




Keith Mayes
Konstantinos Markantonakis
Editors

Smart Cards, Tokens, Security and Applications

 Springer

Keith E. Mayes
Information Security Group
Smart Card Centre
Royal Holloway
University of London
Egham Surrey
TW20 0EX
United Kingdom

Konstantinos Markantonakis
Information Security Group
Smart Card Centre
Royal Holloway
University of London
Egham Surrey
TW20 0EX
United Kingdom

Library of Congress Control Number: 2007939275

Smart Cards, Tokens, Security and Applications
by Keith E. Mayes and Konstantinos Markantonakis

ISBN-13: 978-0-387-72197-2
e-ISBN-13: 978-0-387-72198-9

Printed on acid-free paper.

© 2008 Springer Science+Business Media, LLC.

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, LLC, 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

9 8 7 6 5 4 3 2 1

springer.com

Chapter 7

Introduction to the TPM

Allan Tomlinson

Abstract The Trusted Platform Module (TPM) and smart card devices have many features in common. Both are low cost, tamper resistant, small footprint devices used to provide the basis of a secure computing environment. This chapter presents an introduction to the security mechanisms provided by the TPM highlighting those not typically found on a smart card. The concept of “ownership” is one of the major differences between the TPM and a smart card and this is described in some detail before concluding with a review of some of the security services uniquely provided by the TPM.

Key words: Smart Cards, Trusted Computing, TPM, Security

7.1 Introduction

Smart cards provide a wide range of functionality from simple storage media to complex processors. A common design goal across this diversity however, is the provision of some degree of secure processing, implemented in secure hardware. The Trusted Platform Module, or TPM, is similar to a smart card device in that it is a small footprint low cost security module typically implemented as a tamper resistant integrated circuit (IC). The TPM however, has been specifically designed to be a building block for trusted computing. So although there are many similarities, the different design goals have resulted in a number of differences between the two types of device. One major difference is that the TPM is considered to be a fixed token bound to a specific platform, whereas a smart card is a portable token traditionally associated with a specific user across multiple systems. This is not to

Allan Tomlinson,
Information Security Group,
Royal Holloway, University of London,
e-mail: Allan.Tomlinson@rhul.ac.uk

say that the two technologies are mutually exclusive, but rather that they may be considered as complementary.

The details of smart card architectures and applications are discussed elsewhere in this book so this chapter will focus on the TPM and expand on the “complementary” aspects of this device, describing the specific mechanisms where the TPM differs from a typical smart card IC. In order to fully understand the rationale behind these mechanisms it is important to understand the underlying design goals of trusted computing that guided the development of the TPM.

This chapter therefore begins with some background on trusted computing to explain the features that the TPM is designed to support and the rationale behind the design decisions that were made. The fundamental features of the TPM that emerged from these design constraints are described in section 7.3. This section focuses on the functionality of the basic TPM building blocks to expose some of the differences between the design of smart card ICs and the TPM. Many of the differences that appear, arise from the contrasting requirements for ownership and management between the two types of device. Section 7.3 therefore looks at these concepts in some detail. The differences between the two types of device result in differences in functionality, and section 7.4 looks at some of the security services not found in typical smart card devices, which are fundamental to the operation of the TPM.

7.2 Trusted Platforms

Where smart cards may be considered as general purpose security processors, the TPM has been designed specifically to support trusted computing platforms. Therefore, in order to understand the TPM design requirements, it is first necessary to understand what the desirable features of a trusted platform are. To do this, a definition is required as to exactly what is meant by the term “trusted platform”.

The concepts of trusted computing, and a Trusted Computing Base, or TCB, are not new and are described in many publications ranging from the Orange Book [3] through to more recent material that describe these ideas within the context of contemporary developments in computer security [2, 5, 15]. One such development is the emergence of the Trusted Computing Group¹ (TCG) that has attempted to define what is meant by a trusted computing platform and that has produced a series of standards that can be used to design trusted platforms. The TCG defines trust to be “the expectation that a device will behave in a particular manner for a specific purpose” [16]. The TCG’s definition of a trusted platform therefore, is a platform that behaves in such a manner. This is, necessarily, a rather high level definition: but what this means is that any entity that interacts with such a trusted platform can be given some degree of assurance that the platform (system or computer) will behave in the way that entity expects it to. Providing assurance of expected behaviour does not in

¹ <https://www.trustedcomputinggroup.org>

itself provide any security. To achieve that, the entity relying on this assurance still has to ascertain that the “expected behaviour” is indeed secure. However, assuming that the relying entity is satisfied that the expected behaviour is secure then he may, for example, be assured that any data given to the system is kept confidential, or that no malware is running on the platform.

There are many ways to design trusted computing platforms that allow statements about expected behaviour to be made [15], but the approach taken by the TCG is to have some physically secure trusted component that can be used as a foundation upon which trust in the rest of the system can be built. The purpose of the TPM is to provide this foundation. For this chapter, a more appropriate definition of a trusted platform is provided by Pearson [13] who states that:

“A Trusted Platform is a computing platform that has a trusted component, probably in the form of built-in hardware, which it uses to create a foundation of trust for software processes”

or Balacheff et al. [1] who say:

“A trusted platform (TP) is defined as a computing platform that has a trusted component, which is used to create a foundation of trust for software processes.”

It is perhaps appropriate at this point to make a subtle distinction between what is meant by a *trusted component*, such as the TPM, and a *trustworthy component*. One generally accepted definition of these terms is given by Anderson [9, 10] who states that:

“The proper definition is that a *trusted* system or component is one whose failure can break the security policy, while a *trustworthy* system or component is one that won’t fail”

By implementing this trusted component, the TPM, as a tamper proof IC; and binding it to the platform, usually on a printed circuit board containing a more powerful processor capable of running software applications; the TPM can be used as the foundation of trust for higher level processes that run on the main processor.

7.2.1 Fundamental Features of a Trusted Platform

In order to establish this foundation of trust, the TPM is expected to provide a fundamental set of security features which have been defined by the TCG. The minimum set of features that a trusted platform should have are: protected capabilities; integrity measurement; and integrity reporting [16]. Providing support for these features leads to the definition of the security requirements of the TPM.

Protected Capabilities

To meet the requirements of a trusted platform, according to the TCG [16], the system should provide some form of protected capabilities. In the TPM design prin-

metrics can be signed by the trusted platform and that the challenger has a certificate that can be used to verify the signature.

7.2.2 Additional Features

While the above features would allow the implementation of a basic trusted platform as defined by the TCG, there are other features that would be required to create a more flexible implementation of a trusted platform. These include: confidentiality and integrity protection; secure storage; and process isolation.

Confidentiality and Integrity Protection

In addition to the protection of integrity metrics, a trusted platform should provide both confidentiality and integrity protection to any data as required by the user. This could include both user data, and application code. Moreover, these security services should be available to protect this data while it is being stored, and during the execution of any process.

Secure Storage

The provision of confidentiality requires that a trusted platform is able to encrypt any user data for secure storage. Access to the data is controlled by a securely stored cryptographic key. Similar to encryption is the concept of sealing. In this case access to the data is controlled by platform state, the use of a cryptographic key is optional. This means that data can be sealed to a set of integrity metrics that reflect the platform state. The data will then only be accessible when the platform is in the desired configuration. The sealing mechanism can be used, for example, to ensure that no rogue applications are running before access is granted to sensitive data.

Process Isolation

The secure storage and integrity checking mechanisms will protect data during storage. To protect data during execution the provision of process isolation is necessary. The concept here is to provide an isolation kernel between the hardware and any operating system. This isolation kernel is used to create and manage multiple secure compartments. These compartments exist in parallel, on the same machine, and access is managed exclusively by the isolation kernel. Each compartment can then run its own operating system and applications in isolation from any other processes that are executing in parallel. In this way, application code and data can be protected during execution. Furthermore, the use of an isolation kernel greatly sim-