

# Combining User and Platform Trust Properties to Enhance VPN Client Authentication

**Patrick George, Gérald Maunier**  
**Gemplus Card International**  
**Avenue des Jujubiers**  
**La Ciotat, F-13705, France**

## **Abstract**

*With PC manufacturers aggressively pushing trusted architectures in their new models, Trusted Platforms are quickly becoming a major component of the IT landscape. These platforms embed a security chip, the Trusted Platform Module (TPM), that is primarily used to attest the integrity of the system but that can also accurately identify the platform. While platform identification raises privacy issue in the consumer space, it represents a major requirement for corporations.*

*Platform identification is particularly important in the case of remote access to corporate resources. Today, Virtual Private Network (VPN) client authentication mostly focuses on end-user identity without addressing the trust properties of the platform the end-user is operating.*

*Starting with the benefits that TPM-based platform identification and authentication can bring to VPN client authentication this paper demonstrates how platform and end-user trust properties can be combined in a standard VPN authentication framework, such as the Extensible Authentication Protocol, and discusses possible implementations.*

## **Keywords**

Trusted Computing, Smart Cards, VPN, Authentication

## **1. Introduction**

Trusted Computing is the answer brought by the IT industry to the growing concern, expressed by users and the media, over on-line security issues. This initiative has been promoted by most of the major players of the IT industry through the Trusted Computing Group (TCG) and technology development such as Intel's LaGrande or Microsoft's Next-Generation Secure Computing Base (NGSCB). With PC manufacturers aggressively pushing these architectures in their new models, Trusted Platforms

are quickly becoming a major component of the IT landscape.

Platform identification is a key function of the TCG architecture. Of course, a unique platform identity has brought back the ghost of the Pentium III identification number and has raised flags with privacy advocates. As a result platform identification features have often been downplayed, even within TCG.

Even though platform identification may be a very sensitive issue and a concern for the consumer market, this is definitely less important in the corporate environment. Corporate IT departments legitimately require the ability to track the platforms they issue to corporate employees. For liability and support reasons, they also must be able to make sure that only authorized software run on corporate platforms especially when critical tasks are performed.

One of corporate IT most critical security concerns is certainly remote connection to the corporate network. Today, Virtual Private Networks (VPN) products offer excellent solutions for confidentiality and user authentication. The authentication frameworks they are based on, such as the Extensible Authentication Protocol (EAP), can efficiently communicate end-user trust properties to the back-end servers enforcing access control policies.

While existing VPN solutions covers the end-user side of trust, new risks are emerging on the platform side. Virus, Trojan Horses and other spywares can be responsible for security breaches or sensitive information leakage when corporate data are remotely accessed, even when the platform is operated by a trusted end-user. Remote communication will more and more rely on platform trustworthiness, i.e. on a dependable reporting of the platform trust properties (e.g. this is a genuine corporate platform, it has booted correctly or it is in a "clean" state). The establishment of a reliable platform identity is a prerequisite for the assessment of platform trustworthiness.

This paper introduces VPN client authentication and identifies threats linked to the lack of platform authentication in existing VPN solutions. The mechanisms

defined by TCG to authenticate and identify platforms are briefly presented. Then we propose an EAP method that combines platform with user authentication and that applies to the TPM1.1b-compliant platforms already deployed. The security of this combined authentication scheme relies on a personal cryptographic token, such as a smart card, to perform the user authentication and authorize the TPM operations.

## 2. VPN Client Authentication

A Virtual Private Network is a network that is constructed by using public wires to connect nodes, set up solely for the users of a single company. These networks use encryption and other security mechanisms to ensure that only authorized users can access the network and that data cannot be intercepted. Therefore, it is generally understood that a VPN solution provide, as a minimum, client user authentication and data encryption.

In VPN protocols such as Point-to-Point Protocol (PPP), the client computer sends the user credentials to a remote access server. A secure user authentication scheme provides protection against replay attacks and remote client impersonation. Most PPP implementations provide limited authentication methods. They generally support:

- Password Authentication Protocol (PAP), a simple clear-text authentication scheme,
- Challenge Handshake Authentication Protocol (CHAP), an encrypted authentication protocol that avoids transmission of the actual password on the connection.

Beyond these basic client authentication schemes the Extensible Authentication Protocol (EAP) [1] allows for an arbitrary authentication method. Sophisticated authentication schemes, such as public key infrastructure or token-based methods, are usually built on this framework.

Other VPN solutions are based on IPsec. IPsec is a framework for security that operates at the network layer by extending the IP packet header. The Authentication Header protocol provides connectionless integrity and data origin authentication.

VPN products primarily tackle client user authentication issues, but new requirements for client device authentication are emerging. EAP versatility is key to enable new authentication methods that take into account not only client user, but also client device authentication.

### 2.1 EAP

EAP is an IETF standard extension to PPP that allows for arbitrary authentication mechanisms for the validation of a PPP connection. EAP was designed to provide the dynamic addition of authentication plug-in modules at both the

client and server ends of a connection. Vendors can supply new authentication methods at any time.

EAP allows for an open-ended exchange between the remote access client, also called supplicant, and the authenticator. The exchange consists of authenticator requests for authentication information and responses by the supplicant. The exact authentication method to be used is negotiated by the supplicant and the authenticator. For example, EAP can support authentication methods based on public key infrastructure (EAP-TLS), GSM (EAP-SIM) or one-time passwords (EAP-OTP).

EAP is often associated to Remote Authentication Dial-In User Service (RADIUS), a security service for authenticating and authorizing users. EAP-RADIUS consists in the passing of EAP messages, of any EAP type, by an authenticator to a RADIUS server for authentication.

EAP is now pushed as the standard legacy authentication protocol in IETF. The most significant new use of EAP is with IEEE 802 wireless networks, but it can also be supported by IPsec. The Pre-IKE Credential provisioning protocol [2], a method to bootstrap IPsec authentication protocol and user authentication mechanisms, allows for the transport of EAP payloads.

### 2.2 VPN devices authentication

VPN security mostly focuses on end-user authentication. While the existing techniques, especially cryptographic tokens and biometrics, provide effective identification and authentication of an individual user, VPN vendors are realizing that a comprehensive security policy must also identify the devices operated by authenticated users.

In this context, device authentication is a second level of authentication that ensures that only authorized platforms can access the network. It assures that unauthorized devices are not allowed onto a network, even when operated by an authorized user. The foundation of device authentication is, of course, a reliable identification of the platform.

Intel tried to tackle device identification in 1999 and introduced a processor serial number (PSN) in the Pentium 3 microprocessor family. A PSN was a software-readable unique serial number stamped into the microprocessor. But Intel's PSN caused concerns to privacy advocates [3], on grounds that it may undercut individual user efforts to maintain their anonymity, and Intel had to drop this feature in further products.

Today, there are many techniques to identify connecting devices. The most basic ones are based on unique platform identifiers that are sent to the authentication server while the connection is established. Information such as computer serial number, asset tag number, computer manufacturer or computer model, are generally used as unique identifiers. A variant consists in using networking information like the

device Media Access Control (MAC) addresses. Each Network Interface Card (NIC) has a unique MAC address that is used as the identity of the device the NIC is attached to. To authenticate a connecting device, the authentication server simply checks the platform identifier against a list of authorized identifiers.

More robust solutions are based on cryptographic mechanisms. For example, the device is initialized with pre-shared keys, which can be tied to a specific IP address, used by a cryptographic authentication protocol between the connecting device and the authentication server. Other alternatives rely on device certificates and public key cryptography. A device certificate is a public key certificate, issued by a device manufacturer, tying the identity of the device to the corresponding private key. In this model a private key is inserted into the device at manufacturing time. To authenticate a connecting device, the authentication server engages in a cryptographic authentication protocol with the client device.

## 2.3 Threats and requirements

User authentication only is clearly not enough. Security policies should not just keep unauthorized individuals from gaining improper network access, but they should also prevent non-conforming devices from tampering with other resources. It is essential to get information about the identity and trust properties of every device accessing to the network. Neither user nor device should be granted network access separately, and both elements must be authenticated.

The most understandable set of threats concerns compromised user credentials, such as guessed passwords or cracked key containers. Once a credential is stolen, it can be used anywhere, on any device, to get improper access. But the corporate resources are still protected as long as an authorized machine is not used. Platform authentication mitigates the risk by forcing the intruder to operate from a genuine corporate platform, which might be harder to get access to.

There are other circumstances where an authorized user, with genuine credentials can put the security at risk because he, or she, is attempting to connect from an unauthorized platform. It is rather easy to transfer user credentials (password or certificate) from an IT department managed platform to, for example, a home PC shared by other members of the employee household, or to an Internet Café PC that can be accessed by anyone. These platforms are, by definition, outside of the control of the corporate IT department and may be loaded with viruses or spywares.

Countermeasures are hard to build especially because existing device authentication mechanisms are unreliable. There is no easy, built-in way of identifying what device is actually connecting to a network.

Identification based on inventory information, device serial or asset tag number for example, cannot be trusted because these identifiers can easily be captured and spoofed. An identification based on physical components, such as the MAC address, is not reliable either. NIC can be reinstalled on other platforms and MAC addresses can be modified and spoofed without difficulty.

While device certificates offer a more robust alternative, they require the device to be tamper-resistant to be really suitable. Otherwise the private keys can easily be extracted and reused in a different environment.

Another class of threats concerns authorized devices operated by authorized user, or not, but running an unauthorized environment. These untrusted environments include systems not running the necessary security software, like anti-virus, and operating system booted from an external drive or a live CD. When connecting, devices should be checked for configuration, patches and security software updates, and access should be denied if they do not conform to the device security policies issued by the corporate IT.

Therefore the ideal device identification mechanism should have the following characteristics:

1. Non removable, to prevent credentials from being reinstalled on unauthorized platforms. A hardware-based mechanism is highly recommended to offer tamper-proof storage.
2. Capable of cryptographic processing to engage in authentication protocol with a network access server, to avoid replay attacks and spoofing.
3. Able to report its software environment integrity and securely convey trust properties to the authentication server.

## 3. Trusted Platforms Authentication

While TCG specification enables accurate platform identification, their members have been very careful with privacy issues. They have designed platform identity features in such a way that “analysis of the aggregated activity could [not] reveal personal information that a user of a platform would not otherwise approve for distribution”.

In the following sections, a basic knowledge of TPM and TCG Software Stack (TSS) architectures is assumed. Readers are referred to [4] [5] for more information.

### 3.1 Trusted platforms identities

Each Trusted Platform Module (TPM), the hardware security module defined by TCG, holds a unique RSA key pair called the Endorsement Key (EK). The EK is either generated by the TPM itself or inserted by the TPM manufacturer, and is certified by a trusted third party, such

as the TPM manufacturer. The TPM is physically attached to a particular platform and the public part of the platform's EK (PUBEK) can then be considered as a unique identifier of the platform.

However this globally unique identifier is never divulged directly and is not even supposed to be used in communication with third parties. Instead, TCG specifies that the TPM can hold an arbitrary number of uncorrelated identities. Each TPM identity is associated with a RSA key pair, called Attestation Identities Key (AIK) that is certified by a trusted third party, the so-called privacy certification authority (Privacy CA), and that can be used to identify the trusted platform in different contexts. AIK are bound to the TPM and cannot be migrated to another platform. The process of obtaining a TPM identity has three main phases:

- Create the AIK,
- Build and send the AIK certificate request to the Privacy CA,
- Activate the AIK, and retrieve the AIK certificate.

The AIK certificate will bind the AIK to an arbitrary identity chosen for the new TPM identity. Once activated, the AIK can be used to sign platform attestations.

It is important to notice that the endorsement credential, a globally unique platform identifier, is sent to the Privacy CA with the AIK certification request. This means that the Privacy CA is the only entity able to associate the different arbitrary platform identities to the platform unique identity. This represents a potential privacy breach [6].

### 3.2 Trusted platforms integrity measures

Trusted platform must contain a core of inherently trustworthy services. TCG defines a Root of Trust for Measurement (RTM) that:

- Accurately measures at least one integrity metric indicating the software environment of the platform,
- Accurately records measured integrity metrics to the TPM.

The RTM program must be the first program to execute on the platform, otherwise the programs that execute before the RTM must be trustworthy. In PC, the RTM can be implemented as BIOS instructions that cause the main processor to do RTM work. This set of RTM instruction is generally called the Core Root of Trust for Measurement (CRTM).

In most implementations, the CRTM doesn't perform the full platform software environment measurement. The CRTM does a few simple operations, measures the next software to be executed in the boot process, stores the result in the TPM and passes control to that next software component, and the process goes on until the platform has completely booted.

The measurement itself is achieved by computing SHA-1 hashes of the software components. The resulting hash values are stored inside the TPM Platform Configuration Registers (PCR). These PCR values signed by a TPM identity (i.e. an AIK) will allow a challenging entity to verify whether the platform has booted correctly and if the expected software components are running.

### 3.3 Trusted platforms attestation

Attestation is the mechanism defined by TCG to allow a challenging entity, such as a remote server, to verify the platform integrity measures stored in the TPM PCRs and to fulfill the platform authentication requirement.

Remote attestation is based on the *quote* operation. A quote is a cryptographic reporting of PCR values. To compute a quote, the TPM signs a statement, which names the current value of chosen PCRs, and externally supplied data with a signing key. Platform identity is also attested when the signing key is a TPM identity (i.e. an AIK). The quote is returned to the remote challenging entity together with a locally aggregated measurement log.

The remote challenging entity can validate the quote by:

- Verifying the quote digital signature with the AIK certificate. This proves that the quote has been computed by a genuine trusted platform and which authenticates that platform,
- Comparing the integrity values to the locally aggregated reference measurement log. This proves that the platform has not been compromised.

Compromised systems can modify the measurement log, but they cannot alter the TPM-protected measures to fit the modified log. Fraudulent manipulations are thus recognized when validating the measurement log against the signed measurements.

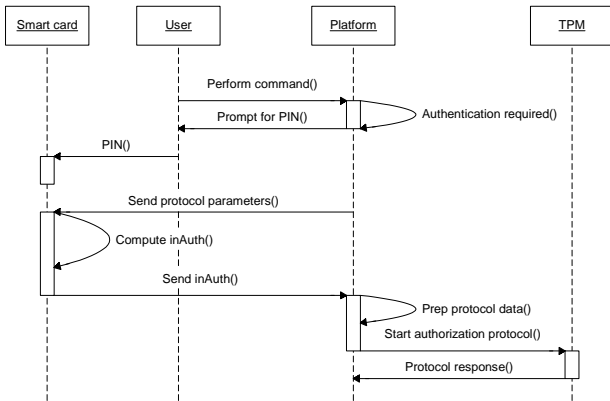
### 3.4 User authentication in TCG

Access to TPM-protected objects is a major issue. Like every sensitive data, access to the AIK is controlled. The TPM user must be authorized before performing any action requiring the use of a specific AIK. While TCG specifications are ambiguous with regards to the concept of TPM user, authorization protocols are clearly defined [4]. They are based on the proof of knowledge of a shared secret: the *Authorization Data* (AD). Authorization data is a 20-byte value -the same size as the output of a SHA-1 cryptographic hash- and is shared between the TPM and authorized users. Knowledge of an object AD is a complete proof of ownership of this TPM-protected object.

The most critical piece of information in TCG authorization protocol is the AD, which is either stored locally on the platform or computed from an external seed secret such as a password. This model raises many

concerns [7]. Since smart cards and other hardware tokens, are widely used to address this kind of user authentication issues, smart card-based authentication can be the answer to the threats linked to the TCG authorization model based on shared secrets. In such a scheme the AD is never exposed and never gets out of the smart card; the smart card computes the parts of the protocol requiring the AD.

**Figure 1 Authorization workflow**



As described in Figure 1, the platform delegates to the smart card the computation of the input authentication value (inAuth). This value is then returned to the platform to perform the authorization command. This solution is highly secure because it brings two-factor authentication to TCG user authentication but, more important, the AD is never exposed to the external world during the processing of the authorization protocol. This approach requires a smart card with computation capabilities. In addition, the verification of the result authentication value, computed by the TPM, can also be performed by the smart card.

## 4. Proposed Solution

Attestation does not only prove that a TPM is genuine, but it also indirectly provides TPM identity information. This mechanism fulfills most of the platform authentication requirements identified in section 2.3.

This section describes a first approach to the combination of platform and user authentication in a VPN client authentication protocol. This solution is a simple authentication protocol that mixes a regular (and rather weak) username and password user authentication with a TPM attestation-based device authentication into an EAP supported method.


EAP plug-ins can also be used in 802.1x with minor efforts. The solution presented in this section also applies to laptop access to corporate WiFi access points.

Our goal was to design a solution that could be easily implemented using existing products, like the Microsoft EAP API, and leverage from the already deployed base of TCG 1.1b-compliant trusted platforms.

It is also important to mention that the choice of username and password user authentication was motivated by didactic reasons. This method requires a user action that emphasizes the user role in the authentication process. Therefore, the smart card is used for platform-related operations only and not for user authentication.

The general architecture for a typical scenario is described in Figure 2. A corporate user, identified by a user name and authenticated by a password, connects to the corporate network using a trusted laptop. Ownership of the laptop's TPM has been taken by the corporate IT department and an AIK, dedicated to VPN authentication usage, has been generated and certified by the IT department Privacy CA.

**Figure 2 General architecture**



In order to protect from unauthorized usage of the AIK, the computation of platform attestation is protected by the user's smart card as described in section 3.4.

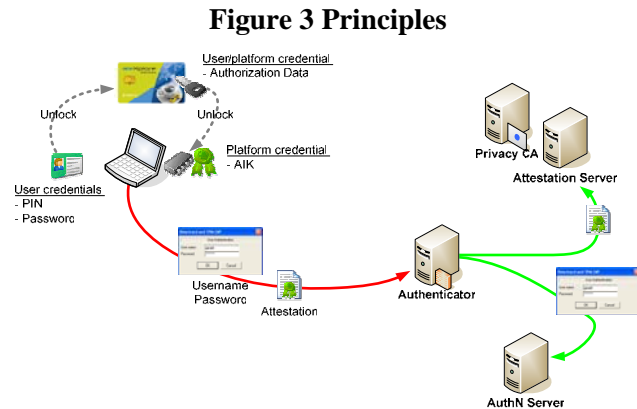
When the user is connecting to the corporate network, a VPN tunnel is established between the laptop and the network access server that controls the access to the corporate LAN. VPN client authentication is performed by the combined authentication protocol within the EAP framework.

While user authentication relies on existing authentication infrastructure (RADIUS server) the attestation part requires two new specific items of infrastructures:

- A Privacy CA that certifies the AIK when the platform's TPM is configured by the corporate IT department, but also acts as platform credential repository,
- An attestation server that performs the attestation verification on behalf of the authenticator.

## 4.1 Client authentication

Client authentication principles are introduced in Figure 3.

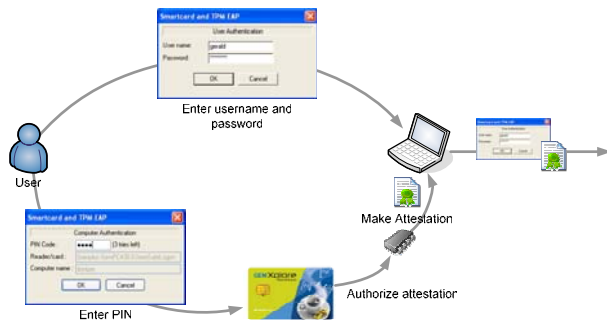


In our proposed scheme, the user holds two types of credentials:

- A password, associated to the username, used by the RADIUS authentication server,
- A PIN, associated to the smart card, used to unlock the AIK authorization data.

The platform holds only one credential: a VPN authentication dedicated AIK, used to sign platform attestations.

**Figure 4 Components interaction**



Client authentication is a 3-step process described in Figure 4. The sequence is the following one:

- The user authenticates with their preferred authentication method. In our scenario this is done with a standard login/password mechanism,
- The user enters a PIN to authenticate to the smart card. The smart card is then used to prove that the user is authorized to the platform and provides explicit proof of consent to the platform authentication,
- The platform's TPM generates an attestation that "quotes" a pre-defined set of PCRs.

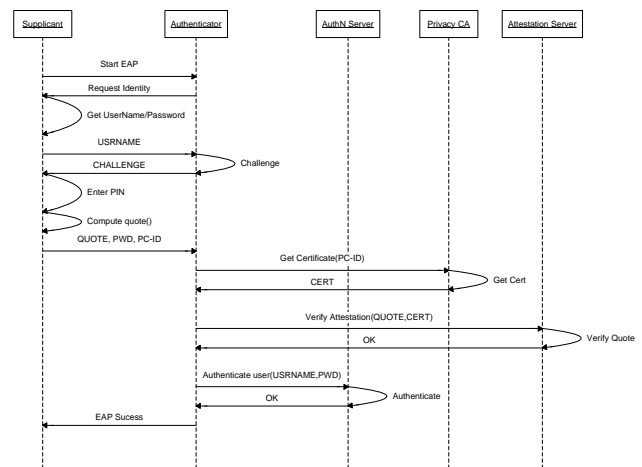
The authenticator forwards user name and password to a RADIUS authentication server, while the attestation is forwarded to the attestation server. If both servers return positive authentication status, the authenticator opens the VPN channel.

## 4.2 EAP integration

EAP message flow of client authentication process is described in Figure 5.

The supplicant provides user's identity to the authenticator, which sends back an anti-replay challenge. The encrypted user's password and a platform attestation - both diversified by the challenge - are sent to the authenticator, together with the corporate platform identity.

**Figure 5 Client Authentication protocol**



The authenticator asks the Privacy CA for AIK certificate associated to the corporate platform identity, and then requests the attestation server to verify the platform attestation to authenticate the supplicant. If the platform is successfully authenticated, user authentication starts, e.g. with a RADIUS server.

## 5. Conclusion and Future Work

This paper establishes the requirements for a combined authentication of both users and platforms applied to the context of VPN connection to a corporate network. It introduces the sensitive issue of platform identification and describes the mechanisms proposed by TCG.

A first practical solution to the combination of platform and end-user trust properties is then demonstrated with EAP, a standard VPN authentication framework. Further research may open the discussion to improvements in order to provide a robust combined authentication protocol.

The solution presented in this paper has been implemented and is fully functional. However some limitations and inconsistencies have been identified and this first proposal is more to be seen as a proof-of-concept for combined authentication.

Two main directions for improvement have been identified:

- User authentication is probably the weakest point of the current solution. Password-based authentication weaknesses are well known and leveraging from the user smart card cryptographic capabilities would be a major improvement.
- TCG specifications are getting more mature and key aspects that were not addressed initially are now defined and stable. At the same time recommendations, made by TCG working groups, apply to TPM attestations and should be taken into account, as well as related research work [8].

All these improvements would lead to the definition of a more robust protocol.

## 5.1 Improve user authentication

The password-based user authentication was chosen for didactic reason, to differentiate between user and device authentication mechanisms. But it is clear that, in a real world environment, a password solution is not robust enough to support a trustable user authentication mechanism.

The most natural solution to improve user authentication would be to leverage on the smart card security. Our solution uses the smart card to protect the AIK authorization data and performs the client part of the authorization protocol. But the smart card could also be used to authenticate the user. For such an approach, strong user authentication, based on digital signature, is probably the most suitable solution. The smart card performs the user authentication by signing an authentication challenge sent by the authenticator. Impact at the authentication server is minimal as RADIUS already supports cryptographic challenge response authentication protocols.

The current model binds the user to the platform. This is achieved via the storage of the AIK AD in the user smart card. This method could be used to implicitly authenticate the user. The fact that the attestation has been computed is indirect proof of user identity. Only the user with the smart card and the PIN is able to compute the attestation. This approach would simplify the protocol and allow the attestation to carry two identities at the same time.

## 5.2 Take TCG evolutions into account

A new version of the TCG specification has been released, which addresses new requirements, brings new mechanisms and provides details about key points.

For example the details of attestation and credential management are being addressed by TCG Infrastructure Working Group. We anticipate some assumptions made in our current proposal, such as using the Privacy CA as an AIK certificate repository, may become irrelevant.

TCG also introduces a new protocol called Direct Anonymous Attestation (DAA), based on a zero-knowledge group-signature scheme. In DAA the AIK certificate is not required but instead a cryptographic proof that the platform has one [9]. This scheme offers a very high level of privacy. Further research would consider the relevance of such a scheme to the corporate IT environment, and then study the impact on implementation.

## Acknowledgements

The authors would like to thank Philippe Leblanc for his contribution to this work, and James Langley for his helpful remarks.

## References

- [1] L. Blunk and J. Vollbrecht, RFC2284, “*PPP Extensible Authentication Protocol (EAP)*”, IETF, March 1998.
- [2] Y. Sheffer, H. Krawczyk, B. Aboba, “*A Pre-IKE Credential Provisioning Protocol*”, IETF ispra Working Group draft, October 2002
- [3] J. Weinberg, “*Hardware Based ID, Rights Management, and Trusted Systems*”, 52 STAN. L. REV. 1251, 1274 (2000)
- [4] Trusted Computing Group, “*TPM Main Specifications – Part 1 Design Principles*”, Version 1.2, October 2003.
- [5] S. Pearson et al, “*Trusted Computing Platforms – TCPA Technology in Context*”, Hewlett Packard Company, Prentice Hall PTR, 2003
- [6] J. Camenisch, “*Better Privacy for Trusted Computing Platforms*”, Proceedings of 9th European Symposium on Research in Computer Security (ESORICS 2004), Springer-Verlag, 2004
- [7] P. George, “*User Authentication with Smart Cards in Trusted Computing Architectures*”, Proceedings of the International Conference on Security and Management (SAM'04), June 21-24, 2004, Las Vegas, Nevada, USA
- [8] R. Sailer, T. Jaaeger, X. Zhang, and L. van Doorn, “*Attestation-based Policy Enforcement for Remote Access*”, in Proceedings of the 11<sup>th</sup> ACM Conference on Computer and Communications Security (CCS), pages 308–317, October 2004
- [9] J. Camenisch, “*Direct Anonymous Attestation: Achieving Privacy in Remote Authentication*”, June 2004