

Trusted Computing: From Theory to Practice in the Real World

Dipl. Math. Alexander W. Koehler

Utimaco Safeware AG, Hohemarkstraße 22, D-61440 Oberursel,
Fed. Republic of Germany
alexander.koehler@utimaco.de

Abstract

Trusted Computing technology is now able to provide solutions to today's enterprise IT security issues. There is a need to increase control over the corporate network. Mobility increases productivity, but with new risks, which have to be handled by technology that is able to ensure security with acceptable costs. Creating security needs well-thought-out concepts, technology, infrastructures, and people who develop, communicate with, and use security products. This paper is intended to deliver one full line of a picture of Trusted Computing, starting with the objectives, looking at the status of Trusted Computing today, and some existing applications, and ending with a view of what will come next.

Broadly-accepted standards are an important prerequisite for any new concept or technology. The Trusted Computing Group (TCG) is the industry-accepted standards body which combines the IT security expertise in Mobile Security of vendors such as Utimaco Safeware AG. Case studies help integrate the views of customers in the presentation on products and vision

Disclaimer

The intention of this paper is to build the bridge between theory and reality. In consequence there is a need to use real products to relate to, and explain, Trusted Computing. This should not be considered as surreptitious product promotion.

1 Why Trusted Computing?

Trusted Computing is a highly-complex matter, and so is the TCG. There are no doubt many definitions of trusted computing. I have chosen this one: It is an IT environment, or an area within it, in which a user, IT administrator or business partner can have "trust". The term "trust" should be considered as a general term that incorporate other concepts such as "system integrity", "reliability", and "predictability". These concepts can be defined as those concepts on which well-known threats such as malicious code, badly-written code, or hardware design weaknesses, could have an influence, as soon as somebody tries to exploit them. This paper, which is simply intended to provide the reader with an introduction to Trusted Computing, does not define the terminology involved in any more detail. There are plenty of technical details, abbreviations and assumptions. Before we begin describing the technology and other details, it is important to agree on some basics. It will then be much easier to understand the TCG and all its ramifications. I would like to begin by classifying these assumptions.

1.1 Assumption 1:

The use of additional hardware components increases security dramatically

If you store a secret, such as a key, on your PC's hard drive, then there is a good chance that somebody else can read it. Several tools are available for detecting key patterns on the hard drive or in RAM. To avoid exposing the key to such attacks it is better way to store the key in a separate location that no malicious software can access. This does not solve the problem that the key has to be transferred to RAM to be processed by the CPU. The solution to that is to have a separate processor at the location where the key is stored. The key will be processed there, and will never leave it. That's secure. This kind of storage and dedicated processing unit package is implemented as smart card, if a removable hardware component is required, or as a Trusted Platform Module, in which case it is integrated with the processor board.

1.2 Assumption 2:

Ultimate protection for data is only given by encryption.

Access control mechanisms may work, but past experience has shown that access control is a barrier that can be overcome without great effort. Usually access control means nothing else than an attribute for a set of data. If the managing software, be it the operating system itself or an application, respects this attribute, the access control mechanism is secure. If other software or operating system is used to access the data, then the access control mechanism no longer provides protection. Improvements in access control system programming does not improve the level of security: there is a need to change technology. The solution to this security task is encryption. Encryption can be considered as the ultimate level of data protection. Encryption does not protect against destroying the content by erasing a file. This is still in the scope of access control. We have seen that the content is well-protected against being read, as it is encrypted. What we need to protect with the ultimate level of security are the keys used for encryption and decryption. Here we refer back to assumption 1, and we can then conclude: Encryption keys need to be stored and processed in dedicated hardware such as smart cards or the TPM (Trusted Platform Module). The TPM is described in detail in chapter 2. TPM (Trusted Platform Module).

To summarize: Ultimate data protection has two prerequisites: 1. Use of encryption technology 2. Use of dedicated hardware for key storage and processing, such as smart cards or a TPM.

So far we have discussed how to provide data with the best-possible protection. But what about the software, the operating system? How do we know that the PC platform is not being manipulated by malicious code? For example, if a thief with a stolen identity tries to use another PC device to penetrate a corporate network. How does the network server know that the PC device which is trying to connect to it is the one it claims to be? The answer to all that is given by the TPM (Trusted Platform Module) which is described in the next chapter.

2 TPM (Trusted Platform Module)

The TPM is a fairly simple, mass-produced silicon chip. It consists of a silicon print similar to the one used on a smart card, such as the Random Number Generator, Key Generation and RSA Engine, Non-Volatile Storage, Program Code and Execution Engine.

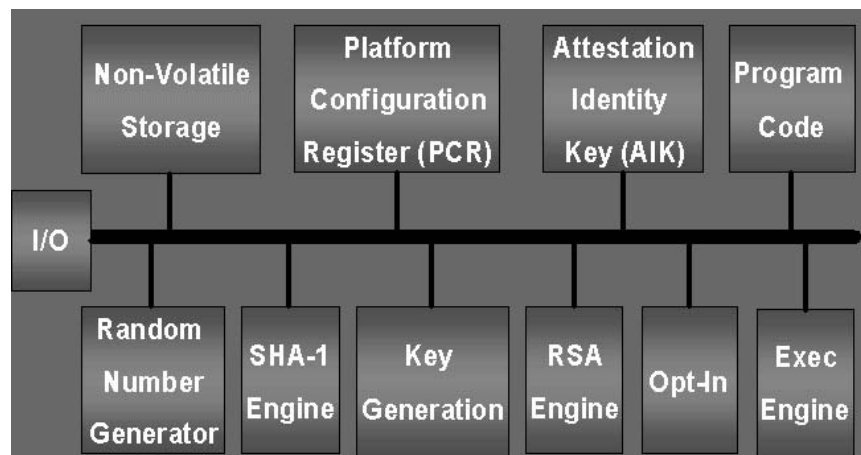


Figure 1: Trusted Platform Module (TPM)

This means that the user use the TPM to perform the same tasks as a smart card, apart from the fact that the TPM "smart card" does not live in a smart card reader, from which it can be removed. The TPM is attached to the motherboard of the PC or other platform. The TPM is much more than a smart card. We will first look at how its additional features will help the user or the IT administrator. The TPM offers two fundamental values to the device:

2.1 "Who am I?"

The TPM can be used to pass on a unique identity to the larger computing device. This identity is based on the ability of a TPM to store an asymmetric private key and perform basic RSA operations using that key. The TPM stores a private key or digital certificate that is associated with a private key stored in the TPM. Using this approach, TPMs can be used to assert identity.

2.2 "Can I be trusted?"

The TPM can be used to evaluate the integrity of software running on a PC or any other computing device. A TPM can be used to store measured integrity information about software for future reference. It does so using so-called HASH functions (algorithms). This creates a way to test whether the software in use is exactly the same as the software installed by the administrator. The assertion is that if the software is the same, the device can be trusted. If the software has been modified, perhaps the device can no longer be trusted. One way to determine if a unit of data has been modified is to hash the original version, preserve the hash value, and use that value as a reference in the future. If the hash value of the unit of data changes, the reason is that the data has been changed. Another capability of the TPM is that it contains Platform Configuration Registers (PCRs) to provide persistent and protected storage for hash values. When an authorized individual configures the device, that person can hash critical components of the software suite in the device and store those values in PCRs. Later, when the device is in use, these critical software components can be hashed in real time and the calculated values can be compared to the stored values. If they match, those components have not been modified since last accessed by the administrator. If they do not match, an unauthorized agent has changed a critical element of the software suite. In this case, it may be prudent for the corporate network security staff to treat the device as compromised.

The tasks of the TPM are:

- Keeping secret information such as passwords and keys out of the reach of software attacks
- Generating high-quality keys, using a hardware random number generator
- Processing private keys inside the unit
- Storing measured data

We have seen before that there are some similarities between the TPM and a smart card. I think everybody agrees that a smart card inserted in the PCMCIA slot of a notebook would neither replace nor control the main CPU (Pentium etc.). Nevertheless, it is sometimes claimed that the TPM controls the PC. This is incorrect. The need for the TPM to store data passed to it from the CPU is based on a simple physical rule: If the data which is measured from an object is stored in the object itself, then the measured data would be exposed to the same kind of attack as the object itself. It would be easy for an attacker to manipulate the measured data in the same way that they manipulated the object's data, and it would be impossible to detect. Hence, any kind of measurement would be worthless. This is the reason why the TPM and CPU need to be two physically-separate units.

The TPM is the most tangible element of the Trusted Computing Group concept. The Trusted Computing Group has defined the standard to which TPMs chips are to be designed. Leading chip vendors already offer TPMs and the most important PC desktop and Notebook vendors already offer multiple product lines that include the TPM chip. However, to create security, much more is needed than this piece of hardware.

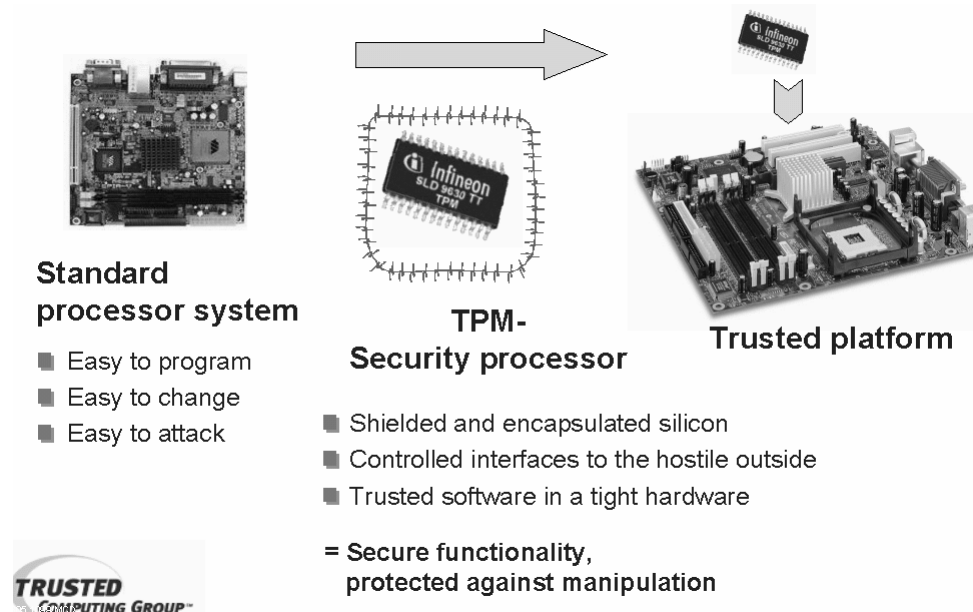


Figure 2: Standard Processor vs. Dedicated Security Controller

3 TCG: Trusted Computing Group

Security in a networked world is a task of worldwide partners. There are many interfaces: from software engineers to special interest groups, government and legal bodies, from industry analysts to the press, and so on. To organize security, there is a need for a central body. The Trusted Computing Group is this body.

The Trusted Computing Group, which was launched on April 8, 2003, is incorporated as a not for profit corporation with international membership and broad industrial participation. The purpose of the TCG is to develop, define, and promote open industry standard specifications for embedded hardware-enabled trusted computing and security technologies, including hardware building blocks and software interfaces, across multiple platforms, peripherals, and devices. By using the building blocks and software interfaces defined by the TCG specifications, the industry can address a range of security needs without compromising functional integrity, privacy, or individual rights.

TCG was created with an organization structure and governance model, as defined by the TCG bylaws, which is similar to many other computing industry standards bodies. These include the following:

- An open membership model with multiple membership levels
- A Board of Directors consisting of Promoters and elected Contributors
- Multiple Work Groups that are open to Promoter and Contributor members, and seek active participation by these members
- A reciprocal reasonable and non-discriminatory (RAND) patent licensing policy between the members
- Supermajority voting at Board and Work Group level to facilitate progress. This structure is designed to enable the rapid development of open, industry standard specifications with broad industry participation, and to foster widespread adoption of the organization's specifications.

The key deliverables of TCG will be hardware and software interface specifications, white papers and other materials that facilitate understanding and adoption of the specifications, and marketing programs that promote awareness and customer adoption.

For more information about TCG, please visit www.trustedcomputinggroup.org

4 TCG, operating systems and hardware platforms

The TPM provides tremendous improvements to IT security when it is integrated into, for example, application software solutions. The TPM could do even more if it could be used by the operating system. During February 2004 Microsoft announced that the TPM will be part of the successor to the Windows XP operating system, code named Longhorn. The secured operating system component has been named NGSCB (Next Generation Secure Computing Base). Developing an NGSCB for Longhorn is still on-going and subject to major changes during the lifetime of this paper. For this reason, this paper will not describe the NGSCB.

If an operating system used the TPM, the real benefit would only come together with redesigned hardware platforms. To get the full picture for Trusted Computing, I propose to examine Intel's announcement in February 2004 of a new project, code-named LT for LaGrande Technology.

Remember that today's PC (micro)processor design goes back to the 1980s. At that time, these microprocessors were designed to do their job despite the limitations inherent in their hardware design. This hardware design was quite different to the one used for minicomputers or even mainframes. For instance, protected execution is a classic must in computer science. It prevents DoS (Denial of Service)-type attacks from being successful.

Today we need to bring back the benefits of a fully-featured hardware design. Improvements to the hardware design need to be compatible with the most common platform in the world, which is the PC platform, so backwards compatibility is mandatory. For this reason Intel is enhancing the PC platform not only with this functionality, but even more. The additional functionality, such as protected graphics, arises from today's threats and needs. (20 years ago, graphics were not an issue on most client devices). This is LT:

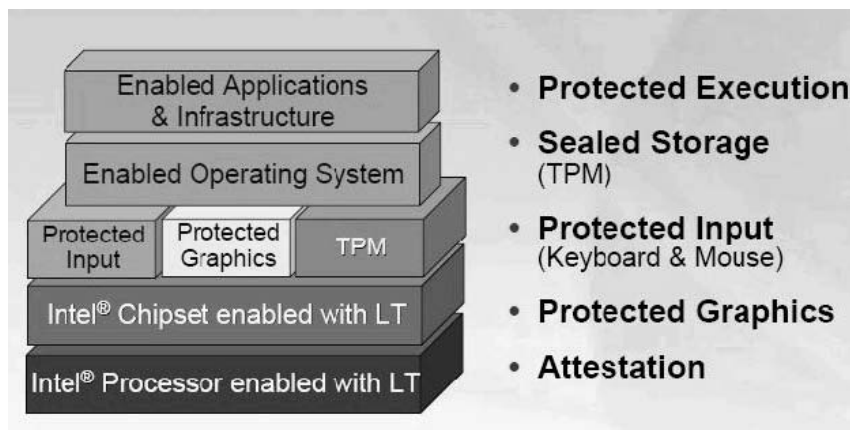


Figure 3: Summary of LT Capabilities

LT improvements over the standard PC platforms are implemented on functions, which are provided by the TPM.

LT provides a kind of "vault", which is a protected partition. Security-sensitive tasks are performed inside this protected partition. The vault has some limitations on its functionality. Software developers are requested to only use the vault for security-sensitive parts of their code. The main part of the code should be executed in the standard partition, as it is today.

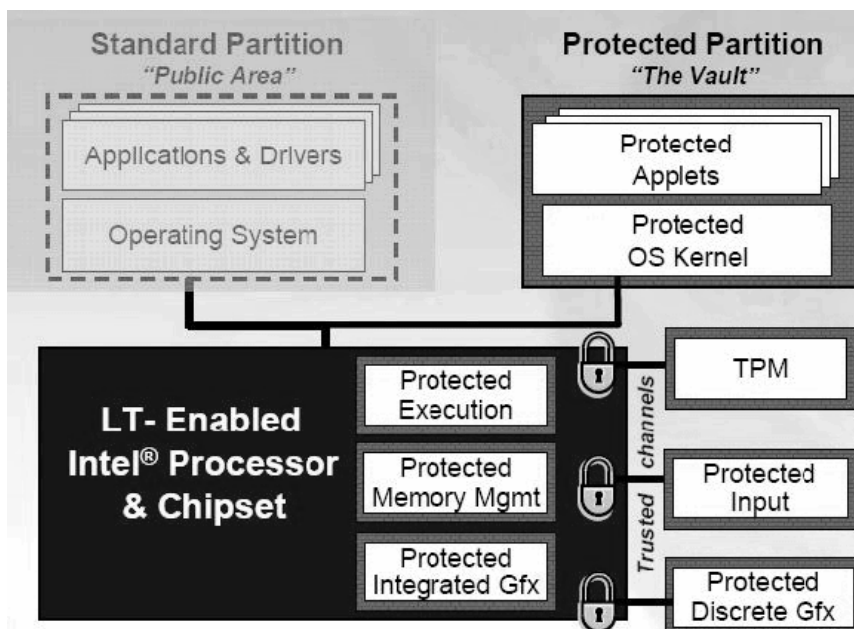


Figure 4: Protected LT Environment

Improvements in operating system design, coordinated with changes in the design of the PC hardware platform, are a great step forward – in theory. But in the real world it will take several years, as millions of today's PC platforms will continue to operate as stand-alone ma-

chines or, most of them, linked to a network. What we will see are several incremental steps by the hardware vendors and operating systems vendors to implement the next stages of IT security.

One advantage of the TCG concept is that there is no need to wait for a rosy future. Today, all major PC vendors already offer Notebooks and PCs with a TPM, so users can already benefit from increased security by using TCG technology today. The next chapter describes one real-life example of security software and shows how TCG technology is beneficial to today's security software.

5 Real-life security products: Increased benefit by using TCG technology

Notebooks are the preferred tool of corporate users. Usually Notebooks carry a large amount of important data. Notebooks are also at threat from several identifiable types of attack.

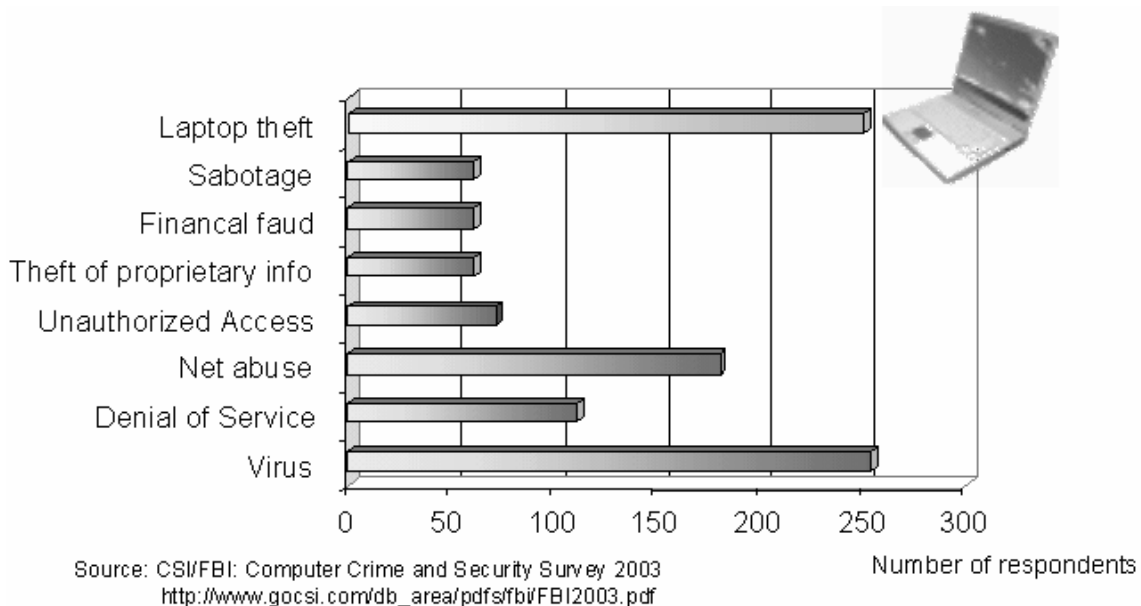


Figure 5: Types of Attack on Mobile Devices

Number one is the same as for any other PC, it is a virus attack. It is essential for a networked device such as a Notebook to have antivirus software installed on it. But what is about the number 2 threat, theft? Theft means that you are no longer in possession of your property. Let me quote what Microsoft says about that. There are the ten immutable laws of security. Law number 3 says: If a bad guy has unrestricted physical access to your computer, it's not your computer anymore. Microsoft is synonymous with today's Windows operating system technology. Does this mean we have to live with the fact that notebooks represent a threat with which we are obliged to live? Not at all. If we widen the scope beyond the operating system, there is an answer.

We classify the status of the notebook into two simple categories. The first is "power off". This is usually the case when the attackers gets hold of the Notebook. Let's have a look at power-off protection.

Imagine the Notebook is stolen, or just left in a taxi. It can be protected. The method is called Bulk Encryption. Bulk Encryption means that every sector on the hard drive is encrypted,

which includes all operating system files. The individual encryption of selected files and folders requires a sophisticated security policy and disciplined users. A disciplined user is a user who is willing, and has the time, to care about where to save files, delete the ones which are no longer needed, and so on. This means the user must comply with some sort of security policy, point by point. Often this is not feasible in the real world. The solution is Bulk Encryption. It frees the user from such tasks and makes them more productive, and makes the system secure. Bulk Encryption runs in the background without any kind of user interaction.

Bulk Encryption means full protection on power-off. It protects both system files and the SAM (Security Account Manager) database in which password hashes are stored. An attacker could swap the SAM with a rogue version. Bulk Encryption protects the operating system, programs, data, etc.: it protects everything. It even prevents the thief from reading the hibernation file in which even individually-encrypted files would be stored as plain (unencrypted) text, if they were being used at the moment when hibernation started.

Bulk Encryption is a technology which is already in use on millions of clients around the world. It has never been cracked for more than 10 years. However, security is not static: threats and defense scenario evolve continuously. For this reason every vendor should go far beyond today's threats when considering modern security technology. They should for example include hardware protection for encryption keys. Another issue to consider is that, in theory, a dictionary attack is always possible against any software-only solution. It is extremely difficult if somebody wants to run it against a certified product. Nevertheless, dictionary attacks simply become impossible if they have to be run against hardware. This is the case with the TPM: if a TCG hardware component, the TPM, is added to a proven IT security software solution, this improves security even more at the present time, and makes security better prepared for future attacks, which are certain to come. Now let's consider what the TPM can do when the Notebook is "on", known as "power-on" protection.

As soon as the Notebook running it is subject to fairly similar threats to a desktop PC. The TPM hardware chip protects credentials such as private keys or secure passwords. Passwords could be generated automatically, and long and complex, as they would be handled internally within a machine. This will make it extremely difficult to crack them. Passwords protected by the TPM will enable Single-Sign-On to the operating system and applications, which means increased security combined with increased user productivity.

The TPM offers another benefit when it comes to insider attacks. Imagine somebody removes a hard disk drive from a corporate desktop and installs it somewhere else. As they know the passwords, they will be able to read the data. The TPM will prevent them from doing so. The password the user knows is only the one used to open the TPM to access other secret information about the operating system, programs or the network. If there is no TPM or if there is another TPM, the attempt to access this other secret information will fail. This concept is called Machine Binding. Data can be tied to the machine.

At the beginning of this paper the True Random Number Generator was described, as one component of the TPM. The application shown here uses this hardware key generator to increase the level of security, as described above.. An example of how these keys can be used in the product referred to here is that, not only is the client authenticated at the server, but vice versa: mutual authentication. This prevents rogue servers cheating an honest client. To end this part of the paper, I would like to quote a famous mathematician: "Anyone who considers arithmetic methods of producing random digits is, of course, in a state of sin", said John von Neumann.

SafeGuard Easy, combined with the TPM and IBM's ESS (Embedded Security Subsystem), is already in use in several European financial services companies. These companies typically use Notebooks as real fat clients. They use custom application software to provide their consulting services to their customers. It would reduce their competitive advantage if this software were available to anybody else. In addition, customer data is also processed on these clients. The service company is responsible for the privacy of this data. The consequence is that security has to be state-of-the-art, yet place no limitation on productivity. "SafeGuard Easy combined with IBM's ESS technology is the only acceptable solution for us," said one of the customers. Obviously the technology is of real benefit to customers.

We have seen here how a real-life high-quality security product can use the benefits of TCG technology. TCG technology means increased security, improvements in user-friendliness and productivity, or reduced costs, and in some cases even a combination of all of these factors. Obviously this matches TCG's declared design goals: Delivering robust security with user control and privacy. In the next chapter we will examine the prospects for the near future. It is another endorsement of the thesis that security is much more than just hardware and software. It is about concepts, infrastructures and so on: the mirrored picture of a networked world.

6 The Near Future of Trusted Computing:TNC (Trusted Network Connect)

So far this paper has focused on the PC platform and Trusted Computing as a sort of point solution. For the sake of completeness, it should be mentioned that Trusted Computing relates to every kind of device that has a certain level of data processing power, and is able to link to a network. This includes PDAs, smart phones, intelligent printers and whatever devices will be developed in the near future. If we look beyond devices, there is the network itself.

The networks, systems, software applications, and data of many enterprises and organizations form a critical foundation and essential structure for their daily operations. Without a reliable and functional network, the business is not secure.

The issue is that point solutions are simply not adequate, and that an end-to-end, comprehensive approach to the security problem is a good way to push security to a higher level. An industry standards-based solution for securing the endpoints of host connections is a critical step on the path to this comprehensive approach, and it is critical for an acceptable solution to our growing network security problem.

The Trusted Network Connect specification for multi-vendor networks will provide a common architecture for vendor solutions that will:

- Ensure endpoint integrity by establishing a level of "trust" in the state of an endpoint. Specifically, solutions based on the specification will ensure the presence, status, and upgrade level of mandated applications, revisions of signature libraries for anti-virus and intrusion detection and prevention system applications, and the patch level of the endpoints, operating system and applications.
- Maintain access policy by helping ensure that the endpoint machine and/or its user authenticate and establish a level of trust before connecting to the network.
- Provide quarantine measures for endpoint machines that do not meet the security policy requirements for "trust" and, if possible, apply appropriate remedial measures, such as upgrading software or virus signature libraries to enable the endpoint to comply with

security policy. TCG is developing an open specification to improve network security and integrity.

The Trusted Network Connect specification, due to be available at the end of 2004, will assist in protecting networks from viruses, worms, denial of service attacks and host software vulnerabilities by allowing users to enforce security policies to prevent vulnerable or untrusted systems from connecting to the network.

7 Summary

IT security is not a static topic but an on-going process. The concept and content of Trusted Computing is an important building block to create robust security products which will provide substantial improvements to IT security. TCG concepts already reach to the future, and an additional strength is that Trusted Computing already works with today's products, infrastructures and business processes.

Abbreviations and References

TCG	The abbreviation has two meanings: Trusted Computing Group; the organization as a standards body www.trustedcomputinggroup.org Technology based on standards provided by TCG
TPM	Trusted Platform Module: a dedicated security chip bound to a PC or other hardware platform, manufactured by Infineon, for example. http://www.infineon.com/cgi/ecrm.dll/ecrm/scripts/prod_ov.jsp?oid=29049
TNC	Trusted Network Connect www.trustedcomputinggroup.org
SafeGuard Easy:	Bulk encryption product by Utimaco Safeware AG www.utimaco.com
ESS	Embedded Security Subsystem by IBM http://www.pc.ibm.com/us/think/thinkvantagetech/security.html
LT	LaGrande Technology by Intel http://www.intel.com/technology/security/
NGSCB	Next Generation Secure Computing Base by Microsoft http://www.microsoft.com/resources/ngscb/default.mspx

Acknowledgments:

First and foremost I would like to thank all the members of the TCG who provided inputs in several formats, from informally discussions to management presentations. Special credit goes to Stacy Cannady (IBM), David Grawrock and Monty Wiseman (Intel), Thomas Rosteck and Hans Brandl (Infineon). I also owe the mentioned gentlemen thanks for the contribution of diagrams as presented in this paper.

In this document names and brands are properties of their respective owners.