WIRELESS WORLD
RESEARCH FORUM

A Mobile service architecture utilising trusted computing

Nicolai Kuntze and Andreas U. Schmidt
Fraunhofer Institute for Secure Information Technology SIT
Darmstadt, Germany

***Abstract*— Convergence of access technologies potentially enables new mobile business scenarios. To enable them, trust between the stakeholders is a prerequisite, e.g., in the form of a cross-domain authentication infrastructure. Service architectures provide solutions for some of these needs. This paper proposes a use case combining the well understood techniques from the mobile domain with the upcoming new possibilities of trusted computing.**

***Index Terms*— authentication, service architecture, trusted computing**

INTRODUCTION

The horizontal integration of on-device access technologies (like 3G, WLAN, Bluetooth, WiMax, RF, etc.) could principally foster new business scenarios which are agnostic with respect to network and access technology. Some of these usage scenarios in the mobile domain, in particular for service access, depend on the underlying trust model. Especially charging requires a high trust into the basic technologies and devices. Trusted Computing (TC) offers new techniques in this area, to establish trustworthy mobile devices and thus enable new use cases. The challenge here is to enable trust between technological and entrepreneurial domains [1, 2].

A TC enhanced mobile device includes a hardware token, the Trusted Platform Module (TPM), which is the root of trust. Inside this module keys can be stored, managed and used without the possibility to tamper them. Moreover the TPM offers trust measurement enabling the device to testify its integrity to a third party. This is done by an attestation process.

Service architectures provide distinct functionalities based on machine to machine (M2M) communication. Based on the M2M paradigm various tasks can be addressed like maintenance, metering of consumption, and charging of goods. We describe a collaboration scenario of a trusted mobile phone in a service architecture. We show that by introducing TC a transitivity of trust can be gained in the sense of providing referral of authentication between different domains [3]. We specifically propose an architecture implementing a point of sales (POS) setting. A salient feature of the scenario is that the POS does not require any network communication capabilities.

A trustworthy point of sales

Figure 1 shows the service architecture in the context of payment interactions between various parties. The owner of the mobile device wants to pay for a good at the POS. The POS equipped with a TPM uses the TPM-enabled mobile device to communicate with the charging unit which could be located at the mobile network operator (MNO), at an external service provider, or at the POS owner/commodity vendor.

Two essential properties of this proposal have to be highlighted. First, the POS uses the mobile device as a toehold in the communication with the MNO. The mobile device enables another device – the POS – to utilise its connection to the MNO for

authentication, authorisation, and accounting (AAA). The mobile device can either ask on behalf of the POS or just relay AAA requests and related data. From the viewpoint of cost-effectiveness it is very advantageous that no GSM or UMTS module is required in the POS. The communication can be performed using a near range communication module like Bluetooth or IRDA.

Second, and most importantly, the authentication of the mobile device is performed by the well-known SIM-card authentication. The authentication between the mobile device and the POS system is based on secrets offered by the mobile device's TPM. This secret is used to authenticate and ultimately authorise the customer at the side of the charging unit. This transition between SIM-card based and TPM based authentication schemes requires a protected forwarding of authentication credentials, enabled by basic TPM features.

The mobile network operator offers in such a scenario services both to his own clientele as well as to the POS proprietor, and can thus profit from his established infrastructure and large subscriber base.

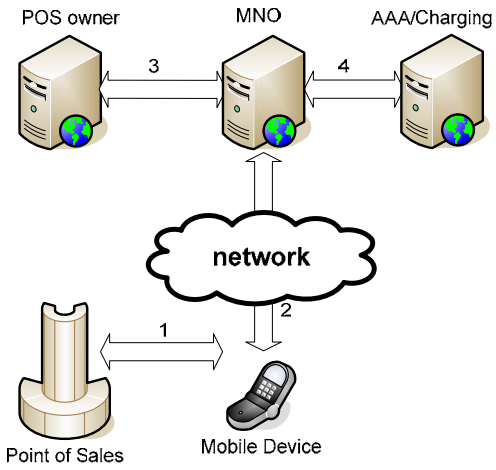

Figure 1 Point of sales scenario, all involved parties and their relations

By separating the duties [4,5] between a charging unit and the MNO, additionally certain privacy options arise could enable minimal need to know policies. For instance, the MNO need not know the identity of the single POS involved or its sales volumes.

Trust relationships

Let us briefly describe the requirements of the envisaged POS scenario with respect to establishment of trust between the involved parties, i.e., authentication, and its maintenance, i.e., data protection throughout a purchase process.

Establishment of trust

Establishing trust between two communication entities requires two distinct attributes to be proved. First, the identity of each entity has to be shown by an authentication process. Based on a successful authentication, the authorisations of the user can be determined and the actions of the user can be logged as necessary. Second, each communication partner requires an assertion about the trustworthiness of his vis-à-vis. This assures him, that the partner behaves as expected.

The first requirement can be solved by genuine authentication, authorisation and audit (AAA) architectures. The required authentication data can be provided, as said, by the MNO user base.

The second requirement describes the trust which is laid into the platforms performing the transactions of the POS scenario. Especially the mobile devices and POS instantiations are exposed to various attacks by third parties. Therefore it is necessary for the POS owner and the charging provider to assure that the infrastructure is in a trustworthy state.

In a distributed scenario, trust can be established in various ways and with different instantiations of duty separations for the authentication and authorisation functions. In the scenario of access to a POS by a mobile user, and mediated by the MNO, the essential trust relation is a mutual one. It operates between the domains of authentication of the MNO, i.e., his user base, respectively the base of trusted mobile devices, and the POS owner's, identifying the trustworthy POS vending machines. While the MNO needs to know that the POS is an uncompromised one from the POS owner's domain to process payment, the POS owner will only allow for purchases from authenticated, trustworthy devices of the MNO's domain.

Data protection

In the economic environment it is often required to implement a certain kind of minimal need to know principle to protect each communication partner from disclosing

corporate secrets or informational assets. For example the charging service provider may not be meant to know the list of goods, the stock, price list, or other details of a certain POS owner. Furthermore, the MNO should not know how many points of sale are owned by a certain company, where these POS' are located, and generally how the POS owner's business is doing.

With respect to these data protection targets and independently of implementation details and variants, the POS scenario can be analysed as follows. During a purchase operation

(i) the mobile device authenticates itself at the POS (see below).

(ii) To verify the authentication token offered by the mobile device the POS connects to the POS owner infrastructure where the decision of acceptation is made. This decision is made based on a trust relationship between POS owner and an authentication provider (e.g. a charging provider). Alternatively the POS requests this directly at the authentication provider. In both cases the POS owner gets no knowledge about the identity of the customer.

(iii) After the authentication the purchase process is performed at the POS. The resulting data containing the billing information like authentication token, good identifier price, are transferred to the POS owner, where a special data package is generated for the charging provider. This package needs in principle only to contain the authentication token and the grand total to charge. Hence the charging provider does not gain any information about the good, quantities, and qualities sold by the POS owner, nor essentials of his infrastructure.

(iv) After the confirmation of the charging the POS owner acknowledges the purchase and the POS vending machine delivers the good.

Steps (iii) and (iv) can be decentralised in an alternative approach, and therefore as well be performed by the POS. The POS in this case requests the confirmation from the charging provider and afterwards requests the acknowledgement by the POS owner.

Out of the view of the charging provider relations between costumers and POS owning companies can arise, e.g., in the form of special offers and rebates. These marketing strategies can remain closed to the other involved parties. In particular there is no possibility to connect individual consumers to the purchased products even for the charging provider who cannot, e.g., build a customer profile using his data alone.

From the point of view of the network operator all traffic is encrypted so it is not possible for her to distinguish between a POS and any other device. Privacy concerns could arise if the network operator and authentication provider are integrated in one party. In this scenario it could be possible to compile user profiles and reconstruct buying patterns.

Trusted assertions

Credentials that can be constructed basing on the functionalities of a trusted platform module (TPM [11]) play a special role in our concept. They are used to relate an agent to the outside world, i.e., to authenticate agent **a** with respect to an authority **A** and to protect the communication between agents and between agents and authorities. Such a credential is denoted by $\mathbf{c}_{\mathbf{a,A}}$.

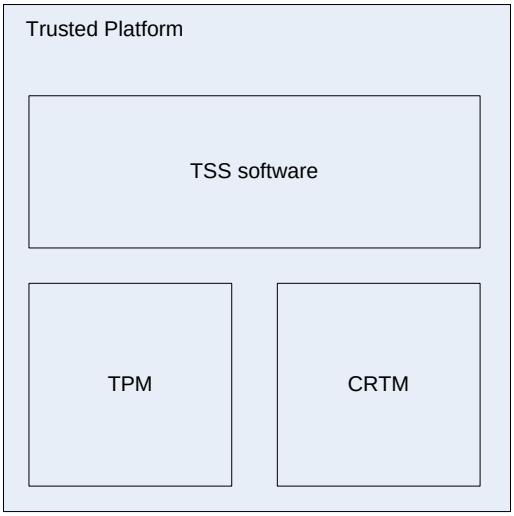

Figure 2 A view of a trusted platform subsystem

A TPM is part of a trusted platform as it is specified by the Trusted Computing Group. Figure 2 shows the main components of a trusted platform. Such a platform should offer the following set of features: Memory curtaining, sealed storage, secure I/O, and remote attestation. Memory curtaining is a hardware enforced memory feature which makes it impossible for concurrent software to read or write in each other's memory area. By this each software is isolated and protected from other malicious running software. Sealed storage offers a secure place for storing and using keys. Secure I/O protects the communication lines of a device, e.g. the communication between the

keyboard and an application. Remote attestation enables a user at the other end of a communication to determine if the requesting application is altered or not.

TPMs provide a number of features that can be used to securely operate a system. Methods for the secure generation, storage, and usage of asymmetric key pairs are the foundation for encrypted and authenticated operation and communication. Trust measurements on the system environment exerted at boot- and run-time allow for trustworthy assertions about the current system state and a re-tracing of how it was reached. The system state is securely stored in platform configuration registers (PCR) tamper-resistantly located inside the TPM. Memory curtaining and sealed storage spaces are enabled by pertinent TPM base functions. Trustworthy system and application software can build on this basis to establish authenticated communication with the exterior and transmit data maintaining integrity and confidentiality.

In particular, the root of trust for reporting (RTR) of a TPM establishes a context for attesting to reported values. The RTR is a cryptographic identity which makes the TPM unique and must be instantiated before the first use of the TPM. In the context of the TPM the RTR is called Endorsement Key (EK). This EK is generated inside the TPM using a built-in hardware (physical) random number generator, and generation takes place when TPM or containing platform is manufactured. The usage of this EK is restricted by the TPM for two cases: (i) establishing a TPM owner and (ii) creating Attestation Identity Key (AIK) values and credentials.

The private part of each key generated by the TPM is stored inside the TPM or in a special protected storage area. This area is encrypted with the Root of Trust for Storage (RTS) Key. This key also is not accessible outside the TPM.

Establishing a unique ID for a system creates privacy concerns as the use of that identity could result in aggregation of activity logs. Solving this problem requires an indirection introducing a third party which offers aliases for the EK. These aliases are referred to as Attestation Identity Keys (AIK). They provide signatures, but not encryption.

Based on the PCRs and the AIK a Remote Attestation can be performed. In preparation the user requests for the mobile device a certificate for an AIK from a trusted third party. This third party, called privacy CA, verifies that the AIK belongs to a TPM which is in a trustworthy configuration and then issues a certificate stating this. This certificate is later used to anonymise the identity of the mobile phone. The certificate of the privacy CA assures that this AIK belongs to a trustworthy TPM. Without the certificate of the privacy CA the user has to offer his identity to the communication partner. During the process of remote attestation the mobile device performs a quote operation which is a cryptographic reporting of PCR values. To compute a quote, a set of PCRs and externally supplied data are signed with a signing key, the AIK. The resulting quote is then returned to the challenging entity together with a measurement log. Based on this data two attributes can be verified. First, that the quote was produced by a genuine trusted platform by verifying the digital signature of the quote with the AIK certificate. Second, the integrity of the platform by comparing the integrity values of the measurement log. A compromised system can tamper the log but cannot change the PCR values as these are protected by the TPM.

In practice the verifier has to know a reference value for the integrity values to match the measurement log. In the PC domain this reference is rather large as every possible version of existing software and hardware has a different hash value and so a different integrity value. The mobile domain distinguishes from the PC domain as the number of hardware combinations is rather small and updates of the used software are also rare. In fact the remote attestation therefore seems to have higher practical feasibility in the mobile domain.

The measurement of the system integrity is implemented as a chain of trust. The Core Root of Trust for Measurement (CRTM) initially measures itself and reports to the TPM. On this trust base it moves up the boot hierarchy. BIOS is measured and booted. The BIOS measures the bootloader (i.e. trusted grub [10]). And the boot loader in turn measures and starts the respective OS. The OS in turn has the access to the TPM to report modifications in the software. This is done by using the Trusted Platform Support Service (TSS) where e.g. TrouSerS [12] is available.

A novel alternative method responding to remaining privacy concerns of the centralised approach embodied in a privacy CA and also

promising to be very efficient is Direct Anonymous Attestation (DAA), a method put forward in [6] and specified by the trusted computing group (TCG). It enables the establishment of trust relationships of a trusted system with external entities. A central goal of DAA is to cover privacy issues related to previous versions of the standards [7]. DAA is based on involved cryptographic protocols known as zero-knowledge proofs.

It has to be noted that DAA has higher demands in terms of computing power and a real-world implementation is still pending. Both TC-based attestation methods can be used as and are referenced to as attestation protocols.

One important step in the lifecycle of a TPM is the "take ownership" process. Basically this is the process of inserting a shared secret into a TPM shielded location. Any entity who knows this secret is a TPM owner. The proof of ownership is done by using a challenge-response protocol. During take ownership the RTS is created and by this it is guaranteed to be unique for each owner of this TPM. If the ownership is revoked the RTS is to be destroyed. A TPM ships with no owner installed.

Certain commands of the TPM require a physical presence of the TPM owner.

Although certain flaws are known in the TCG standards (e.g. [8] points to a flaw in the OIA Protocol an authorisation protocol which represents one of the building blocks of the TPM) that exist currently future versions are likely to remedy them. We assume for the purport of our applications that the functions used are at least secured against common attack vectors in the scenario below.

Using the described functionality, a trusted system **a** can establish what we call a trust credential $\mathbf{t_a}$. Specifically, we assume that the trust credential can be used to attest the validity of three fundamental security assertions of a system to the exterior.

1. The presence of a live and unaltered TPM. This can for instance be carried out using a challenge-response method using the TPM's endorsement credential. Endorsement credentials are pre-installed by the TPM's manufacturer.

2. The integrity of the system and its components. This property is ascertained through trust measurements and communicated via an attestation protocol.

3. That an existing credential $\mathbf{c_{a,A}}$ is unaltered. This must be established by trusted system software and components used to access the credential's data. Again, this assertion is forwarded to other parties using an attestation protocol, and secure communication channels established therewith.

These properties are not independent but build on each other, i.e, to prove 3. one needs first attestation of 2. and 1., etc. The TPM is capable of creating, managing, and transmitting own cryptographic credentials which can convey the described assertions 1.–3.

The basic operation for creating trust between agents in our POS scenario relies on referral trust in the parlance of [9]. That is, on the ability of a trusted agent through assertions 1.–3., to make recommendations to trust another agent or even himself in a special, functional role.

Purchase Process

A coarse description of the purchase process at the operational level is as follows. A user with a TPM-equipped mobile device wants to purchase a soft drink from a likewise trust-enabled vending machine, the point of sales (POS). While the user still makes up her mind on her taste preferences, device and POS initiate a trusted communication session using attestation and transport layer encryption. Device and POS thus achieve mutual assurance that they are in an unaltered, trustworthy state, and begin to exchange price lists and payment modalities (1 in Figure 1).

After the user selects a good and confirms his choice at his device, signed price and payment processing information is transferred to the MNO (2 in Figure 1). After verifying the signatures and optionally informing the good's vendor and a payment service provider (3 and 4 in Figure 1), the MNO sends a signed acknowledgement to the mobile device, which relays it to the POS, where it is verified and the good is delivered.

How are the requirements on trust and data protection listed above fulfilled in the purchase process? In the mutual authentication of the mobile device and the POS, end-to-end security is achieved. In particular, remote attestation of the device toward the POS yields to the latter the assurance that full price and lists of available goods will not leave the device. This is possible since the device is proved to be in a trustworthy state, belongs to the MNO's domain, and, optionally, end-to-end

encryption between POS and device is established (e.g., using TC functionality).

The individual identities of POS and device (and therefore its owner) need not be revealed in the purchase process. Authentication is carried out and trust is established by referral. That is the respective providers of identities of the two domains in question – the MNO and the POS vendor – vouch for the identities of their agents. In more detail this can be implemented as follows.

For the authentication of the POS, $\mathbf{c_{pos,OWNER}}$ is transferred to the mobile device over a protected channel. Authentication of the POS can be achieved by at least two ways protecting the anonymity of the POS. First, the mobile device connects direct the POS owner and verifies the $\mathbf{c_{pos,OWNER}}$. This requires an initial trust of the mobile device into the POS owner. Making it more convenient for the user the MNO can take place offering a trust relation between the user and the POS owner as a service offer. The mobile device asks in this second scenario the MNO about the identity of the POS. The MNO can verify the signature of the offered token and acknowledge the request.

The identity of the POS is revealed to the MNO by this operation. Using the TC concept of a privacy CA on side of the POS owner the POS can change its identity after a certain time or use some identities in parallel providing at least pseudonymity. This prevents compilation of purchase data by the MNO. A similar implementation would by possible using DAA without necessitating the online requests to a privacy CA.

Not only the trust credentials $\mathbf{t_{POS}}$ and $\mathbf{t_{device}}$ with which POS and device carry out remote attestation can be based on AIKs but also the genuine credentials of POS and device can make use of this TC functionality. An AIK is created inside the TPM and its private portion is protected by the shielded capabilities of the hardware instantiation of the TPM. The privacy CA issues a certificate base on the certificates of the platform stating authenticity of the AIK. The AIK can be used in combination with this certificate as a pseudonym of the platform. In our scenario the privacy CA of the POS is under control of the POS owner. Therefore each POS can acquire as many AIKs as necessary thus protect its own identity (duration of validity of the certificates can be chosen as small as possible to avoid the necessity of revocation lists).

The authentication of the mobile device works similar. The combination TC features with standard authentication and authorisation methods allows for a multitude of implementation variants of the POS scenario. A high level protocol realising it is shown in [3].

Conclusions

We have shown that trusted computing can be efficiently used to carry out one time transactions like single purchases. Key features of the presented concept are that it uses and adds security to existing infrastructures and that it establishes trust between only indirectly connected communicating entities. This is a novel usage of trusted computing outside the domain of digital rights management. The advantage of this usage scenario is that no high economic values are depending on the security of single instantiation of a specific trusted platform. While in the case of DRM for broadcast media a broken TPM means uncontrolled proliferation of the no longer protected good, individual broken devices cause only damage which is limited in time and space.

If a broken trusted platform is detected for example by comparison of accounting data with actual purchases the respective privacy CA will no longer certify AIKs of the detected illicit device. This requires a log in the privacy CA.

Many other application scenarios of TC can be envisaged, some of which are outlined in [3]. We believe that the mobile domain is very amenable to novel applications and solutions based on TC and that TC can become an enabler for new service architectures and mobile business models which bring together stakeholders who formerly were not able to establish technical trust and corresponding business relationships.

REFERENCES

[1] Li, F., Whalley, J., *Deconstruction of the telecommunications industry: from value chains to value networks*. Telecommunications Policy 26 (2002) 451–472

[2] M. Marhöfer, A. U. Schmidt, *Trusted Integration of Mobile Platforms into Service-oriented Networks*, 11th German-Japanese Symposium "Security, Privacy and Safety in the Information Society" Tokio, Japan, 13th-16th September 2005.

[3] N. Kuntze, A.U. Schmidt, *Transitive trust in mobile scenarios,* to appear in Proceedings of the

International Conference on Emerging Trends in Information and Communication Security (ETRICS 2006), LNCS, Springer.

[4] Michael J. Nash, Keith R. Poland, Some conundrums concerning the separation of duty. Proceedings of the IEEE Symposium on research in security and privacy, 7-9 May 1990, Oakland, CA

[5] R. A. Botha, J. H. P. Eloff, Separation of duties for access control enforcement in workflow environments. IBM Systems Journal, 40 (2001) 666-682.

[6] Brickell, E., Camenisch, J., Chen, L.: Direct anonymous attestation. In: Proc. 10th ACM Conference on Computer and Communications Security, Washington DC, ACM Press, 2004

[7] Camenisch, J.: Better Privacy for Trusted Computing Platforms. In: Proc. 9th European Symposium On Research in Computer Security (ESORICS 2004), Sophia Antipolis, France, September 13-15, 2004, Springer-Verlag, 2004, pp. 73–88

[8] Bruschi, D., Cavallaro, L., Lanzi, A., Monga, M.: Attacking a Trusted Computing Platform. Improving the Security of the TCG Specification. Technical Report RT 05-05, Dipartimento di Informatica e Comunicazione, Università degli Studi di Milano, Italy, 2005

[9] Jøsang, A., Gray, E., Kinateder, M.: Simplification and Analysis of Transitive Trust Networks. Web Intelligence and Agent Systems, to appear. http://security.dstc.edu.au/papers/JGK2005-WIAS.pdf

[10] Trusted Grub: http://www.prosecco.rub.de/trusted_grub.html

[11] Trusted Computing Group: TPM Specification Version 1.2 Revision 94. http://www.trustedcomputinggroup.org

[12] TrouSerS, the open source TCG software stack http://trousers.sourceforge.net/