrical signal will start the authentication process.
2. *Collect*: Measure sensor data (picture, fingerprint) and calculate a biometric representation (attribute).
3. *Discover*: Start a network discovery in the Digidow network and find the PIA corresponding to the present person. It may be necessary to interact with more than one PIA within this and the next steps.

4. *Transmit*: Create a trusted and secure channel to the PIA and transmit the attribute.
5. *Reset*: Set the state of the system as it was before this transaction.

Since the BS handles biometric data—which must be held confidential outside the defined use cases—a number of potential threats must be considered when designing the BS.

3.2 Attack Vectors and Threat Model

As mentioned before, the BS will work in an exposed environment. Neither the user providing biometric data nor the network environment should be trusted for proper function. There should only be a connection to the Digidow network for transmitting the recorded data. This assumption of autonomy provides independence to the probably diverse target environments and use cases.

In addition to autonomy, the BS should also ensure proper handling of received and generated data. The recorded dataset from a sensor is *sensitive data* due to its ability to identify an individual. Due to its narrow definition, it is affordable to protect sensitive data. Besides that, *metadata* is information generated during the whole transaction phase. Timestamps and host information are metadata as well as connection lists, hashes and log entries and much more (What? Where? When?). There exists no exact definition or list of metadata which makes it hard to prevent any exposure of it. Metadata does not directly identify an individual. However huge network providers are able to combine lots of metadata to traces of individuals. Eventually an action of those traced individuals might unveil their identity. Consequently, a central goal of Digidow is to minimize the amount to minimize the risk of traces.

Privacy defines the ability of individuals to keep information about themselves private from others. In the context of the BS, this is related to the recorded biometric data. Furthermore, to prevent tracking, any interaction with a sensor should not be matched to personal information. Only the intended and trusted way of identification within the Digidow network should be possible.

When a BS is working in production, it will usually be in an exposed environment. There may be an instance of a verifier next to the BS. The connection into the Digidow network may, however, be based on untrusted networks. Furthermore the physical environment may

not be trustworthy. Given this environment, there are a number of threats that need to be considered when building a BS:

- *Rogue Hardware Components*: Modified components of the BS could, depending on their contribution to the system, collect data or create a gateway to the internal processes of the system. Although the produced hardware piece itself is fine, the firmware on it is acting in a malicious way. This threat addresses the manufacturing and installation of the system.
- *Hardware Modification*: Similar to rogue hardware components, the system could be modified in the target environment by attaching additional hardware. With this attack, adversaries may get direct access to memory or to data transferred from or to attached devices.
- *Metadata Extraction*: The actual sensor like camera or fingerprint sensor is usually attached via USB or similar cable connection. It is possible to log the protocol of those attached devices via Man-in-the-Middle attack on the USB cable.
- *Attribute Extraction*: Similar to metadata extraction, the adversary might directly access the attributes via wiretapping the USB cable. The adversary might be able to identify an individual with those attributes.
- *Modification or aggregation of sensitive data within BS*: The program which prepares the sensor data for transmission could modify the data before sealing it. The program can also just save the sensitive data for other purposes.
- *Metadata extraction on network*: During transmission of data from the sensor into the Digidow network, there will be some metadata generated. An adversary could use these datasets to generate tracking logs and eventually match these logs to individuals.
- *Replay of sensor data of a rogue BS*: When retransmitting sensor data, the authentication of an individual could again be proven. Any grants provided to the successfully identified individual could then be given to another person.
- *Rogue Biometric Sensor blocks transmission*: By blocking any transmission of sensor data, any transaction within the Digidow network could be blocked and therefore the whole authentication process is stopped.

- *Rogue Personal Identity Agent*: A rogue PIA might receive the sensor data instead of the honest one. Due to this error, a wrong identity and therefore false claims would be made out of that.

Although all of these attack vectors should be mitigated when used in production, we will address only a subset for the prototype. First, we assume that only authorized personnel has access to the hardware itself. Any other person should only interact with the hardware sensor. Therefore any threat attacking communication between internal system components will not be addressed. Furthermore, we will assume an already established bidirectional channel between BS and PIA. Any algorithms on how the BS finds the corresponding PIA exceed the focus of this work.

On the other hand, any hardware modification including firmware upgrades should be detectable by the prototype system. In addition to this detection the BS should generate trust by having a provable system state before a Digidow transaction is executed. This includes mitigations against attacks on replaying attributes, blocking attribute transmission or aggregating them while running.

3.3 Prototype Concept

Given the threat model and the use cases described in section 3.1, we will introduce a prototype which will address mainly the cryptographic aspects of trust. Furthermore we will discuss system integrity approaches and their limitations.

3.3.1 Integrity and Trust up to the Kernel

We decided to use the PC platform as hardware base for the prototype. There are lots of different form factors available and you can extend the system with a broad variety of sensors. Furthermore, the platform provides full TPM support to enable cryptographic and integrity features. Finally, the platform can run almost all Linux variants and supports relevant pieces of software for this project. A flavour of Linux supporting all features described in this chapter, will be used as OS platform. The ARM platform seems to be capable of all these features as well. However, the support for TPMs, the amount of available software and the ease of installation is better on the PC platform.

As described in section 2.1, the TPM functions can be delivered in three different flavors: As dedicated or mounted device and as part of the platform firmware. The fTPM is part of larger proprietary environments from AMD and Intel which introduces, besides implementation flaws, additional attack surfaces for the TPM. Hence, we will use plugged TPM chips on the platform. Then we are able to deactivate the TPM for demonstration purposes by simply unplugging it.

Any recent PC platform supports TPMs and consequently trusted boot as mentioned in section 2.2. The system will describe its hardware state in the PCRs 0–7 when the EFI/BIOS hands over to the bootloader. We use these PCR values to detect any unauthorized modifications on hardware or firmware level. It is important to also include *empty* PCRs to detect added hardware on the PCI bus with an Option ROM, for example.

With these PCR values we can seal a passphrase in the TPM. The disk, secured with Full Disk Encryption (FDE), can only be accessed when the hardware underneath is not tampered with.

To further reduce the attack surface, the prototype will not use a bootloader like GRUB. Instead, the kernel is run directly from the UEFI/BIOS. Therefore, the kernel is packed directly into an EFI file, together with its command line parameters and the initial file system for booting. This *Unified Kernel* is directly measured by the UEFI/BIOS and is also capable of decrypting the disk, given the correct PCR values.

This setup starts with two sources of trust that are formally defined:

- *TPM*: The TPM acts as certified Root of Trust for holding the PCRs and for the cryptographic function modifying those.
- *RTM*: The Root of Trust for Measurement is part of the mainboard firmware. The tiny program just measures all parts of the firmware and feeds the TPM with the results. However, the program is maintained by the mainboard manufacturer and the source is not available to the public. We have to trust that this piece of software is working correctly.

We implicitly assume that the CPU, executing all these instructions and interacting with the TPM, is working correctly.




Figure 3.1: Extending trust from the Roots of Trust up to the kernel

All parts contributing to the boot phase will be measured into one of the PCRs before any instruction is executed. Decrypting the disk can then be interpreted as authorization procedure against the encrypted disk. Consequently, only a *known* kernel with a *known* hardware and firmware setup underneath can access the disk and finish the boot process in the OS.

The disk encryption is, however, only an optional feature which can be omitted in a production environment when there is no sensitive data on the disk that must not be revealed to the public. The system needs to check its integrity on the OS level and summarize that by publishing an attestation message, before any transaction data is used.

Figure 3.1 illustrates how above processes extend the trust on the system. The TPM is the cryptographic root of trust, storing all measurement results and the target values for validation. Since the RTM is the only piece of code which lives in the platform firmware and is executed *before* it is measured, it is an important part in the trust architecture of the system. An honest RTM will measure the binary representation of itself, which makes

the code at least provable afterwards. Finally, the CPU is assumed to execute all the code according to its specification. Proving correctness of the instruction set cannot be done during the boot process.

When the roots of trust are honest, the trusted environment can be constructed during booting the platform with the PCR measurements. We get a trusted boot chain from firmware up to the kernel with its extensions and execution parameters as a result.

3.3.2 Integrity and Trust on OS Level

With the trusted kernel and IMA, we can include the file system into the trusted environment. According to section 2.3, every file will be hashed once IMA is activated and configured accordingly. By enforcing IMA, the kernel allows access to only those files having a valid hash. Consequently, every file which is required for proper execution needs to be hashed beforehand, i.e. before IMA is enforced. The IMA policy in place should be `appraise_tcb`, to analyze kernel modules, executable memory mapped files, executables and all files opened by root for read. This policy should also include drivers and kernel modules for external hardware like a camera for attached via USB.

3.3.3 Proving Trust with DAA

The features described above take care of building a trusted environment on the system level. DAA will take care of showing the *trust* to a third party which has no particular knowledge about the BS. In the Digidow context, the PIA should get, together to the biometrical measurements, a proof that the BS is a trusted system acting honestly.

To reduce the complexity of this problem, we consider two assumptions:

1. *Network Discovery*: The PIA is already identified over the Digidow network and there exists a bidirectional channel between BS and PIA
2. *Secure Communication Channel*: The bidirectional channel is assumed to be hardened against wire tapping, metadata extraction and tampering. The prototype will take no further action to encrypt any payload besides the cryptographic features that come along with DAA itself.




Figure 3.2: The DAA attestation process requires 5 steps. The PIA may trust the BS afterwards.

For the scope of this thesis, the DAA protocol should be applied on a simple LAN, where all parties are connected locally. The BS will eventually become a member of the group of sensors managed by the issuer. During signup, issuer and BS (member) negotiate the membership credentials over the network. By being a member of the DAA group, the issuer fully trusts that the BS is honest and acting according the specification. The issuer will not check any group members, since they can now act independently of the issuer.

When the BS is then authenticating an individual, the process illustrated in Figure 3.2 will be executed.

1. The PIA gets the public key of the BS group once and independently of any transaction.
2. During the transaction, the PIA will eventually ask the BS for attestation together with a nonce.
3. The BS will collect the PCR values, the integrity log and the nonce into an attestation message signed with the member's secret key (SK).
4. The attestation message will be sent back to the PIA.




Figure 3.3: Overview of the Chain of Trust of the BS

5. The PIA checks the signature of the message, checks the entries of the integrity log against known values, and proves the PCR values accordingly.

Figure 3.3 shows how the sources of trust will be represented in the final attestation message. The four sources of trust are defined as groups which deliver parts of the prototype, but cannot be verified on a cryptographic level. Hence, suppliers must be manually added to these groups by using a well defined check for trustworthiness. For example, any TPM manufacturer has to implement the well defined standard from TCG. There exists, however, no such exact definition for hardware and firmware parts of the platform. Consequently, these parts should undergo a functional analysis before they receive a trust certificate. Trust means that, when the platform is defined trustworthy, the corresponding PCR values should be published to be verifiable when Digidow transactions occur.

The same procedure should be done for the kernel and the used OS environment and of course, the used software. There, only the kernel with its parameters have a corresponding

PCR value. Furthermore, a hash value should be published for any relevant file on the file system.

We can then build a cryptographic representation of the chain of trust in Figure 3.3. The TPM has a signed certificate from its manufacturer, where it derives the endorsement key (EK) from. When all of the above checks against platform, OS and TPM are good, the DAA issuer will assign the platform to the group of trusted BS. The BS has now a member SK for signing its attestation message.

The verifier can now check the valid membership by checking the signature of the message against the issuer's public key (PK). Furthermore, it can check the state of the platform by comparing the PCR values against known values. Finally, it can check the integrity of the running software by checking the hashes in the IMA log against known values. PCR 10 represents the end of the hash chain fed by the IMA log entries.

If all values are good, the BS can be trusted and the Digidow transaction can be continued at the PIA.

4 Implementation

The concept described in chapter 3 will be implemented as a prototype to demonstrate a working implementation and to analyze the speed of those parts of a transaction. Although the goal is to put all these features on a highly integrated system, we decided to start with widely available hardware based on Intel's x86 architecture.

Figure 4.1 shows the setup on a connection level. To show the features of DAA, it is necessary to have three independent systems which are connected via a TCP/IP network. Every host is connected via ethernet to the other systems. To keep the setup minimal, the IP addresses are static and internet is only required during installation. Hence, Service Discovery is done statically, every host knows the IP addresses and functions of each other directly.




Figure 4.1: Prototype setup to show DAA features and the Dataflow from BS to PIA

4.1 Hardware Setup

For demonstrating remote attestation via DAA over a simple network infrastructure, we use three systems with similar configuration. Table 4.1 shows the specification of these systems. We decided to order one system with an AMD processor in it to find differences in handling the TPM between Intel and AMD systems. All features used in this thesis were available on both platform types, so there were no differences found.

The used mainboards come with a dedicated TPM2.0 header which may differ from board to board. A 19-pin header is available on the older platform of *System 2*. As long as TPM and mainboard have the same 19-pin connector they will be compatible to each other. The newer Gigabyte mainboards come with a proprietary 11-pin connector which is only compatible with Gigabyte's TPM2.0_S module. All other modules are however electrical compatible since only unused pins of the full size connector are removed. With a wiring adapter any TPM board would work on any mainboard supporting TPM2.0 even when coming with a proprietary header.

Table 4.1: Systems used for demonstration prototype

	<i>System 1</i>	<i>System 2</i>	<i>System 3</i>
Processor	AMD Athlon 240GE	Intel Pentium G4560T	Intel Pentium G5400T
Mainboard	Gigabyte B450I Aorus Pro Wifi	Gigabyte GA H110N	Gigabyte GA H310N
Memory	8GB DDR4	8GB DDR4	8GB DDR4
Storage	NVMe SSD 128GB	NVMe SSD 128GB	NVMe SSD 128GB
TPM module	Gigabyte TPM2.0_S	Gigabyte TPM2.0	Gigabyte TPM2.0_S
TPM chip	Infineon SLB9665TT2.0	Infineon SLB9665TT2.0	Infineon SLB9665TT2.0

4.2 Select the Operating System

The OS needs to fulfill three requirements for this prototype. First, the TPM must be supported by the kernel. Second, the OS has to support a recent version of the TPM software stack (TSS 3.0.x or newer at the point of writing) for using the Xaptum ECDA[14] project with enabled hardware TPM. Similarly, the `tpm2-tools` must be available in a version newer than 4.0.0. Finally, the support for the Integrity Measurement Architecture (IMA) must be activated in the kernel and supported by the OS. This feature is available in the mainline Linux kernel. However, the corresponding kernel compile parameters must be set.

Ubuntu 20.04 LTS does fulfill above mentioned requirements by default. Ubuntu is also supported by the Xaptum ECDA project, although it was tested with an older version (18.04). When installing Ubuntu on the prototype, we used *Full Disk Encryption* (FDE) which leads to the disk allocation described in Table 4.2. Ubuntu installs Grub by default, and it is used as a fallback bootloader.

<i>Partition</i>	<i>Size</i>	<i>Mountpoint</i>	<i>Comment</i>
nvme0n1p1	512M	/boot/efi	EFI boot partition
nvme0n1p2	1G	/boot	Bootloader partition (Grub)
nvme0n1p3	118G		lvm on dm_crypt
ubuntu-vg-ubuntu-lv	118G	/	root partition on lvm

Table 4.2: Disk layout of the BS prototype

4.3 Trusted Boot

By default, every mainboard with support for TPM2.0 must support trusted boot. When a TPM becomes available, the UEFI/BIOS itself takes all required measures until the boot process is handed over to the OS bootloader (e.g. GRUB). Since Ubuntu uses GRUB 2.04 as bootloader which has TPM support by default, trusted boot needs just to be enabled in the GRUB configuration. In this case, GRUB will be measured from the BIOS to the PCRs 4 and 5, as shown in Table 2.1. According to the documentation[9], Grub itself uses PCR 8 for executed commands, the kernel command line and all commands forwarded to kernel modules. PCR 9 is used to measure all files read by GRUB.

Embedded systems like a productive version of the BS do not need several boot options, making Grub unnecessary. Therefore we can replace Grub's bootloader EFI file with a blob containing all required information to load the kernel directly. This kernel decrypts the disk and boots the remaining system autonomously. Pawit Pornkitprasam [19], [20] and the Tevora company [23] introduced the concept of a *Unified Kernel* for Ubuntu and Arch respectively.

This large EFI file contains the initramfs, kernel command line and the kernel itself. Table 4.3 shows the content of the EFI blob with the corresponding offset addresses as well as the sources in the file system.

All binary resources are available as blobs which can be imported directly. Only the command line parameters need to be customized.

<i>Address</i>	<i>Source path</i>	<i>Comment</i>
0x00000000	/usr/lib/systemd/boot/efi/linuxx64.efi.stub	Linux EFI Stub
0x00200000	/usr/lib/os-release	Linux OS release information
0x00300000	/boot/kernel-command-line.txt	Kernel command line parameters
0x00400000	/boot/vmlinuz	latest kernel image
0x30000000	/boot/initrd	latest initial ramdisk

Table 4.3: Memory layout of the Unified Kernel EFI file

4.3.1 Install and use Trusted Boot

The following shell scripts are available in the folder `trustedboot`. They are tested on a Ubuntu 20.04 server minimal installation on all three devices. These packages need to be installed beforehand on an Ubuntu minimal instance to install trusted boot:

- `binutils` for `objcopy`, generating unified kernel, and
- `tpm2-tools` to interact with the onboard TPM,

Installing trusted boot is done in three steps, assuming being root at the target system:

1. *Execute `install.sh`*: It installs the shell scripts into `/usr/sbin` and adds a new random passphrase to LUKS. It furthermore adds TPM support to the `initramfs` and creates the unified kernel described above.
2. *Reboot*: During reboot, the new PCR values are generated.
3. *Execute `update-luks-tpm.sh`*: The new PCR values are used to seal the LUKS passphrase into the TPM.

After finishing these steps, the system should be able to decrypt the FDE without user interaction and end up at the login prompt. The automatic boot will now work as long as there are no updates affecting the unified kernel. If a kernel update happens during upgrading the system, a new EFI blob must be generated and the TPM seal is to be renewed. Similar to the installation procedure, this can be done in the following procedure:

1. *Execute `update-kernel.sh`*: This includes the latest updates of kernel, its parameters and the `initramfs` into the EFI blob.
2. *Reboot*: Again, the new PCR values need to be calculated.

Listing 4.1: kernel-command-line.txt: Command line for the Kernel

```
1 /vmlinuz ima_appraise=fix ima_policy=appraise_tcb ima_policy=tcb ima_hash=sha256 root
   =/dev/mapper/ubuntu--vg-ubuntu--lv ro rootflags=i_version
```

Listing 4.2: passphrase-from-tpm.sh: Initramfs-script to ask the TPM for the LUKS key

```
1 #!/bin/sh
2 echo "Unlocking_via_TPM" >&2
3 export TPM2TOOLS_TCTI="device:/dev/tpm0"
4 /usr/bin/tpm2_unseal -c 0x81000000 -p pcr:sha256:0,1,2,3,4,5,6,7
5 if [ $? -eq 0 ]; then
6     exit
7 fi
8 /lib/cryptsetup/askpass "Unlocking_the_disk_fallback_${CRYPTTAB_SOURCE_(
   $CRYPTTAB_NAME)}\nEnter_passphrase:_"
```

3. *Execute update-luks-tpm.sh*: Similar to the installation, the new seal is generated and replaces the old one.

The result is a recent kernel with up-to-date configuration.

4.3.2 Detailed description of the scripts

Listing 4.1 shows the used command line which will be saved on `/boot/kernel-command-line.txt`. Relevant for trusted boot is the parameter where the root filesystem is located. This example works for the default Ubuntu LVM on LUKS configuration. The parameters activate also IMA which will be discussed later in this chapter.

If FDE is installed, the boot process need to be aware of how to decrypt the disk. Therefore, the initramfs needs the luks binaries as well as the TPM software stack to unseal the passphrase with the PCR registers. The unseal operation itself is then done with Listing 4.2, which also needs to exist in the initramfs. We copy the script of Listing 4.3 to `/etc/initramfs-tools/hooks` to enable TPM access during boot after the next initramfs update. Next, the script of Listing 4.5 creates a new key for FDE by using the random number generator of the TPM. It is saved in clear text in `/root/keys` to be able to update the sealing operation when new PCR values are used. Using a clear text password file in the `/root` directory is not per se considered insecure as the disk must be decrypted beforehand, meaning that the passphrase also exists unencrypted in memory. Furthermore

Listing 4.3: tpm2-hook.sh: Script copying required TSS files into the initramfs

```

1 #!/bin/sh -e
2 if [ "$1" = "prereqs" ]; then exit 0; fi
3 . /usr/share/initramfs-tools/hook-functions
4 copy_exec /usr/bin/tpm2_unseal
5 copy_exec /usr/lib/x86_64-linux-gnu/libtss2-tcti-device.so.0
6 copy_exec /usr/sbin/passphrase-from-tpm.sh

```

Listing 4.4: update-kernel.sh: Script for updating the unified Kernel

```

1 #!/usr/bin/bash
2 set -e
3 PARTITION_ROOT=/dev/mapper/ubuntu--vg-ubuntu--lv
4 DISK=/dev/nvme0n1
5
6 mkdir -p /boot/efi/EFI/Linux
7 update-initramfs -u -k all
8 LATEST=`ls -t /boot/vmlinuz* | head -1`
9 VERSION=`file -bL $LATEST | grep -o 'version_[^_]*' | cut -d '_' -f 2`
10 objcopy \
11 --add-section .osrel="/usr/lib/os-release" --change-section-vma .osrel=0x20000 \
12 --add-section .cmdline="/boot/kernel-command-line.txt" --change-section-vma .cmdline
   =0x30000 \
13 --add-section .linux="/boot/vmlinuz-$VERSION" --change-section-vma .linux=0x40000 \
14 --add-section .initrd="/boot/initrd.img-$VERSION" --change-section-vma .initrd=0
   x3000000 \
15 "/usr/lib/systemd/boot/efi/linuxx64.efi.stub" "/boot/efi/EFI/Linux/Linux.efi"

```

only root has access to the target directory. The backup is needed for updating the system and when issues with the TPM appear. When an updated kernel is booted, the seal cannot be opened anymore since the old PCR values cannot be restored anymore. It is, of course possible to generate a new passphrase every time when the PCR values change intendedly. In this case no passphrase need to be stored on the file system.

Finally, Listing 4.4 creates the unified kernel according to Table 4.3 using the command `objcopy` and copies it on the EFI disk partition. The offset addresses need to be chosen according to the size of the included blobs. All steps described above are summarized in Listing 4.6.

When the unified kernel is installed, the system needs to be rebooted to generate the PCR values for the new boot chain. The FDE decryption with the TPM will of course fail, since there is no sealed passphrase available yet. This step is done now since the new unified

Listing 4.5: create-luks-tpm.sh: Script to create a new LUKS key

```

1 #!/bin/bash
2 set -e
3
4 CRYPTFS=/dev/nvme0n1p3
5
6 echo "creating_secret_key"
7 mkdir -p /root/keys
8 tpm2_getrandom 32 -o /root/keys/fde-secret.bin
9 chmod 600 /root/keys/fde-secret.bin
10 cryptsetup luksAddKey $CRYPTFS /root/keys/fde-secret.bin
11
12 # /usr/sbin/update-luks-tpm.sh #not required before reboot

```

Listing 4.6: install.sh: Script to install Trusted Boot on Ubuntu

```

1 #!/bin/bash
2 set -e
3
4 cp -vf ./passphrase-from-tpm.sh /usr/sbin/ || exit 1
5 cp -vf ./update-luks-tpm.sh /usr/sbin || exit 1
6 cp -vf ./update-kernel.sh /usr/sbin || exit 1
7 cp -vf ./create-luks-tpm.sh /usr/sbin || exit 1
8
9 cp -vf ./tpm2-hook.sh /etc/initramfs-tools/hooks/ || exit 2
10 awk -i inplace '/luks/{print_$0_",discard,initramfs,keyscript=/usr/sbin/passphrase-
    from-tpm.sh"}' /etc/crypttab
11
12 cp -vf ./kernel-command-line.txt /boot/ || exit 3
13 /usr/sbin/create-luks-tpm.sh
14 /usr/sbin/update-kernel.sh
15 efibootmgr --create --disk /dev/nvme0n1 --part 1 --label "ubuntu_unified" --loader "\
    EFI\Linux\Linux.efi" --verbose
16 echo "Installed_successfully!_Please_reboot_and_execute_update-luks-tpm.sh_
    afterwards"

```

Listing 4.7: `update-luks-tpm.sh`: Script for updating the Sealing of the TPM Object with new PCR values

```

1 #!/usr/bin/bash
2 echo "Updating_TPM_Policy_with_current_available_PCrs"
3
4 set +e
5 tpm2_evictcontrol -c 0x81000000
6
7 set -e
8 tpm2_flushcontext -t
9 tpm2_createprimary -C e -g sha256 -G ecc256 -c /root/keys/e-primary.context
10 tpm2_createpolicy --policy-pcr -l sha256:0,1,2,3,4,5,6,7 -L /root/keys/pcr-policy.
    digest
11 tpm2_create -g sha256 -u /root/keys/obj.pub -r /root/keys/obj.priv -C /root/keys/e-
    primary.context -L /root/keys/pcr-policy.digest -a "noda|adminwithpolicy|
    fixedparent|fixedtpm" -i /root/keys/fde-secret.bin
12 tpm2_flushcontext -t
13 tpm2_load -C /root/keys/e-primary.context -u /root/keys/obj.pub -r /root/keys/obj.
    priv -c /root/keys/load.context
14 tpm2_evictcontrol -C o -c /root/keys/load.context 0x81000000
15 # tpm2_unseal -c 0x81000000 -p pcr:sha1:0,1,2,3,4,5,6,7 -o /root/test.bin #proof that
    the persistence worked
16 rm -f /root/keys/load.context /root/keys/obj.priv /root/keys/obj.pub /root/keys/pcr-
    policy.digest
17 tpm2_flushcontext -t

```

kernel is measured the first time. Listing 4.7 shows how the passphrase is sealed into the TPM with all relevant PCR registers. Line 15 of the script shows the unseal operation which is used to get the passphrase out of the TPM. It should result in an exact copy of the original passphrase in `/root/keys`. This finishes the installation of trusted boot resulting in an unattended disk decryption process during system boot. The result is a trusted boot chain which ensures, that the system only has access to the encrypted disk when the kernel with its parameter is known—and therefore trusted.

4.4 Integrity Measurement Architecture

The result of trusted boot is a measured and therefore trusted kernel with its command line parameters and modules. IMA extends this trust to the file system as described in section 2.3. All features of IMA are already implemented in the kernel sources and do not

need additional packages. The Gentoo wiki page about IMA[11] describe which kernel compiler flags are related to IMA:

- CONFIG_INTEGRITY=y
- CONFIG_IMA=y
- CONFIG_IMA_MEASURE_PCR_IDX=10
- CONFIG_IMA_LSM_RULES=y
- CONFIG_INTEGRITY_SIGNATURE=y
- CONFIG_IMA_APPRAISE=y
- IMA_APPRAISE_BOOTPARAM=y

According to their blog, Redhat¹ supports IMA since Enterprise Linux 8 and so would the Redhat clones like CentOS, Rocky or Alma Linux. Ubuntu enabled the required kernel compile flags by default on version 20.04 in the ubuntu kernel repository². However, we found no official information on when IMA was introduced or whether IMA support will continue in future releases.

Every kernel supporting IMA creates a virtual file at:

```
/sys/kernel/security/ima/ascii_runtime_measurements
```

When IMA is disabled, which is the default, this file has only one entry representing the boot aggregate. By enabling IMA via kernel command line parameters, this virtual file gets filled according to the policies defined. The first four parameters used in Listing 4.1 define the behavior of IMA and how the measurement log should look like.

- `ima_appraise=fix` appends the filehash as an extended file attribute for every accessed file.
- `ima_policy=appraise_tcb, ima_policy=tcb` analyzes resources owned by root and opened for execution.

¹<https://www.redhat.com/en/blog/how-use-linux-kernels-integrity-measurement-architecture>, last visited on 9.7.2021

²<https://kernel.ubuntu.com/git/kernel-ppa/mirror/ubuntu-5.4-focal.git/tree/security/integrity/ima/Kconfig>, last visited on 9.7.2021

- `ima_hash=sha256` sets the hashing algorithm.
- `rootflags=i_version` must be enabled when mounting the filesystem since IMA is checking that number before re-hashing the resource.

Unfortunately, the resulting IMA log is between 2000 and 3000 lines long when the system is freshly booted. The log can blow up to several 10000 lines when uptime exceeds several days. IMA has therefore a policy language rules for hashing files can be customized. It might, for example, be useful to exclude log files from being hashed.

However, a third party can comprehend the state of the attesting system, when parsing the IMA log. Together with the corresponding value of PCR 10, where the hashes of all entries are chained, the TPM can certify correctness of the log.

Besides the log entries, every hashed file gets an extended file attribute called `security.ima` holding current hash value in base64 encoding. The apt package `attr` contains the binary `getfattr` which can show these additional attributes. Now, the file attributes are shown with:

```
getfattr -m - -d trustedboot/install.sh
```

Having attributes for all relevant files allows IMA to enforce appraisal. This means that the hash of every file is checked before it is accessed. Before enforcing IMA, the root filesystem should be parsed to create attributes for every file. Touching the file is not enough, the file must be opened for read, as shown in the following command:

```
time find / -fstype ext4 -type f -uid 0 -exec dd if='{}' of=/dev/null count=0 status
  =none \;
```

This command opens every file in an ext4 filesystem for read, causing the kernel to create/update the file attributes, but does not read any data. It takes about fifteen minutes to parse the file system on a Ubuntu minimal server system given the hardware described above.

The final step is to enforce IMA appraisal via kernel command line parameter. By changing the parameter `ima_appraise=enforce` in `/boot/kernel-command-line.txt`, IMA enforces the hash and will deny access if the hash is not correct. The new parameters are activated by using the workflow for updating the kernel described in subsection 4.3.1.

The result is a system which allows only resources with correct hash attribute to be read or executed. However, an adversary may still create the file hash itself and overwrite the file with the hash. Hence these hash values must be verified in some way before they can be trusted. Two options are therefore available:

- *Attest*: Let the complete IMA log be part of remote attestation and verify the hash values against a trusted known value database. This makes the verification part very complex but the state of the attesting system is well documented.
- *Sign*: Sign hash values for static files like executables or libraries. The host can directly check whether the file was tampered with and act accordingly. Linux kernels support this feature with the *extended verification module* (EVM).

We use with DAA an attestation variant where the IMA log is part of the message. Details about the message design will be discussed in subsection 4.8.1.

4.5 Dynamic System Analysis

IMA is a comprehensive tool for checking the integrity of a file or executable or library before it gets executed. When access is allowed, IMA's job is done. An executable can however change its behaviour during runtime, either intended with remote procedure calls or unintended with a remote code execution vulnerability.

One approach for runtime analysis is to create a log of syscalls for a certain process. Syscalls are the only interface from the program to the OS when managing resources needed by the program. Hence logging the trace of syscalls show which resources it needs and which files are accessed or modified. Any unusual behavior could be detected immediately given a good knowledge of the used resources when executing a good version of the program.

The trace logs are created with `auditd` which requires the corresponding apt package to be installed.

4.6 Prove TPM2 Certificate Chain

Every TPM has a corresponding certificate which is part of a certificate chain maintained by the TPM manufacturer. In our case, Infineon certifies its TPM with a number of intermediate CAs which itself are certified with Infineon's root CA. The TPM certificate is available for RSA and ECC cryptofamilies respectively. Since the verification workflow is the same on all machines and for both cryptofamilies, we demonstrate on system 1 how the process works. Note that this works for Infineon TPMs. Other Vendors like STM, AMD or Intel may provide certificates via download of their website.

1. Read the certificate from the TPM NVRAM. The RSA certificate is located at address 0x1c000002, that for ECC on address 0x1c00000a:

```
root@amd1:~# tpm2_nvread -C o 0x1c00000a -o amd1_ecc.crt
```

2. Download the certificates from the intermediate and root CA from infineon's website:

```
root@amd1:~# wget https://www.infineon.com/dgdl/Infineon-TPM_ECC_Root_CA-C-
v01_00-EN.cer?fileId=5546d46253f65057015404843f751cdc -O
infineon_ecc_root_ca.crt #Infineon root CA
root@amd1:~# wget https://www.infineon.com/dgdl/Infineon-OPTIGA-ECC-
Manufacturing-CA_036-Issued_by_RootCA.crt-C-v01_00-EN.crt?fileId=5546
d46262475fbc0162486417b73cbe -O infineon_ecc_intermediate_ca_036.crt #
Infineon intermediate CA
```

3. Convert all certificates into PEM format. OpenSSL can only verify a chain in PEM format.

```
root@amd1:~# openssl x509 -inform DER -outform PEM -in infineon_ecc_root_ca.
crt -out infineon_ecc_root_ca.pem
root@amd1:~# openssl x509 -inform DER -outform PEM -in
infineon_ecc_intermediate_ca_036.crt -out
infineon_ecc_intermediate_ca_036.pem
root@amd1:~# openssl x509 -inform DER -outform PEM -in amd1_ecc.crt -out
amd1_ecc.pem
```

4. Check the certificate chain with OpenSSL. The option `-untrusted` is required since the provided root CA is not in the OS' trust store and hence of unknown trust.

```
root@amd1:~# openssl verify -verbose -CAfile infineon_ecc_root_ca.pem -
untrusted infineon_ecc_intermediate_ca_036.pem amd1_ecc.pem
amd1_ecc.pem: OK
```

When OpenSSL returns OK, the certificate chain is intact and the TPM is indeed one from Infineon. To be correct: The website, probably hosted by infineon, provides a certificate chain which matches and the links to the corresponding parent certificate is correct. Unfortunately, Infineon do neither provide any website certification nor any checksums of the provided certificates. So, if the above described check fails, no source of trust can ensure that the root certificate is correct.

We found, that both, ECC and RSA chains of all TPMs are intact. For documentation reasons, we provide the OpenSSL SHA256 checksums of all used certificates in Table 4.4. These checksums were generated with:

```
openssl x509 -noout -fingerprint -sha256 -inform pem -in amd1_ecc.pem
```

Similar to trusted boot, Ubuntu requires two installed packages to support the features discussed in this section:

- `auditd` to analyze system calls of processes, helpful when using IMA, and
- `attr` for showing the file IMA hashes, which are saved as extended file attributes.

IMA-Settings erklären. Available on Ubuntu, RedHat and optionally Gentoo. The kernel has the correct compile options set.

Table 4.4: ECC certificate chain of used TPMs

<i>Certificate</i>	<i>SHA256 checksum</i>
ECC root CA	CF:EB:02:FE:CD:55:AD:7A:73:C6:E1:D1:19:85:D4:C4: 7D:EE:24:8A:B6:3D:CB:66:09:1A:24:89:66:04:43:C3
ECC intermediate CA 036	2C:BF:6D:6C:39:94:47:0A:EC:D2:B1:F8:7F:5D:21:11: 08:D1:E0:A4:A8:1E:D9:CB:A2:6B:D2:98:96:5D:E0:1C
ECC intermediate CA 011	B5:59:A8:4E:63:32:A7:B4:E3:2B:4D:37:39:E3:72:E3: C5:17:BA:5A:4C:CF:FE:E1:DA:AF:80:BC:64:16:28:EE
ECC System 1 TPM	D8:C3:21:3E:BB:C2:CB:96:EA:14:2F:D5:57:61:79:04: BB:DE:8A:F9:21:2A:11:7D:4B:2E:FC:64:BD:35:C1:6E
ECC System 2 TPM	FE:1B:EF:42:8B:68:35:D3:FC:5A:13:A0:AE:12:19:BA: A1:60:D6:59:38:1D:79:8E:76:50:48:BE:5C:BD:83:A5
ECC System 3 TPM	92:29:68:6D:50:EE:34:08:30:DF:E7:30:D8:F3:C0:C7: 13:3C:DF:F0:6D:9E:2B:2E:0F:54:76:AE:B8:D6:1A:DA
RSA root CA	89:9E:35:47:4C:98:07:EB:4C:7F:2F:7A:12:DA:00:28: FB:25:0C:D0:21:54:D0:00:9F:CA:7D:9C:66:57:4F:3B
RSA intermediate CA 036	21:6C:47:D2:77:FC:40:CE:90:F0:86:83:21:CB:5E:F5: 91:FC:1D:D8:D0:E4:FD:A1:A2:C8:3C:17:BE:01:B0:7E
RSA intermediate CA 011	A8:33:79:F9:2A:34:1B:EB:61:B6:7F:03:50:66:44:94: 0F:EB:4B:85:EA:50:4A:9D:22:13:BC:A5:2C:88:9F:89
RSA System 1 TPM	F1:C7:6A:00:CF:2B:63:4D:38:C0:2E:73:3C:84:BF:30: 5C:C2:D3:61:DF:34:D8:95:BB:F1:0F:FB:6B:0C:79:E2
RSA System 2 TPM	CB:1F:7D:20:FE:B2:11:C4:2B:20:6B:4F:66:A6:14:1A: 37:94:5F:85:93:6D:2E:92:85:57:BF:3A:BF:9E:DA:DD
RSA System 3 TPM	BF:0B:4E:77:80:18:86:9A:EF:09:06:96:E2:4D:72:A3: 47:B6:E3:8F:AA:F9:9C:2E:C0:13:AB:70:E3:E4:5D:93

4.7 Using the DAA Protocol

Direct anonymous attestation is a group signature scheme which uses the TPM as crypto-processor and key store. The feature of identifiable instances of sensors is not required when interacting with the Digidow network. Only the *trusted* state of the sensor system and the membership in the corresponding group is relevant. Hence, the group membership is an essential part to provide trust to the users, requiring a deep knowledge on what hardware and software is installed and which vulnerabilities it might have.

The DAA group membership states that the system is provisioned from a trusted party, namely the DAA issuer. The level of trust is ultimate since the used version of DAA is only partly dynamic by lacking support of membership removal. During a Digidow transaction, the sensor attests its state by signing a message containing the IMA log and the PCR registers. Any party interacting with the sensor is then able to check trustworthiness via integrity and valid membership of the sensor.

We describe in the following which programs need to be installed and what configuration is required to demonstrate a working implementation of DAA.

4.7.1 Provision Hosts of Test Setup

The demonstration setup, shown in Figure 4.1, consists of three independent hosts which are connected together via TCP/IP. Every host represents one party in the DAA scheme, each requiring additional software to support the DAA protocol. Xaptums ECDAA library needs to be installed on all three hosts but only the sensor requires TPM support. Similar to that, the ECDAA network wrapper is required to support the network communication part.

The member needs, besides DAA protocol support, software to capture and process an image of the USB webcam. We developed a small Rust program called `sensor-capture` for capturing a face image from a webcam. For biometric processing, we transform the image into an embedding with the face recognition prototype of Digidow³.

³<https://git.ins.jku.at/proj/digidow/prototype-facerecognition>

4.7.2 Installing Xaptum ECDAAs Library

Xaptum's ECDAAs Library provide the cryptographic functions and the protocol primitives for DAA. A file based demonstration of the protocol is provided within the project. We have to build the ECDAAs library from source since the provided deb packages do not support Ubuntu 20.04. Therefore we need the C build environment as follows:

```
root@amd1:~# apt install gcc cmake build-essential doxygen doxygen-latex parallel
```

The sensor host requires TPM support which is enabled with the additional package `libtss2-dev`:

```
root@amd1:~/ecdaa/build# apt install libtss2-dev
```

Download the repository from GitHub and create the build folder on the filesystem:

```
root@amd1:~# git clone https://github.com/xaptum/ecdaa.git
root@amd1:~# mkdir -p ecdaa/build && cd ecdaa/build
```

The next step is to build and install the required dependencies from source. Cmake uses the environment variable `CMAKE_PREFIX_PATH` as installation target. Xaptum provided a shell script for the complete routine:

```
root@amd1:~/ecdaa/build# export CMAKE_PREFIX_PATH=/usr
root@amd1:~/ecdaa/build# ../../travis/install-amcl.sh ./amcl /usr FP256BN
```

Finally install the build of the project with `cmake`. Set the variable `DECDAAs_TPM_SUPPORT` respectively:

```
root@amd1:~/ecdaa/build# cmake .. -DCMAKE_BUILD_TYPE=Release -DECDAAs_CURVES=FP256BN
-DMAKE_INSTALL_PREFIX=/usr -DECDAAs_TPM_SUPPORT=ON
root@amd1:~/ecdaa/build# cmake --build . --target install
```

Now, all prerequisites are installed to build and install the ECDAAs network wrapper which is a contribution of this thesis.

4.7.3 Installing ECDAAs network wrapper

Copy the folder `ecdaa-network-wrapper` to the build directory and change to this directory:

```
root@amd1:~# git clone https://git.ins.jku.at/proj/digidow/ecdaa-network-wrapper.git
root@amd1:~# cd ecdaa-network-wrapper
```

Initialize Cmake with the following command:

```
root@amd1:~/ecdaa-network-wrapper# cmake .
```

Then build the preferred targets, depending which host is used. For example, to build the member with TPM support, use:

```
root@amd1:~/ecdaa-network-wrapper# cmake --build . --target ecdaa_member_tpm -- -j 2
```

The following targets are available:

- `ecdaa_issuer`: Creates the binary for the issuer.
- `ecdaa_member`: Builds the member executable without TPM support. This should only be used for testing purposes.
- `ecdaa_member_tpm`: The member binary with TPM support.
- `ecdaa_verifier`: Creates the verifier binary.
- `ecdaa_all`: Builds every binary listed above at once.

When all above steps are finished successfully, the host is capable of taking its role in the DAA protocol.

4.8 DAA Demo Application

For demonstration purposes, we use an USB webcam to take a photo of the person being in front of the sensor. This photo is then processed to generate a face embedding, which is small enough to be sent with the DAA attestation message.

The first part is done with a small Rust program called `bs-capture`. It uses the libraries from the `video4linux` project to capture a still image and saving it to disk. Ubuntu 20.04 requires the following packages to be installed:

```
root@amd1:~# apt install rustc cargo libv4l-0 libv4l-dev v4l-utils
```

Then the program can be downloaded and executed:

```

1 root@amd1:~# git clone https://git.ins.jku.at/preisach/bs-capture.git
2 root@amd1:~# cd bs-capture/
3 root@amd1:~/bs-capture# cargo run

```

The program assumes that a webcam is available at `/dev/video0`. It takes a still image which is saved as `frame.jpg` in the working directory, which is in this example `/bs-capture`.

This image needs then to be processed to generate the face embedding data. Therefore we use the project *Prototype Facerecognition* which uses a trained tensorflow network to generate embeddings. The Git repository contains an install script which cares about installing all dependencies for setup:

```

1 root@amd1:~# git clone https://git.ins.jku.at/proj/digidow/prototype-facerecognition
  .git
2 root@amd1:~# cd prototype-facerecognition/
3 root@amd1:~/prototype-facerecognition# ./install.sh

```

4.8.1 DAA Network Protocol

The network protocol provided by `ecdaa-network-wrapper` adds to the cryptographic implementation of Xaptum's `ecdaa` project a network communication layer. It is designed to match the workflow of a Digidow transaction, affecting the decision which party is defined as listener and which as sender.

- *Start issuer listener*: During startup of the issuer server, the program loads the public/private key pair if present. Otherwise, a new key pair will be created. The issuer listener is always active to manage group membership and queries for the issuer's public key.
- *Broadcast issuer's public key*: The DAA group public key created by the issuer is necessary for group enrollment and for verification of any messages signed by a DAA group member. Consequently, verifier and (potential) member must get this key first. Figure 4.2 shows the two steps that are visible on the network. Since this communication contains only public data, no additional privacy measurements are required.




Figure 4.2: Protocol to get the issuer's public key




Figure 4.3: Protocol to add a new member to the issuer's group

- *Enroll member to issuer's group:* This step extends the DAA group and transfers the trust of the DAA group to the new member. This protocol requires the issuer's public key to be present at the member. If this is not the case, the member asks for it automatically. The procedure is a four-way handshake, as shown in Figure 4.3. We describe in subsection 2.4.2 that the key exchange is cryptographically secure, meaning that no adversary can extract any private keys when getting access to these messages.
- *Send signed messages:* The diagram in Figure 4.4 shows that the verifier listens and the member initiates the communication. This implementation reflects the Digidow transaction workflow in mind where the sensor (=member) sends a signed message to the PIA (=verifier). Hence, sending attestation messages with biometric information of the user will happen once per transaction in the Digidow network.




Figure 4.4: Protocol to send an attestation message to the verifier

5 Testing

With the setup described in the previous chapter, we created a system which is able to capture and process biometric data. The system encapsulates this data into an attestation message and sends it to the PIA which is the DAA verifier. We show in the following section how well each part of the setup work. Furthermore we analyze the footprint in memory and on the disk as well as timing if applicable. The tests are only applied to System 1 and 3 since System 2 has a comparable hardware configuration but has a processor which is one generation older. Furthermore, only two TPMs support the cryptographic parts required for ECDA. We discuss this issue in further detail in section 5.4. Consequently, System 2 was used as DAA verifier since this host does not require a TPM.

5.1 Trusted Boot

The first part of the setup is trusted boot which is well integrated in recent releases of kernel and GRUB bootloader. Furthermore, unlocking the optional disk encryption with the TPM works seamless with the kernel, even when using the manually generated unified kernel without GRUB. Only when performing an update for GRUB, it will modify the entries in the EFI boot manager. Consequently, we recommend a check of the boot entries after a system upgrade.

Having a backup boot option with GRUB is useful for maintenance tasks. Especially for modifying the IMA configuration via kernel command line, it may be necessary to boot with a backup kernel. Hence, a backup boot process is strongly recommended for test setups. This requires a backup password for the disk encryption since which allows to bypass the TPM during booting. Otherwise there are no updates possible with the current setup since the affected PCRs are used by the EFI bootloader and cannot be precomputed.

GRUB already supports trusted boot and activation requires a line in the corresponding config file. Unlike GRUB, the unified kernel does not perform any measurements. Only when asking the TPM for the disk encryption key, the initramfs must have the TPM stack available. These files use about 62 KB of space in initramfs which is negligible compared with the complete image using about 80 MB. Furthermore, only in this case exists a measurable difference in the boot performance since asking the TPM takes around 4 s. This is the time when the boot process posts unlocking via TPM and stays there until the disk is unlocked. The used amount of memory was not tested here, but it should be comparable to the results when generating a TPM key for the member. According to these results, we assume that less than 10 KB are used to hold the TPM stack in memory.

5.2 IMA

IMA appears to have an easy setup but a complex set of consequences. When setting IMA in fixing or enforcing mode, the system's performance drop is significant. We show in Table 5.1 the boot performance of the used test systems given a setup for a biometric sensor described in chapter 4 with TPM backed disk encryption enabled.

The boot procedure is an example for a reproducible file intensive process. It starts with pressing the power button when the system is in off state and ends when the CLI login mask appears. Rebooting starts when the `reboot` command is entered in the shell. Again, the CLI login mask ends the rebooting process. Further timing tests comparing the different IMA states are shown in the next section.

When enabling IMA, the file hashes will be stored as extended file attributes. According to the `xattr` man page[15], the ext4 file system creates therefore a single block per file where all attributes must find place. The block size of the used filesystem is 4 KB. The number of files is determined with the following command

```
find / -fstype ext4 -type f -uid 0 | wc -l
```

System 1 holds after setup as Biometric Sensor and several system updates 156947 files. This results in estimated additional disk usage of 613 MB when adding the `xattr` block for each file. Further attributes, however, will not require more disk space on ext4 since every extended attribute must exist in this single block.

Table 5.1: Booting and rebooting times with IMA (n=10)

Performance		System 1			System 3		
		IMA off	IMA fix	IMA enf	IMA off	IMA fix	IMA enf
<i>Boot</i> [s]	min				0.03	0.03	
	avg				0.03	0.03	
	max				0.06	0.06	
<i>Reboot</i> [s]	min				0.35	0.35	
	avg				0.37	0.37	
	max				0.37	0.41	

IMA holds all runtime information as virtual files in memory, including the IMA log. In context of memory usage, it is clear that the size of the virtual files without the redundant information (like PCR number and log format) is a rough estimation for the memory usage within the kernel. Both memory and disk usage do not change between IMA's fixing and enforcing mode since enforcing only adds file access control.

Saving the file to disk, takes about 2 MB of size per 10000 lines in the log. This indicates a linear dependency to the number of accessed files for the log and the kernel's memory usage. The log can easily get over 100000 entries when the system is running long enough. System 1, for example, had a log file with 214561 lines after about 15 days of uptime, resulting in about 40 MB of size. Saving that file on disk took several seconds which indicates that extracting the log from the kernel is very inefficient. The log file showed that a major impact on the size was the performed system update. Generating the iniramfs blob and compiling programs from source generates combined several 10000 lines in the log.

5.3 Processing and Sending Biometric Data

Capturing and processing biometric data from the user is quite seamless and the cryptographic part of ECDA is reliable enough for this prototype. During the following tests,

Table 5.2: Memory usage measured with Valgrind

<i>Task</i>	<i>System 1</i>	<i>System 3</i>
DAA TPM key generation	10,160	10,160
DAA TPM join	23,864	23,864
DAA join without TPM	19,296	19,296
Image capturing	93,703	93,703
Image processing	1,318,722,747	1,385,416,573
Compiling DAA message	1,115,639	1,115,597
Sending DAA message	36,072	36,072

all software and hardware parts worked as expected. Neither TPM nor software errors were encountered.

Analyzing disk and memory usage is only meaningful for the DAA member. The implemented prototype of the DAA issuer does only negotiate the membership key. Revocation lists and group management are not implemented yet, although the ECDAA library provides datastructures and functions for that. Similarly, the DAA verifier only checks the signature. In production use, both entities must hold the revocation list and perform further checks to trust the DAA member.

The application's memory usage is measured with `valgrind`. It measures the allocated heap space in memory which is shown in Table 5.2.

The memory usage is constant over all procedures but creating the DAA message. This step depends on the size of the files which it summarizes, especially when taking the IMA log into account. Therefore the memory usage is measured while IMA is set off, representing a lower bound of memory usage for this part. Besides calculating the face embedding of the captured image, the whole transaction can be executed using less than 100 Kilobytes of heap memory. This would be affordable even for embedded devices with respect to memory consumption. On the other side, the face embedding algorithm uses over 1.3 GB. The slight difference between the two systems at the processing part seems to be consistent over several runs, whereas the result at compiling the message may result of slight differences in the input file.

Table 5.3 shows each relevant step for the Digidow sensor. Each part shows the time

Table 5.3: Performance results of joining a DAA group and sending a Digidow transaction message (n=10000)

		<i>System 1</i>			<i>System 3</i>		
<i>Task</i>		IMA off	IMA fix	IMA enf	IMA off	IMA fix	IMA enf
<i>Join without TPM [s]</i>	min				0.03	0.03	
	avg				0.03	0.03	
	max				0.06	0.06	
	first				0.04	0.13	
<i>Join with TPM [s]</i>	min				0.35	0.35	
	avg				0.37	0.37	
	max				0.37	0.41	
	first				0.40	0.42	
<i>Capture [s]</i>	min				0.91	0.91	
	avg				1.06	1.06	
	max				1.12	12.48	
	first				1.34	1.46	
<i>Process embedding [s]</i>	min				4.07	4.09	
	avg				4.12	4.14	
	max				4.74	4.46	
	first				5.99	40.21	
<i>Collect [s]</i>	min				0.09	0.19	
	avg				0.10	n/a	n/a
	max				0.11	n/a	n/a
	first				0.11	0.24	
<i>Send with TPM [s]</i>	min				0.26	0.27	
	avg				0.28	0.27	
	max				0.28	0.29	
	first				0.28	0.40	
<i>Transaction sum [s]</i>	min				5.38	5.50	
	avg				5.56	n/a	n/a
	max				6.14	n/a	n/a
	first				7.72	42.31	

Table 5.4: Number of additional entries in the IMA log

<i>Additional log entries</i>	
Boot	~2000
DAA TPM enrollment	
Image capturing	
Image processing	
Compiling DAA message	
Sending DAA message	

of execution of the first attempt and a basic statistical numbers of the 9999 remaining executions. The first two blocks represent generating the private key of the sensor and executing the DAA join handshake over the local network. The difference between these two programs is not only the usage of the TPM. Joining the DAA group using a TPM is in this test a two-step job. The first step creates key within the TPM whereas the second step uses the join protocol. This is not required in the other case since the private key is created on the fly at the beginning of the join handshake.

The following 4 blocks represent the 4 steps for starting the Digidow transaction. Capture takes a picture of the user

When IMA is in fixing or enforcing mode, the corresponding log will be filled according to Table 5.4.

This means—given that the (very slow) hardware TPM had to extend PCR 10 for every line in the log—the slowdown is mainly caused by the interaction with the TPM itself.

Since the IMA log file is also essential for remote attestation, the information of this file must be transmitted to the DAA verifier.

- payload without IMA log about 15KB
- No encryption for payload, but doable – depends on the way how Sensor and PIA can communicate together
- IMA log much too large
- Test results how long the process of capturing takes – with and without IMA

5.4 Limitations

- older TPM does not support ECDSA
- Documentation available for TPM APIs, but no changelog for tpm2-tools.
- Trusted boot and IMA can just handle static resources like files, kernel modules and firmware of hardware components. Code transmitted over network or otherwise dynamically generated can not be recognized. This is an open door for non-persistent attacks.
- Documentation on IMA is mostly outdated and so are some tools. Further customization of rules may be useful to reduce log size. However major Linux distributions support IMA by default on recent releases.
- Complexity of verifying system state is too high and is connected to system complexity. Reducing number of dependencies and relevant file count is key for this problem.
- Implemented DAA does not support a full dynamic group scheme. This might be useful in the future, maybe with a custom implementation of a recent DAA version.

During these memory tests, valgrind showed a large number of possible memory leaks in the python binary itself. The following example is executed:

```
root@amd1:~/jetson-nano-facerecognition# valgrind python3 img2emb.py data/test-
images/test2.jpg
```

Valgrind ends with the report as follows:

```
==1648== HEAP SUMMARY:
==1648==   in use at exit: 32,608,730 bytes in 227,287 blocks
==1648== total heap usage: 810,162 allocs, 582,875 frees, 1,385,416,573 bytes
         allocated
==1648==
==1648== LEAK SUMMARY:
==1648==   definitely lost: 3,144 bytes in 28 blocks
==1648==   indirectly lost: 0 bytes in 1 blocks
==1648==   possibly lost: 523,629 bytes in 12,842 blocks
==1648==   still reachable: 32,081,957 bytes in 214,416 blocks
==1648==                of which reachable via heuristic:
==1648==                 stdstring      : 537,414 bytes in 11,917 blocks
==1648==                 newarray      : 8,920 bytes in 5 blocks
```

```

==1648==      suppressed: 0 bytes in 0 blocks
==1648== Rerun with --leak-check=full to see details of leaked memory
==1648==
==1648== Use --track-origins=yes to see where uninitialised values come from
==1648== For lists of detected and suppressed errors, rerun with: -s
==1648== ERROR SUMMARY: 58173 errors from 914 contexts (suppressed: 0 from 0)

```

This report shows that the Python binary (here Python 3.8 from Ubuntu 20.04) is not memory safe, which is a significant drawback for the system and software integrity. Any binary which is directly involved in the DAA protocol frees every allocated block. Furthermore any binary in the TPM2 software stack is memory safe according to valgrind. The used shell commands may not free every allocated block, however valgrind still finds no errors in these programs.

When IMA is set to enforcing, some unexpected problems appeared during updating Ubuntu. During apt upgrade, the package manager downloads the deb packages into its cache folder at /var/cache/apt/. These files, however, do not have the security.ima attribute when the download is finished. The kernel prevents due to theses missing attributes any access and breaks the upgrade process. It is not clear why the files are not hashed although apt is run by root and every file created by root must be hashed with the active IMA policy. Creating a text file via text editor shows exactly this behaviour.

The missing attributes are also a problem when attempting to turn IMA off again. This requires an updated command line and hence an updated unified kernel. Generating the new kernel work fine but moving the blob into the EFI partition fails due to the missing support of extended file attributes. The copy process seems to create the file, but it fails when trying to write the first chunks. As a result, the only way to update the kernel on this system is to boot a backup kernel and moving the file without IMA set to enforcing. Applying customized IMA rules might solve both problems.

Another relevant part for attestation is to use the IMA log file do recalculate the produces hashes. This means on one hand recalculating the hash in the IMA log entry. On the other hand the chain of hashes should result into the value held in PCR 10. A small script shown in Listing 5.1 tries to recalculate this value. It uses the debug register PCR 16 which is resetable without reboot and has the same initial value as the first 10 PCRs. When IMA is off, the log holds only one entry of the boot aggregate. Then the SHA1 value can be computed with that script. However, comprehending the PCR 10 value with IMA enabled was not possible. Our tests took into account that PCR and log file could be modified

Listing 5.1: Attempt to recalculate the value of PCR 10

```
#!/usr/bin/bash
set -e
tpm2_pcrreset 16
cut -d '_' -f 2 /sys/kernel/security/ascii_runtime_measurements > ima.hashes
while read i ;
do tpm2_pcrextend 16:sha1=$i ;
done < ima.hashes
tpm2_pcrread sha1:10,16
```

when loading programs to read these resources the first time. Loading the log several times eventually ends up in a stable log and PCR value (it does not change anymore even when reading the log another time). The value of PCR 10 was still not reproducible.

Furthermore the documentation of calculating these values did not mention how the sha256 hash in PCR 10 is calculated. `tpm2_pcrextend` requires a sha256 hash as input for the corresponding PCR bank, but the IMA log only provides sha1 hashes. Any verification procedures regarding the sha256 bank of PCR 10 are currently not implemented.

6 Conclusion and Outlook

6.1 Future Work

- Remove building tools on target device - just deliver binaries
- Remove complex runtime environments like Java, Python, etc. to reduce bloating the integrity logs
- Set file system read only, just use e.g. a ramdisk for working files
- Integrate USB sensors into the trusted/integrity environment, including device firmware.

6.1.1 Closing the chain of trust between TPM manufacturer and DAA issuer

Activate a credential with to certify that the Membership key is in the Endorsement hierarchy, which can be verified with the TPM certificate.

- Theoretical concept in the *Practical Guide to TPM 2.0*, pp 109 ff
- Practical approach: with EK, AK and AIK to show validity of EK:
https://ericchiang.github.io/post/tpm-keys/?utm_campaign=Go%20Full-Stack&utm_medium=email&utm_source=Revue%20newsletter#credential-activation

Further integration in the Digidow environment if DAA is useful for that.

6.2 Outlook

Hardening of the system beyond IMA useful. Minimization also useful, because the logging gets shorter.

Bibliography

- [1] FIDO Alliance. *FIDO ECDA Algorithm Implementation Draft*. 2018. URL: <https://fidoalliance.org/specs/fido-v2.0-id-20180227/fido-ecdaa-algorithm-v2.0-id-20180227.html> (visited on 07/07/2021) (cit. on p. 26).
- [2] Will Arthur, David Challener, and Kenneth Goldman. *A Practical Guide to TPM 2.0*. Jan. 2015. DOI: 10.1007/978-1-4302-6584-9 (cit. on p. 11).
- [3] Jan Camenisch, Liqun Chen, Manu Drijvers, Anja Lehmann, David Novick, and Rainer Urian. “One TPM to Bind Them All: Fixing TPM 2.0 for Provably Secure Anonymous Attestation”. In: May 2017, pp. 901–920. DOI: 10.1109/SP.2017.22 (cit. on p. 19).
- [4] Jan Camenisch, Manu Drijvers, and Anja Lehmann. “Universally Composable Direct Anonymous Attestation”. In: *Public-Key Cryptography – PKC 2016*. Ed. by Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang. Berlin, Heidelberg: Springer Berlin Heidelberg, Mar. 2016, pp. 234–264. ISBN: 978-3-662-49386-1. DOI: 10.1007/978-3-662-49387-8_10 (cit. on pp. 19, 21, 22, 25).
- [5] Jan Camenisch and Anna Lysyanskaya. “Signature Schemes and Anonymous Credentials from Bilinear Maps”. In: vol. 3152/2004. Aug. 2004, pp. 56–72. DOI: 10.1007/978-3-540-28628-8_4 (cit. on p. 21).
- [6] Jan Camenisch and Markus Stadler. “Efficient Group Signature Schemes for Large Groups”. In: *CRYPTO ’97* 1296 (Jan. 1997) (cit. on p. 20).
- [7] TPM2 Software Community. *TPM2 Tools*. 2020. URL: <https://github.com/tpm2-software/tpm2-tools> (visited on 05/15/2020) (cit. on p. 10).
- [8] MITRE Corporation. *Search Results for “tpm” in the CVE Database*. 2021. URL: <https://cve.mitre.org/cgi-bin/cvekey.cgi?keyword=tpm> (visited on 05/15/2021) (cit. on p. 9).

- [9] Free Software Foundation. *GRUB 2.04 User Manual: Measuring Boot Components*. 2019. URL: https://www.gnu.org/software/grub/manual/grub/html_node/Measured-Boot.html (visited on 03/29/2021) (cit. on pp. 15, 40).
- [10] Dann Frazier. *Secure Boot*. 2020. URL: <https://wiki.ubuntu.com/UEFI/SecureBoot> (visited on 07/24/2021) (cit. on p. 15).
- [11] Inc Gentoo Foundation. *Integrity Measurement Architecture/Recipes*. 2019. URL: https://wiki.gentoo.org/wiki/Integrity_Measurement_Architecture/Recipes (visited on 07/07/2021) (cit. on pp. 19, 46).
- [12] Trusted Computing Group. *TCG PC Client Platform Firmware Profile Specification Revision 1.04*. 2019. URL: https://trustedcomputinggroup.org/wp-content/uploads/TCG_PCClientSpecPlat_TPM_2p0_1p04_pub.pdf (visited on 08/01/2020) (cit. on pp. 13, 14).
- [13] Trusted Computing Group. *The TPM Library Specification*. 2019. URL: <https://trustedcomputinggroup.org/resource/tpm-library-specification/> (visited on 05/16/2020) (cit. on p. 8).
- [14] Xaptum Inc. *Source repository for the ECDAA C Library*. 2021. URL: <https://github.com/xaptum/ecdaa> (visited on 07/07/2021) (cit. on p. 39).
- [15] Michael Kerrisk. *Xattr Man Page*. 2020. URL: <https://man7.org/linux/man-pages/man7/xattr.7.html> (visited on 09/19/2021) (cit. on p. 59).
- [16] Microsoft. *Secure Boot Overview*. 2014. URL: [https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-8.1-and-8/hh824987\(v=win.10\)](https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-8.1-and-8/hh824987(v=win.10)) (visited on 07/24/2021) (cit. on p. 15).
- [17] Daniel Moghimi, Berk Sunar, Thomas Eisenbarth, and Nadia Heninger. “TPM-FAIL: TPM meets Timing and Lattice Attacks”. In: *29th USENIX Security Symposium (USENIX Security 20)*. Boston, MA: USENIX Association, Aug. 2020. URL: <https://www.usenix.org/conference/usenixsecurity20/presentation/moghimi> (cit. on p. 9).
- [18] Matus Nemec, Marek Sys, Petr Svenda, Dusan Klinec, and Vashek Matyas. “The Return of Coppersmith’s Attack: Practical Factorization of Widely Used RSA Moduli”. In: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. CCS ’17. Dallas, Texas, USA: Association for Computing Machinery, 2017, pp. 1631–1648. ISBN: 9781450349468. DOI: 10.1145/3133956.3133969. URL: <https://doi.org/10.1145/3133956.3133969> (cit. on p. 9).

- [19] Pawit Pornkitprasan. *Full Disk Encryption on Arch Linux backed by TPM 2.0*. July 2019. URL: <https://medium.com/@pawitp/full-disk-encryption-on-arch-linux-backed-by-tpm-2-0-c0892cab9704> (visited on 02/27/2020) (cit. on p. 40).
- [20] Pawit Pornkitprasan. *Its certainly annoying that TPM2-Tools like to change their command line parameters*. Oct. 2019. URL: <https://medium.com/@pawitp/its-certainly-annoying-that-tpm2-tools-like-to-change-their-command-line-parameters-d5d0f4351206> (visited on 02/27/2020) (cit. on pp. 10, 40).
- [21] David Safford, Dmitry Kasatkin, and Mimi Zohar. *Integrity Measurement Architecture (IMA) Wiki Page*. 2020. URL: <https://sourceforge.net/p/linux-ima/wiki/Home/> (visited on 03/20/2021) (cit. on p. 16).
- [22] Nabil Schear, Patrick T. Cable, Thomas M. Moyer, Bryan Richard, and Robert Rudd. "Bootstrapping and Maintaining Trust in the Cloud". In: *Proceedings of the 32nd Annual Conference on Computer Security Applications*. ACSAC '16. Los Angeles, California, USA: Association for Computing Machinery, 2016, pp. 65–77. ISBN: 9781450347716. DOI: 10.1145/2991079.2991104. URL: <https://doi.org/10.1145/2991079.2991104> (cit. on p. 19).
- [23] Tevora. *Configuring Secure Boot + TPM 2*. June 2019. URL: <https://threat.tevora.com/secure-boot-tpm-2/> (visited on 06/19/2020) (cit. on p. 40).