Transitive trust in mobile scenarios

Nicolai Kuntze, Andreas U. Schmidt

Fraunhofer Institute for Secure Information Technology SIT,
Rheinstrasse 75, 64295 Darmstadt, Germany,
{Nicolai.Kuntze,Andreas.U.Schmidt}@sit.fraunhofer.de,
WWW home page: www.sit.fraunhofer.de, www.math.uni-frankfurt.de/~aschmidt

Abstract. Horizontal integration of access technologies to networks and services should be accompanied by some kind of convergence of authentication technologies. The missing link for the federation of user identities across the technological boundaries separating authentication methods can be provided by trusted computing platforms. The concept of establishing transitive trust by trusted computing enables the desired cross-domain authentication functionality. The focus of target application scenarios lies in the realm of mobile networks and devices.

1 Introduction

Current information technology imposes on users a multitude of heterogeneous authentication mechanisms when they want to access networks, services, or content. The technical access channels to these desiderata are, however, undergoing a continual process of convergence. The mobile domain provides a striking example [1,2]. The access to services through mobile devices shows a trend to become network-agnostic. Driven by the horizontal integration of technologies, users will soon be able to consume services seamlessly from a single device via a variety of channels and transport methods such as 2G, 3G, WLAN, Bluetooth, WiMAX, MobileIP, or the upcoming Zigbee. Accordingly, end users' attention will shift away from the pricing of bandwidth to that of content and services. Custom must then be attracted by offering applications and content with good price to quality ratio. Little room is left for returns generated by charging for network access and data transport. Business models necessarily undergo drastic changes, of which the mushrooming of virtual network operators is the salient epiphenomenon. Research has long forseen this evolution toward 'value networks' [3,4].

Thus, information networks are becoming ever more service oriented. On the application layer, identity management (IDM), as embodied, e.g., in the Liberty alliance standard suite, has proved to be a successful foundation for the user-centric integration of service access [5]. Mobile networks with millions of users and even more identities are already using IDM for essential services like roaming [6]. Yet, arguably, these top-level methods require infrastructural support of some kind [7]. In particular, it is desirable to overcome the boundaries between logically, technically, or even physically separated domains and their respective authentication methods. This signifies a second layer of technological convergence, namely convergence of authentication methods and the domains of trust defined by them. This is the subject matter of the present paper.

We argue that trusted computing (TC) can be a means to the above mentioned ends. In fact, two systems or devices can assure each other of their being in a trustworthy state through TC methods like direct attestation. If the devices carry credentials from various trust domains, they can then use TC-secured communication to exchange them. This assignment of credentials by trustworthy transmission between carriers yields *transitive trust relationships*. This allows for the mediation of trust between domains and user or device identities, and in fact, some of the concepts we present are rather similar to logical identity federation. However, transitive trust by TC enables the traversal of authentication domains hitherto separated by technical or even physical boundaries. The concept of transitivity of trust relationships was recently analysed in [8].

The paper is organised as follows. Section 2 explains the basic notions behind transitive trust, in particular the three most primitive operations supported by it. The exposition, while theoretical, is not completely formalised in view of the intended application scenarios. Three of the latter scenarios are described in ascending level of detail in Section 3.

Not by coincidence are these applications chosen from the mobile realm. In fact we show that mobile devices equipped with TC are not only good carriers for credentials but also excellent links between trust domains, when applying the methods of transitive trust. As will become clear from the few scenarios we consider, potential business models, enabled by transitive trust, abound. Needless to say, the newly conceived trust relationships that we describe in concrete business scenarios must be supported in the real world by contractual relationships.

2 Transitive trust by trusted platforms

A completely formalised definition is outside of the scope of the present paper, since we aim at rather specific application scenarios. Nevertheless we want to provide a theoretical descriptions that allows to assess the generic character of the transitive trust relationships supported by trusted platforms, i.e., systems secured by TC as described below. A more formal treatment, e.g., along the lines of [8,9] or [10], is certainly possible. Yet, it would not contribute much to the present topic since we are more interested in pinpointing the properties and functionalities of trusted platforms involved in the establishment of transitive trust.

We use a simple model for actors in trust domains consisting of trust *principals* and *agents*. Trust principals are the subjects defining an authentication domain by issuing credentials to users or enrolling them to their devices. They control domain membership and applicable authentication methods, and therefore define a domain of trust like an identity provider. Trust principals are denoted by capital letters A, B, C, ... Agents asking for access to services provided in a certain domain are denoted by a, b, c, ... The notion of agent signifies *classes* of individuals, i.e., groups of agents who enjoy the same access rights in a certain application context when authenticated using their respective (individual) credentials. A subgroup of agents is written as $a' \subset a$ as usual. Our terminology is different from that in [9] in order to clearly separate the party issuing an authentication request (the agent) from the one answering it (the principal).

Credentials $\gamma_{a,A}$ are objects or data which authenticate agents a with respect to a principal A. We do not specify the particular kind of credentials used, nor the accompanying authentication methods. This notion is very generic and comprises classical examples like SIM/USIM, Hardware tokens, Smartcards, PKI-based certificates, PIN/TAN-based methods, or even personal credentials, e.g., Machine Readable Transfer Documents or a health (professional) card.

It should be clear that the overall security of the authentication assertions of transitive trust that are described below depend on the 'weakest link' in the trust chain. These assertions can in particular not be stronger than those provided by the original credentials. Furthermore, the trust scope implicated by a successful authentication, i.e., the specific type of trust assumed in a given principal-agent relationship, may vary from domain to domain. As already mentioned, risks arising from these complexities must be assessed and mitigated in the context of the specific application scenario at hand. Common instruments for that are contracts between principals and their agents and bridging contracts between principals.

2.1 Trust credentials

Credentials that can be constructed basing on the functionalities of a trusted platform module (TPM [11]) play a special role in our concept. 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, Direct Anonymous Attestation (DAA), a method put forward in [12] and specified by the trusted computing group (TCG), 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 [13].

Although certain flaws are known in the TCG standards (e.g. [14] 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 scenarios below.

Using the described functionality, a trusted system, viewed as an agent a, can establish what we call a *trust credential* τ_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 DAA.
3. That an existing credential $\gamma_{a,A}$ 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 direct attestation 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.

We now describe three basic, independent operations for creating trust between agents and principals. These methods represent the essence of transitive trust enabled by trusted platforms. They all rely on *referral trust* in the parlance of [8]. 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.

2.2 Restriction

By the method of restriction, a subgroup of agents $a' \subset a$ belonging to the authentication domain of principal A can be defined. Agents of class a authenticate themselves in the conventional way associated to their credential $\gamma_{a,A}$. This establishes an authenticated channel, over which agents of subclass a' transmit an additional trust credential $\tau_{a'}$ identifying them as members of a'. Since by this method the trust and original credentials are used independently, only assertions 1. and 2. are needed.

The additional security and in effect higher trust in agents of a' provided by them allows to ascribe to a' more service access rights than to a-agents. In particular, the integrity of client software can be attested by 2. Those clients can access content or services only available to the privileged subgroup. This is in fact the classical scenario used to enforce copyright protection through digital rights management (DRM). A higher security level is provided by restriction in a very generic way. The possibility for A to check the consistency of the trust credential $\tau_{a'}$ with that of $\gamma_{a,A}$ makes at least the subclass a' more resilient against cloning attacks on the credential $\gamma_{a,A}$. This kind of attack is not uncommon in the mobile sector [15].

This raised resilience against cloning is the main reason why the usage of a trust credential is advantageous for the definition of the subclass a'. The latter definition can be implemented in various ways. The first-best approach is restriction under the authority of the principal. She can manage access control lists based on *individual* trust credentials identifying a single TPM. Or, e.g., she can use individual trust credentials to establish a secure channel with a'-agents and distribute a shared secret to them. This secret can reside in the part of the system protected by the TPM and thus become part of $\tau_{a'}$. In turn it may be used in subsequent authentication requests toward A, keeping an agent's individual identity secret.

A proper choice of enrolment method and time for the trust credential is essential for the validity of the additional trust provided by the restriction operation. If the credentials γ and τ are impressed on the agents independently of each other, i.e., not both under the control of the principal A, then, e.g., resilience against cloning attacks is restricted. Since A cannot associate the two credentials belonging to an individual agent, she can at best avoid to grant two agents with identical γ service access by using a first-come-first-served approach. It is possible to improve on this by forcing an activation of $\tau_{a'}$ at an early stage, e.g., the time of roll-out of a mobile device. Higher cloning-resilience can only be achieved if the principal individualises both credentials and controls their deployment to the agent.

It may be more the rule than the exception that the trust credential $\tau_{a'}$ provides stronger authentication than the original one $\gamma_{a'\subset a,A}$. Conventionally, τ would then be the preferable credential to authenticate agents of class a' with. It is essential for the understanding of the present concepts to notice that this is often not practical. Namely, the communication channel through which τ is conveyed to the principal is only available after authentication by γ. A paradigm is the access to mobile networks as described in section 3.1.

2.3 Subordination

By subordination an agent a in principal A's domain can enable the access to this domain, or certain services of it, for another agent a'. By this, a' is effectively included in A's domain of trust, respectively, A's domain is extended to a'. As for restriction, a authenticates himself using a generic credential $\gamma_{a,A}$ and then produces a specific trust credential σ_a identifying those agents of A's domain who are allowed to dominate certain other agents. The subordinated agent a' shows a trust credential $\sigma_{a'}$ to a, who in turn mediates the access to A's services, either by forwarding authorisation requests, or granting them himself. Furthermore, the authentication of a and a' can also be mutual rather than one-sided.

Implementation variants of this operation and authorisation based on it are manifold, despite its simplicity. The most restrictive approach would be to use the secure communication channels between a and a' (mutually authenticated by σ_a, $\sigma_{a'}$), and a and A to forward every single authorisation request from a' to A including the trust credential $\sigma_{a'}$. Independently of the degree to which A takes part in authorisation, the act of authentication for subordination is generically between a' and a. Nevertheless, in many scenarios $\sigma_{a'}$ is controlled and enrolled by A, and the principal can in implementation variants also partake in authentication, e.g., by facilitating steps in a challenge-response protocol.

If genuine trust credentials are used for subordination, the operation employs only TPM functions 1. and 2. above. TPMs provide user functions for the revocation of keys, which is a point of failure in this case. Thus one might use a dedicated credential $\gamma_{a',A}$ for subordination. Such a credential should then live in the trusted part of the subordinated system and be secured in the authentication by function 3. to mitigate forgery.

A subordination scenario is outlined in 3.2.

2.4 Transposition

Transposition operates between the trust domains of two principals A and B. The authentication of an agent b of B's domain is mediated by an agent a of A's domain and the principal A. This can make sense for instance if direct communication between b and B is not possible as in the scenario of Section 3.3.

We assume that authentication of a to A is done as above. Trust credentials τ_a and τ_b are used for (mutual) authentication of b to a (or between them). Here, the third TC function of τ_b is used to prove the integrity of a credential $\gamma_{b,B}$ with which b is ultimately authenticated with respect to B. The generic situation for the latter authentication is as follows. The credential $\gamma_{b,B}$ is forwarded to A. This bears the assurance that an authentic (by $\gamma_{a,A}$) and untampered (by τ_a) agent has handled the latter credential. In effect a establishes a trusted path for the transmission of $\gamma_{b,B}$. Whether or how $\gamma_{b,B}$ is transferred from A to B to finally authenticate b depends on communication means and contractual relations. The transposition concept leaves this open.

Again, transposition can be implemented in numerous variants. In particular, part or all of the functionality necessary for authentication of b can be deferred to A or a. From B's perspective, efficiency gains by such an outsourcing or even decentralised approach to authentication must be balanced with the protection of secrecy of his business data and processes, which, to a certain extent have to be turned over to A.

On the other hand, in the generic transposition operation where $\gamma_{b,B}$ is forwarded to B who in turn completely controls the authentication of b. Then, additional cryptographic means can be applied to render any sensitive information about the relationship of b and B inaccessible to a and A. In particular, B might want to keep his agents anonymous to A, and even the mere size of B's domain of trust might be an informational asset worth of protection.

3 Scenarios

This section outlines three concrete application scenarios of economical relevance, corresponding to the three operations explained above. The first two are sketched on a rather high level, while the third and most complex one is used to detail processes and protocols. A detailed description of the first two scenarios would be very similar.

3.1 Functional discrimination of mobile devices

As already said, the paradigm for restriction scenarios is DRM. We want to pursue a slightly different direction and take a look at the relationship between network operator and customer in the mobile domain. The standard form of customer retention exerted by the mobile network operator (MNO) is SIM-lock, a crude form of functional restriction of mobile devices bonding mobile devices to SIMs of a certain MNO. Based on transitive trust restriction, a finer grained functional discrimination of mobile devices becomes possible. Depending on the device vendor's and MNO's business models, various client functions of the device can be restricted to certain, more or less privileged customer groups. The management of mobile devices, of which functional discrimination is an important instance is viewed by the industry as a fundamental application area of TC [16].

A multitude of benefits accrue to MNO and customer in this kind of scenario. First, it is cost-efficient to produce a single product line with many appearances to the end-user, rather than marketing a multitude of makes and models as customary today. Second, the up- and downgrading of functionalities can be implemented dynamically, without physical access to the device. To the user, the relative seamlessness with which device control operates is an ergonomic benefit and allows for better customisation and even personalisation.

The efficient means to implement functional restrictions of mobile devices is provided by the trusted boot process and operating system of the trusted platform it represents. Thereby, the trust credential can attest two properties via DAA. First, that the device belongs to a certain, restricted group defined explicitly or implicitly by a list of enabled functions. Second, that the device actually is in a state where only the allowed functions are enabled. The set of functions to be managed could be pre-configured and the dynamic control effected via simple changes of parameters, e.g., for values in PCRs.

The enforcement level of this approach is stronger as compared to SIM-lock precisely because the trusted platform's base operation software is tamper resistant. Based on this assurance, the MNO can deliver specific services or content only to the restricted group privy to it. Thus functional restriction provides the foundation on the client side for further service discrimination, policy enforcement, and DRM proper.

As a simple instance using the transitive trust restriction operation, a prepaid mobile phone can be implemented. The phone carries in its trusted storage area a running total which is decremented by a trusted software. While the initial access to the mobile network is still established using SIM authentication, DAA and the trust credential then yield assurance to the MNO that the running total is nonzero, upon which access to the network's communication services can be granted. This releases the MNO from operating (or paying for) a centralised accounting.

3.2 Bonding of mobile accessories

For the mobile domain, an application of subordination which suggests itself is to extend the authentication of devices toward an MNO to devices not equipped with SIM cards or even physical access to the mobile network. A commercial application is the extension of SIM-lock to such devices. For the purpose of customer retention, such a scheme can for instance be combined with loyalty programmes. Just as SIM-locked mobile phones are highly subsidised, an MNO can give away technical accessories such as digital cameras, media players, or high quality headsets. The functioning of those subordinated devices is then dependent on authentication toward a mobile device or any device in a specific MNO's network.

In effect, the accessories can be given away for a very low price or even for free on the condition that they work only within the subsidising MNO's network. The devices are bonded to the MNO. As an additional benefit for the MNO, the traffic generated by subordinated devices is bound to his own network (as traffic volume is a traditional economic value indicator for MNO businesses). Of course, advanced service provisioning can be based on accessory bonding, e.g., the MNO or another provider can offer storage, organisation, and printing services for photographs taken with a bonded camera.

3.3 Point of sales

We now come to scenarios employing the transposition operation, and here present the related technical processes and communication protocols in some detail.

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 DAA 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. After the user selects a good and confirms his choice at his device, signed price and payment processing information is transferred to the MNO. After verifying the signatures and optionally informing the good's vendor and a payment service provider, 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.

The benefits for the vendor that arise in this scenario basically stem from the transitive trust relationship that is mediated between MNO and POS by the mobile device. That it is economically attractive is a view shared by prominent market researchers [1]. The scenario entails in particular that no network communication is required during the initiation of a trusted session, that no transaction

[1] As John Curtis, head of the department Information, Communications & Entertainment of KPMG Germany put it: "Doch permanent subventionierte Handys auf den Markt zu werfen, bringt langfristig keinen Geschäftserfolg. Sinnvoller ist es, sich mit Hilfe attraktiver konvergenter Dienstleistungen [. . .] eine stabile und loyale Kun-

data needs to be stored in the POS, and that, ultimately, the POS does not need to be equipped with networking capabilities — at least for the sales process. In this way the MNO provides payment services as well as authorisation control for the vendor. This requires little more than a TPM and a short-range communication module in the vending machine. In extended service scenarios, the customer's mobile devices can as well be utilised to transfer valuable information to the POS, e.g., updated price and commodity lists, or firmware.

A similar example regards home automation and lets a user and her mobile device become part of the maintenance service of, say, the heating system of her home. Based again on their respective TPMs, heating system and mobile device establish a secure communication channel to exchange maintenance data, or data used for metering. This can be done both at specific user requests or even seamlessly during normal operation of device and heating system, every time the machine-to-machine communication module of the device gets in the range of the one in the heating system. In this way, the mobile device can notify user and a maintenance chain about necessary repairs and also support accounting and billing. Here, a trusted computing approach not only ensures the protection of personal data, it also enables a simple means of remote maintenance and home automation in non-networked homes by efficiently utilising the mobile network.

Returning to the POS scenario, we now describe one possible implementation in more detail. We concentrate on the authentication processes and leave selection, purchase, and payment aside.

The variant of transposition we consider is that of *maximal mutual trust*. That is, both principals A, the MNO, and B, the POS' owner, can trust the involved agent of the other domain, i.e., the POS b, respectively the mobile device a. The raised level of security ensuing from this may be desirable in particular from B's perspective, depending on the sensitivity of business data handled by a and A as mediators, for instance if accounting and charging services of B are transferred to A. The process to achieve this kind of transposition can be divided into two principally independent steps.

A) Establishment of trust of a and A in agent b.
B) Establishment of trust of b and B in agent a.

These two steps are in fact equivalent to two subordination operations with exchanged roles. A sequence diagram for both steps is shown in Figure 1. Note that A) and B) can be interchanged or even overlap.

denbasis aufzubauen. Damit wird man für Werbekunden und Partner im digitalen Handel attraktiv und eröffnet sich neue Einnahmequellen [. . .] Verrechnungsmanagement wird deshalb künftig zu einer Schlüsselkompetenz." (Throwing subsidised handsets on the market is not a sustainable strategy for success. It makes more sense to build a stable and loyal customer base with attractive and convergent services [. . .]. In this way, new revenue sources open up and attractiveness for advertising customers and partners in digital trade is increased [. . .] Charging management will therefore be a future key competency). [KPMG Germany press release, 19th March 2006. http://www.kpmg.de/about/press_office/13609.htm]

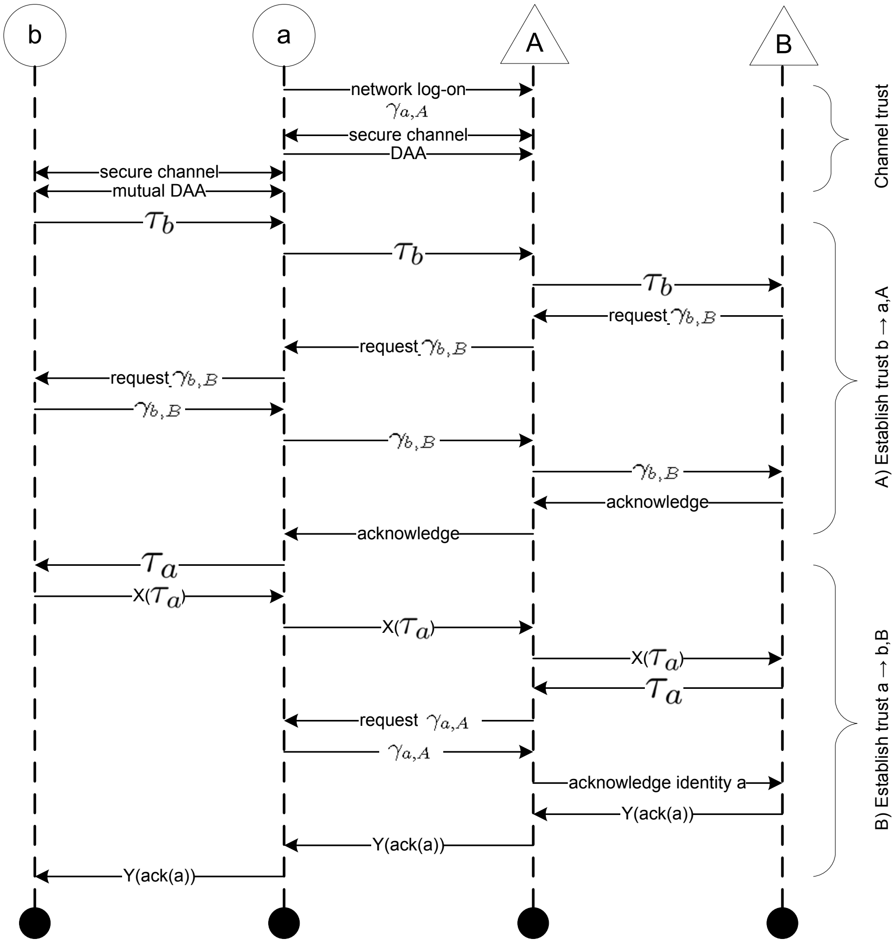

Fig. 1. Sequence diagram for the transposition operation from POS b via mobile device a, MNO A, to POS owner B. The notation $X(\cdot)$, $Y(\cdot)$ means protection by secrets X, Y shared between b and B.

The two main steps must both be preceded by an establishment of a secure communication channel between b and a and between a and A, respectively. For the latter, the usual log-on of the mobile device to the network based on $\gamma_{a,A}$ is augmented by attestation of the trusted platform a via DAA toward A over a secured channel based on, say, encryption on the transport layer. For the former, mutual platform attestation over an encrypted channel is carried out between b and a.

A) The trust credential τ_b of b is passed on to B, attesting to B that there is one of his untampered POS down the communication line. B then requests and receives proper authentication from b with $\gamma_{b.B}$. The underlying assumption that B can associate trust and generic credentials of agents in his domain is a central anchor for trust in the present variant of transposition. In effect B is an identity provider for trust credentials of his domain.

B acknowledges successful authentication of b to A who passes it on to a. The trust relationship between the two principals and A and his agent a assures the latter two actor of the authenticity of b.

B) Agent a initiates his authentication toward B and b by handing his trust credential to b. This credential cannot be utilised by b directly to authenticate a, but is rather used as a pledge which is then redeemed by b at the principals. To that end, b uses some secret X he shares with his principal to protect τ_a. X can for instance be established using the Diffie-Hellman method [17]. The protection of τ_a by X prevents a and A from tampering with the authentication request that is embodied in the message $X(\tau_a)$ passed on to B.

It should be noted that, apart from transport and addressing information, a and A need not know for which of A's agents authentication is requested, if X comprises encryption. Thus, the identity of the authenticated agent a could be kept secret from A in an advanced scenario. This could be used to protect the privacy of agents in the domain of A, e.g., with respect to their purchasing patterns.

B sends τ_a to A and with that requests from A the authentication of it. If A does not have a registry of all valid trust credentials in his domain or any other means of authenticating them then A has to exert a secondary authentication of a by the generic credential $\gamma_{a,A}$ (again assuming that association of τ_a to γ_a is possible). A acknowledges the identity of a to B. This acknowledgement is passed on from B to b, again protected by a shared secret Y to prevent tampering with it on its way.

4 Conclusions

We introduced the notion of transitive trust for a pragmatic purport. It is intended as a conceptual blueprint for the systematic construction of concrete, TC-based application scenarios. The examples exhibited show that transitive trust has a potential to be a fertile concept to that end. In particular, new application and business scenarios are enabled by transitive trust as well as more efficient and/or more secure implementations of old ones. Protection of privacy is not in opposition to the use of TC in those scenarios. It can, on the contrary, be supported in carefully constructed implementation variants of transitive trust.

Returning to the possibility of formalising our concepts, let us briefly sketch how they relate to those of [9, Sections 2 and 3]. Firstly, for a full treatment not only principals and agents must be taken into account, but also the trusted computing base (TCB) which can issue the assertions 1.–3. about an agent. That is, in every of the three operations the agent speaks for its TCB in transmitting, as a channel, the trust credential and the assertions to the principal. In this way the TCB *hands off* authority to the agent to speak on its behalf. It must then be shown that this procedure establishes a credential (in the sense of [9]) for the agent. The joint authentication with respect to the agent's principal using the ordinary and the trust credential is stronger than the original one. In the case of subordination for instance, it allows a secondary hand-off from the subordinated to the subordinating agent, allowing the principal to infer the authenticity of the former. This proceeding would establish a formal description of the most basic technical operation used in all three transitive trust operations. It must be accompanied and complemented by a formalisation of transitivity *per se* on a higher conceptual level, as in [8]. Both formal levels must be coherently interweaved to produce a full formalisation of our concepts. This task, which is well beyond the scope of the present paper, should be treated elsewhere.

Economically the prospect to federate the identities of millions of subscribers of mobile networks with other providers of goods and services, is rather attractive. As said, transitive trust is very similar to identity federation, but TC has additional application potential due to the possibility to transgress boundaries of authentication domains that are closed to IDM on the application layer. The standard way of remote attestation using trusted computing [11] is analogous to classical ID federation in that every TPM is assigned to the domain of trust of its manufacturer *and* the device vendor *viz* OEM via the platform endorsement keys. Thus every TPM-equipped device has an unique identity which can be resolved by these principals who in turn can act as identity providers. Demand for provisioning of identity federations in the mobile domain is confirmed by the recent market survey [1, Section 6] — whether MNOs or other parties are ready to take on that role remains to be seen.

While this way of TC-based remote attestation does not provide for anonymity the novel feature DAA in the TCG specification version 1.2 is qualitatively very different. Resting on involved methods of zero-knowledge proofs [13] it enables a trusted platform in principle to convince the outside world of any of the assertions 1.–3. without revealing its identity, with cryptographic security. In particular, this could be used to issue trustworthy assertions about the membership of a trusted device in a certain domain while staying fully anonymous. Namely, the system state asserted by DAA can include information about the presence of a generic domain credential on the device, without revealing it. This potentially opens up a new area of research and applications centred around methods of anonymous, privacy-protecting, methods for the establishment of trust.

A particular trait of transitive trust mentioned above is the enabling of de-centralised authentication through the trusted agents. A benefit of such approaches can be enhanced resilience and availability of service access. They can also be a base for de-centralised authorisation and ultimately de-centralised business models, such as super-distribution of virtual goods from agent to agent, cf. [18,19,20].

As a further example, in an advanced scenario for the restriction operation, it can be envisaged that a group of agents defines itself in a manner similar to building a web of trust [21] of which PGP is a well-known instance [22]. To that end, the transposition operation could be used to establish mutual trust between agents, extend it to trust paths in a community, and eventually define the subgroup as the resulting web of trust.

References

1. KPMG: Consumers and Convergence - Challenges and opportunities in meeting next generation customer need. 2006. `http://www.kpmg.de/about/press_office/13611.htm`
2. Marhöfer, M., Schmidt, A. U.: Trusted Integration of Mobile Platforms into Service-oriented Networks. Contribution to the 11th German-Japanese Symposium "Security, Privacy and Safety in the Information Society" of the Münchner Kreis, Tokio, Japan, 13th-16th September 2005
3. Li, F., Whalley, J.: Deconstruction of the telecommunications industry: from value chains to value networks. Telecommunications Policy **26** (2002) 451–472
4. Ulset, S.: Mobile virtual network operators: a strategic transaction cost analysis of preliminary experiences. Telecommunications Policy **26** (2002) 537–549
5. Clauß, S., Köhntopp, M.: Identity management and its support of multilateral security. Computer Networks **37** (2001) 205–219
6. Rannenberg, K.: Identity management in mobile cellular networks and related applications. Information Security Technical Report **9** (2004) 77–85
7. Lopez, J., Oppliger, R., Pernul, G.: Authentication and authorization infrastructures (AAIs): a comparative survey. Computers & Security **23** (2004) 578–590
8. 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`
9. Lampson, B., Abadi, M., Burrows, M., Wobber, E.: Authentication in distributed systems: theory and practice, ACM Transactions on Computer Systems (TOCS) **10** (1992) 265 - 310
10. Maurer, U.: Modelling a Public-Key Infrastructure In: Proc. 1196 Symposium on Research in Computer Security (ESORICS' 96), E. Bertino (Ed.), Lecture Notes in Computer Science, vol. 1146, Berlin: Springer-Verlag, 1996, pp. 325–350
11. Trusted Computing Group: TPM Specification Version 1.2 Revision 85. February 2005. `www.trustedcomputinggroup.org`
12. Brickell, E., Camenisch, J., Chen, L.: Direct anonymous attestation. In: Proc. 10th ACM Conference on Computer and Communications Security, Washington DC, ACM Press, 2004
13. 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
14. 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
15. Cheney, P.: How a terror group cloned Ted Rogers' cellphone. The Globe and Mail, Toronto, Canada, December 17, 2005
16. NTT DoCoMo, IBM, Intel Corporation: Trusted Mobile Platform Protocol Specification Document — Revision 1.00. 04/05/2004. `http://www.trusted-mobile.org`

17. Diffie, W., Hellman, M. E.: New Directions in Cryptography. IEEE Transactions on Information Theory **22** (1976) 644–654
18. Schmidt, A.: Incentive Systems in Multi-Level Markets for Virtual Goods. In: [23] 134–141
19. Schmucker, M., Ebinger, P.: Alternative Distribution Models based on P2P. In: [23] 142–149
20. Rajasekaran, H.: An Incentive Based Distribution System for DRM Protected Content Using Peer-to-Peer Networks. In: [23] 150–156
21. Khare, R., Rifkin, A.: Weaving a web of trust, World Wide Web Journal **2** (1997) 77–112
22. Zimmermann, P.: The PGP user's guide, the International PGP Home Page, October 1994. `www.pgpi.org`
23. Nesi, P., Ng, K., Delgado J. (Eds.): Axmedis 2005, Proceedings of the 1st International Conference on Automated Production of Cross Media Content for Multi-Channel Distribution, Volume for Workshops, Industrial, and Application Sessions, Firenze University Press, 2005.