

Article

On Using TPM for Secure Identities in Future Home Networks

Holger Kinkelin *, Ralph Holz, Heiko Niedermayer, Simon Mittelberger and Georg Carle

Faculty of Informatics, Technische Universität, München, Boltzmannstraße 3, 85748 Garching, Germany; E-Mails: holz@net.in.tum.de (R.H.); niedermayer@net.in.tum.de (H.N.); mittelberger@net.in.tum.de (S.M.); carle@net.in.tum.de (G.C.)

* Author to whom correspondence should be addressed; E-Mail: kinkelin@net.in.tum.de; Tel.: +49-(0)-89-289-18006; Fax: +49-(0)-89-289-18033.

Received: 28 October 2010 / Accepted: 28 December 2010 / Published: 7 January 2011

Abstract: Security should be integrated into future networks from the beginning, not as an extension. Secure identities and authentication schemes are an important step to fulfill this quest. In this article, we argue that home networks are a natural trust anchor for such schemes. We describe our concept of home networks as a universal point of reference for authentication, trust and access control, and show that our scheme can be applied to any next generation network. As home networks are no safe place, we apply Trusted Computing technology to prevent the abuse of identities, *i.e.*, identity theft.

Keywords: next generation networks; future home networks; identity management; trusted computing; authentication; trust

1. Introduction

It is still an open question what the network of the future will look like. Although a plethora of proposals have been submitted and discussed, there is no clear contender for the first prize. Some proposals even call for thousands of future networks in parallel [1]. However, one of the issues that any future network will have to address is access control, to both the network and network services offered therein.

It is probably fair to say that the general consensus is that this kind of mechanism should

- be easy to understand and use, *i.e.*, user-friendly,
- work from different geographic locations (roaming),

- work with different devices (versatility and portability),
- allow only legitimate users to access the network or a given service.

Controlling access will become an ever greater issue in future as a multitude of devices will be connected to the Internet (possibly leading to an “Internet of Things”), and quite likely via different access technologies.

We believe now is a good time to design new access control mechanisms for the network that, unlike today, fulfill the requirements listed above and can be applied to any type of new network. Our proposal is based on the idea of using home networks as a natural point of reference and trust anchor.

The remainder of this article is organized as follows: first, we present our authentication and identification scheme for home networks in next generation networks and describe two candidate technologies that can be applied together (Section 2). In Section 3, we elaborate on a way to increase the security of home networks, and thus the security of the system as a whole, by using Trusted Computing technology. We discuss benefits and disadvantages of our concept in Section 4. We conclude with a summary in Section 5.

2. The Home Network as The Universal Center of Reference

A secure scheme for access control needs a point of reference that the participants trust. We argue that a user’s home network is a natural point of reference, with properties that are useful in the context of future networks.

Given the current development towards more networked devices, it is plausible to assume that home networks will gain in importance dramatically. A future home network will be much more than the mere collection of Internet-capable devices that can be found in today’s homes. In future, the home network will offer its users advanced services like control over a number of home appliances, act as a universal storage, backup and maybe even streaming solution [2] and provide services to share and work on documents with others. In short, a future home network will be at the center of a user’s digital life. Such networks are, e.g., investigated in the project *AutHoNe* [3,4] (Autonomous Home Networking).

Because of the multitude of devices, services and users (all of which are later referred to as *entities*), these home networks will require secure and flexible mechanisms for authentication and access control. The first step in achieving this quest is to find suitable identifiers for entities and to design schemes for authentication and access control based on these identifiers. The second step is to harden the home networks against identity theft as it may be caused by, e.g., malware running in a home network.

2.1. Why to Leave the Beaten Path

We argue that user-centric approaches are well suited to the following requirements for secure identifiers in future home networks:

- **R1: Autonomy of Home Network Domains:** The management of identities must remain in control of the users of a home network, which we define as a domain.
- **R2: Hierarchical Identities for Entities:** The identity must reflect the membership of the entity to a certain home network as well as the individual entity itself.

- **R3: Provable Identities:** Entities must be able to prove their membership in the home network domain to a) other entities in their own home network and b) to entities in another home network.
- **R4: Offline Authentication:** Entities must be able to authenticate to other entities without the help of a third service that has to be online all the time for the scheme to work.
- **R5: Scalability:** An identifier scheme must scale to the huge numbers of devices in networked homes that all need to authenticate to each other.
- **R6: User-Friendliness:** The system must be as easy to use as possible.

The most widely-used scheme to identify users is simply a password in a challenge-response protocol. This is also the scheme that is still predominant in today's home networks. While it provides good control to users, works offline and enables authentication on a per user-basis (R1, R3, R4), it suffers from lack of scalability (think of configuration) (R5) and does not fulfill R2 nor R6, either. Given the usual issues with passwords, this scheme does not seem to be a good step forward and should be discontinued in home networking.

Every now and then, the claim is made that smart cards might be a good alternative instead. We view them as a useful complementary tool within our own concept but believe that they need a proper framework to really exploit their possibilities. They are also a rather expensive option as one card per entity is needed and must be configured. The additional hardware that has to be purchased (today) is another drawback.

Our proposal is thus to leave this well-trodden path once and for good and design a scheme that is based around turning the home network into an authority for the devices that are part of it.

2.2. The Home as an Authority

In our scheme, we thus introduce a Certification Authority (CA) that is resident inside the home network, as depicted in Figure 1. This *Home CA* issues certificates to entities within the home that a) identify the entity as a member of this home and b) bind the identity to the corresponding public key, which is (in general) also issued by the Home CA. The Home CA can be implemented as a piece of software running on a home network server, but we will explain later how this design should be hardened.

The identifier of an entity consists of two parts: the identifier of the home network plus an identifier for the entity that remains strictly local to this home network. We let the identifiers be self-certifying. The identifier of a home network thus consists of the hashed public key of the Home CA:

$$\text{HomeID} = H(\text{PubKey}_{\text{HomeCA}})$$

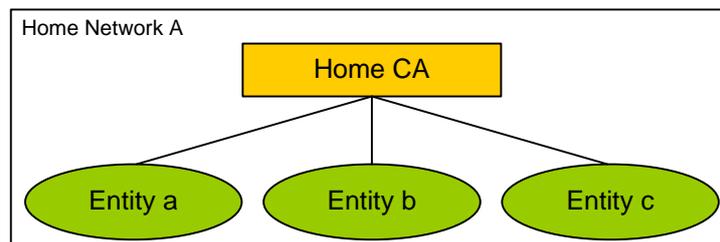
The identifier of an entity is then generated by hashing the public key of the Home CA and concatenating it with the hashed public key of the individual entity:

$$\text{EntityID} = H(\text{PubKey}_{\text{HomeCA}}) \cdot H(\text{PubKey}_{\text{Entity}})$$

This design fulfills R1-R5. As for user-friendliness, we have shown in [5] how an easy-to-use assistance system can be created that helps the owner of the home network and owner of the new identifier in the certification process.

We have so far not elaborated on how the Home CA can identify itself to other entities. In the following, we describe two flavors of our approach. They can be used alone or together.

Figure 1. Simple certificate hierarchy inside home networks: the Home CA issues certificates to entities.



2.3. Self-Certified Home Certificate

The first and more straight-forward option is to let the Home CA create a certificate for itself, *i.e.*, it signs its own public key with its private key.

As a home network is a reasonably controlled environment, it is feasible enough to introduce devices (entities) to this certificate and accept it. To this end, one can use PIN-based techniques to secure the exchange of certificates. Another convenient option is near-field communication together with a trust-on-first-use policy [6].

Entities within the home network will then trust the Home CA's certificate. Every entity in the local home network can use the thus acquired certificate to verify the certificate of another entity that belongs to the *same* home network.

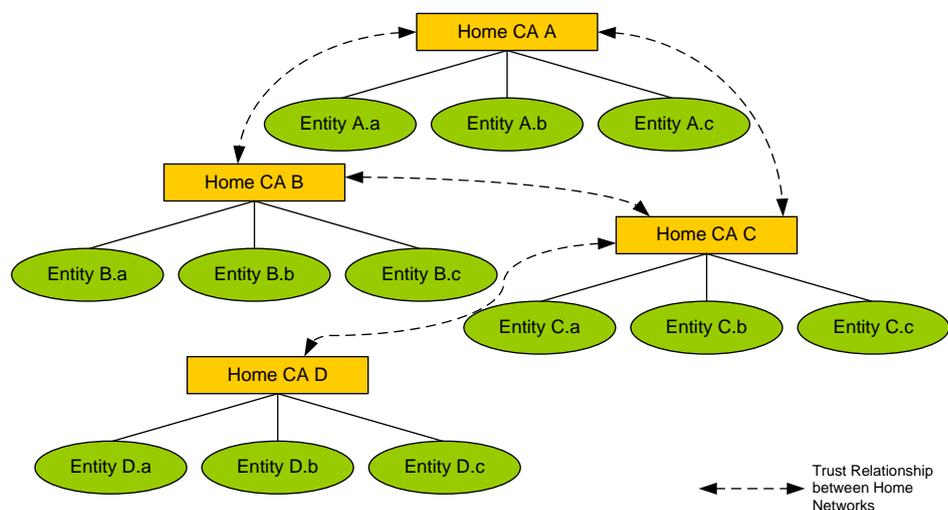
When an entity from *another* home network must be verified, this is not directly possible as entities only know and trust the own Home CA's certificate.

One way to overcome this is to build on the social relations that are likely to exist between the owners of home networks. For instance, members of different home networks can meet up in person and exchange their home certificates, using, e.g., an application installed on their smart phones. This can be again done with near-field communication (possibly strengthened by setting up a secure Bluetooth connection *etc.*).

We view this key exchange as a valid approach because trust relationships would likely need to be established between the homes of family members, friends, work colleagues and so on. The assumption of social relations is thus likely to hold, and the result would probably be a Web-of-Trust-like structure between home networks that reflects real-life relationships (Figure 2).

Within the simplicity of the key exchange mechanism lies its beauty: both partners are mutually identified in person, which allows for a high level of assurance with respect to the correctness of the other Home CA's certificate. Access control is now easy to implement: Home CAs can store different access rights for entities in their own domains as well as in other domains. This can be done with any granularity that the user requires (e.g., enable file-sharing for all entities in home network B, enable it only for some, ...). It also allows to store very different trust profiles for other home networks (e.g., enable a service only for home networks that are grouped under 'friends' *etc.*). This authorization can also be automated to some degree (if desired) and simplified for ease of use. We have built a framework based upon the *eXtensible Access Control Markup Language (XACML)* for the management of authorization policies within home networks [7].

Figure 2. Home Networks use self-certified Home Certificates. After a key exchange—the exchange of home certificates between ‘befriended’ home networks—foreign home networks and entities within can authenticate to each other.



2.4. Home CA with a Certificate Issued by a Trusted Third Party

In this section, we describe a complementary alternative to the self-certified approach above. Our idea is again based on an observation of social links between entities. Whatever a future network will look like, it is reasonable to assume that Internet Service Providers will continue to play the role they have today. In particular, they will provide the infrastructure that users need to connect their home networks to the Internet. Our idea is to extend the contractual relationship: when a user signs a contract with an ISP, he lets the ISP sign the certificate of his Home CA. The ISP can also provide the user with (root) certificates of other ISPs which this ISP knows to be authentic.

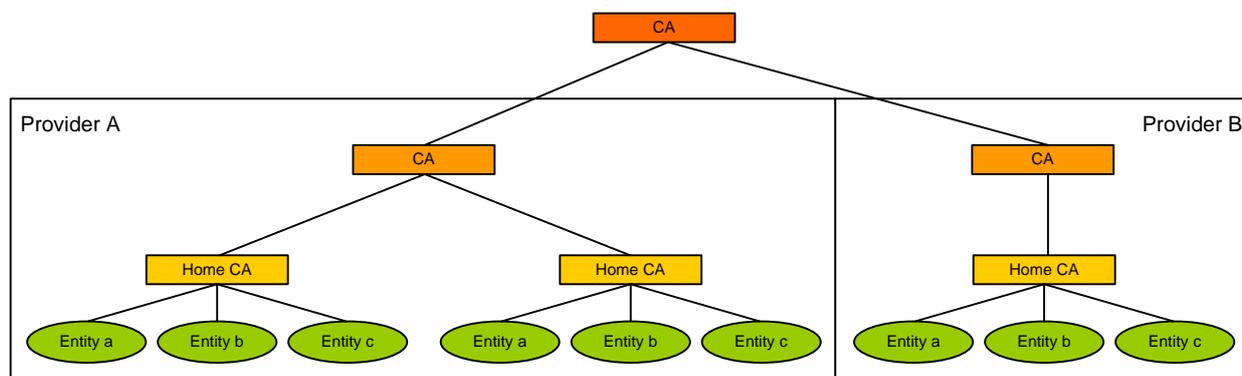
The result is a structure as depicted in Figure 3. Entity certificates still will be issued by the Home CA, but certificate chains between home networks can now also exist because of (possibly contractual) relationships between ISPs. It is of course always possible to have ISPs not only cross-sign each other, but introduce a common CA further up the certificate chain that issues certificates to ISPs.

There are drawbacks to this certification solution, however. The resulting PKI structure is likely going to be a multitude of ISP root certificates that are stored on the home network’s server, with the user having little control over the introduction of new certificates. In essence, the ISP would decide which other ISPs are trusted to be authentic and provide authentic identity information about their users. This removes control from the user. In effect, it resembles the collection of root certificates that we find in today’s Web browsers. The problems this incurs have recently been highlighted in [8], and the dangers of commercially driven PKI structures have been addressed in [9,10]. Furthermore, this PKI does not reflect the social structures of the people in a home and, thus, it is at first unclear how to derive access rights from the authenticated identities.

We thus see the benefit of this solution in providing a first indication of another’s entity authenticity. In contrast to the previous solution, access rights to the services of a home network would have to be set considerably more restrictive by default. Users would be recommended to update

access rights only afterwards, after they have been able to determine the other entity's identity and trustworthiness themselves.

Figure 3. Home CAs obtain their certificate from a trusted CA provided by the home network's ISP. An implicit path-of-trust between many home networks exists now that can be used to authenticate home networks and entities within.



2.5. Building up Trust between Home Networks

Networked homes are a very dynamic environment: new devices become part of a home, keys change, users lose their keying material and need to establish it again etc. We have already emphasized the control that users should have over their infrastructure. We will now show that it is possible to provide them with means to reflect the level and change of social relationships while providing a good level of security.

Our approach is built on a protocol that we have described in formal terms in [11]. The protocol—Peer Domain Protocol Authentication (PDP-A)—is a four-party authentication protocol with key establishment for domain-structured systems. Authentication servers of a domain aid their clients by participating in the authentication process. Their task is two-fold: firstly, they establish contact to other domains and store information about them, and secondly, they provide their clients with relevant information during the authentication run to aid in the authentication decision. Where two servers have a strong preexisting security association, their clients can authenticate to each other at the same level of assurance.

In our scenario, the Home CAs also act as the Authentication Servers. PDP-A allows to carry additional information between the Home CAs and the client entities. This can aid in building additional trust between home networks, which is useful in determining access rights in a semi-automated way.

We explain this by three examples. We consider the case of two devices from different home networks wishing to authenticate to each other.

Keys exchanged between domains We have shown in Section 2.3 how two devices or entities can exchange their Home CAs' public keys. Once this has happened, authentication between other entities of the same home networks benefits from the information that PDP-A conveys to the users: the other entity is known to belong to a domain to which a high level of assurance exists. Access rights can be set immediately in this case (although the user should probably still confirm first).

No keys exchanged between domains, but ISP assurance The situation is slightly different where two domains have not had previous contact, and Home CAs' keys have not been exchanged, but ISPs

have signed the certificates. In this case, the authentication would be based on the thus established certification chain. PDP-A would signal the different (lower) level of assurance to the user and access rights can be automatically set to be more restrictive. The user can then communicate with the other entity, making use of the authentication and the established encryption. If he believes the other entity to be authentic or otherwise trustworthy, he can still correct the access rights later.

No previous contact between domains, no ISP assurance In this situation, no keys have been exchanged between the home networks and no certification chain via ISP certificates can be found. PDP-A would in this case conduct a first key exchange between domains, but warn the user that the communication is not authenticated and fraud might be possible. The user can either stop the authentication process or continue with it and allow communication to happen (with no access rights set). If he becomes convinced of the other entity's trustworthiness or authenticity (for example, because communication was by VoIP), he can set access rights. In this case, a certain (very low) level of assurance would also be set for the other domain. This information will be stored on the home network server and be re-used in later authentication processes with other entities from that domain. However, as a man-in-the-middle might have been present during the first contact, it is not advisable to increase access rights much higher without exchanging keys between the domains (out-of-band) at some point.

3. A Trust Anchor for a Home Network's CA

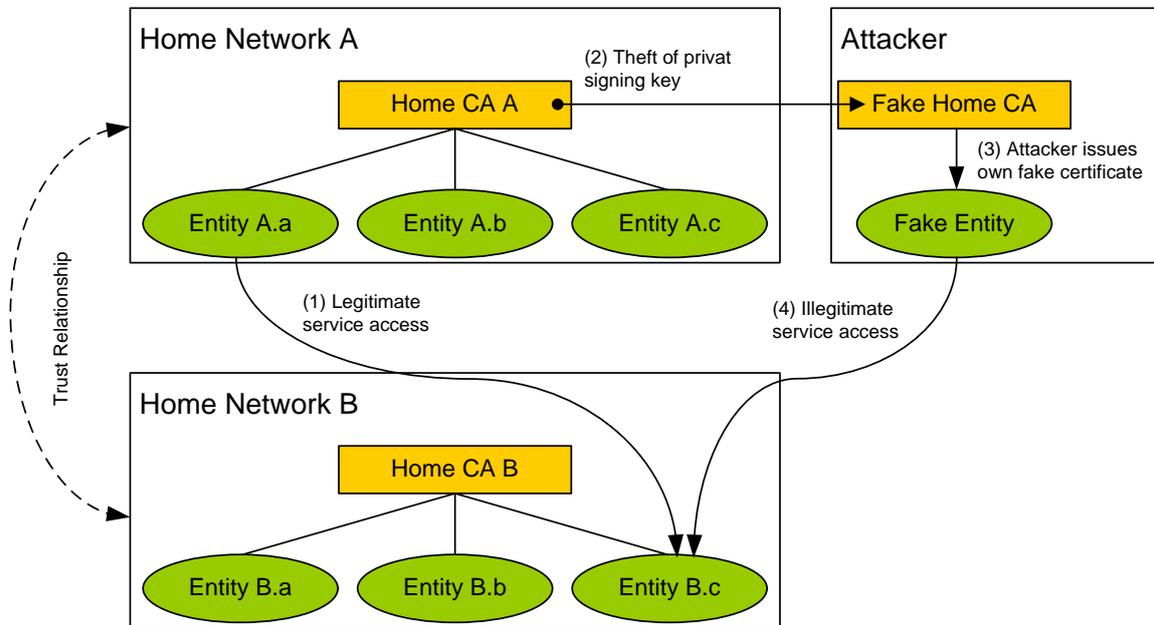
We have so far only elaborated on the certification scheme and the establishment of trust. However, one major problem inside home networks lies within the fact that home networks are no safe place *per se*. This is because their users are normally inexperienced and do not care much about security. Therefore, the chance that malware is resident inside the home network is not to be neglected. Whenever a certification authority is run in such an unsafe environment, the private signing key of the CA is endangered.

For instance, if malware succeeded in stealing a Home CA's private signing key (later referred to as the *Home Key (HK)*), the attacker would be able to sign certificates for his own entities that look as if they had been signed by the legitimate Home CA. If default access rights to services within another home network are granted to members of the compromised home network, the attacker can access those services. Theft of data, abuse of services, *etc.* can be the consequences, see Figure 4. Without protection against malware, identifiers issued by the Home CA are less trustworthy as it can not be guaranteed that the identifier is legitimate.

A technique that makes it difficult or even impossible that the Home Key can be stolen is needed. Again Smart Cards would be a candidate technology, but because of the reasons listed above, we believe that they need a strong framework to be useful. A technology often used is the so called *Trusted Platform Module (TPM)* which has been specified by the *Trusted Computing Group*.

Of course, malware could also steal the private key and certificate of an end entity. The protection methods we describe can also be applied with minor changes to their private keys. However, in this article, we focus exclusively on the protection of the Home CA's key.

Figure 4. Between home network A and B exists a trust relationship, *i.e.*, entities from home network A may access services in home network B (1). The attacker succeeds in stealing the Home Key of Home CA A (2). The attacker now is able to issue a certificate to its own entity (3) and later is able to access and abuse the service within home network B.



3.1. Trusted Computing and the Trusted Platform Module

In this paragraph, we give a high-level introduction into aspects of the *Trusted Platform Module (TPM)* that are most relevant in our context. For further details, please refer to the TCG standards [12–15].

A TPM is a low-cost cryptographic chip, which can be embedded into many computing devices like notebooks or desktops. One of the most important abilities of a TPM is to generate, use and manage asymmetric keys. A private key is securely stored inside the TPM’s storage and never exposed to the memory of the computing device. Properties of keys can be specified during their creation, *e.g.*, to prohibit that a private key ever leaves the TPM. A program using the keys to encrypt or sign data needs to hand over the piece of data to the TPM and the TPM hands back the encrypted blob or signature. Keys managed by the TPM can be protected against unauthorized usage with an authorization secret (PIN). Every time the user application desires to use the key, this PIN needs to be provided to the TPM.

During the production process of the TPM-equipped computing device, various credentials are issued by the manufacturer that can be used to prove the compliance of the TPM-chip and other platform components with the TCG standard to other parties. The most important credential is the EK Certificate which is issued by the manufacturer of the TPM for the *Endorsement Key* of the TPM. The EK can be seen as the TPM’s master key. The EK certificate asserts that the EK is resident inside a standard-compliant TPM.

Trusted Computing guidelines do not allow to use the EK directly to encrypt or sign data. Instead, so-called *Attested Identity Keys (AIK)* are generated, which can be seen as anonymized aliases for the

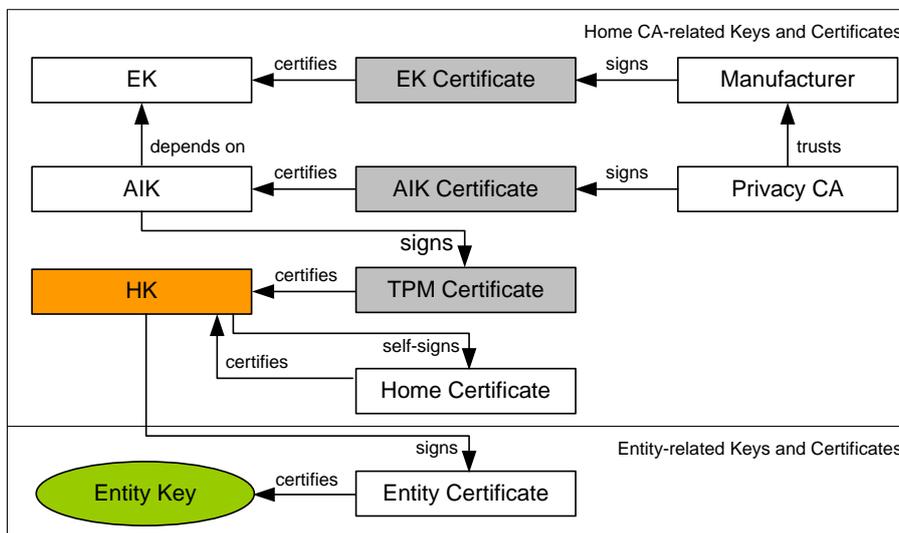
EK. The generation process of an AIK involves an entity called *Privacy Certification Authority (PCA)* which issues an AIK Certificate. This certificate attests the usage of a genuine TPM.

An important property of TPM keys is their migratability. Every key in a TPM, except the EK, AIKs and keys signed by one of them, can be specified as migratable. A migratable key can be exported after providing a migration secret. The advantage of a migratable key is the possibility to backup the keying material.

3.2. Home TPM CA

When using a TPM to secure the Home CA’s private signing key (Home Key, HK), the key hierarchy inside the home network needs to be enhanced by the keys and their certificates that enable and prove the HK’s protection by a TPM. As explained above, an AIK needs to be created as an alias for the EK first. Due to constraints of the TPM’s specification, the AIK may not be used as HK, but it can be used to sign another non-migratable TPM-protected key, the Home Key. This leads to the following key and certificate hierarchy within the home network: EK → AIK → HK. The HK can then be used to sign the certificate of an entity (see Figure 5).

Figure 5. Changed key and certificate hierarchy inside the home network. Certificates depicted with grey boxes are TPM-related and used to attest that the HK is stored within a TPM. For compatibility with non-TPM systems, a self-signed Home Certificate is still created. The certificate for an entity’s public key is signed by the HK.



Malfunction of a TPM chip can lead to the loss of the HK. This means that the Home CA is not able to issue certificates to new entities anymore. The home network needs to create a new HK and Home Certificate and finally re-establish trust relationships to other home networks. A higher level of usability and a backup in case the TPM is malfunctional can be provided by making the HK migratable. In order to sign a migratable HK, another non-migratable key, referred to as *Intermediate Key (IMK)*, has to be added to the key hierarchy. Finally, the following key chain is built: EK → AIK → IMK → HK. Again, the HK is used to sign the entity certificate.

Both key hierarchies can be used to prove to others that the HK is stored within a standard compliant TPM. There is practically no difference in the security of the key's protection. But Home Networks that trust a Home Certificate that belongs to a non-migratable HK should assign a higher level of confidence to this certificate than into a Home Certificate of a migratable HK. The reason is that the backup of a migratable HK could be stored insecurely.

The HK is protected against unauthorized use by applications or malware by using a PIN. The owner of the home network needs to provide this PIN when access to the HK inside the TPM is needed to sign a new certificate of an entity.

3.3. Advanced Attacks on the TPM CA

Although the TPM-protected HK of a Home CA cannot be stolen, different attacks are possible that can lead to identity theft. For example, an attacker who has introduced malware into the Home CA can eavesdrop on the PIN. The possession of the PIN enables the attacker to create a certificate. Alternatively, malware could trick the Home CA during the certification process to certify a public key that belongs to the attacker. Although this kind of attack is difficult to perform, it is still possible. Therefore we consider the trustworthiness of a Home CA's execution environment as crucial, at least during the certification process.

Trusted Services within Home Networks We propose the following countermeasure: we run the TPM Home CA within a *trustworthy execution environment*. Trustworthy means that only well-known and verified software is installed within the execution environment.

The TC technology provides a mechanism referred to as *Remote Attestation (RA)* which is able to assess the status of an execution environment. For RA, integrity measurement values (SHA-1 fingerprints) are recorded by the TCG's *Trusted Boot* process and IBM's *Integrity Measurement Architecture (IMA)* [16]. Both technologies make use of properties of the TPM that allow for the tamper-proof creation and reporting of measurement values. The fingerprints are sent to a verifying external entity for assessment. This verifying entity needs to be equipped with reference fingerprints of trustworthy software and can compare the fingerprints. When all reported fingerprints match the reference fingerprints, the verifying entity can be confident about the trustworthiness of the remote system, *i.e.*, about the absence of malware. The result of this verification can be sent to the owner of the home network via a secure channel before the certification process starts, *i.e.*, before the owner enters his PIN for unlocking the Home Key (see Figure 6).

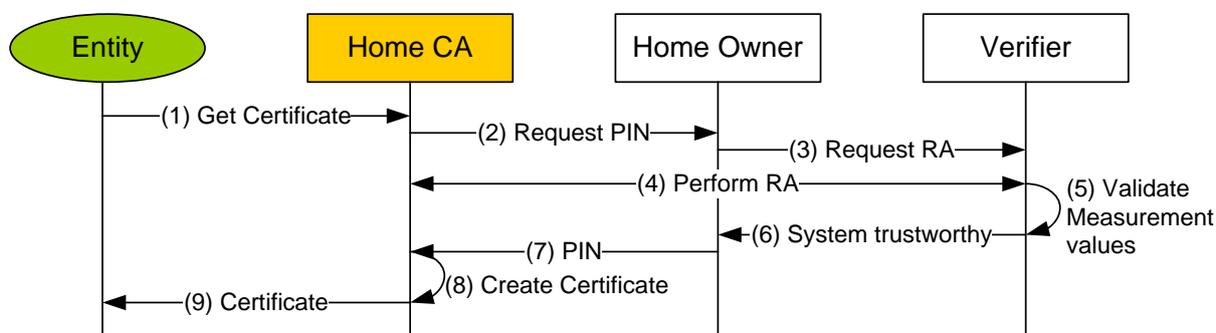
From practice, we know that the crux here is complexity: the trusted execution environment must be small and well defined in order for RA to be feasible. The extreme diversity of operating systems and software components is a major obstacle for RA services as no verifier would be able to know fingerprints of all software used.

(Virtualized) Trusted Execution Environments Our proposal is to use small and strictly predefined operating system images that only run the TPM CA. For instance, such operating systems could run on dedicated low-cost devices or within virtual machines on XEN-like architectures. For XEN VMs, early implementations exist that provide access to a virtualized TPM. Virtualized TPMs make use of the hardware TPM of the machine and basically can be used within a virtual machine just like a hardware TPM [17].

The needed operating system images, either for a dedicated device or the virtual machine, can be distributed by trusted entities on the Internet. This distribution process could be realized like the distribution of signed software packages of today’s Linux operating system packet managers. If Home Certificates issued by the CA of an ISP are used, the ISP could also distribute the operating system image.

For both options, dedicated device or VM, the small scale and predefined operating system image allows to establish the trustworthiness of the TPM CA during the certification process of an entity. Note that this works with either of the two methods we have described in Section 2.

Figure 6. Remote Attestation of the Home CA. When a certificate needs to be created (1), the Home CA requests the home owner to enter the PIN (2). The owner triggers the RA process (3–5) and enters the PIN (7) if the Verifier assesses the Home CA as trustworthy (6). The Home CA now is able to unlock the HK and create and deliver the certificate (8 + 9).



4. Discussion

We discuss benefits and disadvantages of our concept in this section.

Benefits of our Identification and Authentication Scheme Our concept has the advantages that it is user-centric, that it can make use of information from different sources, and that it helps in establishing trust relationships with other home networks. It is also very extensible: information about authenticity and trustworthiness could be derived from other sources, like, e.g., social networks where fingerprints of home networks can be stored in a profile.

Users’ social relations are reflected by the exchange of certificates in a convenient way. At the same time, it is possible for users to contact previously unknown domains (homes) and access rights can be determined in a semi-automated way. The scheme is also useful for the common case where domains have no exchanged keys or even no previous contact. It scales with the available information and aids users in making informed choices while not annoying them with unnecessarily technical background.

Benefits of TPM Usage in Home Networks Using the TPM as a hardware safeguard has the advantage that stealing the Home Key becomes much less feasible. The attacker would need to break the security mechanisms provided by the TPM. To our knowledge, no software-based attack that succeeds in the theft of a key from a TPM is known. However, once such systems are employed and there is enough financial incentive for identity theft of home networks, we expect that such attacks will appear. Our solution cannot prevent this, and software patches that fix exploits will continue to be needed. Migratable or non-migratable keys can be used as the Home Key. The first option makes it possible to

backup the Home Key which prevents the loss of the HK in case the TPM is defective. The drawback is that a migratable Home Key is less trustworthy than a non-migratable one.

Malware on the TPM-enabled Home CA could still interfere during the certification process. Remote Attestation can be applied to fight such attacks. The Remote Attestation mechanism is only feasible for small scale operating systems. The first option would be to use a dedicated device as a Home CA. Dropping hardware prices make this possible. The second option would be to use virtualized TPM-enabled machines. Here, trusted services can run in different virtual machines with a high level of isolation. This will provide a high level of flexibility. But here the system security depends heavily on the security of the hypervisor used. If malware running in a virtual machine succeeded in compromising the hypervisor, it might be able to tamper with the communication between the virtualized TPM and the hardware TPM. This could lead to wrong measurement values and a remote attestation without effect.

Although there are some unsolved problems when using TPM technology, and TPM is certainly not the silver bullet that it is sometimes touted to be, we believe that on the whole it can be used effectively to enhance security.

5. Summary

In this article, we have shown that identities and authentication schemes for next generation networks can have a trust anchor within future home networks. We have designed such a scheme that has the advantage that it works independently of what a future network will look like. We have described mechanisms that issue certificates/identities to entities within the home network. By establishing trust chains that reflect the social relationships of users we generate a Web-of-trust-like security structure for authentication and access control. Internet Service Providers can reinforce this structure by certifying home networks with whom they have contractual relations. A supporting authentication scheme supplies users with information about the trustworthiness of an entity in an authentication process. These mechanisms can help to determine access rights for authenticated entities.

However, security mechanisms must work within an unsafe home network environment. We have shown how different technologies taken from the Trusted Computing Group can be applied to mitigate attacks on the system, e.g., different kinds of identity theft attacks.

References and Notes

1. Völker, L.; Martin, D.; Khayat, I.E.; Werle, C.; Zitterbart, M. A Node Architecture for 1000 Future Networks. In *Proceedings of IEEE International Workshop on the Network of the Future*, Dresden, Germany, June 2009.
2. Fouquet, M.; Niedermayer, H.; Carle, G. Cloud Computing for the Masses. In *Proceedings of 1st ACM Workshop on User-provided Networking*, Rome, Italy, December 2009.
3. AutHoNe is funded by the German Federal Ministry of Education and Research.

4. Carle, G.; Kinkelin, H.; Müller, A.; Niedermayer, H.; Pahl, M.O.; König, A.; Luckenbach, T.; Scholl, K.; Schuster, M.; Thiem, L.; Petrak, L.; Steinmetz, M.; Niedermeier, C.; Reichmann, J. Autonomic Home Networks in the BMBF project AuthoNe. In *Proceedings of 8th Würzburg Workshop on IP (EuroView 2008)*, Würzburg, Germany, July 2008.
5. Müller, A.; Kinkelin, H.; Ghai, S.; Carle, G. An Assisted Device Registration and Service Access System for Future Home Networks. In *Proceedings of IEEE IFIP Wireless Days*, Paris, France, December 2009.
6. Stajano, F.; Anderson, R. The Resurrecting Duckling: Security issues for ubiquitous computing. *IEEE Comput.* **2002**, *35*, 22–26.
7. Müller, A.; Kinkelin, H.; Ghai, S.K.; Carle, G. A secure service infrastructure for interconnecting future home networks based on DPWS and XACML. In *Proceedings of 1st ACM SIGCOMM Workshop on Home Networks (HomeNets 2010)*, New Delhi, India, September 2010.
8. Eckersley, P.; Burns, J. An observatory for the SSLiverse. In *Proceedings of Talk at DefCon 18*, Las Vegas, NV, USA, July 2010.
9. Gutmann, P. PKI: It's not dead, just resting. *IEEE Comput.* **2002**, *35*, 41–49.
10. Ellison, C.; Schneier, B. Ten Risks of PKI: What You're not Being Told about Public Key Infrastructure. *Comput. Security J.* **2000**, *16*, 1–7.
11. Holz, R.; Niedermayer, H.; Hauck, P.; Carle, G. Trust-Rated Authentication for Domain-Structured Distributed Systems. In *Proceedings of 5th European PKI Workshop: Theory and Practice (EuroPKI 2008)*, Trondheim, Norway, June 2008.
12. Trusted Computing Group. TCG Specification—Architecture Overview, Revision 1.2, 2004.
13. Trusted Computing Group. TPM Main Specification Level 2 Version 1.2, Revision 103, Part 1: Design Principles, 2007.
14. Trusted Computing Group. TPM Main Specification Level 2 Version 1.2, Revision 103, Part 2: Structures of the TPM, 2006.
15. Trusted Computing Group. TPM Main Specification Level 2 Version 1.2, Revision 103, Part 3: Commands, 2006.
16. Sailer, R.; Zhang, X.; Jaeger, T.; van Doorn, L. Design and Implementation of a TCG-based Integrity Measurement Architecture. In *Proceedings of 13th Usenix Security Symposium*, San Diego, CA, USA, August 2004.
17. Berger, S.; Ceres, R.; Goldman, K.; Perez, R.; Sailer, R.; van Doorn, L. *vTPM: Virtualizing the Trusted Platform Module*; Technical Report; IBM Research: Hawthorne, NY, USA, 2006.