

On the deployment of Mobile Trusted Modules

Andreas U. Schmidt,
Nicolai Kuntze, Michael Kasper
Fraunhofer Institute for Secure Information Technology SIT
Rheinstraße 75, 64295 Darmstadt, Germany
Email: {andreas.schmidt,nicolai.kuntze,michael.kasper}@sit.fraunhofer.de

Abstract—In its recently published *TCG Mobile Reference Architecture*, the TCG Mobile Phone Work Group specifies a new concept to enable trust into future mobile devices. For this purpose, the TCG devises a trusted mobile platform as a set of trusted engines on behalf of different stakeholders supported by a physical trust-anchor. In this paper, we present our perception on this emerging specification. We propose an approach for the practical design and implementation of this concept and how to deploy it to a trustworthy operating platform. In particular we propose a method for the take-ownership of a device by the user and the migration (i.e., portability) of user credentials between devices.

I. INTRODUCTION

As a first deliverable within their scope and work programme, the Mobile Phone Work Group of the Trusted Computing Group (TCG MPWG) has published a specification [1], which offers new potentials for implementing trust in mobile computing platforms by introducing a new, hardware-based trust anchor for mobile phones and devices. This trust anchor, called a Mobile Trusted Module (MTM), has properties and features comparable to a Trusted Platform Module (TPM, see [2], [3]). Concurrently the MPWG issued a much more universal security architecture for mobile phones and devices on a higher abstraction level. The pertinent specification is called *TCG Mobile Reference Architecture (RA)* [4] and abstracts a trusted mobile platform as a set of tamper resistant trusted engines operating on behalf of different stakeholders. This architecture offers a high degree on flexibility and modularity in design and implementation of the trusted components to all participants in hard- and software development.

An important aspect of the *TCG Mobile Reference Architecture* is the potential to virtualise significant parts of a trusted mobile platform as trusted software applications and services. The trusted execution chain for this rests on the MTM. The implementation of this chip depends on the security requirements of its specific use-case. For high levels of protection and isolation, an MTM could be implemented as a slightly modified Trusted Platform Module (TPM). This enables cost-effective implementation of new security-critical applications and various innovative business models, in both the mobile and generic computing domain [5]–[7].

The present paper discusses the main structural features of the RA, highlighting the capabilities of the MTM as the main functional building block. After this technology review, we propose two basic methods for usage of the RA, namely the

set-up of a trusted subsystem on a device by a remote owner, and its migration from one device to another.

This paper is organised as follows. In Section II, we explore the significant parts of the *MPWG Reference Architecture*. It is divided into four parts. Subsection II-A gives an overview of the security architecture, and subsection II-B details the concepts of the proposed architectural approach for an MTM and the requirements to virtualise its functionality, whereby a high security and isolation level is maintained. Furthermore, we propose a model for remote stakeholder take-ownership in II-C and migration of trusted subsystems in II-D. In Section III, we show how such an architecture can be implemented on trustworthy operating platforms.

II. TCG MPWG REFERENCE ARCHITECTURE

The TCG MPWG has developed an architecture on a high level of abstraction for a trusted mobile platform, which offers numerous variations for design and implementation. In this section, we reflect essential parts of this architecture and an overview of significant platform components in terms of our objective.

A. Architectural Overview

A trusted mobile platform is characterised as a set of multiple tamper-resistant engines, each acting on behalf of a different stakeholder. Broadly, such a platform has several major components: trusted engines \mathcal{TE} , trusted services \mathcal{TS} customised by trusted resources \mathcal{TR} . A general trusted mobile platform is illustrated in Figure 1.

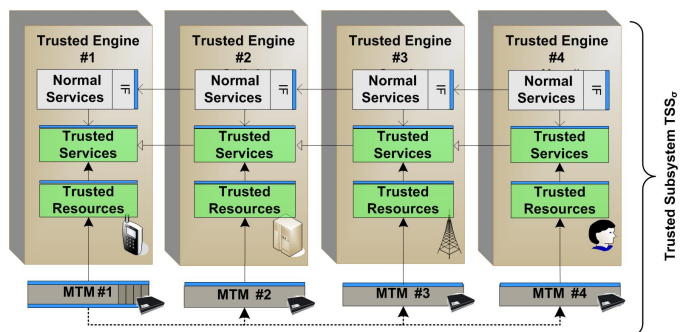


Fig. 1. Trusted Mobile Platform Architecture

We define a trusted subsystem \mathcal{TSS}_σ as a logical unit of a trusted engine together with its interrelated hardware compartment. A \mathcal{TSS} of a stakeholder σ can formally described

by a tuple

$$TSS_{\sigma} = \{TE_{\sigma}, TS_{\sigma}, TS_{\epsilon}, TR_{\sigma}, SP_{\sigma}, SC_{\sigma}\}$$

In each trusted subsystem TSS either a remote or local entity acts as a stakeholder, who is able to configure its own subsystem and define his security policy SP_{σ} and system configuration SC_{σ} within an isolated and protected environment. The *MPWG Reference Architecture* specifies the following principal entities: the local stakeholders *Device Owner DO* and *User U*; and the remote stakeholders *Device Manufacturer DM*, and more general *Remote Owners RO* (e.g. a communication carrier, or service provider). The functionality of a TSS is either based on dedicated resources of an embedded engine TE_{σ} , or may be provided by trusted services TS_{ϵ} of external engines.

Each subsystem is able to enforce its security policy SP_{σ} and subsystem configuration SC_{σ} . As a consequence, the functionality of a trusted subsystem TSS_{σ} is constrained by the available resources TR_{σ} with their derived trusted services TS_{σ} , by the offered functionality of external trusted services TS_{ϵ} , by the security policy SP_{σ} , and finally the system configuration SC_{σ} of an engine's stakeholder.

All internal functions executed inside TSS_{σ} are isolated from other subsystems by the underlying security layer and is only accessible if a proper service interface is defined and exported. A TSS_{σ} relies on the reputation of the stakeholder σ as basis for that trust. Therefore, each stakeholder issues a security policy SP_{σ} and a set of credentials belonging to embedded trusted components of its subsystem TSS_{σ} . This policy contains reference measurements (RIM), quality assertions and security-critical requirements.

1) *Trusted Engines*: The most important concept within the *MPWG Reference Architecture* is that of trusted engines. The

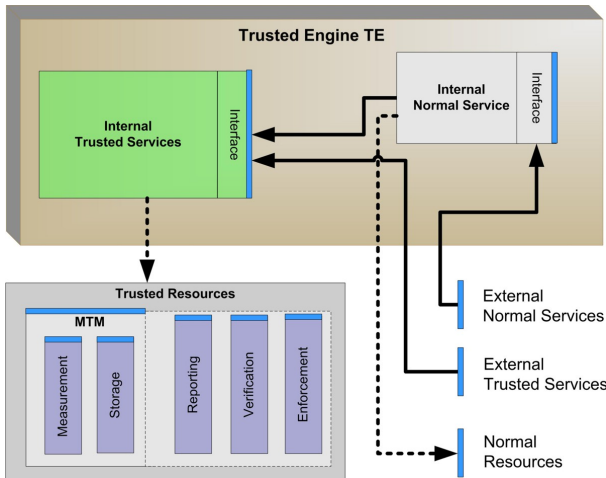


Fig. 2. Generic Trusted Engine

purpose of a trusted engine is to provide confidence in all its embedded services, which are internally or externally provided by the engine. It is a protected entity on behalf of a specific stakeholder that has special abilities to manipulate and store

data, and to provide evidence of its trustworthiness and the current state of the engine. Figure 2 shows a generic trusted engine. In general, each engine has at least following abilities:

- implement arbitrary software functionalities as trusted and/or normal services,
- provide the evidence for its trustworthiness,
- report the evidence of its current state,
- obtain and use Endorsement Keys (EK) and/or Attestation Identity Keys (AIK),
- access a set of trusted resources, and
- import and/or export services, shielded capabilities and protected functionality.

In order to establish a definite categorisation, the MPWG differentiates engines according to their functional dispensability. Therefore, an engine is either dedicated to a mandatory (of DO or DM) or a discretionary domain (of DO). Engines inside a mandatory domain are permanently located on a trusted platform and hold security-critical and essential functionality. All essential services of a trusted mobile platform should be located inside the mandatory domain, which does not permit a local stakeholder to remove a remote owner from the engine. Mandatory engines have access to a *Mobile Remote owner Trusted Module (MRTM)* to guarantee that a valid and trustworthy engine state is always present.

Non-essential engines and services are replaceable by the device owner DO and should be located inside the discretionary domain. Engines inside the discretionary domain are controlled by the device owner DO . Discretionary engines are required to be supported by a *Mobile Local owner Trusted Module (MLTM)*.

2) *Trusted Resources*: As illustrated in Figure 2, an internal trusted service has access to several trusted resources. The TCG calls these resources *Root-Of-Trusts (RoT)* representing the trusted components acting on base of the trusted execution chain and providing functionality for measurement, storing, reporting, verification and enforcement that affect the trustworthiness of the platform. The following RoTs are defined for the mobile domain:

- Root of Trust for Storage (RTS),
- Root of Trust for Reporting (RTR),
- Root of Trust for Measurement (RTM),
- Root of Trust for Verification (RTV), and
- Root of Trust for Enforcement (RTE)

Each RoT vouches its trustworthiness either directly by supplied secrets (EK, AIK) and associated credentials, which are only accessible by authenticated subjects of the stakeholder, or indirectly by measurements of other trusted resources. These resources are only mutable by authorised entities of a stakeholder.

In this paper, we group several logically self-contained RoTs to simplify the presentation of interfaces and the communication layer. In a typical arrangement, the RTS and RTR represent one unit, while the RTM and RTV build another unit within an TSS_{σ} . However, note that the RTV and the RTM depend on protected storage mechanisms, which are provided

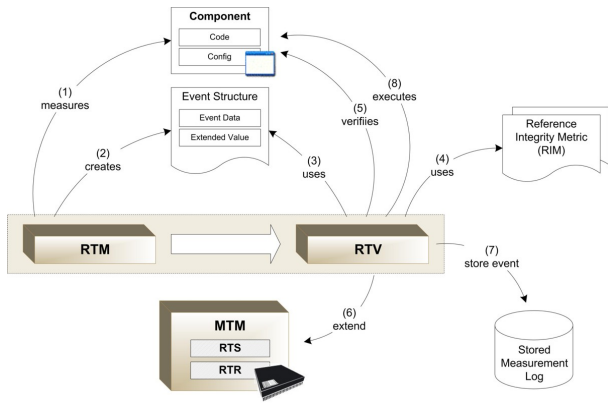


Fig. 3. Measurement and Verification Process

by the RTS. Thus, it is also plausible to implement all RoTs together as a common unit within an engine.

RTS/RTR are the trusted resources that are responsible for secure storage and reliable reporting of information about the state of trusted mobile platform. An RTS provide PCRs and protected storage for an engine and stores the measurements made by the RTM, cryptographic keys, and security sensitive data. An RTR signs the measurements with cryptographic signature keys of TSS_σ .

RTM/RTV In general, an RTM is a reliable instance to measure software components and provide evidence of the current state of a trusted engine and its embedded services. In the mobile domain, to avoid communication costs, this functionality is extended by a local verifier, which checks the measurements against a given *Reference Integrity Metrics* (RIM). This process can be done instantly as the measurements are performed employing a combination of RTM and RTV. Figure 3 depicts such a *Measure*→*Verify*→*Extend* process.

An RTE is required if an engine uses allocated resources and services. In this case, such RoT acts as a trusted boot loader and ensures the availability of all allocated trusted resources and services within that trusted subsystem.

3) *Services of a Trusted Engine*: A trusted engine integrates all functionality by customising available platform resources as software services. Such a service offers computation, storage, or communication channels to other internal or external services and applications based on dedicated or allocated resources. The MPWG categorises them into: trusted, normal, and measured services.

A trusted service customises trusted resources. Thus, a trusted service is implicitly supplied with an *EK* or *AIK* in order to attest its trustworthiness. Trusted services are intended to provide reliable measurements of their current state and to provide evidence of the state of other normal services or resources.

Normal services are customising normal resources and implement functionality, but are not able to provide evidence of their trustworthiness by own capabilities. However, normal services can access internal trusted services to use their provided functionality. Therefore, an internal normal services

is able to vouch its trustworthiness by associated integrity metrics that have been measured by a trusted service.

B. Mobile Trusted Module

The generic term *Mobile Trusted Module (MTM)* refers to a dedicated hardware-based trust-anchor. It is typically composed of an RTS and RTR and has characteristics comparable to a TPM. According to their design objective the MPWG distinguishes between MRTM and MLTM. Both must support a subset of TPM commands as specified in [4]. Additionally, an MRTM has to support a set of commands to enable local verification and specific mobile device functionality.

The *TCG MPWG Reference Architecture* does not exclude to utilise a TPM v1.2 (or even a TPM v1.1) as an MTM, if an appropriate interface consisting of a set of commands conforming to the MPWG specification and associated data structures are provided. Although it is possible to implement this architecture upon a standard TPM, we here focus on a more sophisticated solution based on a Trustworthy Computing Platform such as EMSCB/Turaya [8]. In this context, we expect three different solutions for isolation, key management and protection of TSS_σ .

A **Standard TPM-based Model** uses a non-modified standard TPM to build the trusted computing base. The secret keys are stored into a single key-hierarchy on behalf of \mathcal{DO} as specified in [1]. In this case, an adversary or malicious local owner may be able to access the secret keys of a remote stakeholder and take control of a remote owner compartment. A \mathcal{DO} can also disable the whole MTM or corrupt mandatory engines of remote stakeholders.

A **Software-based MTM-Emulation Model** uses a software-based allocated *MTM*-emulation with an isolated key-hierarchy. All sensitive and security-critical, such as *EK* or *SRK*, are only protected by software mechanisms outside of the tamper-resistant environment of a dedicated MTM [9], [10].

Generic MTM-based Model supporting multiple stakeholders and virtual MTMs. In order to circumvent resulting drawbacks and mitigate attacks, we favour a solution with a higher level of security. For this reason, we adopt the proposed secure co-processor variant of [9] and describe a generic MTM with support for multiple stakeholder environments. In a cost-efficient scenario, the trusted mobile platform is implementable based on a single generic MTM and several virtualised MTMs for each trusted engine. Hence, at least one dedicated MTM has to be available and additionally a unique vMTM has to be instantiated in each trusted subsystem TSS_σ . In such case, a physically bounded MTM in the platform acts as a master trust anchor and offers MRTM and MLTM functionality with respect to its specific use case.

A Trusted Software Layer offers a *vMTM Proxy Service* to all embedded trusted engines \mathcal{TE}_σ . The main task of this service is to route MTM commands from a \mathcal{TE}_σ to its dedicated instance $vMTM_\sigma$. The advantage is that all security-critical MTM commands are tunnelled to $vMTM_\sigma$

D. Migration of a Trusted-Subsystem

If a stakeholder wants to move a source $TSS_{\sigma,S}$ to another MTM-enabled platform, for instance to port user credentials from device to device, all security-critical information including the Storage Root Key (SRK) has to be migrated to the target $TSS_{\sigma,D}$. In our scenario, we assume the same remote owner (e.g. mobile network operator) on both subsystems $TSS_{RO,S}$ and $TSS_{RO,D}$.

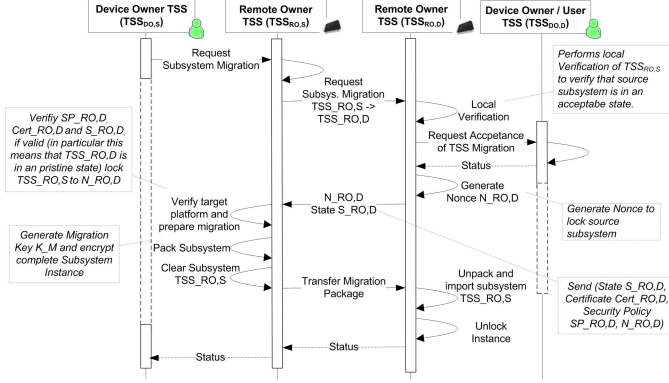


Fig. 6. Trusted Subsystem Migration Protocol

To be able to securely migrate the SRK, we suggest a modification of the current MPWG specification to allow *inter-stakeholder-migration* of a complete isolated key hierarchy. Thus, an isolated key hierarchy is (1) migratable between environments of identical stakeholders, (2) if and only if an entitling security policy on both platforms exists. The advantage of migration between identical stakeholder subsystems is that the migration process doesn't require a trusted third party. We only involve the owner in combination with local verification mechanisms of the TSS_{RO} to migrate the trusted subsystem (including the SRK) to another platform. This enables for instance direct, device-to-device porting of credentials, e.g. using short-range communication. We here propose a complete, multilateral and secure migration protocol, which is illustrated in Figure 6.

Platform and Protocol Precondition. Similar to section II-C, the trusted mobile platform has carried out the same initial steps as mentioned above. Furthermore, the remote owner has performed an remote take-ownership procedure as described in II-C.

Trusted Subsystem Migration Protocol. At the beginning of the migration protocol, the device owner DO_S of the source platform TP_S initialises the migration procedure and requests an appropriate migration service of $TSS_{RO,S}$. Next, the trusted platform TP_S is instructed by TSS_{RO} to establish a secure channel to the target platform TP_D . After the connection is available, $TSS_{RO,S}$ activates the corresponding migration service of $TSS_{RO,D}$ to perform the import procedure. Thereupon, the target subsystem $TSS_{\sigma,D}$ performs a local verification of $TSS_{RO,S}$. If revoked, it replies with an error message and halts the protocol. Otherwise $TE_{RO,D}$ requests an confirmation from the local owner DO_D .

Next, the target subsystem $TSS_{RO,D}$ generates a nonce $N_{RO,D}$. In order to provide evidence of its trustworthiness, $TSS_{RO,D}$ sends all necessary information to the source subsystem $TSS_{RO,S}$. This includes the current state $S_{RO,D}$, a certificate of $TSS_{RO,D}$, security policy $SP_{RO,D}$ and the nonce $N_{RO,D}$. Having received the target subsystem's message, $TSS_{RO,S}$ verifies the state of $TSS_{RO,D}$. If the target system is in a trustworthy state and holds an acceptable security policy and system configuration, the current state of $TSS_{RO,S}$ is locked to nonce $N_{RO,D}$.

The $TSS_{RO,S}$ generates a symmetric migration key K_M , serialises its instance and encrypts it with the migration key, which is bound to an acceptable configuration of $TSS_{RO,D}$. Next, the key-blob and the encrypted instance are sent to the destination $TSS_{RO,D}$. In particular, this includes the whole isolated key-hierarchy $\mathcal{K}_{RO,S}$ with $SRK_{RO,S}$, the security policy $SP_{RO,S}$, and the required subsystem configuration $\mathcal{S}_{RO,S}$.

Finally, the target subsystem $TSS_{RO,D}$ decrypts the received blob and uses $SRK_{RO,S}$ as its own SRK . The subsystem verifies the obtained security policy $SP_{RO,S}$ and the subsystem configuration $\mathcal{S}_{RO,S}$. With this information, $TSS_{RO,D}$ rebuilds the internal structure of the source.

The source system should then be notified of the success of migration and ultimately delete the migrated key hierarchy (or even do it before sending the migration package as indicated for simplicity in Figure 6). Otherwise one obtains replicated trusted subsystems, by themselves indistinguishable to the remote owner. But this may depend on the policies to be enforced in the particular use case.

III. DESIGN OF MOBILE TRUSTED MODULES ON TRUSTWORTHY OPERATING PLATFORMS

A prototypical implementation of the trusted engines and the specified trusted services was realised as an extension to the existing EMSCB / Turaya Computing Platform. Turaya is an implementation of the EMSCB security architecture. It provides fundamental security mechanisms and a protected and isolated execution environment, which meet the requirements of the *MPWG Reference Architecture* [8], [11], [12].

Figure 7 illustrates our model, in which a hypervisor/microkernel executes a legacy operating system in coexistence with a running instance of the EMSCB-based security architecture. The latter controls a virtual machine with several trusted engines and services compliant to the MPWG requirements [1], [4]. In the following paragraphs, we outline the significant platform layers concerning our approach.

The **Hardware Layer** of our model includes a generic MTM as described in Section II-B, in addition to conventional hardware components. This MTM acts as a dedicated master trust anchor for the complete trusted mobile platform.

The **Virtualisation Layer** provides generic hardware abstraction, between the physical hardware of a trusted mobile platform and the *Trusted Software Layer* below. The EMSCB project supports microkernels of the L4-family [14] such as hypervisors [15]. In general, all solutions provide mechanisms

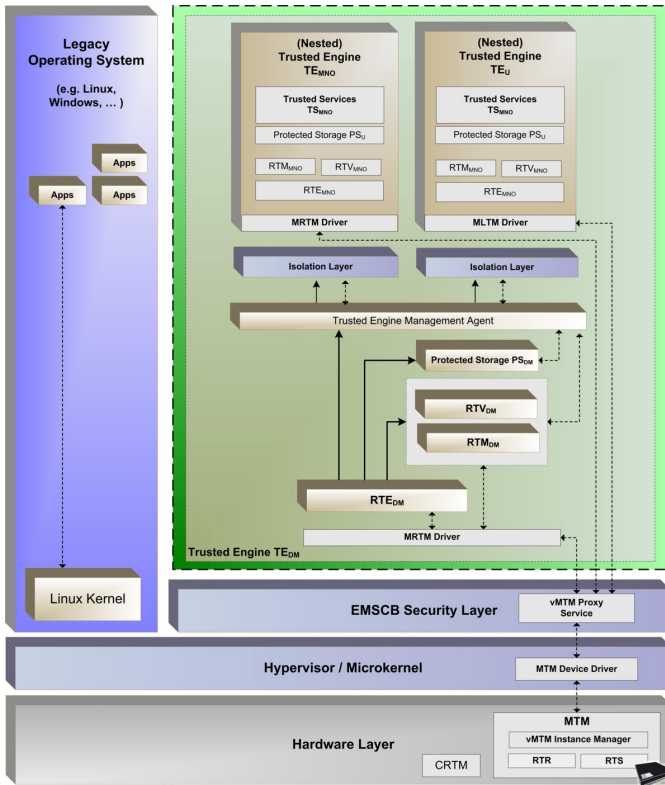


Fig. 7. Trustworthy Operating Platform with multiple Trusted Engines

for resource management, inter-process-communication (IPC), virtual machines, memory management and scheduling. In our specific case, the virtualisation layer includes also a fully functional device driver for a dedicated generic MTM. Furthermore, it is responsible for instantiation of both the trusted software layer and the legacy operating system.

The **Trusted Software Layer** provides security functionality and is responsible for isolation of embedded applications and software compartments. It also implements the *vMTM Proxy Service* as described in Section II-B. Currently, EMSCB/Turaya provide an excellent foundation by its security services (trust manager, compartment manager, storage manager), which are required by the RTR and RTV, Protected Storage and Trusted Engines Management Agent of \mathcal{TE}_{DM} . Therefore, it is reasonable to build the significant parts of the device manufacturer engine \mathcal{TE}_{DM} within this layer.

Trusted engines \mathcal{TE}_{σ} within the **Application Layer** are implemented as parallel and isolated L4Linux compartments [13] on behalf of different stakeholders. Each compartment has access to its *vMTM* instance through an embedded client-side device driver. This driver constrains the functionality with respect to its specific use case (MRTLM or MLTM). Furthermore, \mathcal{TE}_{σ} has an RTE_{σ} , which is responsible for building all required allocated trusted resources and services depending of its specific system configuration SC_{σ} and security policy SP_{σ} .

IV. CONCLUSION AND FURTHER WORK

We have introduced the Trusted Engines and MTMs in terms of our objective. In this context, we have exposed significant parts of the MPWG Reference Architecture and how it can be implemented on a (very slightly modified) TPM trust-anchor. We have shown how to deploy trusted virtualised compartments on devices and exhibited basic operations required in the mobile domain, such as migration.

Using a *vMTM* in lieu of a Subscriber Identity Module (SIM) as a trusted and protected software allows expansion to a much wider field of authentication and identification management systems even on standard PC platforms [16]. Supporting online transactions by authentication via credentials held in a *vMTM* may be one attractive use case. However, there are some privacy and security challenges associated with this implementation on a desktop computer, which require further research. Finally, replacing SIMs/USIMs by multi-purpose *vSIMs* may be attractive even for genuine mobile devices.

REFERENCES

- [1] TCG. TCG MPWG Mobile Trusted Module specification, version 1.0, Revision 1, 12 June 2007. <https://www.trustedcomputinggroup.org/specs/mobilephone/tcg-mobile-trusted-module-1.0.pdf>
- [2] TCG. TCG specification architecture overview, specification revision 1.2, 2004.
- [3] TCG. Trusted Platform Module specification (TPM), version 1.2, revision 94, 2005. <https://www.trustedcomputinggroup.org/specs/TPM>
- [4] Trusted Computing Group, *TCG Mobile Reference Architecture*. Specification version 1.0, Revision 1. 12 June 2007. <https://www.trustedcomputinggroup.org/specs/mobilephone/tcg-mobile-reference-architecture-1.0.pdf>
- [5] Nicolai Kuntze and Andreas U. Schmidt. Transitive trust in mobile scenarios. In Günter Müller, editor, *Proceedings of the International Conference on Emerging Trends in Information and Communication Security (ETRICS 2006)*, volume 3995 of *Lecture Notes in Computer Science (LNCS)*, pages 73–85. Springer-Verlag, 2006.
- [6] Kuntze, N., Schmidt, A.U.: Trusted computing in mobile action. In Venter, H.S., Eloff, J.H.P., Labuschagne, L., Eloff, M.M., eds.: *Proceedings of the Information Security South Africa (ISSA) Conference (2006)*
- [7] Nicolai Kuntze and Andreas U. Schmidt. Trusted ticket systems and applications. In *To appear in: New Approaches for Security, Privacy, and Trust in Complex Systems. Proceedings of the IFIP sec2007. Sandton, South Africa 14-16 May 2007*. Springer-Verlag, 2007.
- [8] EMSCB. European multilaterally secure computing base - towards trustworthy systems with open standards and trusted computing. 2006. <http://www.emscb.com/>
- [9] Stefan Berger, Ramon Caceres, Kenneth A. Goldman, Ronald Perez, Reiner Sailer, and Leendert van Doorn. *vTPM: Virtualizing the trusted platform module*.
- [10] Mario Strasser. *Linux TPM emulator*, 2005.
- [11] Ahmad-Reza Sadeghi, Christian Stübke, and Norbert Pohlmann. European multilateral secure computing base - open trusted computing for you and me. *Datenschutz und Datensicherheit (DuD)*, pages 548–554, 2004.
- [12] PERSEUS. The perseus security framework. <http://www.perseus-os.org>. Applied Data Security Group, Ruhr-University Bochum.
- [13] Michael Hohmuth. *Linux-Emulation auf einem Mikrokern*. PhD thesis, TU Dresden, Fakultät Informatik, Lehrstuhl Betriebssysteme, 1996.
- [14] R. Baumgartl, M. Borriss, Cl.-J. Hamann, M. Hohmuth, L. Reuther, S. Schönberg, and J. Wolter. Dresden realtime operating system (drops). In *Workshop of System-Designed Automation, (SDA'98)*, 1998.
- [15] P. Barham. *Xen and the art of virtualization*, 2003.
- [16] Jane Dashevsky, Edward C. Epp, Jose Puthenkulam, and Mrudula Yelamanchi. SIM trust parameters. *Intel Developer Update Magazine*, 2003.