FIDIS

Future of Identity in the Information Society

Title:	"D3.9: Study on the Impact of Trusted Computing on Identity and Identity Management"
Author:	WP3
Editors:	Ammar Alkassar, Rani Husseiki (Sirrix AG)
Reviewers:	Jozef Vyskoc (VaF Bratislava)
Identifier:	D3.9
Type:	[Deliverable]
Version:	1.1
Date:	Tuesday, 13 May 2008
Status:	[Complete]
Class:	[Public]
File:	fidis_wp3_del3.9.Trusted_Computing.doc

Summary

Trusted Computing (TC) is a key enabling technology adding substantially new security features, making many new use cases possible, which may revolutionize identity management. However, this emerging technology is not undisputed and raises many societal questions related to privacy, rights on ownership etc. This study takes a deeper look into TC concepts like TPMs, Trustworthy Operating Systems etc, and discusses possible use and business cases for TC in the context of identity and identification, pointing out possible risks of this technology in terms of privacy and consumer protection.

The objective of this study is to give an overview of Trusted Computing concepts and its supporting technologies, and to introduce new ideas on how those concepts can support or influence digital identification and identity management systems, including possible privacy and anonymity implications of Trusted Computing specifications defined by the Trusted Computing Group.

This deliverable differs substantially from 33 of ALU-FR, as it addresses mainly the use of TC mechanisms on the client side and focuses on the technology description and its impact on IMS.

The FIDIS NoE receives research funding from the Community's Sixth Framework Program

Copyright Notice:

Members of the FIDIS consortium

1. Goethe University Frankfurt	Germany
2. Joint Research Centre (JRC)	Spain
3. Vrije Universiteit Brussel	Belgium
4. Unabhängiges Landeszentrum für Datenschutz	Germany
5. Institut Europeen D'Administration Des Affaires (INSEAD)	France
6. University of Reading	United Kingdom
7. Katholieke Universiteit Leuven	Belgium
8. Tilburg University	Netherlands
9. Karlstads University	Sweden
10. Technische Universität Berlin	Germany
11. Technische Universität Dresden	Germany
12. Albert-Ludwig-University Freiburg	Germany
13. Masarykova universita v Brne	Czech Republic
14. VaF Bratislava	Slovakia
15. London School of Economics and Political Science	United Kingdom
16. Budapest University of Technology and Economics (ISTRI)	Hungary
17. IBM Research GmbH	Switzerland
18. Institut de recherche criminelle de la Gendarmerie Nationale	France
19. Netherlands Forensic Institute	Netherlands
20. Virtual Identity and Privacy Research Center	Switzerland
21. Europäisches Microsoft Innovations Center GmbH	Germany
22. Institute of Communication and Computer Systems (ICCS)	Greece
23. AXSionics AG	Switzerland
24. SIRRIX AG Security Technologies	Germany

Foreword

FIDIS partners from various disciplines have contributed as authors to this document. The following list names the main contributors for the chapters of this document:

Chapter	*Contributor(s)*
1 (Executive Summary)	Rani Husseiki (Sirrix AG)
2 (General Introduction)	Stefan Köpsell (TU Dresden), Rani Husseiki (Sirrix AG)
3 (Trusted Computing – an Overview)	Stefan Köpsell (TU Dresden), Rani Husseiki (Sirrix AG)
4 (Trusted Computing Group Specifications)	Stefan Köpsell (TU Dresden), Christian Wachsmann (Sirrix AG)
5 (Trusted Computing beyond the TCG specifications)	Rani Husseiki (Sirrix AG)
6 (Application Scenarios for Trusted Computing Technology)	Stefan Köpsell (TU Dresden), Martin Meints (ICPP), Vassiliki Andronikou (ICCS)
7 (Controversial and Legal Aspects of Trusted Computing)	Stefan Köpsell (TU Dresden), Rani Husseiki (Sirrix AG)
8 (Trusted Computing, Identity and Identity Management)	Martin Meints (ICPP), Rani Husseiki (Sirrix AG)
9 TC and Identity Management – a Use Case Scenario	Rani Husseiki (Sirrix AG)
10 (Recommendations and Future Work)	Rani Husseiki (Sirrix AG)
11 (Conclusion)	Rani Husseiki (Sirrix AG)

Table of Contents

1 Executive Summary

In this deliverable, Trusted Computing (TC) technology is introduced to the reader, with an overview of the main concepts, functionalities and features of general TC hardware and software. The overview includes also the TC specifications of the Trusted Computing Group – a specification and standardization body dedicated for Trusted Computing – in addition to other industrial and academic efforts and projects for further development of the technology. After touching on some application scenarios and social and legal aspects of TC, we shed the light on implications of TC on identity and identity management with regard to enhancing digital identification, possible consumer privacy violation, and providing anonymity in TC-based infrastructures. We propose ideas on how the TC concepts and features can influence identification, identity management systems, privacy and anonymity. We then give a brief use case scenario of TC-based identity management across several domains of identification.

TC is a new technology aimed at bringing trust in computing platforms to a higher level by providing evidence about the integrity of a platform to both, the platform's owner and to arbitrary third parties. While the concepts underlying Trusted Computing date back to the 1960s, the technology emerged when adopted by the Trusted Computing Group (TCG), formerly the Trusted Computing Platform Alliance (TCPA).

The three main components of the TCG proposal are under focus, namely the Trusted Platform Module (TPM), a tamper-resistant hardware chip, a kind of (protected) pre-BIOS called the Core Root of Trust for Measurement (CRTM), and a support software called TCG Software Stack (TSS). Those are continuously subject to research, development and implementations by TPM Manufacturers, system integrators and leading industrial and academic open-source projects.

Although TC seems to be a promising technology, some aspects of it still raise problems and reservations from consumers, but also from academic researchers. The fears lie in several aspects like the possibility of restrictive digital content control by content providers by means of technologies such as Digital Rights Management, "lock-in" of specific software on consumer platforms by software providers, restricting the installation of similar software and reducing interoperability, and consumer privacy breach due to some specified protocols requiring disclosure of identification information to Trusted Third Parties. From a legal perspective, some possible implications of TC might not be completely conformant with legislations such as the Copyright Act and the Competition Act.

Nevertheless, TC seems to give ground for new business use cases in various fields such as Distributed Policy Enforcement, secure end-user systems and embedded security. In particular, TC can present advantages in the areas of Identification and Identity Management, and can enable new solutions in those fields. Some features and protocols defined by the TCG can have substantial effect on various aspects of identification and can hence enable new business cases. Functionalities can also affect anonymity aspects using TC-based platforms; other can enhance identification depending on the use of the features.

For this reason, national security agencies, industrial and standardisation bodies as well as the research community are all encouraged to further investigate the advantages of TC

with respect to digital identity, identification and identity management, and to envision more developed use cases of the technology.

2 Introduction

2.1 What is this deliverable about?

This deliverable presents a study on Trusted Computing (TC), a key enabling technology that adds substantial new security features, and its possible impact on identity and identity management. It gives a deep look into TC concepts, features and specifications, which are still developed within and beyond the "Trusted Computing Group", a non-profit organization dedicated for TC specifications that evolved out of the "Trusted Computing Platform Alliance" industry working group.

The deliverable discusses the TCG organization and the specifications it defines. It sheds the light on three main components of the TCG proposal which are the hardware component Trusted Platform Module (TPM), a kind of (protected) pre-BIOS called the Core Root of Trust for Measurement (CRTM), and a support software called TCG Software Stack (TSS).

Moreover, the deliverable discusses TC research and development beyond the TCG specifications. This includes trustworthy operating systems and secure platforms and hardware that are designed and developed nowadays to support the TC specifications. The study also focuses on application scenarios for TC such as distributed policy enforcement, DRM and anonymity services. It also gives an overview of the TC solutions in the market according to TPM Manufacturers, system integrators, industrial and academic open-source projects…

This thorough description of TC concepts, functionalities, specifications and market status is followed by a description of the implications of TC for identity and identity management. Some features and protocols defined by the TCG can have substantial effect on various aspects of identification and can hence enable new business cases. Some functionality can also affect privacy aspects of consumer using TC-based platforms; other can enhance identification depending on the use of TC. The deliverable sheds the light on some TC features that can improve identity management, especially in terms of interoperability and trust establishment.

Readers of this deliverable are assumed to have a fair level of computer science background as it especially focuses on technological advancements that can potentially support identity and identity management, although readers with legal background can also benefit from the legally controversial issues relater to TC, in addition to its privacy implications in the context of identification.

Parts of this deliverable are based on a study to be published by the German Federal Office for Information Security (BSI) [141].

2.2 Relevance of TC for Identity and Identity Management

The reason why TC is addressed in the context of FIDIS is mainly related to the fact that TC introduces new schemes and business cases based on trust establishment between computing platforms, which would definitely have considerable implications on identity based authentication and identity management schemes. On the other side, the specifications and protocols defined within the TCG seem to raise problems regarding privacy since they require

the user to reveal some identification information about the platform he is using. TC seems ambiguous for many people when it comes to its capability for enhancing privacy or anonymity of identity and identification information.

It is therefore important to shed the light on the relevance of TC infrastructures used in identification and identity management systems, and to assess their possible effects on crucial aspects of identity and identity management such as strong authentication, user privacy and anonymity.

2.3 Scope and Structure of the Document

The scope of this document spans over the concepts, features and market status of Trusted Computing according to the TCG specifications and to other industrial research and development in this field. It also tackles the social and legal aspects related to TC before shedding the light on the implications of TC infrastructure on identity, identification and identity management systems. Recommendations and best practices follow this analysis.

Chapter 3 gives an overview of the main concepts defined within the TC context and which are at the basis of the distinguished features of TC implementation.

Chapter 4 introduces the Trusted Computing Group and gives a more or less detailed description of the TC functionalities as specified by the TCG, with a focus on three main aspects: the Trusted Platform Module (TPM) specifications, the TCG Software Stack (TSS) specifications and the Trusted Network Connect (TNC) specifications.

Chapter 5 tackles the industrial research and development and academic open source projects in the direction of enhancing and supporting TC beyond the TCG specification. This includes secure operating systems, virtualization technology and further TC hardware.

Chapter 6 mentions the current application scenarios for TC and their existing use cases in addition to prospective scenarios.

Chapter 7 focuses on the social and legal aspects of TC, especially the controversial issues.

Chapter 8 sheds the light on the implications of TC for identity and identity management with regard to enhancing identification, possible consumer privacy violation, and providing anonymity in TC-based infrastructures.

Chapter 9 describes a use case scenario showing a possible application of TC to identity management.

Chapter 10 gives some recommendations and best practices regarding the use of TC for identification and identity management

Chapter 11 presents the conclusions of this study.

3 Trusted Computing – an Overview

3.1 Introduction to Trusted Computing

Existing networked computing platforms are not able to fulfill the multilateral security requirements of parties like companies, end-users, and content providers. This gets obvious in the huge number of exploits and security updates as well as the high number of attacks through viruses, worms and Trojan horses. Furthermore, the security of existing computing platforms could not be vitally improved in the last years due to the conceptual weaknesses, like their monolithic architecture and thus inherent complexity. This pertains to Windows-based operating systems as well as Linux-based ones.

Most of the currently used IT-systems lack elemental security properties, such as integrity checks or the generation of secure cryptographic keys using appropriate random number generators. Thus, the existing threats thwart the realization of a variety of useful applications and business models, particularly in the area of Digital Rights Management (DRM).
Trusted platform technology should provide evidence about the integrity of a platform to both, the platform's owner and to arbitrary third parties. To take full advantage of trusted platform properties a public key infrastructure (PKI) is required.

The degree of confidence in software-only security solutions depends on their correct installation and execution, which can be affected by other software that has been executed on the same platform. Therefore a trusted platform is a conventional platform containing a hardware-based subsystem devoted to maintain trust and security between machines. It contains a trusted component, probably in the form of a built-in cost effective security hardware that is used to create a foundation of trust for software processes. This extra hardware is roughly equivalent to that of a smart card (with some enhancements) and contains a variety of functions that must be trusted. The trust mechanisms in trusted platforms use selected security mechanisms, but they are ultimately based upon signed statements of "social trust" made by individuals and organizations. In addition, a trusted platform provides hardware protection for keys and other secrets, which may be used to encrypt files or gain access to servers or networks.

Applications and services that would benefit from using trusted platforms include electronic cash, email, hot-desking (allowing mobile users to share a pool of computers), platform management, single sign on (removing the need for the user to be asked to authenticate himself or herself more than once when using different applications during the same work session), virtual private networks, Web access and digital content delivery.

3.2 Main concepts of Trusted Computing

3.2.1 Integrity Measurement, Verifiable and Secure Booting

Integrity measurement is one of the basic and probably most fundamental mechanisms of Trusted Computing. Integrity measurement means to calculate and store the state of a computer or device under secured conditions.

The process of measurement while booting is called *authenticated booting (or authentic booting)*. It is a method that securely *logs* which software is booted on a computing device but does *not influence* the boot sequence (e.g. in deciding which software to execute and which not). After finishing the boot sequence, these logs can be used for checking the state of the system. An important remark is, that authentic booting by no means ensures that the computing device is in a "secure" state.

With the help of the log entries it is also possible to report the state of the computing device to remote entities. By analysing the logs remote entities can decide about the trustworthiness of a system. This approach is particularly suitable for secure *open* platforms, which can be modified in many ways [115].

The technology of authenticated booting is already used besides the TCG specification in other designs. According to Kauer [115] currently available designs of authenticated booting require new hardware or support untrusted applications insufficiently.

There exists another type of booting a system in a trusted way called *secure booting*. Secure booting checks the code before executing it according to a set of security policies and therey avoids that malicious or unauthorized code is running on the computing device. The security policy comprises identifiers of authorized modules. Hash values can be used as identifiers. If an identifier or hash value is not found in the list of the security policy the process of booting is discontinued. Usually this technology is used to ensure locally that the platform is in a secure state, but not to prove this to a remote party. As mentioned in [115] the technique of secure booting can be used to construct secure *closed* platforms, which have a limited number of executable software. Designs for secure booting have been known for over 15 years [113][114].

Trusted booting combines both boot methods mentioned. After measurement of integrity the results are sent to a remote verifier that checks the results [119]. If the computing platform is in an invalid state the remote verifier may initiate a remediation. So the platform has to be updated. After that it starts the integrity measurement again until it is in a valid state.

When a boot sequence can be validated (remote or locally) it ensures that components of the platform are not emulated, so that "a specific hardware with a specific OS with a specific GUI and a specific application is indeed running in the identified device" (Härtig [114]).

3.2.2 Binding, Sealing, and Attestation

The hardware integrated root of trust can provide methods to bind data, licenses and user authentication to a specific platform which consists of specific hardware and software executed on that hardware. This can be realised by cryptographic operations which are executed and stored in a protected hardware environment. Within the TCG specification a unique key, that is certified by the manufacturer and stored in a protected way inside the microcontroller, is essential for verifying the platform as trusted and can also be used to authenticate a special user accurately. This key is the initial point for certificates and further keys.

Sealing allows binding data to a specific computing device and thus sealing sensible data which is stored on its hard disk.

According to [118] a major security problem of computers is storing cryptographic keys securely. Keys or passwords that protect private documents are often retained locally on a hard drive of a PC, side by side with the encrypted documents themselves. Everyone who gains access to the computing system also gains access cryptographic keys and passwords stored. But keys should be kept secure so that only legitimate users can use them.

The technique of sealed storage is based on a key that is partially generated by the identity of the software requiring the key. Furthermore the identity of the computing device that executes the software presents the second part of the key. So, these keys need not be stored on the hard drive but can be created whenever they are required.

If a program different from the program which initiated the encryption or sealing of sensible information would try to decrypt or unseal this data this fails because the generated key is not equal to the original one. That follows from the different identifiers of the software that seal and unseal the data and consequentially the generated keys are different as well. A similar use case is that encrypted data is transferred to another computing device that tries to decrypt the data. This will also be unsuccessful. So, for example emails that you can read on your computer are unreadable on other computer systems.

With the help of sealed storage you cannot prevent that confidential data is copied to another system but you can prevent others from reading it on this system.

By using attestation it is possible to check the hardware and software states of a remote platform. Therefore the results of integrity measurements, which characterise the software and hardware environment of a computing system, are signed with a key by the hardware component. The signed outcomes can be verified by a remote party and needs no physical presence. Together with the signed outcomes certificates are sent that accredit the used key as trusted. A remote attestation can be conducted directly or through a trusted third party that verifies the remote platform as trusted.

A trusted third party (TTP) checks keys and certificates of a computing device. If they are valid, the trusted third party issues a certificate that attests a computing device as trusted.

The relation between a certain computing device and a certificate provided by a TTP should be hidden for anonymous usage.

Direct anonymous attestation (DDA) without using a third party verifier is a further attestation technique. It can be proven that a valid certificate exists without disclosing it. So certificates could be generated anonymously. This technology is based on zero knowledge group signature schemes (see also [109]). Group signature schemes allow that every member of a group can sign messages on behalf of the whole group. This supports the anonymity of the group members and provides a valid signature.

Another method to check that executable code is trustworthy is provided by code signing. Most of the code-signing implementations provide a digital signature mechanism to attest the identity of the author or the producer. A checksum is also attached to the code in order to verify that the code has not been altered. Due to this a user or a system can decide whether the source is trustworthy and whether the code has been modified.

Proof-Carrying Code is a technique that guarantees safe execution of untrusted code. The author of the code adds information (called annotation) about the security policy that is fulfilled by the code. The receiver compares the security policy of the code with a set of safety rules established by him. The code will only be executed by the receiver if the examination of the security policy of the code conforms to requirements of the user. Each tampering of the code or the annotation can be noticed by the receiver.

3.2.3 Process Isolation, Small TCB

In present computing devices it is possible for malicious code to read or alter sensitive parts of programs stored in memory. Process isolation shall increase the security of computing systems by providing full isolation of sensitive areas of memory for all programs expect authorised ones. This prevents programs from being able to read or write the memory of other programs. Access to protected memory of programs could be denied for unprivileged parts of the operating system as well. These protected parts of the memory are also called shielded locations or curtained memory. Even if an intruder controls the whole computing system, he is not able to tamper the shielded locations.

Process isolation can be achieved in hardware and software. As mentioned in [118] software implementation requires rewriting of operating system and device drivers. If necessary even application software has to be rewritten. The software approach does not need new hardware components but requires rewriting lots of software.

Implementing process isolation in hardware can provide backwards compatibility with present software. And the amount of software that has to be rewritten is small compared to the software approach.

A further technique to shield software is called sandbox. Software is insulated from the rest of the system, put into a virtual sandbox. So the software cannot do damage to other parts of the system and it is possible to record its effects. This method provides a kind of test area wherein the user can run unknown or suspect software.

The terms “small Trusted Computing Base” and “small secure platform” describe the idea to make the system core as small and compact as possible. Both include that a minimal set of hardware, firmware and operating system has to be sufficient for fundamental security. All components of a trusted computing base have to be trusted in order to guarantee the reliability of the techniques mentioned in this chapter.

According to Härtig there are some requirements in order to achieve a small secure platform like [114]:

- Flexibility so that it can be used for different kinds of platforms,
- The functionality should be minimised but sufficient for applications that require high security,
- Programs and users are only allowed to access such information and resources that are required for their legitimate purpose.
- Separation of secure and insecure parts of the platform,
- Attestation of its global identitfier to remote or local entities,
- Allowing an accurate evaluation.

As Härtig mentioned a small secure platform is small and simple when it can be completely controlled by about seven people. A lot of today's systems like the Linux kernel are far from achieving this.

There are some approaches to reduce the size of a trusted computing base like the Nizza security architecture in [121].

4 Trusted Computing Group Specifications

In this section, we introduce the Trusted Computing Group which is an industrial work group focused on the specification and standardization of TC technologies, and we briefly review the main functionalities, including binary attestation and binary sealing, of the specifications version 1.1b [9] and 1.2 [10] of the Trusted Computing Group (TCG).
The main components of the TCG proposal are the hardware component Trusted Platform Module (TPM), a kind of (protected) pre-BIOS called the Core Root of Trust for Measurement (CRTM), and a support software called TCG Software Stack (TSS) which performs various functions like communicating with the rest of the platform or with other platforms.

4.1 The Trusted Computing Group

The Trusted Computing Group (TCG) evolved from the Trusted Computing Platform Alliance (TCPA) which was an industry working group focused on the development of trust and security mechanisms in computer platforms. It was formed by Compaq (today part of Hewlett-Packard), Hewlett-Packard, IBM, Intel and Microsoft in January 1999.
In October 1999 the TCPA announced a draft specification and opened the possibility for other companies to join under a non-disclosure agreement.
In August 2000 the first public version of the TCPA Specification was released for comments and has been published as TCPA Specification 1.0 in February 2001. This specification was platform independent and basically defined functions that must be provided by a Trusted Platform Module (TPM) (see section 3.1) from the viewpoint of a hardware manufacturer.
The TPM Work Group, which has been formed in February 2001, revised the specification regarding practical implementation issues and error correction. This led to the TCPA Specification 1.1 which has been published in August 2001. Many specifications of non-hardware functions have been deferred to other TCPA Specifications.
In September 2001 the TCPA PC Specific Work Group published its first specification [97]. This working group was set up to design a special specification for the PC platform.
The next milestone was the TCPA Specification 1.1b [9] which has been released in May 2002.
In April 2003, the TCPA was replaced by a non-profit organization [10], called Trusted Computing Group (TCG). The TCG adopted all TCPA Specifications and continued their development.
In addition to the PC Specific Work Group and the TPM Work Group, which have been adopted from the TCPA, the TCG established several other working groups. These are concerned with the development of specifications for mobile devices, PC clients, servers, storage systems, infrastructure for trusted computing, TCG Software Stack (TSS) (see section 3.2) and Trusted Network Connect (TNC) (see section 3.3).
An overview on the activities of the working groups and the most important specifications will be given in section 2.2.4 and 3.
In November 2003 the last major change to the TCG Specification has been published as TPM Main Specification 1.2 [101]. It essentially describes the platform independent functionality that must be provided by a TPM.

Today the TCG has more than 120 members, including component and system vendors, software developers and network and infrastructure companies.
The TCG is incorporated as a non-profit organization with the goal of "the development, definition and promotion of hardware-enabled trusted computing and security technology, including related hardware and software components, across multiple platforms, peripherals and devices" [10].

The TCG compliant technology has become widely available in the market. We list as TPM Manufacturers: Infineon, Atmel, Winbond/National Semiconductors, Sinosun, STMicroelectronics and Brodcom. As system integrators: Intel, Hewlett-Packard, Lenovo-IBM, Dell, Futjitsu-Siemens Computers, Toshiba, Samsung Electornics, Arcom and Densitron. Firms providing TCG-enabled software: Hewlett-Packard, Infenion, Wave Systems, Softex, Uticamo, NTRU Cryptosystems and Lenovo-IBM.

4.2 Trusted Platform Module (TPM) Specification Overview

The TPM Specification is the main part of the TCG Specifications. It defines all platform independent aspects and functions that must be provided by a trusted platform. All system specific aspects have been sourced out to system specific documents like the PC Specific Specification.

4.2.1 TPM Components and Tamper Protection

The TPM provides an RSA key generation algorithm, cryptographic functions like RSA encryption and decryption, a secure random number generator (RNG), non-volatile tamper-resistant storage, and the hash function SHA-1.
The TCG Specification does not prescribe that TPM devices have to be implemented in hardware but to provide the degree of security claimed by the TCG Specifications with a software implementation may be an infeasible task. Thus most TPM implementations are in hardware.
Hardware TPM devices can be compared to integrated smartcards containing a CPU, some memory, and special applications. The assumption is that the chip is tamper-evident (see section 3.1.2) and mounted on (or integrated in) the motherboard [98] such that removal is evident to visual inspection. The main chip contains a special security controller with some internally shielded memory for the firmware, non-volatile EEPROM for the data and RAM. Furthermore, it contains a cryptographic engine for accelerating RSA encryption and decryption processes, a hash accelerator and a random number generator that is needed to generate secure cryptographic keys. Figure 1 shows the main components of the chip.

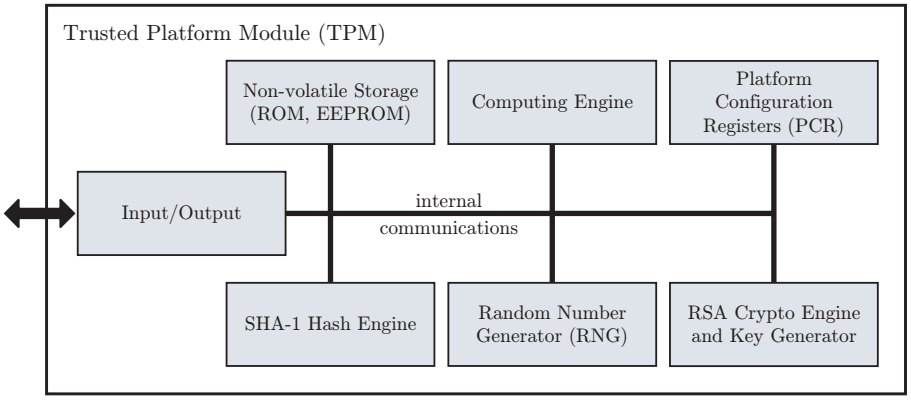

Fig1: Simplified Architecture of the TPM

- The components of a TPM must be trusted to work properly which is intended to be guaranteed by Common Criteria (CC) evaluation and Evaluation Assurance Levels (EAL) [6, 7, 8].

The I/O-Interface of TPMs has not been specified by the TPM Specification but has been left to the platform specific ones. The PC Specific Specification [97] recommends to use the synchronous Low Pin Count-I/O-Interface (LPC-I/O) that is available on most PC motherboards to communicate with the host PC. The connection of the TPM to the motherboard is illustrated in figure 2.
The protocols defining the order of commands and transmissions between the host and the TPM are a command-and-response dialogue, i.e., after every command the host waits for the corresponding response from the TPM before it sends a new request.

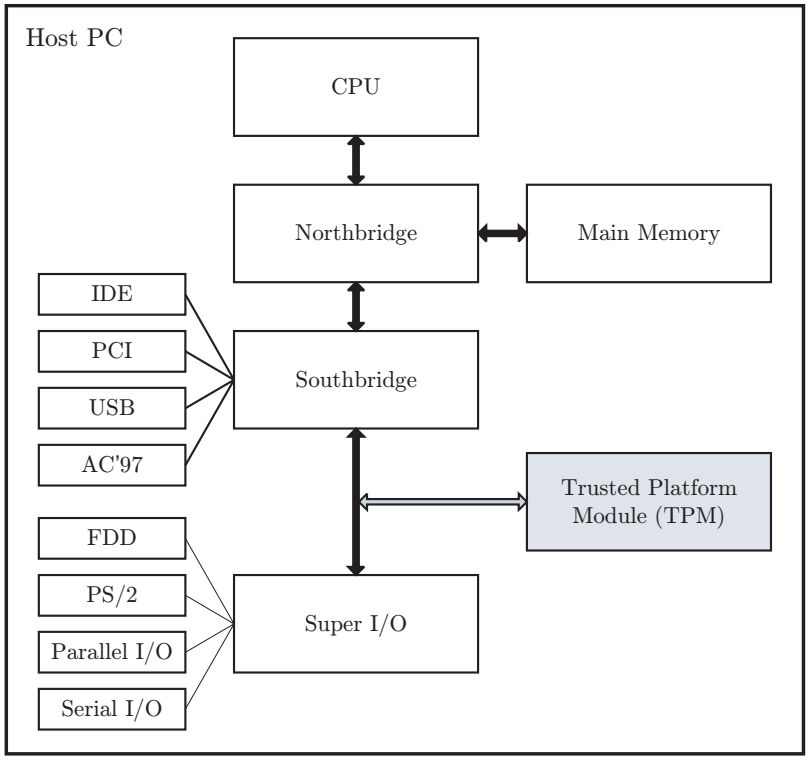

Fig2: Integration of the TPM into a PC platform

Due to the physical properties of the LPC interface, checksums for block protection are not required. To configure the chip, configuration registers can be used to enable or disable functions of the TPM chip, and to configure the I/O addresses for communication with the chip. Data registers are used for data transfer between the host PC and the TPM chip; status and command registers are used to audit and control the performed operations. Depending on the used TPM chip, different layers may exist above the hardware to transport control information, vendor-specific information, or application data (e.g., data to be signed or commands to generate keys).

In terms of compliance to Tamper Protection standards, a TCG compliant TPM should be able to achieve FIPS PUB-140-2 certification [100, 51].

4.2.2 TPM Key Types and Credentials

A TPM contains a Root of Trust of Storage (RTS) which protects data and keys entrusted to the TPM.
The RTS manages a small amount of volatile storage inside the TPM device that is used to hold currently used keys (key slots). Unused keys may be encrypted with a storage key and moved off the TPM chip, e.g., to a hard disk drive. The storage key might be encrypted with another storage key which leads to a key hierarchy as shown in figure 4 with the Storage Root Key (SRK) being the root. The key slots of the TPM are managed by a trusted service outside the TPM which is called Key Cache Manager (KCM).

Key Types

Each TPM protected key is stored with several attributes that identify the type of the key and what it is intended to be used for. These attributes are set during the generation of the particular key and cannot be altered later.

- Storage Root Key: The Storage Root Key (SRK) is used to wrap TPM protected keys which can be stored outside the TPM. This builds a hierarchy of keys represented in figure 4 on an external storage device like a hard disk drive. The SRK is embedded into the TPM and is generated during the process of taking logical ownership of the platform. It can be re-generated by creating a new platform owner which destroys the previous key hierarchy and all the data and keys it contains.
- Endorsement Key: Each TPM device is shipped with an embedded non-migratable Endorsement Key (EK) that has been generated as a part of the manufacturing process in or outside the TPM. Embedded means that the key cannot be removed from the TPM and thus uniquely identifies it and the surrounding platform. The entity that generates the EK issues an Endorsement Credential which should provide evidence that the EK has been properly created and embedded into a valid TPM. Besides the two special keys described above, a TPM can create four different types of asymmetric keys:
- Migratable keys (MK): Migratable keys are cryptographic keys that are only trusted by the party who generated them (e.g., the user of the platform). A third party has no guarantee that such a key has indeed been generated on a TPM.
- Non-migratable keys (NMK): Contrary to a migratable key, a non-migratable key is guaranteed to reside in a TPM-shielded location. A TPM can create a certificate stating that a key is an NMK.
- Certified-migratable keys (CMK): This type of encryption key, introduced in version 1.2 of the TCG specification, allows a more flexible key handling. Decisions to migrate and the migration itself is delegated to two trusted entities, chosen by the owner of the TPM upon creation of the CMK: The Migration-Selection Authority (MSA) controls the migration of the key, but does not handle the migrated key itself. In contrast, the Migration Authority (MA) handles the migration of the key: To migrate a CMK to another platform, the TPM expects a certificate of an MA stating that the key to be migrated can be transferred to another destination. Furthermore, the certificate of the CMK that the owner/user uses to prove that it was really created by a TPM contains information about the identity of the MA resp. MSA.

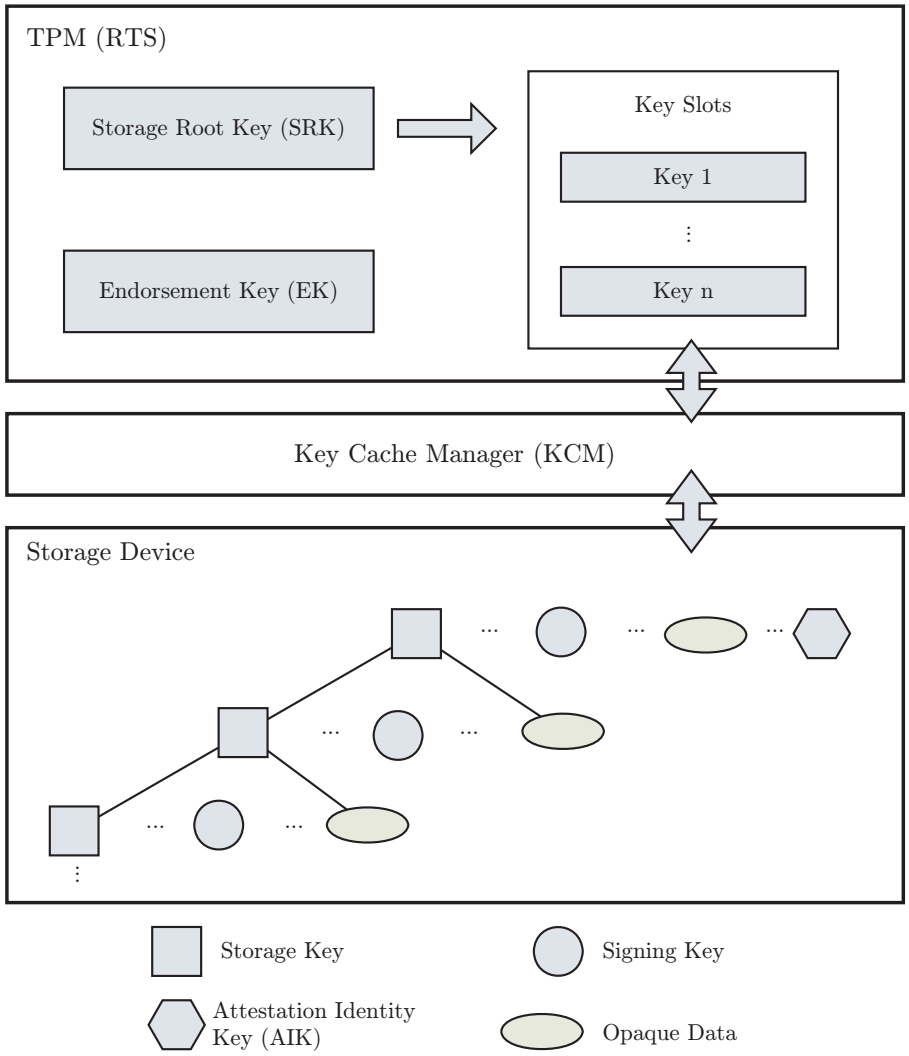

Fig3: TPM Key Management and Hierarchy

- Attestation identity keys (AIK): These non-migratable signature keys provide pseudonymity resp. anonymity of platforms including a TPM. AIKs are locally created by the TPM. The public part is certified by a Privacy Certification Authority (Privacy CA) stating that this signature key is really under control of a secure TPM. In order to overcome the problem that this party can link transactions to a certain platform, version 1.2 of the TCG specification defines a cryptographic protocol called Direct Anonymous Attestation (DAA) [27], eliminating the Privacy CA. AIKs can be used to attest to specific platform configuration states (see section 3.1.4 and 3.1.5). A platform can have multiple AIKs to avoid correlation of platform identities. In order to generate AIKs, the Endorsement, Conformance and Platform Credentials (which are delivered with the platform), the EK and the authorization by the platform owner to use them is required.

TPM Credentials

A trusted platform is delivered with several credentials (digital certificates) that should provide assurance that its components have been constructed to meet the requirements of the TCG Specifications.

Endorsement Credential As already mentioned, the Endorsement Credential should provide evidence that the EK has been properly created and embedded into a valid TPM. It is issued by the entity that generated the EK. This credential contains the name of the TPM manufacturer, the TPM part model number, its version and stepping and the public part of the EK. The EK is a RSA encryption/decryption key pair which is used along with the Endorsement, Platform and Conformance credentials to provide evidence of the platform's identity in a protocol to establish AIKs.

Conformance Credentials The Conformance Credential is issued by an evaluation service (e.g. the platform manufacturer, vendor or an independent lab) with sufficient credibility to properly evaluate TPMs or platforms containing a TPM. It should indicate that the Trusted Building Block (TBB) design and implementation has been accepted according to the evaluation guidelines. A single platform may have multiple Conformance Credentials for multiple TBBs. They typically contain the name of the evaluator, platform manufacturer, the model number and version of the platform, the TPM manufacturer name, TPM model number and version or stepping and a pointer to the location of the TPM and platform conformance documentation. The Conformance Credential does not contain privacy sensitive information or information that can be used to uniquely identify a specific platform.

Platform Credential The Platform Credential is issued by the platform manufacturer, vendor or an independent entity. It should provide evidence, that the platform contains a TPM as described by the Endorsement Credential. The Platform Credential contains the name of the platform manufacturer, the platform model number and version and references to the Endorsement Credential and the Conformance Credentials. The Platform Credential is privacy sensitive since it contains information that can be used to uniquely identify a specific platform.

4.2.3 Integrity Measurement and Reporting

A trusted platform collects information about its current configuration and stores it in a log outside the TPM, called Stored Measurement Log (SML). This enables the detection of modified code and malicious or unwanted software which might compromise the platform's security and thus its level of trust. The information stored in the SML cannot be stored inside the TPM device since it may become very large. Manipulations of the SML will be detected because the digest of the original sequence is securely stored inside the TPM. For this purpose the TPM provides a set of registers called Platform Configuration Registers (PCR) that can be used to store hash values. The TPM hardware ensures that the registers can only be modified as follows: $Ri+1 := SHA1(Ri \,||\, I)$, with the old register value Ri, the new register value $Ri+1$, and the input I. The process of modifying a PCR value is called extending a PCR and ensures that previous will not be ignored and the order of operations is preserved.

The content of the PCRs can be used for verifiable attestation (see section 3.1.5) of the platform's configuration based on Validation Credentials and the chain of trust (see figure 4). Validation Credentials are digital certificates issued by hard- or software manufacturers that provide measurable components (like video and disk storage adapters, memory controllers, processors or software) or other qualified validation entities. They contain the validation entity name, component manufacturer name, component model number, version or stepping and digitally signed reference measurement values taken in a clean-room environment when the component is believed to work properly. The verification of a platform configuration state requires the re-computation of the measurement digest using the reference measurements from the Validation Credentials and a simple comparison of the resulting digest value with the actual content of the PCR. A TPM can attest to a PCR value by digitally signing it with an AIK (see section 3.1.5).

Chain of Trust

As already mentioned, a trusted platform subsequently reports integrity measurement information to the TPM. The idea is that each firm- or software component that is to be loaded or executed is measured before it is started. The result of this check, a message digest, is reported to the TPM in a cryptographically secure manner. Once the value has been submitted to the TPM, it cannot be changed. This means that any change or manipulation of the software state can be recognized since malicious software cannot hide itself by manipulating PCR values or the SML. This implies that the instructions that start the chain of measurements must be trusted which means that they have to function as expected. These instructions are called Core Root of Trust for Measurement (CRTM). Ideally the CRTM would reside in the TPM to profit from its tamper-resistance but due to architectural requirements of the specific platform it might also be located in another device (like the BIOS of the PC platform) which can hardly be manipulated from an remote adversary and should be trusted. After the CRTM measured the system environment consisting of firmware and other components required to give control to the platform's computing engine, which typically consists of the system's CPU, memory and chipset, the CRTM passes control to the Root of Trust for Measurement (RTM). Typically the RTM actually is the platform's normal computing engine which has been previously checked by the CRTM. The RTM inherently generates reliable integrity measurements and reports them to the TPM device building a "chain of trust" as presented in figure 4.

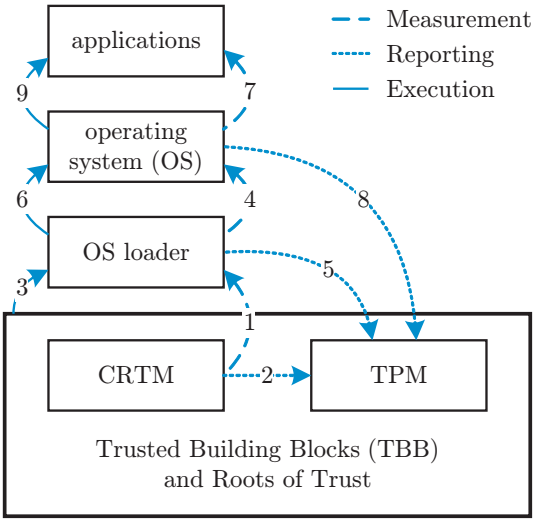

Fig4: Chain of Trust

Trust Assumptions

Trusted Building Blocks (TBB), like the system's main processor (which may function as RTM), BIOS (which may act as CRTM), TPM, memory controller, RAM and the paths between those components that are necessary for integrity measurement and reporting have to be trusted. This means that they have to behave in a way that does not compromise the goals of trusted platforms.

4.2.4 Binding, Signing, Sealing, Sealed Signing and Attestation

In this section, we revisit the basic TC functionalities supported by the TPM, focusing on the specifications of the TCG.

Binding

Binding means that a message can be bound to a certain TPM (and platform) using encryption. When encrypting a message with an asymmetric encryption scheme, the sender uses the public key of the recipient to encrypt a message. The recipient is then able to decrypt the ciphertext with his corresponding private key which can be managed by a TPM. If this private key is a non-migratable key then only the TPM that generated it is able to use the key and thus decrypt the message. Therefore the message is bound to the TPM that protects the corresponding private key.

Signing

By signing a message, the integrity of this message is associated with the key used to generate the signature. This means that a verifier can detect manipulations of a (signed) message and is able to identify its origin by the verification key which might be bound to an identity using a digital certificate. The TPM tags some managed keys as signing only keys. Those keys are only allowed to be used for signature generation. This should prevent them from being used as encryption keys which might comprise security.

Sealing

Sealing is an extension of binding since sealed messages are additionally bound to a set of platform metrics specified by the sender of the encrypted message. These metrics describe a specific platform configuration state that must exist before the decryption of the message is allowed. Therefore Sealing binds a message to a set of PCR values and a non-migratable key protected by a TPM. This provides assurance that protected messages are only recoverable when the platform is in a specific known configuration which is considered to be trusted by the sender of an encrypted message.

Sealed Signing
Signing operations can be linked to specific PCR values and thus a specific platform configuration state. For this reason PCR values are included into the signature. This enables a verifier to inspect a platform's configuration at the time when the signature has been generated. The verifier is then able to decide whether to trust the given platform configuration state and accept the signature or not.

Attestation
Attestation is the process of vouching for the accuracy of information. A TPM can attest to information by digitally signing internal TPM data like PCR values using an AIK. The correctness of this information then can be verified by a third party that checks the integrity measurements and the AIK itself. The AIK can be obtained and verified by using a Privacy CA or a trusted attestation protocol like DAA [27].

4.3 TCG Software Stack (TSS) Specification Overview

The TCG Software Stack (TSS) provides a platform independent software interface for accessing TPM functions [98]. The TSS enables the creation of interfaces for existing cryptographic APIs like MS-CAPI or PKCS#11. This enables TPM support for current and future applications that are using those APIs. In order to take full advantage of a TPM's attestation functions, however, applications will have to support TSS directly.
TSS defines three software interfaces for TCG-enabled software. An overview on these interfaces and some possibilities to make use of them is given in figure 6.
The kernel mode TPM device driver is documented in the TCG TPM Specification. Above the kernel mode driver, a user mode driver, called TPM Device Driver Library (TDDL), provides an operating system independent interface for TPM applications. This separation should ensure that different implementations of TSS are able to communicate with any TPM device and enable the implementation of TPM software emulators as user mode components.
The TSS Core Services (TCS) offers an interface to a common set of platform services like TCG Service Providers or RPC services for communication to a remote TCG Service Provider. The TCS is run as a system process in user mode. It provides services for credential and key management, measurement and event management to handle event log entries and access to PCRs. Additionally it manages access to the TPM device itself since there might run multiple TCG Service Providers in parallel on a single platform.
The TCS must be trusted to manage authorization information which is supplied to the TPM.

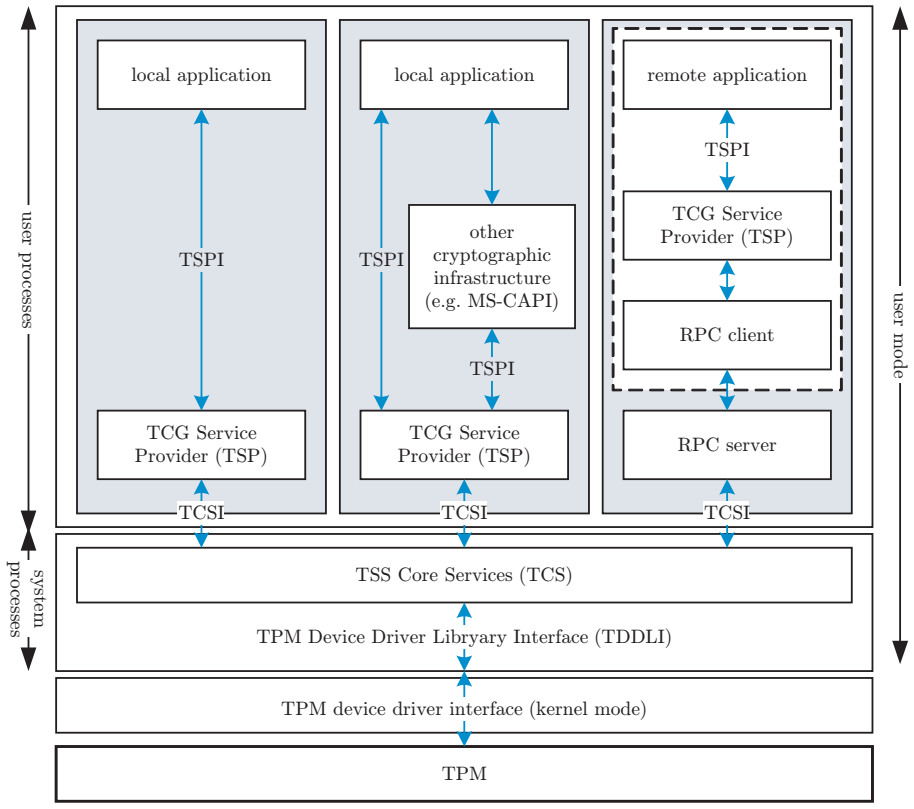

Fig5: TCG Software Stack (TSS) and Interaction Scenarios

The TCG Service Provider (TSP) provides an interface for the C programming language. This interface can be used by applications that make use of TPM features. TSPs provide context management which allows efficient use of application and TSP resources and basic cryptographic functions like the computation of message digests and signature generation. Other cryptographic service providers may use the TSP Interface (TSPI) to communicate with the TPM device. This enables applications not supporting TSP to use TPM functions through their currently supported cryptographic API.

4.4 Trusted Network Connect (TNC) Specification Overview

In May 2004 the TCG published the first specification [98] from the Trusted Network Connect (TNC) Work Group. This specification should be a vendor independent network standard that enhances network security by combining network access control with trusted computing. The goal is to integrate the concepts of trusted computing with existing network access control mechanisms.

4.4.1 Concept of Trusted Network Connect

The overall goal of TNC is to prevent compromise of the hosts that connect to a network or other network resources and thus the network itself. The specification suggests platform authentication through a proof of identity in combination with the integrity status of the platform that wants to connect to a network.

Therefore network access control is based on extended attributes like platform authentication (e.g. by using an AIK), endpoint compliance or software state information which are collected as described in section 3.1.4 and attested to a remote verifier. Based on the integrity information reported by the platform and its proof of identity the verifying instance is able to decide whether it is save to extend the network to that platform.
This kind of network access requires an extension of existing network access policies concerning endpoint compliance, network access, assessment, isolation and repairing. The endpoint compliance policy should establish a level of trust by ensuring that a machine that actually gets network access has a certain software state, e.g., that its operating system, services and applications are of specific versions and running properly. The access policy should also handle machine and/or user authentication before granting network access. A policy for assessment, isolation and remediation should ensure that machines that do not comply with security policies will not be able to get network access and are eventually repaired to meet the security policy requirements again in order to be able to regain network access. This can be done by isolating them in an separate network environment which only allows access to a special server that provides software or virus signature updates for the infected client. Since TNC should integrate with conventional network access control solutions it uses a common three party model consisting of an Access Requester (AR), a Policy Decision Point (PDP) and a Policy Enforcement Point (PEP) as presented in figure 6. The AR might be a VPN Client or 802.1X Supplicant that requests access to a protected network. The AR's request is processed by a PDP which might be a software component of a RADIUS (Remote Authentication Dial In User Service, see RFC 2865 and RFC 2866) server that validates the information provided by the AR (including platform authentication and integrity credentials) against previously defined network access policies. These policies should include integrity and identity aspects of the requesting platform. The PDP reports its decision (access granted or denied) to a PEP which might be a VPN gateway, switch, firewall or 802.11 Access Point that actually allows full or partial network access or denies it completely according to the PDP's decision.

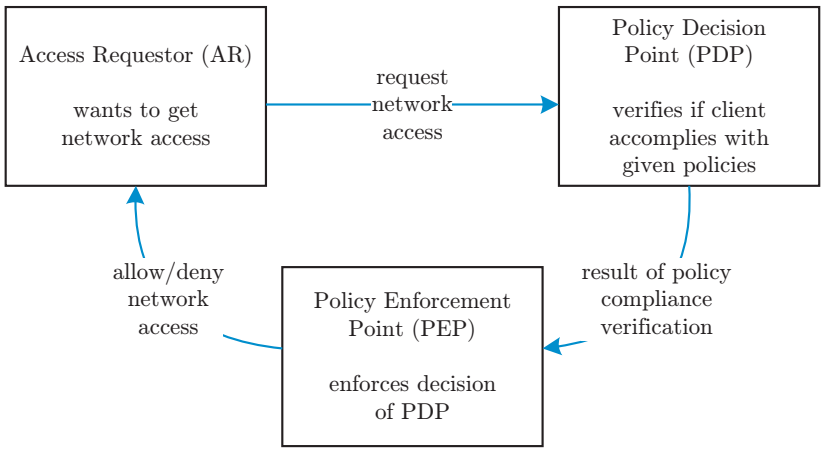

Fig6: The TNC Three Party Model

TNC also aims to enhance existing authentication, authorization and accounting (AAA) protocols developed and standardized by the IETF such as RADIUS or Diameter (RFC 3588)

by adding the ability to measure and report the state of the endpoint platform as part of the authentication and authorization process.

4.4.2 Currently Supported Technologies

There are several technologies available that provide support for the TNC architecture. These include network access, message transport and authentication server technologies of which an overview will be given in the next paragraphs like 802.1x [78].

IEEE 802.1X

IEEE 802 [78] refers to a family of IEEE standards about local and metropolitan area networks. The IEEE 802.1 set of standards (collectively named 802.1X) are mainly concerned with port-based network admission control.
Since TNC has been developed with respect to 802.1X the three party model presented in figure 7 matches to the model of IEEE 802.1X. The TNC AR corresponds to the 802.1X Supplicant which requests access to a network port at an Authenticator which maps to the PDP in the TNC model. The Supplicant will then be authenticated by an Authentication Server that is equivalent to the TNC PEP based on previously defined access policies. TNC will enhance 802.1X in the way that the Supplicant authenticates to the Authenticator by using a TPM protected authentication token (e.g., an AIK) and attested integrity measurements (see section 3.1.4) which should provide evidence of the Supplicant's current system state. The Authenticator then verifies this information and validates it against given policies.

Extensible Authentication Protocol

The Extensible Authentication Protocol (EAP) as defined in RFC 3748 has been originally designed to transport authentication information in the context of PPP and dial-up services. Today EAP is also used in the context of 802.1X and other scenarios. Some EAP methods enable the transportation of integrity measurement information for use with mutual platform authentication. These are especially TLS based EAP methods (like EAP-TLS or PEAP) which can be extended by using the TLS-Attestation Extension protocol.

Virtual Private Networks

Virtual private networks (VPN) are virtual communication networks usually used to communicate securely over a public networks like the Internet. Many VPNs use the IPSec protocol to achieve authenticity, integrity and confidentiality of the transported data. IPSec often uses the Internet Key Exchange (IKE) [48] protocol for authentication and key exchange purposes. IKE can be enhanced by communicating integrity information as part of the mutual authentication and key establishment by adding integrity reporting after the IKE peers have authenticated. The next version of IKE, IKEv2 [65], supports EAP-based messages for peer authentication which can be used in combination with other EAP based approaches proposed by the TNC architecture that will be described below.

Despite of IPSec-based VPNs there are also SSL- or TLS based VPNs. TLS provides security services on the transport layer. This protocol is mainly used to secure the communication between web servers and browsers by offering authentication and confidentiality. TLS can be enhanced to provide endpoint integrity by the TLS-Attestation Extension [99] protocol for delivering integrity measurements between client and server.

Point-to-Point Protocol

The Point-to-Point Protocol (PPP) is commonly used to establish a direct connection between two nodes and is used by most Internet service providers for dial-up access to the Internet. Typically EAP is used over PPP to transport security related parameters which enables the use of the EAP-based approach described by the TNC that will be presented below.

Transport Layer Security

As already mentioned above, TLS can be extended for using integrity measurements for authentication purposes through the TLS-Attestation Extension protocol. TLS and SSL can be used to secure application-related communication including various web services protocols like SOAP. Thus HTTP can also be used to transport integrity measurement information when protected by TLS.

RADIUS and Diameter

RADIUS (RFC 2865) and Diameter (RFC 3588) are standardized authentication protocols. As already described above, RADIUS is a widespread authentication protocol which has been considered in the development of the TNC architecture. Extensions for Radius to support EAP are defined in RFC 3579 which enables the use of EAP for transporting integrity measurement information between the Clients (TNC AR), the Authenticator (TNC PEP) and the Authentication Server (TNC PDP).

Diameter is an alternative to Radius. It uses attribute value pairs (AVP) to deliver various payloads relating to user authentication, service authorization, resource usage, etc. The Diameter protocol can be extended by adding new AVPs. Therefore it can be enhanced by introducing a new AVP for carrying integrity measurement information from the AR to the PDP.

5 Trusted Computing beyond the TCG Specifications

As explained in chapter 3, Trusted Computing comprises a set of functionalities and mechanisms based on which a system can be regarded as "trusted". The TCG aims at defining a set of specifications related to the TPM, which is crucial for providing the Trusted Computing features on a platform. However, some important Trusted Computing concepts are not directly addressed by the TCG specifications. In fact, many scientists argue that adding a TPM to a platform with an unmodified mainstream operating system does not introduce a considerable improvement to the security of the system.

Suitable operating systems support for TCG specifications seems to be necessary. Process isolation and secure input/output paths between applications and users are examples of important features that need to be provided by the system in order for the Trusted Computing TPM functionalities to be efficient. For this reason, security kernels able to run several environments in parallel – based on virtualization – are being considered to provide such features. More over, manufacturers are introducing security enhancements to their CPUs, input and output devices to provide such environments isolation on a same platform.

In the following, we give an overview of the different efforts for providing operating systems and hardware support for the TCG specifications.

5.1 Operating Systems Support and Secure Platforms

5.1.1 Microsoft Next Generation Secure Computing Base (NGSCB)

Microsoft Corporation (http://www.microsoft.com) is an Initial Promoter Member of the TCG. Microsoft planned to integrate in their latest operating system, Windows Vista, (formerly known as Longhorn) security features based on a new software architecture named Next Generation Secure Computing Base (NGSCB) [77]. It would provide native support for TCG 1.2 compliant TPM hardware including TPM management features and TPM based file and folder encryption. Older TPM versions will only be supported through third party TSS implementations.

Originally, Windows Vista was designed in a way to contain a security kernel, named Nexus, which aims at providing a secure computing base that runs in parallel to the regular Windows environment. This enables the execution of traditional Windows software and Nexus security enhanced software. Nexus would provide features such as strong process isolation, sealed storage which enables applications to lock information so that it can only be accessed by themselves and a secure path to and from the user by providing secure channels from the keyboard or mouse to the Nexus enhanced application and from the Nexus enhanced application to the screen. It also offers attestation that provides assurance to third parties that a piece of data has been created and signed by a secure application. In order to provide these features, Microsoft's NGSCB extended the features defined in the TCG Specifications including changes to secure booting, hardware like memory controllers and CPUs to enable strong process isolation and changes to input and output devices that enable a secure path from and to the user.

However, in May 2004, Microsoft decided to reconsider the security features of NGSCB, and Windows Vista was release without most of them. The security features are expected to available in the next years. However, Vista still includes the "BitLocker" which makes use of

TPM to provide secure boot and hard disk encryption. In the same context, Microsoft's "Singularity Project" is an ongoing attempt at providing isolated processes for security and reliability [142].

5.2 Supporting Technologies: Virtualization

Virtualization technology allows several guest machines (i.e Virtual Machines) to share one host machine using a Virtual Machine Monitor (VMM). In contrast to other sharing mechanisms, the virtual machine monitor offers to its virtual machines the same interface that is provided by the host machine or by a different native host. This means that several legacy operating systems can be executed on top the VMM. Therefore, security kernels that provide virtualization are also denoted as secure hypervisor or secure VMM. Virtualization is a very efficient mechanism to build secure IT-system, since it ensures vertical isolation of the several virtual machines.

The combination of Trusted Computing functionalities with virtualization techniques seems to be a promising approach for enhancing the platforms security. The isolation features provided by virtualization techniques is necessary to efficiently provide the Trusted Computing functionalities. In the following, we give an overview of current virtualization methods.

5.2.1 Hypervisor-based virtualization

Hypervisor-based virtualization ensures full virtualization in the sense that the virtual machine simulates enough hardware to allow an unmodified "guest" OS (one designed for the same CPU) to be run in isolation.

An example is Xen [3], an Open Source virtual machine monitor (VMM), which has been developed by the University of Cambridge and is currently distributed under the GNU General Public License (GPL). Xen is able to execute multiple guest operating systems in a virtual machine (VM) by providing an abstract layer above the computer's hardware. It supervises and manages the distribution of resources like CPU time or I/O cycles to the guest systems. In order to be able to work with the Xen architecture, the guest operating system's kernel has to be modified.

IBM Research integrated sHype [79], a hypervisor security architecture, into Xen. sHype provides flexible access control enforcement to strongly isolate and control the sharing of hardware resources and the communication between different VMs. More information about sHype can be found at [79].
Since a TPM is not a device that was designed to be accessed by multiple operating systems at the same time, IBM Research is about to develop a virtual TPM architecture in order to provide TPM support to all operating systems running on Xen. To realize this, the TPM command set defined in the TCG 1.2 Specification has been extended by virtual TPM management commands which enable virtual instances of TPMs that can be transparently used by all guest operating systems. Thus software intended to work with hardware TPMs is supposed to continue to work without any changes when executed on Xen.

5.2.2 VMWare

Instead of modifying legacy operating systems such that they can be executed on top of the underlying abstraction layers, it is also possible to implement a software-based virtual

machine monitor allowing the reuse of unmodified legacy operating systems. A common example of software-based VMMs is VMware [VMWa07].

A VMWare Workstation is able to emulate the complete set of hardware (video adapter, network adapter, harddisk adapters) within the virtual environment to the guest operating system. The aim is to let more than one operating system run simultaneously on the same physical machine by sharing its resources.

However, from a security point of view, VMWare virtualization is not efficient as hypervisor or microkernel based virtualization. It lacks the ability to enforce process isolation on the memory level [webb06] which is usually provided by separation kernels.

5.2.3 Microkernel-based virtualization

A microkernel is a minimized operating system kernel that provides only essential services such as logical address spaces, tread management and inter-process communication (IPC). Processes on top of the microkernel run in their own address space and are therefore strongly isolated from each other. Especially, high-privileged code runs in protected memory, which isolates it from potential intrusions.

In a microkernel-based virtualization, the largest part of the VMM is realized by a process running in user mode. This reduces the complexity of the code running in supervisor mode. Only non-essential services such as networking, display and device-driver processes are run in user-space. The supervisor mode can still perform any operation the hardware can, such as writing to write-protected memory and switching to arbitrary address-spaces. Hence, device-drivers processes for example must invoke the kernel to perform privileged operations for them, allowing the microkernel to check for safety and security.

5.3 Industrial and Academic Open Source Projects

5.3.1 European Multilaterally Secure Computing Base (EMSCB)

The EMSCB project [16, 1], which is partly funded by the German Federal Ministry of Economics and Labour, is an instance of PERSEUS focusing on multilateral security. This means the development of security critical services such that security policies of all involved parties are not violated. EMSCB aims to develop five demonstrators with a micro-kernel-based security architecture using a TCG 1.1b compliant TPM. These five demonstrators will offer hard disk encryption, secure VPN, DRM application, a multilevel security system and a prototype of DRM application that is intended to be used in an embedded automotive platform provided by Bosch/Blaupunkt. Additionally an EMSCB platform consisting of TPM-enabled security critical services like a trusted GUI, trusted storage, application manager, TPM driver, secure bootloader and a security policy manager that uses the trusted GUI which will enable users to define the locally enforced security policy, will be implemented. The source code of the EMSCB platform will be published under an open-source license.

More information is available in [16] or at http://www.emscb.de.

5.3.2 Open Trusted Computing (Open TC)

The Open Trusted Computing (OpenTC) consortium is concerned with the development of a secure computing system for conventional computers and embedded systems based on open source software and trusted computing technology. A major goal is to implement an open trusted computing framework supporting TPMs and security enhanced next generation processors from AMD and Intel. This includes the development of a secure operating system that is able to use the features provided by TPMs and security enhanced processors. It will contain virtualization layers, a Trusted Software Stack for Linux and management software for the TPM. OpenTC will also take care of the development and implementation of protocols for policy management, including distributed policy enforcement, security state monitoring and management and network and configuration management. The project also works on the integration of trusted computing into existing public key infrastructures.

More information is available at http://www.opentc.net.

5.3.3 TrouSerS

TrouSerS [4] is an open-source TCG Software Stack implementation for Linux operating systems which has been developed and released by IBM. It is provided under the Common Public License (CPL).

Currently usable TPM services provided through the TSS interface offered by TrouSerS are RSA key pair generation, RSA encryption and decryption using a PKCS#11 compatible interface, RSA signing and verification, storing data in the TPM's PCRs and logging, sealing of data to arbitrary PCRs, random number generation and secure RSA key storage.

TrouSers also includes the tpm-tools applications (or suite if helper applications), which can be used for command line based TPM management.

5.3.4 Enforcer

The Enforcer project provides a Linux Security Module which improves the integrity of a computer running a Linux operating system by ensuring that there has been no tampering to the computer's file system. The software is able to detect if a file has been changed and to take formerly specified actions when any tampering has been detected. Enforcer is able to store secrets to an encrypted loopback file system which is protected by a TPM. If Enforcer detects a tampered file, the encrypted loopback file system will be dismounted automatically. It is also possible to bind specific files to specific applications so that an application is not able to modify files it is not intended to.

5.3.5 Trusted Linux Client

IBM has implemented a Trusted Linux Client [83]. It utilizes the Linux Security Module (LSM) that mediates all security relevant features of the Linux kernel, kernel modules and a TPM. The Trusted Linux Client implements three kernel modules. These are the TPM kernel module, the Extended Verification Module (EVM) and the Simple Linux Integrity Module (SLIM). The EVM Module provides policy based verification functions based on authenticated attributes which can be seen as a file's security meta data which is checked on open or execution of the file. Additionally, it offers symmetric key-based verification functions which use the TPM as secure key store, and access control based on integrity

containment which uses the SLIM module for access enforcement. The TPM is used to verify the integrity of the EVM by verifying that all its files are authentic. The SLIM uses the results from EVM's file verification in order to give trusted process authority to those files which meet all the file verification requirements, only. Hence, the SLIM manages the state of applications during their execution which means that it classifies them into trusted and untrusted ones.

The TLC offers a kind of "pre-boot" authentication. It is not pre-boot authentication exactly because the system must boot until the initial ram disk has been loaded which might contain the kernel master key which is needed for authentication. The kernel master key is a randomly TPM-generated key which is protected by a TPM. Every user of the system must know this key in order to boot. Multi-user authentication can be realized by assigning a unique authorization password to a unique copy of the kernel master key to each user. The TPM then only releases the kernel master key to the kernel if the PCR values of BIOS, master boot record, GRUB bootstrap loader, Linux kernel and initial ram disk have not been altered. The kernel master key can be stored in the initial ram disk or on an USB flash drive protected by a user authorization password that has to be presented at boot time. The authentication mechanism also supports the use of fingerprint readers.

The Trusted Linux Client has been implemented and tested on Fedora Core 361 and Red Hat Enterprise 462 systems.

5.3.6 tcgLinux

The tcgLinux [80] project is run by IBM Research. Its goal is the development and implementation of a TPM-based Linux run-time attestation. This basically means the generation of verifiable information about the software stack running on a Linux system which can be used by remote parties to determine the integrity of the execution environment of this specific system. tcgLinux therefore implements an integrity measurement feature which is an adapted variant of the Linux kernel that is able to measure each executable, library, or kernel module before it is loaded and executed.

A TPM is used to securely hold an integrity value over the measurement list which is managed by the kernel. The current research mainly considers on how to take advantage of the validated measurement list to justify the security properties of a system's runtime environment.

For more information refer to [80].

5.3.7 PERSEUS Architecture and Turaya security kernel

PERSEUS [15, 2] was developed at the University of Saarland in 1999 in cooperation with IBM Research, Zurich, and is currently pursued by eurobits (European Competence Center for IT Security) at the University of Bochum, Germany. The PERSEUS Project provides an open computing platform offering a basis for the realization of multilateral security based on trusted computing. It is intended to support a wide range of hardware platforms like PC, PDA and embedded systems. PERSEUS uses a micro-kernel which only contains elementary functions like process management, memory management and interprocess- communication and thus minimizes the security-relevant part of a system which allows formal verification of its implementation. The PERSEUS architecture realizes security critical applications like

digital signatures or DRM applications and conventional operating system as separate processes which are only able to communicate to each other or to the system's hardware by involving the PERSEUS security kernel.
PERSEUS will enable the realization of policy enforcement, e.g. enforcing license agreements (if accepted by the user) or permitting access to information services only for payment but prevent content providers to gather more private information about the user as needed to provide their service. Thus PERSEUS may serve as basis for DRM solutions and provides compartment mode security to prohibit access to documents outside a desired workflow. It can also be used to realize a secure multi server system which is able to run different isolated services like a web server, database and firewall in parallel on the same hardware.
The security software layer of PERSEUS consists of a micro-kernel based on the L4 micro-kernels which have been developed at the University of Karlsruhe and the Technical University Dresden. On top of this kernel, a resource management and access control layer has been placed which controls the distribution of the system's hardware resources to the conventional operating system and all security relevant applications above it. PERSEUS also provides a secure bootloader which has been implemented as a TCG-enabled version of GRUB (Grand Unified Bootloader) to assure that a specific operating system configuration is booted. A secure user interface (Secure GUI) will provide trusted paths between the user and secure applications resp. secure applications and hardware. PERSEUS's application manager will ensure the controlled installation and update of software since this offers the possibility to infiltrate malicious code.
Conventional operating systems usually run applications with the rights of the user who started them. This results in running applications with more privileges then they actually need. PERSEUS's application manager will care about privileges of applications since these should be minimal. A trusted computing service will allow virtualization of trusted computing hardware features which will offer attestation and sealing functions to applications.

5.3.7.1 Turaya security kernel

Turaya is the published version of the PERSEUS security software layer. The Turaya security kernel is a small security software layer which can be logically divided into a hypervisor layer and a trusted software layer.

The main task of the hypervisor layer is to provide an abstract interface of the underlying hardware resources like interrupts, memory and hardware devices. Moreover, this layer allows sharing these resources and realizes access control enforcement on the object types known to this layer. Currently a microkernel is used as the foundation of the hypervisor layer.

The trusted software layer builds on the hypervisor layer and offers a Trusted GUI, which controls the graphic adapter and the input devices, i.e., mouse and keyboard, to establish a trusted path between the user and an application. The Trusted GUI labels application windows with unique application names. Moreover, the Trusted GUI enforces a strong isolation between applications on the GUI level. Unauthorized applications cannot, for instance, access the graphical output of other applications or fake their interface to look like the usual password dialog. The Application Manager loads applications and measures the integrity of applications. These integrity measurements can then be reported to local applications as well as to remote applications. In cooperation with Trusted Computing hardware this functionality

constitutes the basis for elaborate Digital Rights Management applications. The advantage of this approach in contrast to other integrity measurement architectures (e.g., [SZJv04]) lies in enhanced end-user privacy protection and improved manageability, e.g., of software updates. The Storage Manager enables other applications to persistently store their local states. It preserves the integrity, confidentiality, and freshness of the managed data such that only the application or the user having produced the data may later re-access it. Finally, the User Manager identifies and authenticates the users and assigns roles to the users. The user management is not part of the insecure operating system to prevent malicious software from "sniffing" user passwords or "stealing" the user identity.

5.4 More Trusted Computing Hardware

5.4.1 Intel LaGrande/TXT Technology

Intel currently develops a set of security enhancements to conventional computer hardware called La- Grande Technology [37]. It basically follows the same ideas as Microsoft's NGSCB and thus probably is intended to become Microsoft's target platform. LaGrande is a set of hardware capabilities that enables to run applications in isolated execution environments. This prevents other applications from unintended observation or compromise of data. An additional extension aims at the protection of input and output to provide a secure path from and to the user by encrypting the corresponding data streams. The system's processor will offer features to realize a protected execution environment whereas the chipset will include mechanisms to enforce memory protection policies, protected channels from input and to output devices and an interface to the TPM device that serves as protected storage and provides platform attestation features.
More information about LaGrande Technology can be found in [37] and at http://www.intel.com/technology/security/.

5.4.2 AMD Pacifica/Presidio

Founded in 1969 and headquartered in Sunnyvale, California, USA, AMD (http://www.amd.com) develops and produces microprocessors, flash memory devices and low-power processor solutions for the computer, communications and consumer electronics industry. AMD is an Initial Promoter Member of the TCG.

AMD develops a security and virtual machine architecture, called Pacifica [17, 105], to be integrated into their next generation processors. Basically this technology will enable a single computer to efficiently run several operating systems in parallel. Therefore a virtual machine monitor (VMM) or hypervisor software is needed to control the execution of the different operating systems and to manage the allocation of hardware resources. With Pacifica AMD provides hardware support for virtual machine architectures. In contrast to conventional VMM solutions which have to emulate the virtual computer in software, Pacifica allows a guest operating system to access the host system's hardware directly which results in higher performance.

Intel offers a similar technology, known as Vanderpool. Due to several design differences between AMD and Intel CPUs these technologies will not be compatible.

Conventional VMM have to manage the virtual machines' main memory in software. Usually this is realized by a table which enables the translation of virtual memory addresses into physical ones. In contrast to Intel, AMD integrates the system's memory controller into the CPU core. Thus Presidio can do this translation in hardware which additionally increases the performance of virtual machines.

Another problem arising from direct hardware access by virtual machines is that DMA enabled devices have the possibility to access the system's main memory without using the CPU. To counterfeit this problem, the memory controller that is integrated into the CPU, has been enhanced by a feature called Device Exclusion Vector (DEV).

Another difference between Intel's Vanderpool Technology and AMD's Pacifica is the support for TPM features to VMs. This enables verifiable startup of trusted software inside the VMs and a secure virtual computing environment.

According to several presentations AMD works on a technology called Presidio which seems to be a security architecture for their next generation CPUs. It is presumably similar to Intel's LaGrande Technology already described above. Actually information about this technology is barely available.

5.4.3 ARM TrustZone

ARM (http://www.arm.com) develops and produces electronic devices for consumer entertainment, wireless and networking solutions including automotive, security and storage devices. ARM's product portfolio includes RISC processors, embedded memories, peripherals and software and development tools. ARM is a Contributor Member of the TCG.

TrustZone

TrustZone [104, 52] is a set of security extensions integrated into ARM's CPU cores that have been designed for mobile phones, PDAs or set top boxes. The TrustZone security solution consists of hardware extensions which provide a secure execution environment in parallel to the normal one. It includes secure software offering basic security services such as cryptography, safe storage and integrity checking.

The basic idea behind TrustZone is the isolation of conventional non-secure applications from secure ones running in a protected trusted environment which can be switched with the normal runtime environment as required. This is handled by a software component called TrustZone Monitor which communicates with the conventional operating system and a secure kernel that provides the secure computing environment. ARM provides this secure kernel together with secure drivers, a secure boot loader and secure software services. These include identification and authentication of a platform, management features for identities, cryptographic keys and certificates as well as I/O access control, secure data storage, basic cryptography functions and code and integrity checking.

Detailed information about TrustZone can be found in [104] and [52] or at http://www.arm.com/products/esd/trustzone_home.html.

6 Application Scenarios for Trusted Computing Technology

Trusted Computing allows the realization of a variety of business models relying on distributed trusted third parties, or a considerably more efficient configuration of some of the existing applications. In the following, some interesting applications of increasing importance will be presented.

6.1 Application Areas

6.1.1 Distributed Policy Enforcement

Existing technical measures of copyright handling on digital content resp. services (see, e.g., [82, 46, 26]) on end-user devices only registered moderate success, since most of the technical solutions can be totally controlled by the end-users due to the lack of appropriate protection in hardware and software. Experiences in the past have shown that hardware solutions (e.g. dongles) cannot be established because of their high complexity, incompatibility, insufficient security, and limited user acceptance [18]. Moreover, a variety of these techniques were treated as trade secrets; a strategy which contradicts the cryptographic principals, because security should not rely on the secrecy of an algorithm but on the secrecy of a secret parameter (e.g., a cryptographic key). In spite of non-disclosure and legal threats by content providers, most of the methods have been broken in the past (see, e.g., [18, 84]).

In contrast to existing insecure solutions, the features offered by the TC hardware combined with an appropriate software security layer provides the appropriate basis for the realization of more secure applications. For instance, license agreements can be enforced if these were accepted by the consumer of digital content: On the one hand, it is ensured that users of online-information (e.g., traveling or navigation information, electronic magazines, etc.) can get access to the desired information only against payments, and that they cannot arbitrarily distribute this information to others. On the other hand it can be prevented that providers get more private information about the user than they actually need for providing their service.

Possible applications with short term potential are copyright protection, eLearning, eBooks, geographical information systems, as well as the area of Telematics in car navigation systems.

Another field of application is the long term high sale expecting area of providing multimedia content, e.g., video and audio data. Here trusted platforms will considerably complicate the unauthorized distribution of digital content.

6.1.2 Compartmented Mode Security

Business processes between companies often require the exchange of sensitive data and documents (e.g., financial accounting, patent motions, technical cooperation), whose usage is regulated by contracts (e.g., through secrecy acknowledgments). Company internal protection measures are essential as well, so that access on documents outside the desired workflow is prohibited. This, for example, shall prevent that employees read sensitive documents, distribute documents (accidentally or purposely) outside the company or perform unauthorized changes.

Existing computing platforms cannot securely handle classified documents (e.g., unclassified, secret, top secret), so that users can circumvent control mechanisms by using available functions for their own purpose or by exploiting known security holes of existing software components.

Business processes between companies often require the exchange of sensitive data and documents (e.g., financial accounting, patent motions, technical cooperation), whose usage is regulated by contracts (e.g., through secrecy acknowledgments). Company internal protection measures are essential as well, so that access on documents outside the desired workflow is prohibited. This, for example, shall prevent that employees read sensitive documents, distribute documents (accidentally or purposely) outside the company or perform unauthorized changes.

Existing computing platforms cannot securely handle classified documents (e.g., unclassified, secret, top secret), so that users can circumvent control mechanisms by using available functions for their own purpose or by exploiting known security holes of existing software components.

Many security problems occur, because companies or public departments are not able to successfully prevent their users to (accidentally or purposely) break the security policies. They are able to install software components on their own or manipulate the IT system otherwise, which leads to potential security lacks, e.g., through viruses, Trojan horses, worms and configuration errors.

TCG enhanced platforms will provide functionalities that allow to securely enforce external and company wide security policies. This is the basis for the realization of a system with Multi Level Security (MLS), which is customized by practical conditions. Existing MLS solutions are not satisfactory up to now because of their high complexity resp. inefficient configuration (strictly separated hardware).

Another important example application, which will be realizable in association with a secure computing platform, are Multi Server Systems (MSS), which run, like virtual machine monitors (VMM), different isolated services (e.g., a database, a web server, and a security gateway) in parallel on a single server.

6.1.3 Secure End-User Systems

Today, a standard personal computer or mobile device, with an off-the-shelf operating system and all the software that one mainly buys for this system, is not secure at all, particularly in the context of digital signatures, eCommerce and eGovernment. Different applications of the same user are not protected from each other and the end-users are confronted with frequent security updates. Moreover, almost all data may nowadays carry executable code and the execution often starts without knowledge of the computer owner. Hence, it is impossible to administer a standard end-user system such that a critical application is protected from all others.

Trusted computing platforms will offer secure booting and authentication mechanisms which are a necessary and sufficient basis for security relevant applications like secure signature generation, home banking or eGovernment and eCommerce applications.

6.1.4 Embedded Security

Another important application area for trusted platforms arises due to increasing integration of computer platforms in different products and devices (embedded systems), e.g., as done by the automotive industry.

The high complexity of the used software leads to higher error probability, which can be compensated by the use of a security kernel. Furthermore, the integration of information- and multimedia systems in cars will play an important role in the future, which will offer new business opportunities for suppliers and manufacturers.

6.2 Existing Uses Cases and Future Scenarios

6.2.1 DRM

Digital Rights/Restrictions Management (DRM) refers to several concepts to restrict arbitrary use of data and to limit it to accordance with a certain defined policy [136]. TPMs provide functionality that can be used by implementations of DRM systems.

Since DRM technologies are considered by many as the most prominent use case of TC technologies, we shortly describe in the following the main types of DRM

The following types of DRM and how TPMs can be of use will be outlined: DRM in companies or administration, DRM for personal files, DRM for media files, and DRM for software products.

6.2.1.1 DRM in companies or administration

In companies (or administration), customer data can be protected using DRM. As one example, EPAL [137] (Enterprise Privacy Authorization Language) offers a mechanism to tag data with a policy that defines what processing is allowed or interdicted. Provided that the company's system enforces such policies, this results in a higher level of privacy protection. For an effective enforcement, Trusted Platforms could be valuable.

Using TC, a company's IT department can also detect when a PC's configuration got changed. A PC detected as compromised could then be excluded from company network communication until system administration had a closer look at it. This would protect the data from attackers within a company.

Apart from protecting customers' privacy, this could also help against other sensitive company information leaking to unwanted targets. Using the cryptographic support of a TPM, even in case a file leaked outside a company, it would still be protected against unauthorised reading. Incidents such as an employee selling customer data to spammers [138], would become more unlikely.

Critics argue that freedom of press [139] would be affected, as it would also be nearly impossible to discover misbehaviour of a company or administration through leaked information.

In case the whole infrastructure of Trusted Platforms of a country would be controlled by its government like within a company, independent information from the Internet could be made

unavailable by defining a policy that allowed only content to be opened that has already been tagged by a certain government authority.

Currently, Trusted Platforms has become common. The TCG industry consortium is pushing it into the market. The Article 29 Data Protection Working Party of the EU in its Working Document concludes, that "the use of TPM […] is likely to become a de facto standard, a necessary feature to participate in the information society. This could have consequences not only in the field of data protection, but also regarding other human rights aspects such as the freedom of speech." [128]

6.2.1.2 DRM for personal files

The cryptographic support of a TPM can enable an application such as e.g. a word processor to secure its files against unauthorised reading. Files could be encrypted so that they can only be opened on the original PC.

Problems will occur when a user finds his TPM being destroyed or malfunctioning. It will be very hard to restore the data, if there is no unencrypted backup of the data (security risk) or no backup of the secret cryptographic keys from inside the TPM. As even in TPM specification 1.2 not all keys are migratable yet (although privacy protection authorities have demanded that users shall have complete and exclusive control [129] of their IT systems), users currently take the risk to irrecoverably lose their data when using TPM-supported DRM for personal files.

Problems might (will) also occur, when a user buys a new PC and wants to transfer his files, as the two PCs will have different TPMs.

6.2.1.3 DRM for media files

Media content providers want to use DRM mechanisms to protect their products against uncontrolled copying and distribution.

On conventional PCs, DRM protection of media files usually fails as there are too many holes for the protected content to slip through. For example, virtual sound or video devices can write unencrypted media data to a file that then is no longer protected. Additionally, analogue copies of the content can not be prevented.

Thus, there is a high interest in turning PCs into tamper proof devices. A PC with a TPM and an OS supporting its functionality (a Trusted Platform) can act as such a tamper proof device. Using a unique public key generated by the TPM to encrypt a media file shall then ensure that it can only be played on the device with the TPM that possesses the unique secret key. In the long run, DRM will work properly only if the control over a platform is taken away from the user to protect its mechanisms from being circumvented or broken.

But even then there might still be some holes left. Therefore another DRM mechanism tags the unencrypted content with a unique identifier (fingerprinting, watermarking [130]). If an unwanted copy of a media file is found, this allows content providers to trace it to the owner. Of course this works only if there is a database that stores *who* has used what unique public key.

This is a big privacy issue, because as to be able to trace back any unwanted copy at any time, the information on who has purchased which media file has to be stored forever. As a result generating detailed profiles of users will be possible.

TPM specification V1.2 introduced a protocol for Direct Anonymous Attestation (DAA [131]), but to ensure to be able to track the source of an unwanted copy, content providers are not interested in offering anonymity to their customers.

DRM for media files therefore does not benefit home users in any way, but – as long as DRM concepts aim to identify them – is a severe threat to their privacy.

6.2.1.4 DRM for software products

As with DRM for media files, DRM for software products is intended to disable uncontrolled copying and distribution. Furthermore, on a Trusted Platform, software can protect itself by restricting the right to manipulate its code or environment.

As an example, cheating in games is both annoying and common. Game developers try to protect their products from cheaters, but as they can not control the environment the game runs in, cheaters usually find ways to manipulate.

While game developers can not directly control a Trusted Platform either, they can let the game use the TPM (the OS has to support this) to detect manipulation and react accordingly.

The same applies to any other software as well, word processors, spreadsheet analysis programs, media players, and even the OS itself. While this implies an increased resistance against malware such as viruses, Trojan horses, or spyware, it also means that once a virus etc. managed to install itself onto a Trusted Platform (e.g. by tricking a user with administrative privileges to click “OK” and authorise himself against the TPM), it can use the same TC mechanisms to protect itself against manipulation, e.g. removal.

End User License Agreements (EULAs) that a user has to approve during software installations currently are not valid in Europe and many other parts of the world [132]. While users of those locations currently can simply ignore those EULAs, software can use TC to enforce such license regulations [133] ignoring the legal situation.

As software can not only use TC to protect itself against unwanted manipulation, it can also do so against manipulations that are intended by the user. For example, users of a certain OS currently can use programs such as xp-AntiSpy to disable mechanisms that automatically connect to the Internet without notice and uncontrollably transfer data to remote servers [135]. Using TC, such manipulation not intended by a software company could be prevented.

Manufacturers of image editing software could protect the product not only from patches against the included Central Bank Counterfeit Deterrence Group (CBCDG) module that prevents users from manipulating images of banknotes [134], but also from other enhancements from independent developers that did not e.g. pay a fee. This would work just as the little chips in some ink cartridges that shall protect the printer aftermarket from third party products [135]. Again, this only works properly if control over the PC is taken away from the user to protect against circumvention.

The ability of manufacturers to protect themselves against competition with third parties can on a technical level lead to the same negative monopoly effects that software patents can cause on a legal level.

6.2.2 Anonymity Services

Despite the controversial discussions about possible further intentions or usage of Trusted Computing to enforce Digital Rights the same mechanisms can also serve to improve confidence in anonymity services. A look on TC related anonymity issues in the context of identification could be found in 8.6, but we shortly describe in the following the implications of TC on general anonymity services.

We refer to an anonymity service as a service which allows its users to access a certain resource in the Internet anonymously. This could be surfing the Web anonymously, sending anonymous e-mails or establishing anonymous TCP/IP connections. We assume a client-server network architecture. The anonymity service itself consists of a number of servers and the end-users need some client software to use the anonymity service. Examples of such kind of anonymity services are *AN.ON* [108], *Tor* [112] and *Mixminion* [111] which are based upon the theoretical concepts of Mixes developed by David Chaum [110].

The security of such anonymity services depends on the fact that not *k* out of *n* anonymity servers collude and that the attacker will not get enough secret information (e.g. cryptographic keys) to deanonymise a user.

Remark: If we are willing to believe that Trusted Computing is absolutely secure, we can just establish the one big centralised anonymity server and use link encryption together with dummy traffic on the lines [107]. This will give us perfect anonymity. But history has taught us that such a perfect security technology does not and will not exist. Therefore we have to carefully decide how much and for what reason/purpose we trust in Trusted Computing. It seems to be hard to give a mathematical sound or formal model for this -- but generally speaking we see Trusted Computing as a complementary technology which can help to make things easier and more secure but we do not want to rely too much on it (like in “the on big anonymity server” scenario).

According to the client-server architecture of the anonymity service one can distinguish the following use cases for Trusted Computing:

- Trusted Computing on the server side
- Trusted Computing on the client side
- Trusted Computing on both sides

These cases will be discussed in the following sections.

6.2.2.1 Trusted Computing on the server side

Trusted Computing on the server side of an anonymity service results in a number of advantages—for the users of the anonymity service as well as the operators of the servers. Moreover it will become easier to deploy and therefore feasible, mostly because the number of servers compared with the number of clients is small and the servers are dedicated for the purpose of the anonymity service so that the resources (hard-/software) could be chosen to

perfectly match the requirements of Trusted Computing. This is more complicate on the client side where typically general purpose machines are wanted.

From a user's point of view the trustworthiness (the "security") is one of the most important properties of the anonymity service. As stated above this often implies that he has to be sure that not *k* out of *n* anonymity servers collude and that the attacker will not get enough secret information (e.g. cryptographic keys) to deanonymise the user. Hence the user has to carefully select the servers he wants to use.

In order to make a substantiated decision one needs to know information about the servers and its operators, e.g. the hard-/software used for the server, who the operator of the server is (private person, company, privacy commissioner), location of the server etc. Naturally this information has to be authentic which is typically realised by the means of certification and digital signatures. The drawback is that this causes a lot of bureaucratic overhead paired with high (monetary) costs. Eventually this leads to deployment problems as volunteering or non-profit server operators cannot afford this.

Trusted Computing can help to solve this problem by making the trustworthiness of an anonymity server *independent* from the trustworthiness of its operator. That means that some of the (organisational) trust in the operators shifts to trust in the technology of Trusted Computing. The reason for this is that Trusted Computing ensures confidentiality even against the owner (operator) of the anonymity server. Note that the client side does not need to implement Trusted Computing. It must only be able to act as a verifier in the remote attestation procedure. This merely means that the client can check digital signatures and cryptographic hash values.

Additionally Trusted Computing offers advantages to the server operator as well. Basically it enhances the security of the server helping the operator to stick to his guarantees regarding the provided anonymity. This is of course of special interest in the case of a commercially offered anonymity service.

If we think of anonymity for the masses, then we need highly available high performance anonymity servers which offer a very good quality of service. Consequently the operation of the server hardware will be outsourced to an Internet service provider (ISP) which has the necessary connectivity and uptime guarantees. Now the operator (who is no longer the owner) of the anonymity server has to be sure that the staff of the ISP will not reveal confidential information. Again Trusted Computing solves this problem by protecting confidential information even against the owner of the machine.

The outsourcing scenario becomes more complicated if we take into account that one will not rent *real* hardware but a *virtual* one. Virtualisation is another emerging technology to load the multicore server machines as much as possible. Typically renting a "virtual" server is much cheaper than renting a "real" one. Trusted Computing has to be designed to deal with virtualisation in a proper way. That means for instance, that not only the virtual machines running on one computer have to be securely isolated but it also has to ensured that the controlling supervisor gains no access to confidential information.

Another advantage for the server operator/owner is that any juridical, political, ethical etc. pressure to deanonymise someone imposed by politicians, law enforcement agencies or

criminals is useless and therefore will not be done as the operator is not able to reveal any confidential information—even if he wants.

As stated above we see Trusted Computing only as a complementary security technology. Therefore we do not propose changes to the anonymity protocols itself. One of such imaginable changes could be to remove nearly all asymmetric cryptographic operations in low latency anonymity services (like "AN.ON" and "Tor"). This could be achieved by establishing only one anonymous communication channel during the login to the anonymity service and when use this channel for all subsequent communication. This introduces linkability between independent communication actions—but that is not a problem because the linkability information is kept confidential within the trusted computing platform.

6.2.2.2 Trusted Computing on the client side

Trusted Computing on the client side can improve the security of the anonymous communication even more. Naturally TC can be the "secure anchor" needed for any kind of confidential communication. Besides this general fact there exist additional advantages.

According to [117] "...anonymity is the stronger, the larger the respective anonymity set is...". Although Trusted Computing cannot directly be used to distinguish or authenticate individuals it can at least provide reliable information about how many distinct machines are connected to the anonymity service. This brings the user in a better position when he tries to estimate the size of the anonymity set as quantity of anonymity. The anonymity service provider can give more valid information about the quality of service (in terms of quantity of anonymity) he offers. If one day "authenticated locations" (i.e. location stamps) become part of the Trusted Computing technologies the calculation of the anonymity set could be further enhanced.

Another interesting point of Trusted Computing on the client side is, that it could strengthen *my own* anonymity if the *other* users have Trusted Computing even if I have not. Like above this arises from the fact that my anonymity depends on the behaviour (cooperation) of the other users. Technically speaking that means that every user can check by itself (using remote attestation) if the others use secure client software, if they send the agreed amount of dummy traffic etc.

Trusted Computing on the client side of the other users is also a helpful tool if one wants to implement the techniques described in [108] to hide online/offline periods making intersection attacks more difficult. The basic idea of [108] is that a user prepares messages which other users send to the anonymity service during his offline times. One of the main problems is, that a user has to be sure that his "proxies" will in fact send the messages he prepared. With the help of Trusted Computing the user can at least be confident that the machines of the "proxies" will send these messages whenever possible (i.e. they are not switched off or disconnected from the Internet).

6.2.2.3 Trusted computing on both sides

Until now we were not able to discover any scenario where the availability of Trusted Computing on the client and server side gives additional benefit compared with just sum up the properties described in the two previous sections. Therefore it is unclear if the use of Trusted Computing on both sides will lead to emergent phenomenons.

7 Controversial and Legal aspects of Trusted Computing

7.1 Controversial issues

Services and transactions supported by computing devices are susceptible to many kinds of risks. Most often the reason for this is that a reliable basis of security is missing. The TCG try to tackle this subject/problem. The solution presented by the TCG is not only favoured but scientists and independent interest groups have expressed concerns.

A large market share of the IT sector is held by the companies which had founded the TCG. The members of the TCG could use Trusted Computing to strengthen their position on the market and restrict competition. Users of computing devices may also be constricted in the way they can use their devices. So the possible use of the various techniques of the Trusted Computing specifications and their effects arise concerns.

Criticisms and concerns were expressed for instance by researchers like Ross Anderson [106] and by organisations like the Electronic Frontier Foundation [118] or the Chaos Computer Club in Germany. Thereby especially the former TCPA engendered mistrust by a lack of openness and a closed-door standardisation process.

In the following we will present controversial issues of Trusted Computing and potential effects of the security, privacy, and customer's position.

Attestation and Sealing

Remote Platform Attestation enables to readout the exact status of a computing device and to detect unauthorized changes to software via a network. For the legitimate user of a remote computer system it is a feature to detect tampering. But also third parties could use this technique to check all software running on the system in order to certificate the system. This third party gets sensible information about the customer's device and is able to influence privacy by linking requests of the customer because of the usage of unique keys like the Endorsement Key.

As mentioned in [115] a remote entity should not know all software installed and should get only a minimal set of information. Otherwise this raises serious issues regarding privacy and market dominance and could be used to limit options of the costumer's device.

For the attestation of a computer device a hash value of all running programs is created and checked against a database to verify the value as correct. According to [115] the hash value could be invalid if an unknown program is running on the computing device and thus a service provider can deny services.

A zero knowledge technique for improved privacy was published in the TPM specification 1.2. The direct anonymous attestation (DAA) enhances the privacy of the computer owner and provides a direct attestation without using a third party. This attestation uses unlinkable digital pseudonyms so that service providers cannot link pseudonyms of the same person or device.

Owner override is a technique proposed by Seth Schoen [118] to combine possible benefits of Trusted Computing with an improved protection of the users' privacy. His suggestion is that an attestation needs not to reflect the actual state of the software environment. But the owner

chooses a picture of his software environment, which can be completely different than reported.

This would support the freedom of choice in software products and the owner of the platform is informed if the software of his computing device has been changed without his knowledge.

A further technique of Trusted Computing is Sealing. As mentioned in chapter 3 Sealing can be used to bind data to a single platform or application. Sealed data is protected against unauthorised access and distribution.

The TPM could be used to enforce e.g. software licenses and support Digital Rights Management (DRM) with the help of attestation and sealing.

DRM can be used to limit the access to all kinds of documents and thus could also provide a kind of censorship. So it is possible to limit the usage of content to a specific platform [118]. The program that has created a file could also prevent any other program from reading it. Thus, the interoperability could be restricted by the techniques of Trusted Computing [118].

Harmful software and certification

Trusted computing could improve the security of today's computing devices. But the concept of the TCG security model does not prevent any kind of software from running. The security model concentrates on software isolation, so it cannot interfere with other programs. Thus it offers only a minor protection against insecure or harmful software like worms or viruses. A way to distinguish between insecure and secure software is to sign software after an extensive evaluation. Hereupon software is only executed if the signature is valid. Thereby the process of certification has to be open and transparent. An independent organisation should conduct the evaluation. Otherwise this power could be misused by a single authority that could decide if a software or hardware gets a certificate or not.

A monoculture of operating systems, evoked by software attestation and automatic updates, should be avoided. Otherwise this would make attacks of trusted systems more profitable because if an attack of one system is successful all other system can be breached, too.

Endorsement key and the TPM

As mentioned in chapter 4.2.3 the Endorsement Key is a unique key that identifies a single TPM and is the main key for all further operations. The central building block of the security model, the TPM, is hardwired on the motherboard and all important keys are stored and used inside the TPM. According to [120] this may support commercial interests rather than increase the security of a customer's computing device. To avoid this, the costumer should control all keys and thus decide the purpose of each key. That means the legitimate computer user should get the possibility to erase the endorsement key and replace it by a key of his choice. In specification 1.2 of the TPM the TCG allows the deletion of the endorsement key. But this means that all credentials and certificates linked with the old Endorsement Key are invalid and the TPM-based keys, credentials and certificates have to be revised, which could be very complex for a normal user.

The linkage of the most important keys to the hardware of a computer device instead to a concrete user is a further deficiency of the specification [120]. The TPM is designed in a way that allows creation, usage and storage of keys only inside it and does not release them normally. That means for instance a software license for a certain computing device could be

bound to hardware integrated keys. Substituting a computer device could require in the future to purchase also new software licenses, if the keys could not be transferred to the new computing device. Binding the keys on a portable device, so that the user can transfer important keys to other computer devices, is a potential solution. Smart Card systems could be used for the most security benefits in a more flexible way.

Implementation and Backdoors

A computer owner cannot verify if the trusted computing hardware has been implemented according to the published specifications. This is an important problem of all cryptographic hardware [118]. Implemented backdoors or undocumented features can endanger the whole security concept. The cryptographic hardware has access to important and secret information and has also opportunities to leak this information through hidden channels. These channels are difficult to identify and thus the security of all involved information is questionable. So third parties would obtain unauthorised access to private information.

The implementation according to the specification only including documented features has to be assured. Only if the operations of a cryptographic hardware are transparent the computing device could be protected. If an accurate examination of the cryptographic hardware is not possible (because of integration or insufficient documentation) the cryptographic component becomes a black box. Weis, Lucks, and Bogk [120] advise that design and production have to be controlled by trustworthy, international institutions.

Implementing real secure components is not a trivial task. For the area of available TC components this was analysed in [116]. It was shown that various bootloaders used for Trusted Computing contain bugs and ways to attack the chain of trust. Moreover they show that with very little effort it was possible to reset the TPM without resetting the whole platform. This would lead to a Trusted Computing base which reports a state of the platform which does no reflect the true state of that platform.

Integration

There are some endeavours by the members of the TCG to integrate the TPM into other hardware components. A combination of cryptographic functions with other hardware building blocks like the CPU (see also LaGrande of Intel) or I/O-components (see also the Super I/O chip set of IBM) complicates evaluation. According to [120] this is not appropriate because functions of the cryptographic component are mixed with other functions. So it is not clear which function belongs to which part of the component and thus it is difficult to examine and verify the cryptographic component. The cryptographic component should be implemented separately in order to facilitate safety-related verification processes.

Cryptographic issues

Cryptographers approve that the TCG uses well evaluated and standardised algorithms like RSA and SHA-1. But techniques like SHA-1 will not fulfil near future security requirements and should be substituted by better techniques [120].

The conversion of the implemented cryptographic chips from the TCG TPM specification 1.1b to 1.2, which has introduced a lot of security-related improvements, is a slow process. So the obsolete version was still integrated much later than the new specification was released.

Because of the hardwired implementation of the cryptographic component an update of cryptographic functions is only possible by changing the cryptographic hardware component.

Hardware Attacks

Recent research has shown that it is very easy to attack the Trusted Computing base if the attacker is allowed to mount attacks against the hardware. Although one can argue that this is not a security breach as the TCG does not claim to protect against hardware attacks. But on the other side the possibility of successful hardware attacks has to be taken into account when designing distributed systems based on Trusted Computing. As history teaches chances are high that this will be not always the case. A possible scenario would be if users are requested to store personal data on a server secured with the help of Trusted Computing. Of course a user can use remote attestation to try to convince himself that the data is protected on the remote system. But as Trusted Computing does not prevent hardware attacks he can never be sure that the operator of the server (or some other person getting physical access to the server) will not learn the confidential data of the user.

Open Source Software and Patents

In order to ensure that only software runs on trusted devices that have no security hole, all computer programs should be evaluated and signed in future. That means an extensive and expensive evaluation of software is necessary before signing them. If an open source program would be signed, only this version gets a signature and thus changing its source code would make the signature invalid. So this kind of certification process is a contradiction to the Gnu Public License, which allows the modification of source code [120].

For example a part of NGSCB technology of Microsoft is covered by patents so it is uncertain if other developers could use this technology. Patents may restrict the development and usage of trusted computing in the sector of open source

General

The concept of trusted computing is an extensive one but it should not be used solely. People might tend to become less sensitive for privacy / secure threats and problems, because they believe (assume), that the technology in itself is highly secure and everything now is safe due to the use of Trusted Computing. If trusted computing is not 100% secure (the highly realistic view) measures to prevent privacy breaches as well as logging and auditing might fail e.g. manipulated by an attacker.

But Trusted Computing for enhancing the security / privacy (of business processes) might force users to eventually use Trusted computing. By this they also get all the negative things arising from Trusted Computing like threats to privacy e.g. due to the identifying endorsement key or lose of control.

On the one hand the TCG still work on the specifications so security and privacy related issues might be improved. Also the process of standardisation is open for large groups of people. On the other hand hardwired cryptographic component based on the TPM specifications are implemented in notebooks and computer systems, which could not be updated.

7.2 Legal aspects in TC in general

The legal aspects related to TC can be categorized in three parts:

1) The legal implications of content control and the possible abuse of TC by software vendors by means of technologies such as Digital Rights Management.
2) The privacy issues stemming from the TC protocols and specifications defined by the TCG, such as the protocol for certifying the AIK by a Privacy CA.
3) The legal liability of failure of TC.

7.2.1 TC, digital content control and privacy issues

The ability of TC to enforce some control on digital content has been widely discussed and researched [106] which has raised a number of controversial opinions regarding the legal implications of TC aspects. Namely, DRM technologies which are heavily based on TC seem to be of great concern to IT lawyers but also for computer scientists.

DRM includes a certain number of technological mechanisms that can together allow a content provider to define rights over his distributed content. The main aim is to reduce unauthorized access, copying and distribution of digital content. One important application of this technology is to control the illegal distribution of media files since enforcement of legal rules and intellectual property legislations such as the Copyright Act seem to have failed so far due to effective piracy.

DRM starts to appear dangerous for consumers when considering the ability of a content provider to "abuse" this technology in a way to impose restrictions on the access and use of the content by consumers and consequently increasing the costs. While content providers have rights to limit illegal distribution of their content, consumers also have current rights on use of certain digital content which should still be preserved. Such a "fair" DRM can only be achieved if the use of the technology is subject to specific legislations addressing technical aspects of DRM, drawing the line between the rights to control and the rights to consume the digital content.

On the other hand, some other TC functionalities such as "Remote Attestation" might be abused by software vendors in a way to force consumers to run their particular software and to perform regular updates in order to be able to open certain files or have access to a certain service. This would limit the compatibility of software from different vendors, and can be abused by making the consumer pay much more for the software and upgrades he was forced to commit to. This means that TC might be incompatible with the Competition Act which regulates "anticompetitive acts" and "abuse of dominant position" in a market by a supplier or group of suppliers. In fact, strong software vendors in the market would be able to impose on the consumer to run certain kind of software in order for his platform to be considered as "trusted" to be able to access the digital content. The software is said to be "locked-in".

TC raises some privacy issues related to some specifications of the TCG, namely the Privacy CA (cf. 8.6.2). The Privacy CA has to be fully trusted in order for a platform to grant it private identification information. In fact, as will be explained later, the Privacy CA, who is responsible for pseudonym management, would have to obtain credentials from the "attesting" platform revealing some identifiers of the platform. This would enable the Privacy

CA to link pseudonyms to the common, identifying machine credential. The DAA protocol has been introduce in TPM version 1.2.

7.2.2 TC and liability of failure

The issue of legal liability of failure of TC infrastructures is still not thoroughly researched. In [122] the author stresses the point that this particular legal aspect of TC is still not addressed so far as are the content control and privacy issues. However, both academic and commercial researchers seem to agree on its importance. One problem with addressing this perspective is the lack of consensus among academic and industrial partners around the meaning of a "Trusted System".

8 Trusted Computing, Identity and Identity Management

8.1 Trusted Computing for Identity Management

Trusted computing technologies aim at affording secure platforms which behave in a consistent way and which are able to prove their own integrity for both their owners and third parties. TC-based platforms provide elementary security functionalities based on which operating systems and applications can operate securely and consistently, turning down any attempts for tampering with them. Hence, a TC-based platform's secure operation is not only trusted by the platform's owner, but also by other parties interacting directly or indirectly with this platform. The root of trust for such platforms lies in the TPM chip (cf. 4.1).

Among other properties, a TPM chip can store a number of keys and credentials. Some of those are non-migratable i.e. can never leave the TPM unencrypted, and are therefore bound to a specific TPM on a specific platform. Binding and Sealing are two of the functionalities provided by a TC platform which make use of the secure storage of keys and credentials within the embedded TPM. On the other hand, a TPM-based platform has its well preserved unique digital identity which can distinguish it from other platforms. The platform's unique digital identity can be derived from the corresponding Platform Credential or Endorsement Key.

Hence, from a conceptual point of view, correlation between a specific platform's identity and its owner's or user's identity seems possible. With TPMs able to store more than one certificate or key, users' identities, platforms identities and their corresponding transactions can all be bound or correlated to each other in several ways.

While Identity Management (IdM) exists in different forms, it focuses on the establishment, description and destruction of subjects' or objects' identities, by associating attributes to each identity. IdM systems make use of those established identities and corresponding attributes to provide ground for services such as authentication, authorization, behaviour analysis, personalized services, roles and pseudonyms management which are all identity dependents. Therefore, with TC platforms able to associate hardware identity to users identities, TC seems to have big influence, positive but also negative, on the capabilities of IdM systems and hence on the efficiency of the services they provide.

In the following sections, we try to envision how TC can affect some aspects of IdM, and how it can open doors or give ground for a new perspective of IdM and IdM systems.

8.2 Types of Identity Management

In the FIDIS-deliverable D3.1 three basic types of Identity Management Systems (IMS) were introduced (Bauer, Meints, Hansen 2005):

1. Type 1: IMS for account management,
   implementing authentication, authorisation, and accounting
2. Type 2: IMS for profiling of user data by an organisation,
   e.g. detailed log files or data warehouses which support e.g., personalised services or the analysis of customer behaviour,
3. Type 3: IMS for user-controlled context-dependent role and pseudonym management.

Implementations of IMS also may be of hybrid types combining different organisational structures and methods characterising the introduced types. Trusted Computing may play a role in all three types of IMS.

8.3 Related Work and Current Problems

Current systems used for automated provisioning of applications and IT resources make use of IdM systems to provide access control over the resources.

One of the cornerstones of any IMS lies in its security. According to the IMS Types defined in 8.2, the security mechanisms used for an IMS could be different. On the other hand, the protection of a user's identity from identity theft also relies heavily on the security mechanisms whether on the user's specific platform, web interfaces, communication protocols or remote platforms and databases. Unfortunately, trusted infrastructures are not commonly used nowadays for identity management systems. In fact, current identity management solutions lack hardware security support.

For example, digital signatures are an important example for identity management applications. Although the legal prerequisites for digital signatures (at least in Europe) exist, in a big scale applications do not so far.

Digital signatures face several problems, mainly low acceptance due to low security for the user. In case an attacker manages to fake a user's digital signature, the user will find himself having to prove that he did not sign. Therefore, digital signatures need tamper proof devices that require users to authenticate themselves. But as long as this authentication is not safe from e.g. an attacker stealing a password, the user still faces the shifting of the burden of proof. Furthermore, the problem of "What-You-See-Is-What-You-Sign" is not a trivial one, so users are required to have a certain level of expertise to be able to judge whether they are tricked or not, and even experts might fail. One can never be sure if one really signs what one can see, but has to trust the applications, the hardware (and their developers).

Some initiatives for supporting identity management with TC already exist. For example, in [123], the author proposes to use TC to establish a level of trust between different identity domains, in a way to allow one credential provider in one identity domain to issue credentials to be used for authentication in another identity domain based on a pre-defined credential issuing policy.

8.4 TC use for identification

Digital identity is broadly interpreted as "the digital representation of a set of claims made by one digital subject about itself or another digital subject" – *Wikipedia.* While digital subjects can represent either human or non-human (device, computers, digital resources…), the value of a Digital identity pertains only in the context in which it is used: a Digital identity used for identification in an enterprise's network could be irrelevant for identification for an email account, for performing a digital signature on a document, or for accessing a private place. Hence, for a single entity, Digital identity does not need to be unified across different contexts, but rather unique in each domain of application.

However, any Digital identity is meaningful only if the attributes it associates to the corresponding digital subject are authentic. Therefore, authentication mechanisms are crucial to ensure correct identification and to prevent from identity theft. Authentication mechanisms

used for correctly authenticating a digital identity are mainly implemented in software and hardware, but can sometimes rely on physical authentication techniques such as iris scanning for human unique identification.

In this section, we give ideas on how TC can help securing the Digital identity, especially in the context of authentication.

8.4.1 TPM credentials storage

Digital Identity can be presented by several means. In many cases, the authenticity of a Digital Identity is crucial to its value. This is why a username is associated to a password, a Public Certificate is signed by a Certificate Authority's private key, a bankcard is associated to a PIN, a digital signature is associated to a private key… Passwords, Private Keys, PINs are known only to the corresponding digital subject. Their confidentiality to the digital subject is necessary to ensure prevention from identity theft. Since a Digital Identity is only relevant if its authenticity can be proven, the way the corresponding confidential information is protected is crucial. In some cases, the confidential information can be memorized (in the case of a PIN or the answer for a challenge-question), but in other cases it needs to be stored in digital form, mainly on hardware devices such as harddisks or smartcards. In most cases, the confidential information has to be provided to other hardware or software in order to be processed for authenticating the Digital Identity.

For this reason, encryption schemes are widely used nowadays to protect the confidential information corresponding to Digital Identities. Authentication mechanisms rely heavily on those encryption schemes to ensure secure authentication of Digital Identity. The encryption schemes are used to protect the confidential information both in storage and during communication for authentication purposes.

Trusted Computing can contribute to the protection of confidential information corresponding to Digital Identities, and can therefore ensure their authenticity. One of the major advantages of a TPM is its ability to provide tamper-proof protection of data and keys entrusted to it. This feature can be employed to efficiently protect the confidential information in hardware. As mentioned in section 4.1.3, the Storage Root Key embedded in a TPM can be used to wrap TPM protected keys which are themselves used to encrypt sensible information on a storage device, as shown in Figure 4. This means that the confidential information associated to a Digital Identity can be saved encrypted on a harddisk or smartcard, with the encryption key shielded by a TPM attached to the device.

The decryption of the information is controlled by some strict rules that can be pre-defined by the digital subject with help of TC functionalities. This could be based on special authentication mechanisms that allow the digital subject to authenticate himself to its own device in order to access the confidential information corresponding to his Digital Identity. Moreover, the decryption of the confidential information can also depend on integrity checks of the device's software and hardware configurations. In other words, the TPM "binds" the confidential information with the authenticity of the user claiming the Digital Identity, and the platforms configuration on which the confidential information is meant to be decrypted and processed. In the case where this information should be communicated to other platforms, the TPM can also bind it to some pre-defined secure configuration of the target platform in order to ensure that its confidentiality is maintained.

This process greatly ensures the confidentiality of private information associated with a Digital Identity, which in turn improves its authenticity and hence its value since identity theft is considerably reduced.

8.4.2 TPM-based authentication policies

The confidentiality of credentials associated to a Digital Identity is not only important for the digital subject himself, but also for the other digital entities to which the Digital Identity is authenticated. In fact, the authenticity of the Digital Identity is evaluated by the verifying entity. The latter can verify if the provided credentials are truly associated with the claimed Digital Identity. Therefore, it is of great interest for the verifying entity to ensure that the confidentiality of the credentials is well preserved and not subject to leakage, whether in storage or during communication. One important factor for ensuring this feature is the state or configuration of the digital subject's hardware and software used to store and communicate the credentials.

The Remote Attestation feature of Trusted Computing can significantly help achieving this requirement. TPMs allow the attestation of a certain state of a certain system as well as the generation of partial identities for its user in form of signed cryptographic keys or certificates (combined type 1 and 3 identity management).

To ensure security of digital transactions, a Trusted Platform can attest its integrity using cryptography. Earlier versions of TPM specifications required the involvement of a Trusted Third Party (TTP) [124].

In short, a user who wants to initiate a certain transaction over the Internet would go to a TTP, attest to it that he has a Trusted Platform and tell it his identity. The TTP then would issue him a certificate he could show to his transaction partner. As the TTP would still be able to reveal the user's real identity from this certificate, the certificate is a pseudonym and therefore a partial identity. The user would not need to reveal his real identity to his transaction partner while the latter can rest assured that the user has a Trusted Platform. This can be regarded as a privacy-compliant scenario.

The TC Remote Attestation feature can therefore allow a new type of authentication policy that is not concerned with the real identity of the device's user, but rather with the specifications of the device itself. In that case, the Digital Identity in focus is that of the user's platform, and the platforms credentials associated to this Digital Identity are entrusted to the corresponding TPM. TPM-supported devices can therefore be authenticated to a network service without any need to reveal the identity of their users. This type of authentication policy opens the door for a new perspective of networking where any user has access to network services as long as the behaviour of his device is pre-determined.

8.4.3 TC platforms identity management

Identity Management can be improved by TC platforms in several ways depending on the purpose of the IMS. As explained in 8.2, three types of IMS are defined, each with a different purpose.

In the case of type 1 identity management, authentication and authorization can be significantly improved with the deployment of TC infrastructures within an identity domain but also between different identity domains.

The TC attestation protocols can help establish trust between the authenticating party and the Authentication Server (AS) or the Policy Decision Point (PDP). The device used for supplying the identity credentials, supported by TC, will be able to attest its status (software stack and hardware configuration) to the AS by means of an attestation certificate. This would prove to the AS that the platform used to supply the identity credentials is conformant with some predefined "secure configuration" and would therefore allow the AS to trust the corresponding device and securely authenticate the user. Trust can also be established in the other direction, with the AS providing an attestation certificate to the user, which would prove to the user that it is not a bogus AS.

Moreover, TC can play an important role in establishing trust between different identity domains. A pre-defined credentials issuing policy can be enforced on Credential Providers belonging to more than one identity domain. This would allow a Credential Provider in one identity domain A to issue credentials for a user to access services in another identity domain B. At the moment of issuing the credentials, the identity domain A should be trustworthy, i.e. adhering to some pre-defined issuance policy. Such a policy can only be enforced by the use of TC. The credentials can be verified by identity domain B to have been issued with the enforced issuance policy at identity domain B by means of TC signatures. Hence mutual trust established between different identity domains based on TC can improve identity management in terms of interoperability.

8.4.4 Business use cases

Based on the ideas and perspectives proposed in sections 8.4.1, 8.4.2 and 8.4.3, several business use cases of TC-based identification and identity management can be envisioned.

First, the tamper-resistance nature of a TPM can be employed to give a solid ground for future authentication mechanisms. As a result, crucial identification information can be more confidentially stored and communicated, which would support various use cases such as mobile commerce. For example, users of mobile devices will have more confidence storing their creditcard information on their devices, and supplying them during a transaction. On the other hand, transaction peers will be able to verify, through adequate protocols, that the identification information is entrusted to a TPM, which would enhance the trust of the peers in the identity claimed by the mobile user. Hence, the level of confidence in the authentication and identification mechanisms is brought to a higher level, and that would boost the transactions rate.

Another business use case of TC-based infrastructure, which supports anonymity, relies on the idea of "trustworthy platforms". As explained in 8.4.2, a TC-based platform will be able to attest its configuration and specifications to a network peer. This would enable business use cases requiring confidence in the behaviour of platforms, rather than in the identity of the person using them. Possible applications would be online gaming, or online gambling, where the contributors would require trustworthiness of their peers' platforms without needing any reveal of identity. A strict policy on platforms' external network access, operating systems properties and configuration, applications installation can hence be enforced on any platform

attempting to join. In this case, a contributor to the game would be confident that its peers will not be able of treachery without those peers having to reveal their real identities. The digital identity of the platforms could be relevant in such a use case in order to identify the platforms eligible to join the network whenever they have successfully attested a configuration which is conformant to the policy.

One more possible business use case of TC for identity management relies on the concept described in 8.4.3 where trust is established between different identity domains. The advantage of this kind of schemes lies in the fluidity of business processes requiring identification of users in different identity domain. An example scenario would be services in large scale organizations. An organization having different identity domains for separate services or departments would usually require a user to provide different credentials to each of the corresponding servers. With the deployment of a TC infrastructure, a trust policy defining which identity domain can issue credentials for other specific identity domains can be enforced across the organization's servers. This way, an employee in one department authenticated by the corresponding identity domain would be able to obtain credentials from this domain to be used for authentication in another identity domain, if the policy allows it. This kind of schemes is realizable with the use of TPM signatures to enforce the trust policy.

8.5 TC identity and Consumer Privacy

8.5.1 TPM Unique Digital Identity

The unique digital identity of a TPM is represented by the Endorsement Key (EK). As explained in section 4.2.3, this key never leaves the TPM, and is used to create the Endorsement Credential and the AIK. Since the EK is an RSA key, its public key is contained in the Endorsement Credential. The validity of the EK and corresponding Endorsement Credential are necessary for establishing the AIK. In order to certify the public part of the AIK, the Privacy Certification Authority should check the validity of the Endorsement Key and Credential. At this point, the identity of the TPM is revealed to the Privacy Certification Authority. This would create anonymity problems as will be explained in section 8.6.1.

8.5.2 Privacy Risks

As it is unlikely that large numbers of TTPs will survive the market, and as TTPs were capable of easily linking different pseudonyms to the same person, this would result in a concentration of detailed profiles of people at a small amount of places [125]. Therefore, this can be regarded as a privacy-invading scenario.

Transactions in the above context can not only be the purchase of DRM-protected digital goods but also communication or authentication processes. If authentication processes were protected by TC (preferably using DAA) identity theft could be reduced.

As already mentioned Trusted Computing may be used to support enforcement of policies when processing data such as EPAL (see chapter 6.2.2). EPAL may well be used in future for basic data and mined results in the context of profiling (type 2 identity management).

8.5.3 TC-Requirements for privacy-aware IdM

Depending on the type of identity management introduced already in chapter 8.2 requirements for privacy awareness are different.

One of the basic privacy problems with type 1 IMS is the use of global unique identifiers. They support linkability of communications and transactions crossing the border of communicational contexts. TC should support the use context or sector specific identifiers that are linkable only (a) based on legal norms, (b) agreed policies or (c) consent of the data subject on an individual case basis. In any case transparency of linking activities is required as a general rule taking defined exceptions for example in case of criminal investigations into consideration.

From the perspective of type 2 IMS (Profiling) TC should support Transparency Enhancing Technologies (TETs), allowing the data subject to understand what results (knowledge) was generated how (method) based on which data (basic data) and used for what purpose. In this case a stable link between (a) basic data and corresponding privacy policies and (b) data controller and his policy for usage of basic data and mined results is required to facilitate supporting tools for automated uncovering and comparison of data with the corresponding policies and notification mechanisms about the comparison results.

In the context of type 3 IMS TC should support the unrestricted generation and use of identifiers managed and controlled by the user. Again prevention of linkability is the main goal that should be achieved.

8.6 TC platforms and Anonymity

While TC based anonymity services have been discussed in 6.2.2, we focus in the following on the implications of TC features on identification of individuals and general identity management systems.

8.6.1 TC implementations and anonymity

One of the major advantages of TC is the ability of TC platform to perform attestation. This would help the system to prove its security properties to a remote system. However, current implementations of TC do not take the privacy issues into account. In fact, the credentials sent by a platform to prove its security properties can reveal identifying information about the platform, and maybe its user. The anonymity during transactions, which is important in many scenarios, could hence be compromised. The concepts of “Privacy Certification Authority” and “Direct Anonymous Attestation” (next sections) were proposed by the TCG to enhance the privacy aspect during remote attestation.

TC can hence play an important role both in a network of anonymity (in case correct implementation of DAA is achieved), but also in a network of identity. In fact TC can help achieving a network with a trusted infrastructure and full anonymity. This would be useful for a network where all partners would like to perform transactions with other peers with secure computing platforms without having to reveal any identification information about them or requiring any such information from their peers.

On the other hand, TC can help achieve a network where all partners are considerably sure of the identity of their transaction peers, since the attestation certificate of the peer contains identification information and is generated by a TPM which can be considered truthworthy.

8.6.2 Privacy Certification Authorities (CAs)

Privay CAs are a kind of Trusted Third Parties that have been defined by the TCG. The role of the Privacy CA is to certify to partners exchanging information that each of them is truly a secure computing platform. This requires all partners to fully trust the Privacy CA, which is hard to achieve.

The protocol used for this kind of certification requires the partner to provide some private information about his platform. This information is used by privacy CA in order to verify the trustworthiness of the partner before certifying the security properties of the platform to another partner. The word “Privacy” stems from the fact that the Privacy CA does not need to reveal the real identity and other private information of the partner to another partner, but rather has just to certify its trustworthiness. This means that “privacy” is achieved between communicating partners thanks to the Privacy CA. However, enough identity and private information need to be revealed to the Privacy CA himself.

This led to the design of the “Direct Anonymous Attestation (DAA) protocol” by the TCG (next section) which aims at providing this level of privacy between communicating partners without the need for a Privacy CA. The goal is to avoid revealing private information to a third party.

When considering identity management systems, the need for a Privacy CA is crucial, and at the same time not really problematic. The reason is that IMS are usually of centralized nature, which means that a central management authority at the heart of the IMS has to be trusted in all cases in order to certify the validity of the claimed identity. This means that this central authority can in principle play the role of a “Privacy CA” which can certify for any partner that the peer partner has a secure computing platform.

The requirement for a protocol such as DAA to avoid the presence of a Privacy CA would be relevant for example in the case where the IMS is responsible for identification during inter-organizational transactions. In this kind of scheme, the IMS is independent from the partners, and “fully” trusting it and revealing private platform information to it could be a risk. An IMS would be able to use this revealed information to identify the peers of a certain transaction, and to link certain transactions to certain identities. This could cause both an anonymity and a privacy breach.

8.6.3 Direct Anonymous Attestation (DAA)

The DDA protocol defined in version 1.2 can prevent the identity of the TC platform’s user when carrying on remote attestation. The anonymous remote attestation can help in anonymous communication and enhance identity management.

In TPM specifications version 1.2, the aforementioned Direct Anonymous Attestation (DAA) has been introduced. DAA allows attestation of a Trusted Platform without the involvement of any other party, trusted or not. DAA gives the user certificates that do not reveal his

identity and look different each time. These certificates, called credentials, are unlinkable [126]. Therefore, this scenario can be regarded a privacy-friendly (type 3 identity management).

As mentioned before, content providers are not interested in offering anonymity to their customers; their main interest is establishing a centrally managed identity management (type 1 IMS). It is therefore unlikely, that a privacy-friendly mechanism such as DAA will be able to unfold its privacy potential to the masses, as long as the TTP scheme exists. It should therefore be dropped from future TPM specifications for the benefit of DAA. But this is also unlikely as there are other applications (apart from the privacy context), where the use of TTPs makes sense. So “best practices” as suggested by the Article 29 Group [127] will have not only to be established, but in addition should be made obligatory.

9 TC and Identity Management – a Use Case Scenario

In order to give the reader an overall understanding of how trusted computing can help improve IdM, we consider a use case scenario with its requirements and architecture, and we mention the advantages of the solution. The approach is based on the work in [140].

9.1 Scenario Description

The scenario is based on the notions mentioned in section 8.4.3, but refers also to other parts of the deliverable. It emphasizes the need for cross identification of users between different identity domains.

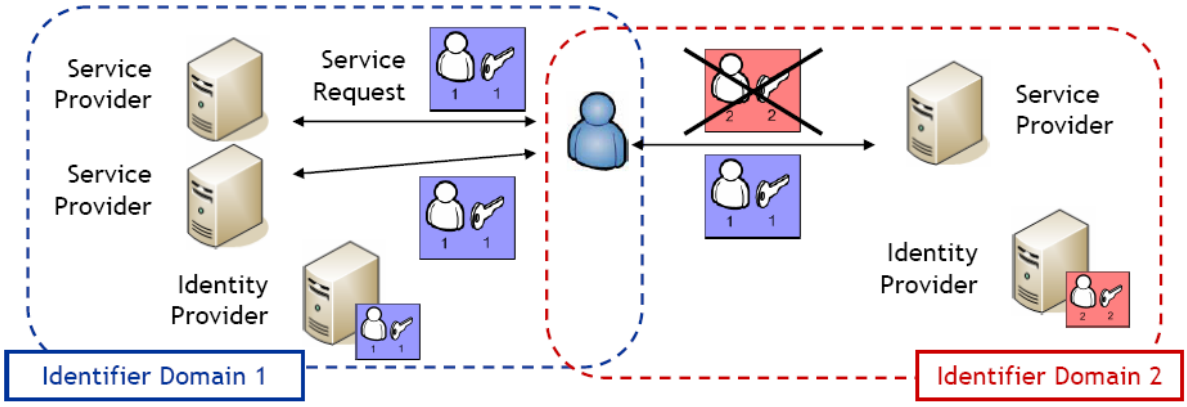

Fig7: Identification across different Identifier Domains [140]

In this scenario, there exist two Identifier Domains, IDom1 and IDom2. In each identifier domain, there is a Service Provider (SP) and an Identity Provider (IP). The IP can issue identity credentials to a user who wants to gain access to the service inside the same identifier domain.

In principle, an SP in IDom1 can validate credentials that have been issued only by an IP inside the same IDom1. However, in certain business use cases, an SP in IDom2 should be also able to validate credentials issued by an IP in IDom2. For that, this SP should be able to verify the trustworthiness of the IP in IDom2, and whether the authentication policy allows this IP to issue this kind of credentials which are valid for both identifier domains.

9.2 Requirements Analysis

In the following, we list the general requirements for the scenario described above:

a) On a local domain level, the identity provider should be able to check the status (e.g. the trustworthiness) of the platform to which a new identity credential has to be granted. This would allow him to authenticate and authorize a new user based on the ability of his platform to preserve the granted identity credential from theft (voluntary handing over can also be considered, but it needs more sophisticated technologies such as support of biometric identification in the device).

b) The identity provider should be able to issue a trusted ticket that can be securely validated by service providers in the same domain. The service providers should be able to check if the identity provider is trustworthy, and if it is allowed to grant such credentials in this domain. This requires the credentials themselves to include information reflecting the status of the identity provider's system at the time when the credential was issued. Typically, such information should be in the form of integrity measurements which could be compared with reference values.

c) For cross-domain validation of credentials, the identity provider should be able to issue credentials that would be valid in other identifier domains (i.e trusted by service providers outside their domains). For that, the service providers should be able to check the trustworthiness of the corresponding identity provider outside the domain, and whether this identity provider is allowed to grant credentials that are valid in the service provider's domain.

9.3 Architecture

The idea behind the architecture is to benefit from the advantages of different TCG specifications in order to establish a trusted infrastructure that allows fulfilment of the requirements mentioned above. The architecture is shown below, together with the protocol needed for establishing the trust, especially across identity domains.

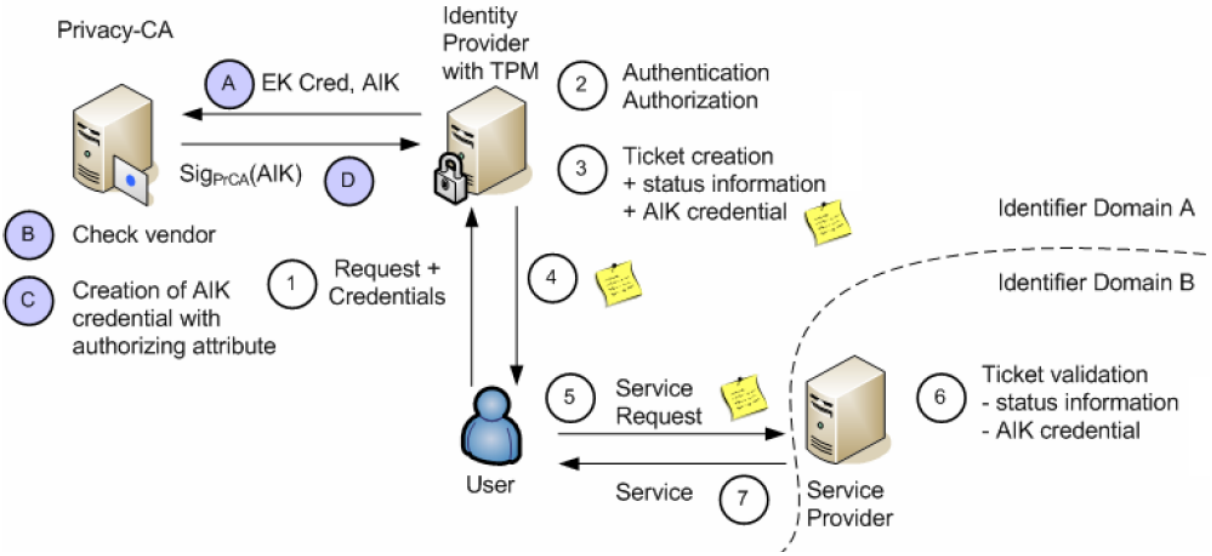

Figure 8: Interoperable credentials across Identifier Domains [140]

- The architecture integrates the role of Privacy-CA explained in section 8.6.2. The Privacy-CA's main role in this case is to check the trustworthiness of the Identity Provider residing in Identifier Domain A, and to issue him an AIK credential with authorizing attribute. For this purpose, the Identity Provider has to communicate his Endorsement Key provided by his TPM vendor, and an AIK generated by his TPM (cf. 4.2.2).

- After regular authentication and authorization of the user, which is usually based on a pre-defined policy, the Identity Provider will then be able to issue an identity credential for the requesting user. This credential should include the AIK credential obtained from the Privacy-CA, in addition to *status information* in the form of integrity measurements performed and stored in the TPM.
- The user will then use his identity credential to authenticate and authorize himself at a Service Provider belonging to Identifier Domain B.
- Since the *status information* is TPM-generated, it can be compared to reference values. The Service Provider will attempt to validate the identity credential of the user based on the AIK credential provided by the Privacy-CA, and the verification of the status information trustworthiness.

With this protocol, a Service Provider residing in Identifier Domain B will be able to accept identity credentials issued by Identity Providers in Identifier Domain A, and that is based on the trust in the Privacy-CA, and the integration of TPM hardware and functionalities.

9.4 Advantages and Disadvantages

The main advantage of the approach described in the previous section is that the changes to the original architecture of the IdM system are not overwhelming. The integration of a TPM as hardware ship in the Identity Provider's system is practically possible. Still, the protocols to be developed should be able to produce a special identity credential with a specific format that could be parsed by the Service Providers. This Trusted Ticket should typically include an attribute statement for the status information. The latter requires the TPM's PCR values, in addition to the AIK credential itself which is used to sign the status information.

The integration of the Privacy-CA role is also possible since it has been specified by the TCG. This will avoid adding a PKI system to provide cross certifications between Identifier Domains. As the infrastructure is already specified by the TCG, the costs of cross certification are surely reduced, despite the need for developed algorithms for handling the special identity credentials both at the Identity and Service Providers sides.

The issue of scalability of the trusted infrastructure remains a problem, since the Privacy-CA should be reachable by all Identity Providers in order for those to obtain AIK credentials. The singularity of the Privacy-CA is essential in order to establish the trust across Domains.

10 Recommendations and Future work

10.1 Trusted Computing Benefits for Identity Management

Trusted computing can enhance the notion of digital identification. With TC being able to protect the confidential information associated to a digital identifier due to the tamper-proof nature of the TPM and its key hierarchy, this ensures the confidentiality of this information and hence improves the authenticity and value of the Digital Identity. Identity theft can be significantly reduced.

Anonymity services can be enhanced by deploying a TC infrastructure. The Remote Attestation feature would enable service providers to provide services to any platform that is able to attest a certain secure-known configuration, without having to reveal his identity. This opens a door for a number of new business use cases.

As for identity management, TC can be used to establish trust between different identity domain, allowing one credential providers in on domain to issue authentication credentials in another domain. This would enhance identity management in terms of interoperability.

10.2 Further Research and Identity Management Considerations

The idea of applying TC techniques to identification and identity management has certainly been already tackled, but one can say that deep research in this direction is still in its early stages. Some perspectives where discussed in chapter 8 but need to be further investigated.

An important aspect to consider is the issue of integrating TC technology in current IMS. The architecture of common IMS might impose some restrictions on the use of TC, mainly implementation issues. It is therefore important to investigate more on the feasibility of TC integration to support current IMS.

Business use cases of TC-based anonymity and identification services, such as the ones proposed in 8.4.4, should also be modelled and their feasibility further investigated.

11 Conclusion

The deliverable gave a description of the Trusted Computing technology status and functionalities, and important specifications of the TCG. It also shed the light on the current and future applications and possible scenarios, in addition to the market status and R&D efforts in the field. It also tackled the controversial and legal aspects of TC before giving a perspective on the use of TC for identification and identity management purposes.

It is clear from the findings that TC has reached an advanced level of detailed specifications, with a fast increasing market and extensive research in the fields. It is also evident that there are many still hidden fields of application for TC, and the analysis given in chapter 8 should raise the awareness to the advantages of TC in the field of identity management, but also in enabling and supporting business cases such as mobile commerce by improving identification and authentication mechanisms.

While legal consideration should also be coupled with the TC technological advancement, it is of great importance to begin a real assessment of the feasibility of the proposed use cases, and the extent to which TC can bring forward the notion of Digital Identity, and improve identification and identity management systems.

12 References

[1] European Multilateral Secure Computing Base (EMSCB). http://www.emscb.de.

[2] The PERSEUS Project. http://www.perseus-os.org.

[3] The Xen Virtual Machine Monitor. http://www.cl.cam.ac.uk/Research/SRG/netos/xen/.

[4] TrouSerS — The Open Source TCG Software Stack. http://trousers.sourceforge.net.

[5] Trusted Computing Group (TCG). https://www.trustedcomputing.org.

[6] Common Criteria for Information Technology Security Evaluation, Part 1: Introduction and General Model, August 1999.

[7] Common Criteria for Information Technology Security Evaluation, Part 2: Security Functional Requirements, August 1999.

[8] Common Criteria for Information Technology Security Evaluation, Part 3: Security Assurance Requirements, August 1999.

[9] TCPA Main Specification, Version 1.1b. https://www.trustedcomputinggroup.org/specs/TPM/TCPA_Main_TCG_Architecture_v1_1b.pdf, February 2002.

[10] Bylaws of Trusted Computing Group. https://www.trustedcomputinggroup.org/about/articles_of_incorporation.pdf, 2003.

[11] Trusted Computing Group Frequently Asked Questions. https://www.trustedcomputinggroup.org/join/levels/, August 2005.

[12] Infineon Technologies AG. Product Brief — TPM 1.2 Hardware. http://www.infineon.com/tpm, May 2005.

[13] Infineon Technologies AG. Product Brief — TPM 1.2 Software. http://www.infineon.com/tpm, May 2005.

[14] Utimaco Software AG. Utimaco enables the Use of Trusted Hardware Platforms for Secure Mobile Computing. http://www.utimaco.com/content_press/p170204.html, February 2004.

[15] Christian Stüble Ahmad-Reza Sadeghi. Towards Multilaterally Secure Computing Platforms — With Open Source and Trusted Computing, 2005.

[16] Norbert Pohlmann Ahmad-Reza Sadeghi, Christian Stüble. European Multilateral Secure Computing Base — Open Trusted Computing for You and Me. Datenschutz und Datensicherheit (DUD), pages 548–554, September 2004.

[17] AMD. Secure Virtual Machine Architecture Reference Manual, May 2005.

[18] Ross Anderson. Security Engineering: A guide to building dependable distributed systems. John Wiley & Sons, 2001.

[19] Arcom. APOLLO EBX format Intel Pentium M or Celeron M based embedded PC. http://www.arcom.com/products/icp/pc104/processors/apollo_datasheet.pdf, 2005.

[20] Arcom. VIPER PC/104 format 400MHz Intel PXA255 XScale embedded controller. http://www.arcom.com/products/icp/pc104/processors/viper_datasheet_v2.pdf, 2005.

[21] Atmel. AT97SC3201 — The Atmel Trusted Platform Module. http://www.atmel.com/dyn/resources/prod_documents/doc5010.pdf, August 2004.

[22] Atmel. Atmel And NTRU Announce An Unbeatable Hardware/Software Security Combination For Trusted Computing Products. http://www.atmel.com/dyn/corporate/view_detail.asp?FileName=ATMLNTRU.html, November 2004.

[23] Atmel. AT97SC3203 Advanced Information Summary. http://www.atmel.com/dyn/resources/prod_documents/5116s.pdf, July 2005.

[24] Atmel. AT97SC3203S for SMBus Protocol Summary. http://www.atmel.com/dyn/resources/prod_documents/5132s.pdf, August 2005.

[25] Atmel. Trusted Platform Module AT97SC3201 Summary. http://www.atmel.com/dyn/resources/prod_documents/2015s.pdf, June 2005.

[26] Stephan Mooney Bill Rosenblatt, Bill Trippe. Digital Rights Management: Business and Technology. John Wiley & Sons, 2001.

[27] Ernie Brickell, Jan Camenisch, and Liqun Chen. Direct Anonymous Attestation. Technical report, IBM Research, March 2004.

[28] Broadcom. Broadcom Revolutionizes LAN Communications by Introducing the World's First PCI Express Gigabit Ethernet Controllers for Server, Desktop and Mobile PCs. http://www.broadcom.com/press/release.php?id=461159, October 2003.

[29] Broadcom. Broadcom Licenses Infineon TPM Management Software for Integration With Broadcom TPM Products to Provide a Complete Trusted Computing Group 1.1b Security Solution. http://www.broadcom.com/press/release.php?id=495640, February 2004.

[30] Broadcom. BCM5752 Product Brief. http://www.broadcom.com/collateral/pb/5752-PB00-R.pdf, 2005.

[31] Broadcom. BCM5752M Product Brief. http://www.broadcom.com/collateral/pb/5752M-PB00-R.pdf, 2005.

[32] Broadcom. Broadcom Controllers Integrate TPM 1.2 enabling OEMs to Offer Hardware-Based Security as a Standard Feature on all PCs. http://www.broadcom.com/press/release.php?id=700509, April 2005.

[33] Hewlett-Packard Development Company. HP ProtectTools Embedded Security — The HP Trusted Computing implementation. ftp://ftp.compaq.com/pub/products/security/embedded_security_-_implementation.pdf, October 2003.

[34] Hewlett-Packard Development Company. HP ProtectTools Embedded Security: Expanding trust within the enterprise computing environment. ftp://ftp.compaq.com/pub/products/security/HP%20ProtectTools%20Embedded%20Security%20WP%20-%20OV.pdf, May 2003.

[35] Hewlett-Packard Development Company. HP ProtectTools: Authentication technologies and suitability tasks. ftp://ftp.compaq.com/pub/products/security/FINAL_5983-1956EN_Security%20Technologies.pdf, June 2005.

[36] Hewlett-Packard Development Company. HP ProtectTools: Firmware security features in HP business notebooks. ftp://ftp.compaq.com/pub/products/security/FINAL_4AA0-0697ENW.pdf, June 2005.

[37] Intel Corporation. LaGrande Technology Architectural Overview, September 2003.

[38] Intel Corporation. Product brief — intel 865g chipset. http://www.intel.com/design/chipsets/865G/865G_PB_8.pdf, 2003.

[39] Intel Corporation. Product Brief — Intel Desktop Board D865GRH. http://www.intel.com/design/motherbd/rh/rh_productbrief.pdf, 2003.

[40] Intel Corporation. Product Brief — Intel 915G Express Chipset. http://www.intel.com/design/chipsets/915G/915G_pb.pdf, 2004.

[41] Intel Corporation. Product Brief—Intel 925XE Express Chipset. http://www.intel.com/design/chipsets/925xe/925xe_pb.pdf, 2004.

[42] Intel Corporation. Product Brief — Intel Desktop Boards D915GEV, D915GUX, D915GAV, and D19GAG. http://cache-www.intel.com/cd/00/00/14/90/149067_149067.pdf, 2004.

[43] Intel Corporation. Product Brief — Intel 945G Express Chipset. http://www.intel.com/products/chipsets/945g/prodbrief.pdf, 2005.

[44] Intel Corporation. Product Brief — Intel 955X Express Chipset. http://www.intel.com/products/chipsets/955x/prodbrief.pdf, 2005.

[45] Intel Corporation. Product Brief — Intel Desktop Boards D945GTP, D49GCZ, D945GNT, D945PSN and D945PAW Classic Series. http://cachewww.intel.com/cd/00/00/21/84/218428_218428.pdf, 2005.

[46] National Research Council. The Digital Dilemma, Intellectual Property in the Information Age. National Academy Press, 2000.

[47] NTRU Cryptosystems. Product Brief — NTRU Core TCG Software Stack. http://www.ntru.com/products/ntru_ctss_brief.pdf, 2005.

[48] D. Carrel D. Harkins. The Internet Key Exchange Protocol (IKE), RFC 2409, November 1998.

[49] Densitron. ConnectBus-II Single Board Computer for Gaming — DPX-114. http://www.densitron.com/computers/pdfs/dpx114.pdf.

[50] Densitron. Pentium 4 ConnectBus-II Board for Gaming — DPX-115. http://www.densitron.com/computers/pdfs/dpx115.pdf.

[51] Federal Information Processing Standards (FIPS). FIPS PUB 140-2: Security Requirements for Cryptographic Modules. http://csrc.nist.gov/publications/fips/fips140-2/fips1402.pdf, January 1994.

[52] Tom R. Halfhill. ARM DonsArmor — TrustZone Security Extensions Strengthen ARMv6 Architecture, August 2003.

[53] Hewlett-Packard. Fact Sheet — HP unveils family of security-enhanced, wireless notebook PCs and high-performance mobile workstation. http://www.hp.com/hpinfo/newsroom/press_kits/2003/telecom/nr_securitynotebooks.pdf, October 2003.

[54] Hewlett-Packard. HP Compaq Business Desktop d530 Series. http://www.hp.com/hpinfo/newsroom/press_kits/2003/nar/ds_desktopd530.pdf, May 2003.

[55] Hewlett-Packard. HP Compaq Business Notebook nc6000. http://www.hp.com/hpinfo/newsroom/press_kits/2003/nar/ds_notebooknc6000.pdf, 2003.

[56] Hewlett-Packard. HP Compaq Business Notebook nc8000. http://www.hp.com/hpinfo/newsroom/press_kits/2003/nar/ds_notebooknc8000.pdf, September 2003.

[57] Hewlett-Packard. HP xw4200 Workstation. http://www.hp.com/workstations/pws/xw4200/xw4200.pdf, April 2005.

[58] Hewlett-Packard. HP xw6200 Workstation. http://www.hp.com/workstations/pws/xw6200/xw6200.pdf, November 2005.

[59] Hewlett-Packard. HP xw8200 Workstation. http://www.hp.com/workstations/pws/xw8200/xw8200.pdf, November 2005.

[60] Hewlett-Packard. QuickSpecs — HP Compaq nc4010 Ultraportable Business Notebook. http://www.hp.com/products/quickspecs/11832_na/11832_na.pdf, August 2005.

[61] Softex Inc. Data Sheet — OmniPass. http://www.softexinc.com/omnipassentdatasheet.asp, 2004.

[62] Toshiba America Information Systems Inc. Protege R200 Series. http://www.toshibadirect.com, 2005.

[63] Toshiba America Information Systems Inc. Protege R205 Detailed Product Specification. http://www.toshibadirect.com, 2005.

[64] Toshiba America Information Systems Inc. Tecra S3 Series Detailed Product Specification. http://www.toshibadirect.com, 2005.

[65] C. Kaufmann. Internet Key Exchange (IKEv2) Protocol, Internet Draft, September 2004.

[66] Computer Systems Laboratory. Secure hash standard. 180-1, April 1995.

[67] Lenovo. ThinkPad R52 Notebooks — Data Sheet. http://www.lenovo.com/thinkpad/, August 2005.

[68] Lenovo. ThinkPad T43 Notebooks — Data Sheet. http://www.lenovo.com/thinkpad/, August 2005.

[69] Lenovo. ThinkPad T43p MobileWorkstations—Data Sheet. http://www.lenovo.com/thinkpad/, August 2005.

[70] Lenovo. ThinkPad X41 Notebooks — Data Sheet. http://www.lenovo.com/thinkpad/, April 2005.

[71] Lenovo. ThinkPad X41 Tablets — Data Sheet. http://www.lenovo.com/thinkpad/, June 2005.

[72] Lenovo. ThinkPad Z60t and Z60m Notebooks—Data Sheet. http://www.lenovo.com/thinkpad/, September 2005.

[73] Lenovo. ThinkVantage Client Security Software 5.4. http://lenovo.com/thinkvantage, August 2005.

[74] Lenovo. ThinkVantage Client Security Solution — Solid security made simple. http://lenovo.com/thinkvantage, August 2005.

[75] Lenovo. ThinkVantage Client Security Strategy and Client Security Solution 6.0. http:/lenovo.com/thinkvantage, October 2005.

[76] Arjen K. Lenstra. Further progress in hashing cryptanalysis. http://cm.bell-labs.com/who/akl/hash.pdf, February 2005.

[77] Microsoft Corporation. Microsoft Next Generation Secure Computing Base — Technical FAQ. http://www.microsoft.com/technet/security/news/ngscb.mspx, July 2003.

[78] Institute of Electrical and Electronics Engineers (IEEE). IEEE Standard 802.1x — Port-Based Network Access Control, June 2001.

[79] IBM Research. sHype — Secure Hypervisor. http://www.research.ibm.com/secure_systems_department/projects/hypervisor/.

[80] IBM Research. tcgLinux — TPM-based Linux Run-time Attestation. http://www.research.ibm.com/secure_systems_department/projects/tcglinux/.

[81] IBM Research. Virtual TPM Architecture for Xen. http://www.research.ibm.com/secure_systems_department/projects/vtpm/.

[82] Ahmad-Reza Sadeghi and Markus Schneider. Electronic payment systems. In Eberhard Becker, Willms Buhse, Dirk Günnewig, and Niels Rump, editors, Digital Rights Management — Technological, Economic, Legal and Political Aspects, volume 2770 of LNCS, pages 113–137. 2003.

[83] David Safford and Mimi Zohar. A Trusted Linux Client (TLC). http://www.research.ibm.com/gsal/tcpa/tlc.pdf, 2004.

[84] Bede Liu Drew Dien Edwar W. Felton Scott A. Craver, Min Wu. Read-ing Between the Lines: Lessons Learnd form SDMI Challenge. USNIX 2001, 2001.

[85] National Semiconductor. Product Brief: PC8374T SafeKeeper Desktop TrustedI/O. http://www.winbond-usa.com/products/winbond_products/pdfs/APC/PC8374T.pdf, August 2004.

[86] V. Shoup, editor. Collision Search Attacks on SHA-1, volume 3621. Springer, 2005.

[87] Sinosun. Sinosun Trusted Computing Solution. https://www.trustedcomputinggroup.org/ShowcaseApp/sh_catalog_files/4b16e66a5d4ad26ea97bc62d52efc792095d0211/SSX35\%20Product\%20Description\%20-\%20Summary\%20Mar.05.pdf, March 2005.

[88] Utimaco Software. SafeGuard Easy — The Ultimate PC security solution. http://www.utimaco.com/content_products/sg_easy.html, 2005.

[89] Data Brief: ST19WP18-TPM-A Trusted Platform Module (TPM). http://www.st.com/stonline/products/literature/bd/10425.pdf, 2004.

[90] Data Brief: ST19WP18-TPM-B Trusted Platform Module (TPM). http://www.st.com/stonline/products/literature/bd/10425.pdf, 2004.

[91] Data Brief: ST19WP18-TPM-B Trusted Platform Module (TPM). http://www.st.com/stonline/products/literature/bd/10425.pdf, 2004.

[92] Data Brief: ST19WP18 Trusted Platform Module (TPM). http://www.st.com/stonline/products/literature/bd/10425.pdf, 2004.

[93] STMicroelectronics. ST's Trusted Platform Module Provides Complete Trusted Computing Group-Enabled Security Solution for Desktop and Laptop PCs. http://www.st.com/stonline/press/news/year2004/p1499m.htm, September 2004.

[94] STMicroelectronics. STMicroelectronics Enters Volume Production of Trusted Computing Solution and Delivers More Than One Million Chips to Motherboard Manufacturers. http://www.st.com/stonline/press/news/year2005/t1655m.htm, July 2005.

[95] Wave Systems. Embassy Trust Suite — Security for the Enterprise PC. http://www.wave.com/products/03-000164-2_ETS.pdf, 2005.

[96] Wave Systems. Wave Systems TCG-Enabled Toolkit — Enabling TCG-Compliant Application Development. http://www.wave.com/products/03-000172_TK.pdf, 2005.

[97] Trusted Computing Group (TCG). TCG PC Specific Implementation Specification. https://www.trustedcomputinggroup.org/groups/pc_client/TCG_PCSpecificSpecification_v1_1.pdf, August 2003.

[98] Trusted Computing Group (TCG). TCG Specification Architecture Overview. http://www.trustedcomputing.org, April 2004.

[99] Trusted Computing Group (TCG). TLS Extensions for Attestation. http://www.trustedcomputing.org, July 2004.

[100] Trusted Computing Group (TCG). TPM Main Specification — Part 1: Design Principles. https://www.trustedcomputinggroup.org/groups/tpm/mainP1DP_rev85.zip, February 2005.

[101] Trusted Computing Group (TCG). TPM Main Specification 1.2. http://www.trustedcomputing.org, February 2005.

[102] Infineon Technologies. Infineon Extends Security Across the Enterprise with HP — New HP Compaq Business Notebooks Include Advanced Security Chip Technology. http://www.infoneon.com, Information Number INFSMS200310.007e, October 2003.

[103] Infineon Technologies. Infineon Helps Bring New Level of Security to Computer Networks; Provides Embedded Security Chip Solution for New HP Compaq Business Desktop PC. http://www.infineon.com, Information Number INFSMS200305.078, May 2003.

[104] Don Felton Tiago Alves. TrustZone: Integrated Hardware and Software Security. http://www.arm.com/pdfs/TZ%20Whitepaper.pdf, July 2004.

[105] Christian Vilsbeck. AMD Pacifica: Virtualisierung von CPU & Speicher, October 2005.

[106] Ross Anderson: Trusted Computing' Frequently Asked Questions – Verson 1.1, 2003; URL: http://www.cl.cam.ac.uk/~rja14/tcpa-faq.html

[107] Paul Baran: On Distributed Communications: IX. Security, Secrecy, and Tamper-Free Considerations. Memorandum RM-3765-PR, August 1964, The Rand Corporation, 1700 Main St, Santa Monica, California, 90406 Reprinted in: Lance J. Hoffman (ed.): Security and Privacy in Computer Systems; Melville Publishing Company, Los Angeles, California, 1973, 99--123. http://www.rand.org/pubs/research_memoranda/RM3765/index.html

[108] Oliver Berthold, Hannes Federrath, Stefan K\"opsell: Web MIXes: A system for anonymous and unobservable Internet access. in Proc. of Designing Privacy Enhancing Technologies: Workshop on Design Issues in Anonymity and Unobservability, Springer Verlag, LNCS 2009, July 2000, 115-129.

[109] Jan Camenisch, Ernie Brickell, Liqun Chen, 2004, Direct Anonymous Attestation

[110] David Chaum: Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. Communications of the ACM 24/2, 1981, 84--88.

[111] George Danezis, Roger Dingledine, Nick Mathewson: Mixminion: Design of a Type III Anonymous Remailer Protocol. in Proc. of the 2003 IEEE Symposium on Security and Privacy, May 2003.

[112] Roger Dingledine, Nick Mathewson, Paul Syverson: Tor: The Second-Generation Onion Router. in Proc. of the 13th USENIX Security Symposium, August 2004.

[113] Michael Gross, 1991, Vertrauenswürdiges Booten als Grundlage authentischer Betriebssysteme. In: Verläßliche Informationsysteme, Tagungsband, Informatikfachberichte 271

[114] Hermann Härtig, Lenin Singaravelu, Calton Pu, Christian Helmuth, 2006, Reducing TCB Complexity for Security-Sensitive Applications: Three Case Studies, EuroSys 2006

[115] Bernhard Kauer, 2004, Authenticated booting for L4

[116] Bernhard Kauer: OSLO: Improving the security of Trusted Computing

[117] Andreas Pfitzmann, Marit Hansen: Anonymity, Unlinkability, Unobservability, Pseudonymity, and Identity Management - A Consolidated Proposal for Terminology. http://dud.inf.tu-dresden.de/Anon_Terminology.shtml, V0.28, August 2006.

[118] Seth Schoen: Trusted Computing: Promise and Risk, 2003

[119] TCG Infrastructure Working Group, 2006, Archicture Part II - Integrity Management

[120] Ruediger Weis; Stefan Lucks; Andreas Bogk: TCG 1.2 – fair play with the 'Fritz' chip?, 2004

[121] Hermann Härtig, Lenin Singaravelu, Calton Pu, Christian Helmuth, 2006, Reducing TCB Complexity for Security-Sensitive Applications: Three Case Studies, EuroSys 2006

[122] Yianna Danidou, Legal Implications of Trusted Computing, BILETA Annual Conference, 2007

[123] Andreas Schmidt, Lecture 7: Trusted in Identity Management Systems, Trusted Computing: Introduction & Applications, 2007

[124] Wikipedia – The Free Encyclopedia, Trusted Third Party, http://en.wikipedia.org/wiki/Trusted_Third_Party

[125] Tätigkeitsbericht, Independent Centre for Privacy Protection (ICPP), http://www.datenschutzzentrum.de/material/tb/tb26/kap11.htm#Tz11.2

[126] Sebastian Clauß, Marit Hansen, Els van Herrwegen, Andreas Pfitzmann, Privacy-Enhancing Identity Management, http://www.jrc.es/home/report/english/articles/vol67/IPT2E676.htm

[127] Working Document on Trusted Computing Platforms and in particular on the work done by the Trusted Computing Group (TCG group), Article 29 Data Protection Working Party, http://europa.eu.int/comm/internal_market/privacy/docs/wpdocs/2004/wp86_en.pdf

[128] Working Document on Trusted Computing Platforms and in particular on the work done by the Trusted Computing Group (TCG group), Article 29 Data Protection Working Party, http://europa.eu.int/comm/internal_market/privacy/docs/wpdocs/2004/wp86_en.pdf

[129] TCPA darf nicht zur Aushebelung des Datenschutzes missbraucht werden, Entschließung der 65. Konferenz der Datenschutzbeauftragten des Bundes und der Länder am 27./28. März 2003 in Dresden, http://www.datenschutz-berlin.de/doc/de/konf/65/top05.htm

[130] Scientific Evaluation of DRM Systems, Hannes Federrath, http://www-sec.uni-regensburg.de/publ/2002/FederrathDRM20020129.pdf

[131] Direct Anonymous Attestation, Achieving Privacy in Remote Authentication, Jan Camenisch, http://www.zisc.ethz.ch/events/ISC2004Slides/folien-jan-camenisch.pdf

[132] Lizenz/EULA, Wikipedia – Die freie Enzyklopädie, http://de.wikipedia.org/wiki/Lizenz#EULA

[133] Tätigkeitsbericht, Independent Centre for Privacy Protection (ICPP), http://www.datenschutzzentrum.de/material/tb/tb26/kap11.htm#Tz11.3

[134] xp-Antispy FAQ at http://www.xp-antispy.org/

[135] EU-Kommission für Banknoten-Kopierschutz, Christiane Schulzki-Haddouti, http://www.heise.de/newsticker/meldung/47083

[136] Markus Hansen, Jan Möller: Digital Rights Management zwischen Sicherheit und informationeller Selbstbestimmung, in: Bundesamt für Sicherheit in der Informationstechnik (BSI) (Ed.): IT-Sicherheit geht alle an!, Tagungsband zum 9. Deutschen IT-Sicherheitskongress des BSI, 2005, 159-171, http://www.datenschutzzentrum.de/vortraege/050510_hansen-moeller_bsi.htm

[137] EPAL 1.2, W3C Member Submission, http://www.w3.org/Submission/EPAL/, and EPAL FAQ, Jan Möller, Independent Centre for Privacy Protection (ICPP) Schleswig-Holstein, http://www.datenschutzzentrum.de/faq/epal.htm

[138] AOL Employee Charged For Selling ISP's Customer List To Spammer, Spam News Ticker,http://www.unspam.com/fight_spam/articles/1410.html

[139] Positionspapier Trusted Computing, CODEattac, ATTAC Austria, http://www.attac-austria.org/apps/ewiki/index.php?id=CODEattac_PosPap_TrustedComputing

[140] Barbara Fichtinger, Eckehard Hermann, Nicolai Kuntze, Andreas Schmidt: Trusted Infrastructures for Identities, Koblenz 2007, http://www.virtualgoods.org/2007/10_VG07_Fichtinger_Hermann_Kuntze_Schmidt.pdf

[141] Ammar Alkassar, Lothar Fritsch, Rani Husseiki: TC-ERTA - Analysis for security techniques complementary to Trusted Computing, Bundesamt für Sicherheit in der Informationstechnik (BSI), 2008

[142] Microsoft Research: Singularity, http://research.microsoft.com/os/Singularity/

Annex 1: Glossary

AAA	Authentication, Authorization and Accounting
AC'97	Audio Codec '97
ACPI	Advanced Configuration and Power Management
AGP	Advanced Graphics Port
AIK	Attestation Identity Key
AMD	Advanced Micro Devices
API	Advanced Programmable Interface
APM	Advanced Power Management
AR	Access Requester
ASIC	Application-Specific Integrated Circuit)
ATA	Advanced Technology Attachment
ATX	Advanced Technology Extended
AVP	Attribute Value Pair
BIOS	Basic Input/Output System
BTX	Balanced Technology Extended
CA	Certification Authority
CC	Common Criteria
CF	Compact Flash
CMK	Certified Migratable Key
CPU	Central Processing Unit
CRTM	Core Root of Trust for Measurement
CSP	Cryptographic Service Provider
CSS	Client Security Solution (Lenovo ThinkVantage Technology)
CTSS	Core TCG Software Stack
DAA	Direct Anonymous Attestation
DDR	Double Data Rate
DES	Data Encryption Standard
DMA	Direct Memory Access
DMVT	Dynamic Video Memory Technology
DNB	Dedicated Network Bus (Intel)
DRM	Digital Rights Management
DTS	Digital Theater System
DVI	Digital Video Interface
EAL	Evaluation Assurance Level
EAP	Extensible Authentication Protocol
EBX	Embedded Board Expandable
EEPROM	Electrically Erasable Programmable Read-Only Memory
EFS	Encrypted File System (Microsoft)
EK	Endorsement Key
EM64T	Extended Memory 64-bit Technology
EMSCB	European Multilateral Secure Computing Base
EPROM	Erasable Programmable Read-Only Memory
EVM	Extended Verification Module

FIFO	First In First Out
FIPS PUB	Federal Information Processing Standards Publication
GPIO	General Purpose Input/Output
GPL	General Public License
GRUB	Grand Unified Bootloader
HD	High Definition
HDD	Hard Disk Drive
HDTV	High Definition Television
HMAC	Hash Message Authentication Code (keyed)
HP	Hewlett-Packard
HT	Hyper-threading Technology
I/O	Input/Output
IBM	International Business Machines
ICH	Input/output Controller Hub
IDE	Integrated Device Electronics
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IKE	Internet Key Exchange protocol
IKEv2	Internet Key Exchange protocol version 2
IPSec	Internet Protocol Security
iSCSI	Internet SCSI
ISO	International Organization for Standardization
KCM	Key Cache Manager
KTM	Key Transfer Manager (Wave Systems)
LAN	Local Area Network
LCD	Liquid Crystal Display
LDAP	Lightweight Directory Access Protocol
LEAP	Light Extensible Authentication Protocol
LGA	Land Grid Array
LoM	LAN on Motherboard
LPC	Low Pin Count
LSM	Linux Security Module
MA	Migration Authority
MBR	Master Boot Record
MCH	Memory Controller Hub
MD5	Message Digest 5
MK	Migratable Key
MLS	Multi Level Security
MMC	Multimedia Card or Microsoft Management Console
MS-CAPI	Microsoft Cryptographic Advanced Programmable Interface
MS-CSP	Microsoft Cryptographic Service Provider
MSA	Migration-Selection Authority
MSS	Multi Server System
NAS	Network Attached Storage
NGSCB	Next Generation Secure Computing Base (Microsoft)

NIC	Network Interface Controller
NIST	National Institute for Standards
NMK	Non-migratable Key
OTC	Open Trusted Computing (a.k.a. OpenTC)
OTPROM	One Time Programmable Read-Only Memory
PATA	Parallel ATA
PCI	Peripheral Component Interconnect
PCR	Platform Configuration Register
PDA	Personal Data Assistant
PDP	Policy Decision Point
PEP	Policy Enforcement Point
PKCS#11	Public Key Cryptography Standard #11
PKI	Public Key Infrastructure
PP	Protection Profile
PPP	Point-to-Point Protocol
PS/2	Personal System 2
PSD	Personal Secure Drive (HP ProtectTools Embedded Security Solution)
RADIUS	Remote Authentication Dial-In and User Service
RAID	Redundant Array of Independent (or Inexpensive) Disks
RAM	Random Access Memory
RF	Radio Frequency
RFID	Radio Frequency Identification
RNG	Random Number Generator
ROM	Read-Only Memory
RPC	Procedure Call
RTC	Real Time Clock
RTS	Root of Trust for Storage
SATA	Serial ATA
SCSI	Small Computer System Interface
SD	Secure Digital
SDRAM	Synchronous Dynamic Random Access Memory
SHA-1	Secure Hash Algorithm 1
SLIM	Simple Linux Integrity Module
SMBus	System Management Bus (a.k.a. SMB)
SML	Storage Measurement Log
SoC	System on a Chip
SRAM	Static Random Access Memory
SRK	Storage Root Key
SSL	Secure Sockets Layer
ST	Security Target
SXGA	Super XGA
TBB	Trusted Building Block
TCG	Trusted Computing Group
TCPA	Trusted Computing Platform Alliance
TCS	TSS Core Services
TCSI	TSS Core Services Interface

TDDL	TPM Device Driver Library
TDDLI	TPM Device Driver Library Interface
TLC	Trusted Linux Client
TLS	Transport Layer Security
TNC	Trusted Network Connect
TOE	Target of Evaluation
TPM	Trusted Platform Module
TSP	TCG Service Provider
TSPI	TCG Service Provider Interface
TSS	TCG Software Stack
UDMA	Ultra DMA
ULV	Ultra Low Voltage
USB	Universal Serial Bus
UXGA	Ultra XGA
VGA	Video Graphics Array
VM	Virtual Machine
VMM	Virtual Machine Monitor
VoIP	Voice over Internet Protocol
VPN	Virtual Private Network
WAN	Wide Area Network
WLAN	Wireless Local Area Network
WSXGA	Wide Super XGA
WWAN	Wireless Wide Area Network
WXGA	Wide XGA
XGA	Extended Graphics Array