Analysis of Trust Properties and Related Impact of Trusted Platforms

Siani Pearson, Marco Casassa Mont, Stephen Crane
Trusted Systems Laboratory
HP Laboratories Bristol
HPL-2005-55
March 18, 2005*

trust, dynamic trust, privacy, trusted platforms, sociology, prime

This paper draws a distinction between persistent and dynamic trust and analyses this distinction within the context of trusted computing technology.

* Internal Accession Date Only
Approved for External Publication

Analysis of Trust Properties and Related Impact of Trusted Platforms

Siani Pearson, Marco Casassa Mont and Stephen Crane

Trusted Systems Laboratory, Hewlett Packard Research Labs, Filton Road, Stoke Gifford, Bristol, BS34 8QZ, UK
{Siani.Pearson, Marco.Casassa-Mont, Stephen.Crane}@hp.com

Abstract. This paper draws a distinction between persistent and dynamic trust and analyses this distinction within the context of trusted computing technology.

1 Introduction

This paper demonstrates how trusted computing can provide both persistent and dynamic trust, and assesses the role of both of these within the context of on- and off-line trust provision. Specifically, it provides:

- background analysis of trust (in the form of a summary of models of trust emerging from the social sciences, cross-disciplinary backgrounds, and consideration of models of trust specifically related to the e-commerce domain)
- reasons why companies might want to be associated with trust
- contrast between persistent v. dynamic, and social v. technological trust
- linkage of this analysis to the real world deployment of Trusted Platforms
- explanation of how Trusted Platforms provide the basis for assurance and assessment of the trustworthiness of services and systems

2 Analysis of Persistent and Dynamic Trust

In this section we analyse the complexities of the notion of trust, and draw a distinction between social and technological trust, each of which may be further subdivided into persistent and dynamic trust.

2.1 Trust: a complex notion

It is hard to pin down the meanings of many words. *Trust* is particularly tricky since it is not a simple notion. Typically, we think in terms of 'entity A trusting entity B for something', which is complex not least for the following reasons:

Not always transitive. If A trusts B and B vouches for C, does A trust C in this case? In other words, is trust a transitive notion? The answer is "not always", although it can be under specific circumstances.

Dynamic. Furthermore, trust is dynamic rather than static – there can be differing phases in a relationship such as building trust, a stable trust relationship and declining trust. Trust can be lost quickly: as Nielsen states [15]: "It [trust] is hard to build and easy to lose: a single violation of trust can destroy years of slowly accumulated credibility".

Varying degree and scope. Trust levels differ both in the sense of varying degree and scope of trust: entities typically trust – or do not trust – each other to fulfill selected obligations or for a particular purpose, rather than for everything. On the other hand, trust in certain areas can transfer to trust more generally, as shown by big brands having an advantage when moving into new areas of business.

However, it is useful to have a succinct definition of trust if at all possible, particularly if you are claiming to provide an increased level of trust in something. If you look in any English dictionary, you will find (at least one use of the word) *trust* to be defined in similar wording to the following: "a firm belief in the reliability or truth or strength etc. of a person or thing". However, this is not the end of the story. To date, we have no universally accepted scholarly definition of trust, although "confident expectations" and "a willingness to be vulnerable" are usually viewed as critical components. Evidence from a contemporary, cross-disciplinary collection of scholarly writing suggests that a widely held definition of trust is as follows [20]: "Trust is a psychological state comprising the intention to accept vulnerability based upon positive expectations of the intentions or behavior of another". Yet this definition does not fully capture the dynamic and varied subtleties considered above.

In general, we can conclude that it is difficult to define trust because there are different facets of trust. In the case where "trust" is applied in an on-line business context relating to people having confidence in enterprises, these facets include:

A technological basis that is the focus of this paper

A contractual side that includes both laws and underwriting or contracts

Customers' image that is built up via previous interactions with a company, brand image, publicity, etc.

In the following sections some of the major attempts to provide social theories of trust are considered and also how such reasoning has been applied more generally and in particular to the e-commerce domain. Such background analysis supports further consideration in Section 3 of the extent to which trust is increased by using trusted computing technology.

2.2 A social science view of trust

There are two main approaches to modeling trust in social science:

Temporal aspect. Trust has been considered to have a temporal aspect for a long time, ever since Aristotle stressed that friendship cannot exist without trust and that trust needs time. In the twentieth century, Niklas Luhmann viewed trust as a representation of the future. This is rather similar to the belief we hold when reasoning inductively that after experiencing a historical pattern of behavior, similar behavior can be expected in the future. For example, even without knowing the laws of physics, we trust that the sun will shine tomorrow in the same way that we have seen in the past.

Risk aspect. Social scientists have strongly stressed that risk is a central aspect of trust. For example, Luhmann believed that trust is an investment that involves risky preliminary outlay, where we accept risk in order to reduce the complexity of what we think about the world. In a similar vein, Georg Simmel believed that trust is an intermediary state between ignorance and knowledge, and the objective of gaining trust may fail [21]. Again, more recently, Nissenbaum in [16] stressed how trust involves vulnerability.

In addition, other interesting properties of trust have been suggested, such as:

Trust is necessary to allow us to function in the world. Luhmann in [13] belives that "trust is the glue that holds everything together in social life". This is because it reduces the complexity of how we think about the world around us so that we are only then capable of action and decision-making.

Trust is a learning process. In the personal sphere, trust is a historical process of individuals learning to trust others without having to give unlimited trust. However, according to [13], we do not really understand the process.

On a larger scale, social order is replaced by legal order. If you look in a dictionary, it is very probable that some of the definitions of trust will mention law, and this is no accident. Indeed, one reason why trust is necessary is because we do not have the resources on a personal level to analyze all the information that we need during our working life. Therefore, as societies become more advanced, such delegation increasingly requires trust in functional authorities and institutions, particularly in the area of knowledge (and technology). However, as mentioned above, if these institutions or powerful individuals (such as politicians) let down the people who trust them, there is the risk of a big change of attitude towards them. This leads us to the following point…

Trust can be fragile. People can tolerate some problems, but when a certain threshold is reached, trust can flip to distrust, and fixing the individual problem will not regain the trust that has been lost.

Trust may be irrational. Many social scientists believe that trust is not a matter of reasons and is unpredictable, in that it involves processes that cannot be calculated in advance. In fact, there is a difference of opinion

regarding whether giving precise reasons generates extra trust, or the opposite.

2.3 A cross-disciplinary view of trust: dynamic trust

Various people have tried to carry across the understanding gained via social scientists' models of trust to other domains. For instance, Rousseau *et al* have attempted to provide a cross-disciplinary model of trust [20]. They identify common elements of trust across disciplines, arguing that the definitions are variants of the same theme, and highlighting that trust is a process rather than being static. Trust is a process because it develops over time, and is *dynamic* insofar as multiple perspectives are necessary in order to explain different aspects of trust. For example, according to the economic view, trust is a cause, whereas the sociological perspective classifies it as an effect. From a social psychological perspective, in contrast, trust is considered as an interaction of the two. Although different scholars may particularly focus on one level of analysis (such as the individual, group, society, or firm level), this does not necessarily mean that they disregard the others.

The implications of this, according to Rousseau and colleagues, are that trust comes in different forms in different relationships. Even in the same relationship, the bandwidth of trust will vary depending on what development stage the relationship is in. This bandwidth ranges from deterrence-based trust, over calculus-based trust and relational trust to institution-based trust.

In conclusion, a multi-level analysis is important in order to try to understand the complex phenomenon of trust.

2.4 Analysis of on-line trust

How do the aspects of trust considered above relate specifically to trust in the domain of e-commerce, and are there additional features that relate to this area? In this section, some of the more important issues that relate to on-line trust are highlighted. For further general discussion related to trust in Information Technology, see [1;3].

Delegation of Trust to Authorities

As considered above, people cannot always be expected to work things out for themselves, particularly when technology is involved. They will look to someone to set an example (for example, the Consumers' Association, or role models). Due to a lack of information and time, together with the huge complexity of IT security, it is impossible for users of IT products to identify the level of security offered by individual products. They need to rely upon the reliability of a product being assessed by experts via evaluation and certification procedures, such as using criteria catalogues. Such criteria catalogues are widely used: for example, the 'orange book', ITSEC, Common Criteria in ISO/IEC.

Trust is extremely difficult to measure as it is fundamentally concerned with an individual's subjective feelings towards another entity. The authors of the ISO/IEC

standards believe that it is possible to measure, test and evaluate the security assurance of a product or system that is to be trusted. The idea is that if a certain assurance level is reached, it is worthy of trust being invested in it. However, although such evaluation and certification should guarantee security that can be quantified and verified, this will not necessarily serve in creating trust by means of reducing complexity in such a way that they can be understood and verified by the user [17].

The Relationship of Security to Trust

As we have already seen, there is a great deal more to on-line trust than security.

Some would argue that security is not even a component of trust. For example, Nissenbaum argues that the level of security does not affect trust [16]. She argues that security is increased in order to reduce risk, and not to increase trustworthiness. However, we would argue that, according to the situation, security may increase the level of trust, decrease the level of trust or indeed be neutral as Nissenbaum suggests. An example of increasing security to increase trust comes from people being more willing to engage in e-commerce if they are assured that their credit card numbers and personal data are cryptographically protected [7].

There can be a conflict between security and privacy. For example, some methods of enhanced authentication can result in privacy concerns (such as manufacturers' issue of identification numbers associated with networked devices). Indeed, in order for users to regard a computing system as trusted, it is important that increased security does not have an adverse effect on privacy.

An interesting point is that visual clues are lacking on-line, and this can have an effect on security and trust. For example, you lose having the immediate suspicion arising from seeing an adult frequenting a children's chat room, or a scruffy man selling an expensive car.

In conclusion, enhancing security will not necessarily increase trust, but it is an important enabler and can do so.

Components of Trust within the E-Commerce Domain

Recent research has been carried out to model trust within the e-commerce domain: the Cheskin Research and Studio Archetype/Sapient study [4] defines 'three key elements of web trust' from six 'primary components of the building block of trust' which break down into a total of '28 different ways in which trustworthiness may be established'. The six primary components are:

- *Seals of approval* information about companies that specialize in assuring the safety of web sites
- *Brand* importance of the company's reputation in choosing to do business with them
- *Navigation* the ease of finding what the visitor seeks
- *Fulfillment* the process one works through from the time a purchase process is initiated until the product is received
- *Presentation* ways in which the look of the site communicates meaningful information

- *Technology* the ways in which the site technically functions

Egger has carried out related work to identify many factors that mediate trust in e-commerce [5].

Friedman *et al* in [6] argue that it is not appropriate to use the language of trust in relation to people interacting with machines – rather, computer technologies provide 'suitabilities' that follow from features of the technology. How can we engineer technology that cultivates the conditions for trust online? The authors of that paper offer 10 trust-related characteristics of online interaction that can be taken into account when designing and implementing systems, of which reliability and security of the technology is just one. Others include anonymity, accountability, status cue markers, insurance and performance history and reputation.

2.5 Benefits for a company in being associated with trust

There are a number of reasons why a company would wish to be trusted, including the following:

Trust is a better strategy than power games. Kumar in [12] argues that the power balance between manufacturers and retailers has shifted over the last few years and has introduced new trust relationships. The traditional model of fear and intimidation, resulting in powerless but resisting victims, is held to be a poor long-term strategy that is unsustainable. Indeed, trying to build a position of power is incompatible with the approach of building trust. Alternatively, if companies decide to try to build trust with customers, both parties are likely to be more committed to the relationship and less likely to seek an alternative. In addition, they can both be more open and collaborative, exchanging more information and even sharing confidential information.

Brand image is associated with trust. Reputation is perhaps a company's most valuable asset [16]. Yet brand image suffers if there is a breach of trust or privacy. For example, banks can lose the reputation and trust of their customers by denying that phantom Automatic Teller Machine (ATM) withdrawals happen and the situation is made worse if they take the customers who have been affected by such withdrawals to court. Indeed, marketing seems to be an exercise of establishing trust. There are even cases where a different brand has been created to protect the trustworthiness of an existing one. Furthermore, reputation may be inherited from the "real world" into the online world, as with the Financial Times online service.

Brand image can be leveraged to sell trusted systems. Someone's willingness to carry out business with a Trusted Platform [23] will depend on the intended use and on the level of trust in the platform and the owner of the platform. In particular, the manufacturer of a platform is visible to a third party communicating with that platform. For example, Original Equipment Manufacturers (OEMs) can exploit their reputation for quality to make their platforms the preferred solution for business-critical services.

Note however that a company's reputation may not be justified: for example, you could build up a bogus reputation by targeting customers who have been given correct stock tips rather than those who were given wrong ones.

2.6 Assessing the impact of computer systems: social and technological trust

When assessing how trust may be increased by computer systems, we see it as helpful in distinguishing between social and technological trust. We also distinguish between persistent and dynamic trust, but there is not always a clear-cut distinction between these categories.

Social trust is trust which arises through social mechanisms, behaviour and values. This includes mechanisms such as:

- Legal contractual relationships
- Liability protection
- Sanctions
- Assurance (of technological mechanisms)

The examples given above are of infrastructural mechanisms that may vary over time, but in general are relatively static.

.Social trust can also arise through more dynamic means, which are liable to substantial change at short notice, such as:

- Brand image
- Look and feel
- Reputation
- History of interactions

Technological trust is elicited through technological means as opposed to social mechanisms. This may include:

(a) **persistent (static) trust in systems**. These are the underlying security infrastructure, well-known practices and the technological features corresponding to the static social mechanisms described above. They can include the following information:

- Certified hardware (for example, tamper-resistant hardware)
- Protocols
- Certified cryptographic techniques
- Assurance
- Other security features
- Audit and enforcement

(b) **dynamic trust in systems**. This is confidence that a particular environment or system state is trusted (at a given time, for a particular purpose). A system's behaviour can change according to a given context, and in particular if it has been hacked, and in some cases system behaviour can be driven by policies (dictated by

people, business needs or even malicious people) that change over time. For example, dynamic trust could be affected by the following information being divulged:

- A particular system has been compromised
- Spyware is running on a platform
- Software is in a certain state
- Policy enforcement has not been carried out

The focus of this paper is on a subset of social persistent and of technological persistent and dynamic trust. Both social and technological aspects of trust are necessary when designing online systems, quite apart from additional social guarantees of privacy and security. Trust in a computer system is underpinned by trust in individuals, in companies, and in brand names who vouch for the system. Protocols and services should be designed in such a way that everyone agrees that they are suitable and do not contain security weaknesses. In addition, you also need to know that a service executes properly, in a dynamic way. As we shall see in Section 3, you can do this by making measurements and checking the results of these measurements against values that have been created and signed by someone that you trust.

2.7 Summary

Trust is a complex notion and a multi-level analysis is important in order to try to understand it. There are many different ways in which on-line trust can be established: security may be one of these. When assessing trust in relation to computer systems, we have distinguished between social and technological trust, and between persistent and dynamic technological trust. All of these aspects of trust can be necessary.

Social trust in a hardware or software component or system is an expression of confidence in technological trust, because it is assurance about implementation and operation of that component or system. In particular, there are links between social trust and technological trust through the vouching mechanism, because it is important to know who is vouching for something as well as what they are vouching.

Mechanisms to provide dynamic technological trust need to be used in combination with social and technological mechanisms for providing persistent trust: as we shall see in the following section, if software processes provide information about the behaviour of a platform, that information can only be trusted if entities that are trusted vouch both for the method of providing the information and for the expected value of the information.

3 Deploying Trusted Technologies

This section considers how trusted computing meets the need for increased confidence in platforms via dynamic and persistent trust.

3.1 Trusted Platforms

Recently, the industry-wide Trusted Computing Platform Alliance (TCPA) (and now its successor the Trusted Computing Group (TCG)) have been designing and developing specifications for computing platforms [22;24] that create a foundation of trust for software processes, based on a small amount of hardware. TCG has adopted all the TCPA specifications, and so for consistency we shall henceforth refer to TCG technology.

These organizations have published documents that specify how a Trusted Platform (TP) must be constructed. A Trusted Platform is a normal open computer platform that has been modified to maintain privacy using a special hardware device. It does this by reporting on the platform integrity and protecting private and secret data and identity information against subversion, and by attesting to the properties of the platform to a challenging party, i.e. to prove that it is a Trusted Platform while maintaining anonymity (if required).

Within each physical Trusted Platform is a Trusted (Platform) Subsystem, which contains a Trusted Platform Module (TPM), a Core Root of Trust for Measurement (CRTM) and support software (the Trusted platform Support Service - TSS). The TPM is a cost-effective security hardware device (roughly equivalent to a smart card chip) that is tamper-resistant and has cryptographic functionality. The CRTM is the first software to run during the boot process and preferably is physically located within the TPM. The TSS functions do not need to be trustworthy, but are required if the platform is to be trusted; they include functions needed for internal and external communication.

Trusted computing technology incorporates this standard hardware solution with little support from the operating system within the TCG-enabled computers which are already commercially available. Allied protected computing environments under development by certain manufacturers and open source operating systems such as Linux can support TCG facilities further and allow their widespread and convenient use. Intel's LaGrande hardware and chipset modifications [9] are designed to provide access-protected memory, and give greater support to secure operating systems. Microsoft's Next-Generation Secure Computing Base (NGSCB), formerly known as Palladium, is a secure computing environment that can run in parallel with Windows [14]. It provides additional protection beyond that defined by the TCG specification, including some protection of the video output and keyboard input. These different trusted computing implementations are all Trusted Platforms since they accord to the same underlying philosophy and basic principles of operation, as espoused in [23].

Trusted computing addresses some central concerns of people using PCs: it protects data that is stored on those machines (even while they are interacting with other machines over the Internet) and it aims to put everyone in the position where they can feel confident that they can:

- Protect their data
- Find out whether their platform is in a trustworthy state
- Have the means to decide whether it is reasonable for them to trust other platforms

As we have discussed above, trust involves a myriad of issues, all of which are important for business. TCG has taken the approach of addressing the issue of trust

(confidence) for businesses rather than just trying to improve the level of information security *per se* – although security improvements do form part of the solution. Trust is considered to be the fundamental concept in the business world and information security is an important (even vital) enabler.

In essence, a genuine Trusted (Computing) Platform is a platform that is trusted by local users and remote entities, including users, software, websites and third parties. In order to enable a user to trust a computing platform, a trusted relationship must be built between the user and the computing platform that can tell the user that an expected boot process, a selected Operating System (OS) and a set of selected security functions in the computing platform have been properly installed and operate correctly. The user then makes his or her own judgment whether or not he or she trusts the boot processing, OS and security functions.

3.2 How Trusted Platforms can provide persistent and dynamic trust

A Trusted Platform's Trusted Subsystem contains fewer functions than the Trusted Computing Base (TCB) (i.e. the set of functions that provide the security properties of a platform) of conventional secure computers, yet also contains some new functions. Rather than taking the usual approach of formal assessment and certification to provide evidence that a computer can operate securely if it is operated in certain tested configurations (see Subsection 2.4), the Trusted Subsystem provides a less formal means of showing that the TCB can be trusted in a wide variety of configurations: the Trusted Subsystem first demonstrates that it itself can be trusted and then demonstrates that the remainder of the TCB in a Trusted Platform can also be trusted. This involves certification from trusted entities that are prepared to vouch for the platform in various configurations, as described further below.

Broadly speaking, the view taken by the proponents of trusted computing (see for example [23]) is that we can think of something as being trusted if it operates in the expected manner for a particular purpose, or can be relied upon to signal clearly if it does not. The TCG definition of trust is that something is trusted "if it always behaves in the expected manner for the intended purpose" [23]. A similar approach is also adopted in the third part of ISO/IEC 15408 standard [10]: "a trusted component, operation or process is one whose behavior is predictable under almost any operating condition and which is highly resistant to subversion by application software, viruses and a given level of physical interference".

Within a platform, a trust hierarchy operates such that, for example, trusting that software running on the platform operates in the expected manner is underpinned by trust that the platform is at that time properly reporting and protecting information (dynamic trust), again underpinned by another layer of trust that that platform is capable of properly reporting and protecting information (static trust). The foundation of these multiple layers of trust is the "root of trust", in the form of an actual device that can be trusted intrinsically, and can report on other aspects of the system. Such a root of trust is missing from existing computers; there is no obvious way to check if the system is running correctly, as expected, and has not been deliberately or

inadvertently tampered with in some way. In TCG technology, the "root of trust" takes the form of the TPM.

We believe that categorizing trust in terms of the analysis presented in Section 2 helps in understanding how *Trusted Platforms* enhance trust.

Dynamic trust

In order to know whether a platform can be trusted at a given time, there are processes in a TP that dynamically collect and provide evidence of platform behaviour. These processes carry out measurement and provide a means for the measurement method to show itself to be trustworthy. When any platform starts, the CRTM (inside the BIOS or the BIOS Boot Block in PCs) starts a series of measurements involving the processor, OS loader, and other platform components. The Root of Trust for Reporting (RTR) – implemented as the TPM – is needed to dynamically store and protect against alteration the results of this measurement process, and to reliably cryptographically report the current measured values.

Persistent trust

The social basis for trust is that trusted third parties vouch (a) for the mechanisms that collect and provide evidence of dynamic trust as well as (b) that particular values of integrity metrics represent a platform that is behaving as it should. In essence, certain third parties are prepared to endorse a platform because they have assessed the platform and others are willing to state that if measurements of the integrity of that platform are of a certain value, it can be trusted for particular purposes.

In order to do the former, an endorsement key (in fact, an asymmetric key pair) is embedded into the TPM. The public endorsement key is signed by the manufacturer and published in the form of a digital certificate; the private endorsement key is known only to the TPM and is used only under the control of the owner of the platform. Social trust is used to recognise a specific genuine TPM: you trust a specific TPM because you can inspect the endorsement certificate, which is a trustable assertion by the manufacturer that produced it. Other elements of a Trusted Platform also have certificates: these vouch for the design of a TP, that a specific TPM was incorporated into a TP, that the design of the RTM and TPM meet the TCG specification and so on.

Proof that a platform is a genuine Trusted Platform is provided by cryptographic attestation identities. Each identity is created on the individual Trusted Platform, with attestation from a Public Key Infrastructure (PKI) Certification Authority (Privacy-CA), which is an organisation chosen by the user to return pseudonymous identities for the platform. To obtain attestation from a CA, the platform's owner sends the CA information (i.e. the certificates described above) that prove that the identity was generated on a genuine Trusted Platform. The platform owner may choose different CAs to certify each TPM identity in order to prevent correlation, or even use a protocol for creating identities that maintains anonymity.

Verifying trustworthiness
Both dynamic and static trust are involved in the decision by an enquirer (either local user or remote entity) whether a platform is trusted for the purpose intended by that enquirer: if the enquirer trusts the judgment of the third parties that vouch for the system components, and if the platform proves its identity and the measurements match the expected measurements, then the enquirer will trust that the platform will behave in a trustworthy and predictable manner. The platform reports information to the enquirer to enable that decision to be made [13], and analysing this requires intelligent application of cryptographic techniques; optionally, use could be made of a third party service to perform or help with this analysis.

By the following means an enquirer (whether local or remote) can decide whether the identity and software state of a platform can be trusted:

1. *The enquirer performs a cryptographic check of the user identity.* The enquirer's trust in the TP is based on trust in the Privacy-CA, since proof that a platform is genuine is given in the endorsement, platform and conformance credentials, which are sent to the Privacy-CA in order for a TCG identity to be issued. The enquirer decides whether to trust that the identity corresponds to a user of a genuine TP based on (1) whether the public Privacy-CA key corresponds to the signature on the user identity and (2) whether the enquirer trusts this Privacy-CA (it may be necessary to refer to a chain of CA certificates to decide this).
2. *The enquirer challenges the platform to obtain integrity metrics.* The enquirer verifies the integrity metrics certificates (that vouch for the expected integrity metrics of platform components) and compares these certified metrics to the reported metrics. If they match, and if the enquirer trusts the issuers of these integrity metrics certificates, the enquirer can trust that these metrics correspond to certified software.

3.3 Building upon platform trust to provide trusted services

The first Trusted Computing Platforms are already available for purchase and several more types will be appearing on the market throughout 2005. The first generation only provides a protected storage capability – it does not expose the full functionality described in the TCG specifications and can only be trusted to protect secrets in a certain way since there is no trusted boot process. Trusted Platforms that do provide the full functionality described in the TCG specifications provide roots of trust for systems, but even so they do not provide a complete trust solution: it is necessary to provide additional trust functionality built on top of them. In the short term, this must include security enhancements at the operating system level (as mentioned above), right up to trust management techniques (see for example [2;8]).

In this section we briefly assess how services can potentially be made more trustworthy using this technology.

Design features: maximising system trust
Measures which may be taken to maximise trust in a system include:

- Where appropriate, using only platform and software components developed under appropriate quality controls and from reliable vendors (for example, checkable via platform integrity metrics)
- Protecting the environment in which critical software resides, by using a TPM to protect functionality that must be trusted (e.g. for reporting on the trustworthiness of the system) and if possible by isolating critical software within OS compartments
- Using hardware components for critical functions, such as the TPM combined with smart cards

Checking trustworthiness of services (including trust management services)
Selected trusted software may be integrity checked to ensure that it is operating as expected and has not been modified or substituted in an unauthorized manner. This process would involve a TTP (usually the vendor of the software) publishing or otherwise making available a signed version of the integrity measurements that should correspond to genuine code. Upon boot, the code may be integrity checked and not be trusted for use if this integrity check fails. The integrity checking is performed as an extension to the platform integrity checking process [23;25], namely by measuring integrity metrics and comparing these with certified correct metrics. The code can be protected further by running within a protected environment such as the TPM (if there is sufficient space) or within a suitably isolated compartment.

The TPM can be used to provide protected storage for logs, digests, etc. via TCG protected storage mechanisms [25] so that such data cannot be interpreted by unauthorized entities. However, if these data are not stored within the TPM itself or within other tamper-resistant hardware, they will not be protected against unauthorized modification or deletion — although alteration to such data can be detected (for example by storing a digest within the TPM).

Trust sustainability
In order to maintain a trust relationship between service requester and provider over time, or at least until a service is completed, the service requester may periodically re-challenge the provider, as discussed above, to check the latest integrity metrics. The analysis required may need to take account of other factors too, such as time and history. Another approach is to have enhanced trusted software on the provider platform that monitors any changes to the platform state against pre-registered conditions provided by the service requester and notifies the requester if the changes impact the conditions [26]. This can be more efficient but requires additional initial setup and infrastructure; moreover, it can potentially lessen trust and security if the root of trust for reporting is no longer the TPM hardware chip.

3.4 How Trusted Platforms can underpin assurance

We are currently involved in deploying trusted technologies in order to provide more trustworthy privacy and identity management systems, within the context of PRIME, an EU Framework VI project on Privacy and Identity Management within Europe [19]. Not only are trustworthy systems more likely to be used, but by developing such

systems we can help maximize the number of EU citizens to use PRIME results, contribute to the other aims of Privacy Enhancing Technologies (i.e. the protection of human rights and the informational self-determination of the citizen) and tap into a business driver (because trustworthy systems foster e-commerce, and for the reasons considered in Subsection 2.5 above).

Trusted Platforms can underpin assurance in several ways:

- There is trusted certification material associated with Trusted Platforms, such as endorsement and platform certificates, as considered above
- The TPM can provide other signed information, in particular integrity metrics
- The TPM can certify other information, such as profiles (see [20])
- TCG non-migratable keys can be used in order to protect personal information and by binding such keys to attribute information within server-side certificates, services can be bound to authorised servers

Delegation of trust to authorities (cf. Subsection 2.4 above) is needed in order to provide the trusted certification material: TCPA/TCG provides conformance requirements such that the manufacturer must obtain a conformance certificate vouching for the correct security-related design and implementation of the TPM, using the Common Criteria. Protection profiles are defined by TCPA/TCG for both the platform and the TPM, and the manufacturer must create a security target that describes how an actual design meets the security requirements of the corresponding protection profile. The manufacturer must then present these documents to an accredited Common Criteria conformance laboratory in order to obtain conformance certification for the TPM and platform (only if the security target satisfies the protection profile).

Providing clients with greater assurance about services and fulfilment of contractual obligations

By use of TCG technology, a trusted server would have a strong identity (which could be pseudonymous if desired, but this would not usually be the case on the server side). Although we do not yet have this capability with current products, and furthermore ta level of infrastructure (including CAs) is required to provide such a strong identity, this infrastructure may be fairly readily provided in intra-company scenarios by binding TCG non-migratable keys to certificates issued by Corporate IT. This identity could potentially be specified in *Service Level Agreements* (SLAs), which would help with the automated fulfilment of contractual obligations. This would also be particularly useful in giving consumers greater trust about the services they are using. In the case of roaming scenarios for appliance or mobile phone providers, the client would be able to know that requested information is really coming from the appropriate server or service, which would give consumers a greater level of trust in engaging in e-commerce. A similar argument applies for e-government – for example, people might be more prepared to vote on-line if they could have greater assurance about the service they were using.

4 Future Work

One aspect of our work within PRIME is that it is in practice often too simplistic for an enquirer platform to provide a trust metric after analysing integrity information from another single platform only, as described above in Subsection 3.2, because often there would not just be one server, but a network of (distributed) servers. If a client wishes to assess the trustworthiness of the back end system, then it is necessary to aggregate trusted values from a network of platforms in a meaningful way. We are currently researching how to aggregate trust measurements from components within an infrastructure, and how to analyse the overall trust of systems that may change over time, such as within adaptive enterprise models.

Usually in practice the situation can be very complex since the 'server side' consists of various servers (e.g. PCs), potentially each of them with TPMs, etc. Users might be completely unaware of the topology and the current set of PC servers that will actually handle their information (e.g. adaptive enterprise or current load balancing techniques where resources are dynamically allocated based on workloads); nor is it in the interest of the enterprise to disclose this topology. We are currently involved in providing a suitable solution in such a case by developing a Trust Aggregator mediator component within PRIME; this provides only a partial solution because it is not always known in advance which exact resources will be used and the mediator may not have complete control over this for the whole of the service provision. In some models the authoriser is not a human, i.e. remote authorisation could be implemented between services and also within a service, such that the identity manager could provide such authorisation internally.

5 Conclusions

In conclusion, answers to questions about technology-mediated trust involve a combination of technology and also (changing) human attitudes and behavior. In order to determine whether a system is trustworthy, we have to ask whether we have assurance that the system will behave as it should and also whether we trust the people behind the technology. Trusted Platforms help in doing this, but note that still we trust the people if we believe that they will not exploit their potential to hurt us. Business with strangers is risky, no less so for business partners online than it is off-line [11]. By the mechanisms described above the next versions of Trusted Platforms will aim to provide a root of trust for other trust service technologies.

Acknowledgements: Our ideas on this topic benefited from useful input and discussions with Graeme Proudler (HP) and Giles Hogben (JRC).

References

1. ACM, Special Issue on "Trusting Technology", *Communications of the ACM*, vol 43, no 12, December 2000.
2. Blaze, M., Ioannidid, J. and Keromytis, A.D., "Experience with the KeyNote Trust Management System: Applications and Future Directions", Proceedings of the First International Conference on Trust Management (iTrust 2003), Crete, Greece, pp. 284-300, May 2003.
3. Castelfranchi, C. and Y.-H. Tan, eds., *Trust and Deception in Virtual Societies*, Kluwer Academic Publishers, 2001.
4. Cheskin, "Research and Studio Archetype", *eCommerce Trust Study*, University College London. January 1999. Available via http://www.sapient.com/cheskin/.
5. Egger, F. N., *Increasing Consumers' Confidence in Electronic Commerce through Human Factors Engineering*, MSc project, University College London, 1998.
6. Friedman, B., P. H. Kahn Jr. and D. C. Howe, "Trust Online", *Communications of the ACM*, 43, no. 12, December 2000, pp. 34-40.
7. Giff, S., *The Influence of Metaphor, Smart Cards and Interface Dialogue on Trust in eCommerce*, MSc project, University College London, 2000.
8. Grandison, T. and Sloman, M., "Trust Management Tools for Internet Applications", Proceedings of the First International Conference on Trust Management (iTrust 2003), Crete, Greece, pp. 91-107, May 2003.
9. Intel, "LaGrande Technology Architectural Overview", September 2003. Available via http://www.intel.com/technology/security/downloads/LT_Arch_Overview.pdf.
10. ISO/IEC 15408 (all parts), "Information technology – Open Systems Interconnection – Evaluation criteria for information technology security", *International Organization for Standardization*, Geneva, Switzerland, 1999.
11. Jupiter, *Trust Online: Barrier Today, Strength Tomorrow*, Research Report, 4 April 2001.
12. Kumar, N., "The Power of Trust in Manufacturer-Retailer Relationships", *Harvard Business Review*, Nov-Dec 1996, pp92-106.
13. Luhmann, N., "Trust as a Reduction of Complexity", *Trust and Power: Two works by Niklas Luhmann*, New York: John Wiley & Sons, 1979, pp. 24-31.
14. Microsoft, Next-Generation Secure Computing Base home page, http://www.microsoft.com/resources/ngscb.
15. Nielsen, J., "Trust or Bust: Communicating Trustworthiness in Web Design", *Jacob Nielsen's Alertbox*, 1999. Available via http://www.useit.com/alertbox/990307.html.
16. Nissenbaum, H., "Can Trust be Secured Online? A theoretical perspective", *Etica e Politica*, no. 2, Dec 1999.
17. Osterwalder, D., "Trust Through Evaluation and Certification?" *Social Science Computer Review*, 19, no. 1, Sage Publications, Inc., Spring 2001, pp. 32-46.
18. Pearson, S., "A Trusted Method for Self-profiling in e-Commerce", *Trust, Reputation and Security: Theories and Practice*, R. Falcone et al. (eds.), LNAI 2631, pp, 177-193, Springer-Verlag, Berlin, 2003.
19. PRIME: Privacy and Identity Management for Europe, European RTD Integrated Project under the FP6/IST Programme, http://www.prime-project.eu.org/, 2004
20. Rousseau, D., S. Sitkin, R. Burt and C. Camerer, "Not so Different after All: a Cross-discipline View of Trust", *Academy of Management Review*, 23, no. 3, 1998, pp. 393-404.
21. Simmel, G., *Soziologie*, 5th ed., p. 263, Berlin, 1968.
22. Trusted Computing Platform Alliance, *TCPA Main Specification*, Version 1.1, 2001. Available via www.trustedcomputing.org.

23. Trusted Computing Platform Alliance, *TCPA Design Philosophies and Concepts*, Version 1.0, 2001b.
24. Trusted Computing Group, *TCG TPM Specification*, Version 1.2, 2003. Available via https://www.trustedcomputinggroup.org.
25. Yan, Z. and Cofta, P., "A Mechanism for Trust Sustainability Among Trusted Computing Platforms", S. Katsikas, J. Lopez and G. Pernul (eds): TrustBus 2004, LNCS 3184, pp.11-19, Springer-Verlag Berlin Heidelberg, 2004.