Reliable Identities using off-the-shelf hardware security in MANETs

Nicolai Kuntze, Andreas Fuchs, Carsten Rudolph
Fraunhofer Institute for Secure Information Technology (SIT)
Rheinstrasse 75, 64295 Darmstadt, Germany
{*nicolai.kuntze*|*carsten.rudolph*|*andreas.fuchs*}@*sit.fraunhofer.de*

Abstract

Application scenarios for mobile ad-hoc networks (MANETs) impose a variety of non-standard security requirements. furthermore, in many scenarios owner and user of devices do not always have physical control over the device. Therefore, security in MANETs should be rooted in hardware security anchors. For current PC architectures a relatively cheap hardware anchor is readily available, the so-called trusted platform module TPM as standardized by the Trusted computing Group. This paper shows that TPMs can provide the basis for rather complex security mechanisms that can support a variety of security properties in MANETs. In addition to straightforward requirements like authenticity or confidential storage of data on the device, also more complicated requirements like unlinkability of multiple identities or restrictions to the validity of identity certificates are discussed.

1. Introduction

Mobility and high-speed access to voice and data links are highly desirable in various situations concerned with the public safety and highly dynamic environments. Public and military organizations depend on secure and reliable communication systems in order to strategize, command, control, and operate their resources in their respective environments. But also in public applications like IPTV as it is developed in the Nanodatacenters Project. Authentication can be considered as the keystone of network security because it is the first step toward prevention of unauthorized access to network resources and sensitive information.

MANETs are considered one promising approach to face the challenges of such environments due to their ability to provide ad hoc connectivity to and between mobile devices. MANETs are defined as nodes with an included routing capability forming autonomous, self-forming, and self-maintaining networks that are capable of adapting to changes and merging dynamically into hierarchies [4].

Most of the research on security for MANETs has been on the routing algorithms and takes the authenticity of the messages as an assumption [13]. Also the trustworthiness of the devices is often considered as given as the user is considered as friendly. These underlying assumptions are not necessarily true for different kinds of scenarios. For example the the military is interested to provide for its forces communication systems able to allow for secure communication in ad hoc scenarios [11]. There, each participant resp. node may take over different roles in the network like relaying messages or deciding on changes in the routing tables due to territorial changes and the impact on the physical network. Similar scenarios are also used by aid organisations or the civil defence. Each node in these networks has therefore access to different sensitive information on the participants of the network and may also be able to inject data relevant to the availability of the offered service. This is due to the fact that the underlying routing protocols are created in a collaborative approach.

This paper presents security building blocks for security in MANETS using inexpensive off-the-shell components, namely Trusted Platform Modules (TPM). Among others, the paper shows how TPMs can be used to authenticate nodes in a MANET or manage multiple identities and roles within. Part of the presented concepts is to provide proof on the authentication of the user but also to reveal the identity of him. A second aspect is the lean implementation of the key and identity management allowing for complete offline scenarios without any central entity available.

The rest of the paper is organised as follows. In section 2 a high level security analysis on the scenario is presented. Section 3 introduces the basics on Trusted Computing (TC) with the focus on the previously defined requirements. Section 4 reveals our concepts for secure identities and is followed by section 5 presenting a short conclusion.

2. Problem Statement

Devices in MANETS have to face a large variety of threats. In the past, the majority of research was concentrated on the communication, i.e. on the efficient transportation of messages. One of the most challenging issues is secure routing. The prevention of network layer attacks has to deal with attacks like wormholes, blackholes, or flooding. Further, monitoring and traffic analysis can reveal information on the devices present in a MANET and on the structure of the MANET and roles of the devices. Proposed solutions are mostly based on changing the underlying protocols. In contrast to these approaches, we concentrate on the security of the devices themselves and on secure roots of trust for the construction of secure MANET protocols. This section summarizes the security requirements for MANET devices considered in this paper. These requirements are discussed on an architectural level. Perils resulting from e.g. faults in the software or hardware implementation are not covered. The underlying scenarios do not assume that the owner of the node has direct physical control during the whole life-time of the device. Thus, we have to keep in mind that beside software attacks using the regular attack vectors like buffer overflows and attacks on the communication, the node is also threatened by direct physical attacks e.g. excerping information directly from the memory.

Some requirements are concerned with the identity of a device and the user operating the device. The MANET should not be open for foreign devices and also stolen devices should not be able to join the network. External attackers must not be able to impersonate devices in the MANET. Therefore, the identity of the device needs to be authentic when joining the network and when communicating with other nodes. Previous proposals included a securely stored private key as basis of authentication. However, for many scenarios in particular in the context of tactical MANETs a device can have multiple identities depending on the role or the group the device belongs to. It should not be possible to link different of these identities (in some cases even for other legitimate devices in the MANET). Furthermore, identities can have a limited validity. Possible parameters for this limitation include time or the location of the device. Note that revocation based on lists as in public key infrastructures (PKI) is difficult or impossible in MANET scenarios.

For the protection of integrity and correctness of the data produced and communicated by a device in the it is not sufficient to protect the communication link. Manipulations on the device (e.g. software changes) can also be used to introduce malicious data or to change the behaviour of the device. Further, data stored on the device needs to be protected. Thus, one essential requirement is the protection of the integrity of the device. In practice it is difficult and expensive to build tamper-proof devices. However, in many scenarios it is sufficient to have tamper-evidence, i.e. communication partners can tell, whether a device is still in its original state or has been changed.

Security solutions in MANETs must also satisfy a number of other (non-security) requirements. Among others, solutions must be resource-efficient, not require expensive hardware, reliable (e.g. no smart card interfaces) and it must be possible to run with non-centralized management during run-time.

3. Trusted Computing

As shown in section 2 protection of the node's identity is an important security requirement. Furthermore, the lack of control on the physical access to the node induces strong requirements on the protection level. One possibility is to root security mechanisms in strong hardware security anchors. *TC* [12] offers such a hardware root of trust providing certain security functionalities. In this section these functionalities are introduced.

TC as defined by the Trusted Computing Group (TCG) are computer systems extended by additional components which shall bring trust to the computing environment. Trust means that components of the system always work as implemented. To achieve this goal, the TCG has published and is still working on specifications describing architectures, affecting system components at any level from hardware to the operating system.

Most important hereby is the specification of the Trusted Platform Module (TPM). This module is mostly realized as a hardware chip hard-wired to the computer platform. It implements basic cryptographic functionality like SHA-1 calculation, message digest creation, random number generation, creation of 2048 bit RSA key pairs, and a RSA engine for encryption and signing purposes. Realized as an independent hardware module, it can provide protected capabilities allowing to shield secret data efficiently. This implementation also allows for in depth testing and validation of the soft- and hardware. The TCG defines three different roots of trust. These are components on which the trust to the whole system is built on.

The Core Root of Trust for Measurement (CRTM) [7] is implemented e.g. as an extension of the BIOS. Its duty is to perform measurements of system components involved in the boot process. Measured components then can perform measurements of other components involved in the next stage of booting. Through this principle of transitive trust, trust in the correctness of the measurement values can be passed on to the OS and the software executed.

The second root of trust is the Root of Trust for Reporting (RTR). One of the aims of TC is to enable computer systems to proof to other platforms that it is in a trusted state. Therefore the results of measurements of system components have to be presented to the remote platform. To guarantee the genuineness of these data, they are signed. For this purpose every TPM contains a 2048 bit RSA key pair, the Endorsement Key (EK), which is generated before shipping. The EK, together with an EK Credential, represents the identity of the platform. Pseudonymous representatives of the EK, so called Attestation Identity Keys (AIK) are used for signatures, for example of measurement results used by the remote party to verify the correctness of the desired state (Remote Attestation or RA) [5].

The third root of trust is the Root of Trust for Storage (RTS) with the purpose to establish a secure storage for cryptographic keys and other sensitive data. The RTS is implemented by introducing the Storage Root Key (SRK), a second 2048 bit RSA key pair stored in the non-volatile memory of the TPM. The SRK never leaves the shielded location of the TPM. That allows building a hierarchy of keys, with the SRK as the root, in which direct successors are protected by encryption with the SRK. These keys on their part can protect any number of other keys. These keys allow to bind data to a device or even to a particular state of the device.

4. Security building blocks

In the literature various authentication mechanisms are proposed for MANETs. [1] presents different ways to classify the approaches, such as the node functions (homogeneous or heterogeneous) and the type of credentials used for authentication (identity-based or context-based). However, existing solutions (PKI, pre-shared keys, etc.) can only provide a subset of required properties. Many of the more complex requirements including anonymity or pseudonymity as discussed above are not supported in these approaches. Especially symmetric approaches like pre-shared keys are not scalable and it is hard to refresh or exchange these keys. They also need to be transferred securely to the device. Asymmetric cryptography based on a PKI requires a central CA to store the certificates and to maintain a Certificate Revocation List. In an identity based encryption based system also a centralised server is required that may be a single point of failure. Delivering the private keys to the devices, key escrow, and key refreshing [2] are also challenging.

Software-based solutions are not sufficient because of the possibility of sophisticated physical attacks. Existing hardware solutions are either not flexible enough (e.g. by using protected storage of one private key within the device) or are based on smart-cards. It should be noted that smart-cards are not suitable for heavy duty environments. One reason here is the physical connection of the smart card that is not compliant with all required operational environments. Similar reliability problems exist for other (inexpensive) external interfaces. Some work already exists combining TC with P2P network technology.

TC as described above is available for most PC platforms and is inexpensive to be integrated in current systems. TPMs provide a large variety of security functionality that can be used to build solutions for many of the requirements for MANET devices. It even provides a lightweight public key infrastructure that can provide multiple pseudonymous identities for each device in combination with hardware protection for these keys. The main role of the TPM is usually seen as the provision of hardware roots of trust for measuring and reporting the state of a platform and for secure storage. In the following subsections we give an overview of other possible security mechanisms relevant for MANETs.

Actual work on combining TC and MANETs is on an early step of research. TC is used in [3] to establish a pseudonymous authentication scheme for peers and establishment of secure channels between them. Using an internal attestation of components [6] applies TC methods to protect ad-hoc routing. A trusted reference monitor is introduced by [8] to monitor and verify the integrity and properties of running applications and to enforce policies on them. The topic is addressed in [9] from a more economic point of view by looking on the media industry and possible distribution schemes.

4.1. Provision of Unclonable Identities

The most prominent requirement for nodes deployed in a tactical MANET scenario is the provision of authentic identity information. In this context, we consider an identity to be a certain user operating on a certain device. We assume that the user is authenticated towards the device. The mechanisms for user authentication are out of the scope of this paper. Nevertheless, below we will also consider scenarios where a user is forced to provide her authentication credentials e.g. her password, to an attacker. We assume, that a device can be used by different users, thus the device identity can be different for each user. In the following, we do not distinguish between user and device identity.

Many of the described requirements from Section 2 rely on authentic identity information. Furthermore, secure MANET protocols can also be build on authentic identities. First of all, it is necessary to prevent device impersonation in the scenario. This can be achieved by digitally signing messages with an asymmetric key that can only be used by a certain individual. This also requires, the existence of a proof, showing that the used key belongs to the designated individual, i.e. a certificate. Second, identity information is necessary, in order to secure the confidentiality of a communication. This can be done by encrypting data with an asymmetric key, such that only a certain individual is able to decypher the original message.

The techniques of TC implement several ways to achieve the above mentioned strong identification of a specific individual operating its personal device. In order to achieve this, each TPM is equipped with a key generator and non-volatile secure storage. When the Ownership of a TPM is taken, a unique key (EK) is created within the TPM and stored securely. This EK is used to securely identify a device against a Certificate Authority (CA) in the process of identities being created on a device.

When a new identity shall be created, the TPM generates a key pair (AIK) and an identity request. Such a key is usually secured by a password that will allow only the owner of the key to load it to the TPM and to used it. Further concepts of authentication are developed within the TCG. The identity request carries the public part of the EK and information about the newly generated key for review by a CA. The CA may then evaluate if the generating TPM is trustworthy and issue a certificate, containing the newly generated AIK and information about the owner. The certificate, together with a reference to the AIK is then encrypted with the EK of the platform. Finally, the owner of the platform can retrieve the certificate from the TPM, that provides identification of the AIK's owner.

The advantage of this process is, that the AIK that identifies a certain individual will be generated on the device it will be used on and may never be extracted from it. It can therefore be considered unclonable. Even the key owner cannot voluntarily publish the key or be forced to extract the AIK's private part. Additionally by binding the identity certificate to the issuing TPM, the CA can be assured that a valid TPM generated the AIK.

4.2. Association of Identities to Situations and Groups

The static identification of individuals does not always satisfy the desired properties in the context of tactical MANETs. With each change of context the role of an individual or the membership in a certain group or team may change. It may also be desirable to not disclose membership in one group to the participants of a different group, though they are considered trustworthy within the group's own context.

As introduced in the previous section each TPM provides a unique identifier and mechanisms to securely generate and certify pseudonymous identities. The advantage of this separation between EK and AIKs can be used to generate many AIKs on one device. This means that several identities can be used in parallel on one device and the identities can be changed depending on context regarding time, scenario, or location. This dissociation from the unique identifier provided with the EK in terms of AIKs allows the user to protect its identity by providing pseudonymous identities that cannot be linked. Only the CA that issued the certificate would be able to link these pseudonymous identities.

According to changing situations and group memberships the AIK certificate can be used to indicate the membership of this device and its user to a particular team and define its related role. It allows for multi homed devices participating in different groups at the same time. Also, identities of certain roles or groups can be time limited independently (mission AIK).

4.3. Anonymity towards Unauthorized Parties

The traceability of personal in the field may lead to unwanted side effects. It may be desirable that unauthorized parties shall (1) not be able to identify a certain individual among a set and (2) not be able to trace the path of pseudonymized individuals in the field. The sheer movement pattern may give out unwanted information like group movement or role and position among a single group.

Based on the notion of groups and teams from the previous section, it is possible to restrict the disclosure of the identity of an individual to other members of the same group only. For this use case, the techniques of certified migration from the TC specification can be utilized. When a new group is created, a new key pair, that is associated with this groups, is generated within a TPM. Unlike an AIK, this key is allowed to be transfered from one TPM to another. But this so called migration will be restricted. Only if a Migration Authority (MA) agrees to a specific transfer, the transfer can be performed. This will prevent users of the key and even owners of the platform to transfer the secret portion of the key to other platforms at their will. Each migration needs to be approved by a MA. Note that this approval can happen long before the actual migration takes place.

With such a key being deployed to all members of a group, participants will not start communication sessions in the way we are used to from e.g. SSL, by first disclosing their certificates in clear text. Instead, even the very first connection attempts will be encrypted using the group key. Certificates are not disclosed to non-members and because of the padding with random data in PKCS#11, an outsider is not able to link two connection attempts by the same individual.

4.4. Enforcing Temporal Restriction on Identities

It could be necessary that the device controls the usage of the identities instead of the individual end user. Current TPMs implement tick sessions to measure time. A tick session nonce identifies the current tick session that is started with the last power up of the TPM. One possibility to use this information is the inclusion of tick counter values and a tick session nonce into certificates [10]. The enforcement of the restriction is deferred to the communicating party by including the restriction to the AIK certificate issued.

Actual enforcement of time restrictions by the TPM requires extending TPM functionality. The TPM_Key data structure could be extended to optionally include tick session nonce and a max_ticks value for a key. These values can be checked during loading the key in the TPM or for each use of the key. To establish the restricted AIK inclusion of the tick session nonce and the validity time is required in the AIK creation process. This scheme could also be applied to limit the life time of data stored at the device if the data is encrypted using a key with this special limitation. It is to be noted that in case of a power failure of the TPM the tick session is lost and all data protected is rendered inaccessible. This can also be seen as a protection against physical attacks on the TPM.

4.5. Confidentiality of Stored Data

When data is stored on a device, it should not be disclosed to unauthorized users. The most common technique is to encrypt the stored data either as a whole (e.g. a full disk), or to encrypt certain data that is considered secret. Such protection mechanisms may secure data from offline attacks. If the key can be retrieved from the platform or a password protecting the key is weak encryption is rendered useless. TPMs provide a secured keystore and also provide protection against password dictionary attacks. Similar protection mechanisms are also implemented for smart-cards.

In a pure smartcard-based solution attacks from insiders or through blackmailing or forced interrogation could reveal credentials for access to the smart-card. These credentials then provide full access to encrypted data. Thus, data could either be copied directly to other devices or at least by booting a different OS.

The specification of the TPM provides the means to restrict the usage of a key to a certain platform configuration. Integrity measurement values of components during boot are recorded in platform configuration registers. Afterwards, when access to the TPM's key is requested, the current measurements representing the current platform configuration - including OS and applications - is tested against the conditions associated with the key. Only if the platform can be considered running trustworthy software the access to confidential data is granted. This software configuration could further enforce protection of the data, e.g. restrict access. This approach renders even attacks that involve the operator of a device useless in such a way, that the confidential information cannot be digitally copied from the device to the outside. Retrieval of information as would be possible during regular operation may still be possible and cannot be prevented.

These mechanisms can also be used for cases in which the OS itself or certain applications include data that is considered to be confidential. This may include protection of intellectual property rights theft as well as the disclosure of other relevant information. Export restrictions in many countries do not allow for the export of certain technologies. In some cases these restrictions are not applied on a physical embodiment but on algorithms or data used by the devices. By assuring that a certain information is only accessible if the device is in the desired state (again defined by the producer of the device) restricted technology can be adopted and introduced to markets otherwise not accessible.

In some use cases confidential data needs to be shared between team members within the MANET. One possibility is to use public keys of other team members to encrypt information. However, such an approach requires decrypting data on the device and then encrypting with a different key. Furthermore, the TPM cannot enforce the use of valid keys of other team-members. Another approach is to use key migration mechanisms of the TPM. The encrypted data is directly transferred and the key is migrated to the TPM of the other device. In the case of certified migration keys, this migration can again be controlled by a Migration Authority. Thus, the restriction of migrating data to team members only can be enforced by the TPM.

4.6. Tickets for the Measurement List

On basic functionality of TC is RA. However, RA has a variety of practical problems. During RA a unique identifier for the software state is securely transferred to the verifier. In order to check this identifier the verifier needs to know reference values for all valid identifiers. Thus, a possibly very large database (DB) of all valid states is required. This database needs to be either stored on the device itself or be provided by a central DB server. Both cases are not efficient for MANETs. A ticket scheme can be used to replace the measurement list by a placeholder that vouches for the given value signed by a central authority. The validity can also be bound to the tick session by including the tick session nonce in the ticket and requesting a TPM_TickstampBlob (timestamp) during the protocol.

5. Conclusion

This paper can only give a brief introduction into the variety of security mechanisms for MANETs that can be build on TPMs and TC technology. Most of the solution sketched in this paper can be realized with inexpensive commercial of the shelf hardware as it is used for embedded systems and PCs. However, software support for the TPM needs to be developed. Additional development is needed for nearly all areas, starting from trusted boot processes with proper integrity measurement up to software dealing with AIK certificated or supporting certified migration. However, this paper shows that hardware roots of trust exist for a large variety of security mechanisms and developing MANET applications or MANET protocols using these roots of trust can lead to new security properties for MANET applications. The building blocks described in this paper only use parts of the TPM functionality. Other features, like direct anonymous attestation could provide further interesting security properties.

References

[1] N. Aboudagga, M. Refaei, M. Eltoweissy, L. DaSilva, and J. Quisquater. Authentication protocols for ad hoc networks: taxonomy and research issues. In *Proceedings of the 1st ACM international workshop on Quality of service & security in wireless and mobile networks*, pages 96–104. ACM New York, NY, USA, 2005.

[2] S. Balfe, K. Boklan, Z. Klagsbrun, and K. Paterson. Key Refreshing in Identity-Based Cryptography and its Applications in MANETs. In *Military Communications Conference, 2007. MILCOM 2007. IEEE*, pages 1–8, 2007.

[3] S. Balfe, A. Lakhani, K. Paterson, et al. Trusted computing: Providing security for peer-to-peer networks. In *Proceedings of the Fifth International Conference on Peer-to-Peer Computing (P2P05), Konstanz, Germany, August*, pages 117–124.

[4] L. Genik, M. Salamanian, H. Schotanus, E. Hansson, C. Verkoelen, and P. Mason. Mobile ad hoc network security from a military perspective. Technical report, Defence R&D Canada Ottawa, 2004.

[5] K. Goldman, R. Perez, and R. Sailer. Linking remote attestation to secure tunnel endpoints. In *Proceedings of the first ACM workshop on Scalable trusted computing*, pages 21–24. ACM New York, NY, USA, 2006.

[6] M. Jarrett and P. Ward. Trusted Computing for Protecting Ad-hoc Routing. In *Proceedings of the 4th Annual Communication Networks and Services Research Conference (CNSR'06)-Volume 00*, pages 61–68. IEEE Computer Society Washington, DC, USA, 2006.

[7] S. Pearson. Trusted Computing Platforms, the Next Security Solution. *Bristol UK: HP Laboratories*, 2002.

[8] R. Sandhu and X. Zhang. Peer-to-peer access control architecture using trusted computing technology. In *Proceedings of the tenth ACM symposium on Access control models and technologies*, pages 147–158. ACM New York, NY, USA, 2005.

[9] S. Schechter, R. Greenstadt, and M. Smith. Trusted computing, peer-to-peer distribution, and the economics of pirated entertainment. In *Proceedings of The Second Annual Workshop on Economics and Information Security*, pages 29–30. Springer, 2003.

[10] F. Stumpf, A. Fuchs, S. Katzenbeisser, and C. Eckert. Improving the scalability of platform attestation. In *Proceedings of the 3rd ACM workshop on Scalable trusted computing*, pages 1–10. ACM New York, NY, USA, 2008.

[11] H. Tang, M. Salmanian, and C. Chang. Strong authentication for tactical mobile ad hoc networks. Technical report, Defence R&D Canada – Ottawa, 2007.

[12] Trusted Computing Group. TPM Specification Version 1.2 Revision 103. *Trusted Computing Group*, 2009.

[13] L. Venkatraman and D. Agrawal. A novel authentication scheme for ad hoc networks. In *IEEE Wireless Communications and Networking Conference (WCNC 2000)*, volume 3, pages 1268–1273, 2000.