The Advent of Trusted Computing: Implications for Digital Forensics

Mike Burmester
Department of Computer Science
Florida State University
Tallahassee, Florida 32306-4530
850-644-6410
Email: < burmeste@cs.fsu.edu>

Judie Mulholland
Florida Cybersecurity Institute
Florida State University
Tallahassee, Florida 32306-4530
850-644-7182
Email: <judiemul@cs.fsu.ed

ABSTRACT

The release of computer hardware devices based on “trusted computing” technologies is heralding a paradigm shift that will have profound implications for digital forensics. In this paper, we map out the contours of a trusted environment in order to establish the context for the paper. This is followed by the main components of the TC architecture with an emphasis on the Trusted Platform and the Trusted Platform Module (TPM). The next section presents a synopsis based on three threat models, *viz.*, (i) pc owner-centric, (ii) trusted computing-centric, and (iii) digital forensics-centric and then briefly touches on the implications and unintended consequences of trusted computing for digital forensics. Finally, the last section of the concludes with a recommendation on how to mitigate the negative effects of trusted computing.

Categories and Subject Descriptors

K.5.0 [Legal Aspects Of Computing]: General.

General Terms

Security, Legal Aspects.

Keywords

Cybercrime, data recovery, encryption, file systems, forensics, specifications, Trusted Computing.

“We shape our tools, and thereafter
our tools shape us”—Marshall McLuhan.

1. INTRODUCTION

The Trusted Computing Group (TCG) is a not-for-profit industry-standards organization that was set up to establish specifications for architectures, functions and interfaces that support hardware-based trusted computing solutions. As part of their mandate, the TCG has been developing a set of guidelines [8] that will serve as a baseline for a wide variety of platforms—from personal computers, personal digital assistants, to cellular telephones.

A number of initiatives falling under the auspices of trusted computing (TC) are currently under development. The most notable ones are: (i) hardware-related projects—Intel is developing a new chip called LaGrande Technology (LT) and AMD is working on one called Pacifica. (ii) Software-related projects—Microsoft is releasing a new operating system they have christened *Windows Vista*—originally called Palladium/ Next-Generation Secure Computing Base (NGSCB)/Longhorn. At the time of this writing, a dominant design has begun to coalesce around a single TC architecture.

To establish the context for the paper, we begin by mapping out the contours of a trusted environment. This is followed by the main components of the TC architecture with an emphasis on the Trusted Platform and the Trusted Platform Module (TPM). The next section presents a synopsis based on three threat models, *viz.*, (i) pc owner-centric, (ii) trusted computing-centric, and (iii) digital forensics-centric. Section 5 outlines the implications of trusted computing for digital forensics with respect to file system analysis and evidence recovery. Finally, the last section of the paper offers some recommendations on how to mitigate the negative effects of trusted computing for law enforcement.

2. TRUSTED COMPUTING OVERVIEW

The TCG defines trust as “the expectation that a device will behave in a particular manner for a specific purpose” [8]. To be considered a trusted environment, a minimum of three conditions must be present:

Protected capabilities—are based on a set of commands that have exclusive permission to access shielded locations (e.g., memory and/or registers) where it is safe to work on sensitive data.

Integrity measurement—is the process of obtaining metrics of platform characteristics that affect the integrity (trustworthiness) of a platform.

Integrity reporting—serves two main functions: (i) to expose shielded-locations for storage of integrity measurements, and (ii) to attest to the authenticity of stored value based on trusted platform identities.

SAC’06, April, 23-27, 2006, Dijon, France.
Copyright 2006 ACM 1-59593-108-2/06/0004…$5.00.

3. TC ARCHITECTURE

This section describes the logical layout of the TC architecture as outlined in the TCG documentation [8].[1] At present, the TCG specifications are being designed to provide personal computers with an essential hardware base for client-side security. According to Safford, the TC architecture provides two important security functions: secure storage of signature and encryption keys and system software integrity measurement [7]. It should be noted that the TC architecture includes both hardware i.e., the trusted platform module (TPM) and software components i.e., the trusted support services (TSS). Given the focus of this paper on data recovery, only hardware issues will be dealt with.

The Roots of Trust represent the minimum functionality needed to describe the properties that affect the trustworthiness of a computing environment. The trusted platform is comprised of three Roots of Trust: (i) *a root of trust for measurement* (RTM)—measures integrity and enables transitive trust; (ii) *a root of trust for storage* (RTS)—presents summary values for integrity digests and maintains the sequence of digests; and (iii) *a root of trust for reporting* (RTR)—reports information held by the RTS. The Roots of Trust must be trusted due to the fact that any misbehavior taking place within the confines of the system might not be detected. Each root is expected to function correctly without external oversight. The Trusted Building Blocks (TBB) and the Roots of Trust form a trust boundary where measurement, storage and reporting can be accomplished using a minimal configuration. According to the TCG specifications, "[t]he TBB should be established such that devices containing other measurement code do not inadvertently extend the TBB boundary where trustworthiness of the linkages has not been previously established" [8]. Or, as Stafford points out, "integrity measurement can be used to detect software compromise, such as a rooted kernel, and to lock down use of protected keys and data if a compromise is found" [7].

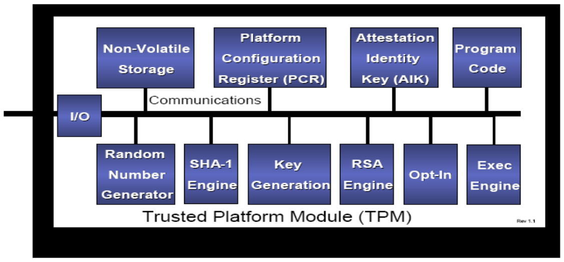

Figure 1. Logical layout of the TPM [8]

3.1 Trusted Platform Module (TPM)

The main components included in the TPM schema that are expected to have the greatest impact on the personal computing environment include: (i) *secure functions* with a focus on the modes of operation and the issuance of credentials, (ii) *TC keys* with a focus on measurement and the protected message exchange protocols, and (iii) expanded capabilities with a focus on secure input/output, memory curtaining, sealed storage, and attestation.

Secure Functions. The trusted chip (i.e., the TPM—*see Fig 1*) manages three main groups of functions: (i) public key functions, (ii) trusted boot functions and (iii) system initialization and management functions.[2] In order to verify that there have been no malicious additions to the hardware or software, measurements (i.e., SHA-1 hashes) are made during the boot process and stored in the Platform Control Registers (PCRs).

Based on the current configuration, the TPM behavior is limited by a combination of three mutually-exclusive modes of operation:

Enabled / Disabled—the TPM may be enabled/disabled multiple times within a boot period. When the TPM is enabled, all features are available; whereas when the TPM is disabled, all operations are restricted except the ability to report TPM capabilities and to accept updates to the Platform Configuration Register (PCR).

Activated / Deactivated—when activated all features of the TPM are available. In a deactivated state, the TPM is similar to disabled except that operational state changes such as "change owner" or "activation with physical presence" are possible.

Owned / Un-owned—a platform is owned when the owner of a platform is authorized to perform all functions including operational state changes.

TC Keys. The main classification for TC keys are non-migratable vs. migratable. Non-migratable keys embedded in the TPM include: (i) the Storage Root Key (SRK) and (ii) the Endorsement Key (EK). Migratable keys may be exchanged (exported/ imported) which enables the TPM to sign application data and enforce usage restrictions. This allows the key pair to follow the user around irrespective of device type. To extend non-migration attributes to opaque data, data are stored with the RTS using a non-migratable storage key. This means that as long as an opaque object is controlled by the TPM, it cannot be decrypted elsewhere.

Within the TCG schema, keys are considered communication endpoints. Therefore, if communication endpoints are poorly configured or keys are improperly managed, a breach in security may result. The TPM advances security by providing both key management and configuration management features (e.g., features such as protected storage, measurement and reporting are combined to “seal” keys and platform configurations making endpoint definition stronger.[3] The TCG defines four classes of protected message exchange:

Binding—is based on the traditional operation of: (i) encrypting a message using the intended recipient's public key and (ii) recovering the message using the intended recipient's private key. If the private key is a nonmigratable key, then only the TPM that created the key may use it.

Signing—is a process that associates the integrity of a message with the key used to generate the signature.

[1] The background material for section 3, unless noted, is drawn from the TCG Specifications [8].

[2] Using the initialization and management functions, the owner can turn functionality on and off, reset the chip, and take ownership.

[3] Protected messaging is based on two principles: (i) that messages intended for one and only one individual can be encrypted using a public key and (ii) the message can be protected from tampering by signing with a private key.

Sealing—binds a set of metrics—a platform configuration state that must exist before decryption can proceed—to a message. The symmetric key used to encrypt the message is associated with a set of PCR register values and a non-migratable asymmetric key. Sealing ensures that **"**protected messages are only recoverable when the platform is functioning in a very specific known configuration" [8:16].

Sealed-Signing—can be used to provide an assurance that the platform that signed the message meets specific configuration requirements.

Any command that affects security and privacy or is capable of revealing platform secrets must be authorized which means that a secret must be supplied as part of command invocation. Commands that do not require authorization include: (i) informational commands (i.e., they contain no security or privacy information) and (ii) privacy relevant meta commands (i.e. they are needed to configure command validation).

Expanded Capabilities. Once the TPM has been activated, new features will be available to pc owners, content providers and law enforcement agents (LEAs). The particular capabilities singled out for our review are the ones generating the most controversy within the end user community:

Secure Input and Output (I/O)—to minimize the type of threat posed by keyloggers and screen-grabbers, secure I/O provides a tamperproof communications route between a user and an application. Under secure I/O, the keyboard and mouse will be protected from physical attacks; screenshots or scrapes will be disabled; and programs that deliberately corrupt, modify or mislead the user will be prevented from running or operating.

Memory Curtaining—memory that has been isolated from other internal processes enables trusted programs to run without interference.[4] Encryption keys locked in a data vault (a chip attached to the motherboard) are used to maintain privacy and integrity. Although process isolation can be achieved using software, the advantages of hardware include: (i) greater backwards compatibility, (ii) less code needs to be rewritten and (iii) fewer changes to device drivers and application software.

Sealed Storage—encryption keys, based on a combination of hardware and software, are used to store data in an encrypted format means the data can be read only by the same combination of software and hardware. If an application other than the one that was used to seal the data attempts to decrypt or unseal the data, the operation will fail. Similarly, if the data is copied in encrypted form to a different machine, attempts to decrypt it will be unsuccessful.

Attestation—is the process of verifying and vouching for the accuracy of information and it works by having the TPM generate a certificate that confirms—NO unauthorized installs, updates or changes to have been made to the user's hardware or software. Attestation is designed to prevent data (e.g., commands, executables, private information) from being sent to/from a compromised or insecure computer.

4. THREAT MODELS

Computer Security concerns the protection of information assets. For personal computers this means the protection of stored data and programs. Protection typically involves integrity, confidentiality and availability.

Scenario 1: The traditional pc threat model. In the traditional security model for personal computers, the threats are external and do not involve the owner of a personal computer (pc). That is, the owner is trusted and has full control over the pc. The owner is identified by a password and/or biometrics. The adversary is an unauthorized user (a hacker)–*see Fig 2*. With networked systems of computers, the computers are only used as platforms and the information assets are stored centrally and managed by a network administrator who enforces the access control policies of the system. With such systems, the administrator is the only trusted party.

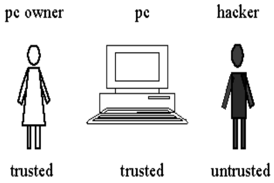

Figure 2. The traditional pc threat model

Scenario 2: The trusted computing threat model. The security model for trusted computing is similar to the personal computers model, except that in this case the trust between the pc and its owner is broken–*see Fig 3*. That is, every user, including the owner of the pc, is untrusted. Only the pc is trusted. The owner has restricted access to the information assets stored on the hard drive of her/his computer. The restrictions are intended to limit and contain the damage that can result from any security flaw in the operating system of the computer, as well as to protect its owner from, inadvertently exposing or corrupting information assets stored on the hard drive (e.g., by importing malicious code), privacy threats (by encrypting stored data on the hard drive with keys generated by the hardware), illegal copying or file sharing, unfriendly behavior to the software and publishing industry, by tethering (preventing files from migrating), lock-ins (only approved software will run), forcing upgrades/downgrades, and possibly other non-disclosed mechanisms (the good, the bad and the evil?). This model can be regarded as a special case of the security model for networks in which the network is replaced by a single computer and the administrator by the operating system of the computer. This is essentially a Big Brother model [4], in which (the hardware of) the computer is designed in such a way so as to protect its owner from "wrongdoings", where the wrongdoings are determined to a large extent by business and corporate interests. This does not benefit the software industry as

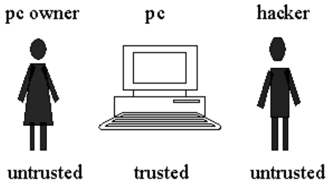

Figure 3. The trusted computing threat model

[4] With curtained memory, even the operating system is denied access.

a whole, because it introduces anti-competitive practices [4, 7] but it enforce Digital Rights Management [23].

Scenario 3: The digital forensics threat model. The security of the models discussed so far, focuses on preventing attacks. For our last model, the model for digital forensics, security focuses on detection. This model is similar to the model for trusted computing, only that in this case the hacker is replaced by a trusted law enforcement agent (LEA). The owner of the computer remains untrusted–*see Fig 4*. The objective of the LEA is to extract incriminating data stored on the computer. The computer is trusted not to corrupt this data, and to make it possible for the agent to decrypt it. The main difference from the model for trusted computing is that in this case the "wrongdoings" are determined by well-established legal procedures, based on the interests of society as a whole, rather than the interests of the software and publishing industry.

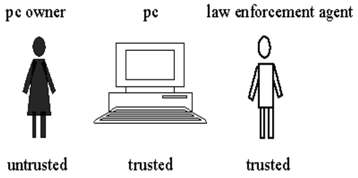

Figure 4. The digital forensics threat model

5. TC IMPLICATIONS

As noted earlier, trusted computing has generated a ground swell of controversy. Without the addition of user-friendly fixes—*viz.*, some type of override mechanism—opposition is likely to continue [2]. Once trusted computing is deployed on a massive scale and the reality of a 'locked down' computing environment starts to sink in, there is bound to be a backlash. However, from a digital forensics point of view, the advent of trusted computing, is not all bad. In fact, the TC-enabled features most feared by the naysayers may become a boon for cyber-investigators. On the other hand, if file-encryption becomes the norm, trusted computing may turn out to be law enforcement's worse nightmare. To get an inkling of the potential impact of TC and its unintended consequences, this section focuses on three key elements in the digital forensics arsenal: acquisition, file system analysis and data recovery methods.

Acquisition. At the scene of the crime, it has become standard practice to "bag and tag" evidence and take it back to a safe environment (e.g., a certified forensics lab) for imaging and analysis [8]. When dealing with servers, to avoid disruption, most forensics examiners—once normal safeguards are in place—will acquire the evidence right on the spot. With trusted computing, it is still unclear what type of acquisition policies should be followed. For example, if it is known *a priori* that a case involves unencrypted data, it will be safe to follow 'standard operating procedures.' Depending on the circumstances, it will be up to the forensics team to decide where and how to acquire the evidence. Alternatively, if a TC-enabled box with encrypted data becomes part of an investigation, cyber-investigators are well advised to approach these machines/devices as if they are mission critical. In any event, forensic teams responsible for data recovery should err on the side of caution. Depending on what type of secure I/O or remote attestation has been set up, these machines may interpret any unauthorized interference as a threat and act accordingly. Not to mention, preventative measures—passphrases/biometrics, curtained memory and sealed storage—may have been set up to thwart unauthorized access. Ideally, LEAs will secure the cooperation of the pc owner who will reveal pertinent information. Most likely, unless some kind of plea bargain or immunity arrangement is worked out beforehand, there will be little or no incentive for the pc owner to cooperate since without the decryption keys, incriminating data will remain protected[5] i.e., unrecoverable. For forensics practitioners, this means that a new generation of intermediary forensics tools will be needed that can to extract data from TC-enabled machines.

File System Analysis. Given the ease with which data can be modified, a major issue confronting all cyberinvestigations is "what type of data can be trusted."[6] When dealing with TC-enabled computers, not only will more system data be stored in tamper-proof logs but data that were previously out-of-bounds will now be routinely signed, sealed and bound to a user. Every time someone who operates a TC-enabled machine comes in contact with a digital object, a unique fingerprint will be created. It is assumed that once critical mass is reached, law enforcement will be able to rely on digital signatures and time stamps derived from authentication procedures to corroborate evidence and rule out suspects—in much the same way that DNA is currently used. Similarly, it is expected that hashes/digests that are generated as a by-product can be used for separating 'known from unknown' file types and data carving purposes. In other words, law enforcement will have at their disposal a historically rich source of metadata they can use to more closely associate individuals with the actions, thereby increasing the likelihood that this evidence will be admissible in court.

Data Recovery. At the time of this writing, details regarding Microsoft's new operating system (Windows Vista) are few and far between. To date, no guidelines, comparable to the TCG specifications, have been published. Therefore, it is difficult to hazard a guess as to how well data recovery efforts will fare under trusted computing. To consider what some of the implications might be, we can conjecture the following:

In keeping with past releases (e.g., Windows 9x/0x, NT, XP), Vista will most likely retain the same layout, data structures (records, signature values, flags, options) and file formats (indexing, journaling) that first appeared in the FAT file system and were later revamped/revised and incorporated into NTFS [3:351-395]. If so, that is good news. Apart from learning new terminology and tweaking some data recovery tools, no significant changes in digital forensics modus operandi will be required for recovering unencrypted data on a TC-enabled machine. It is expected that the Microsoft OS will retain little endian ordering, the Master File Table (MFT), metadata, and file attributes. DOS partitions, clusters, sectors and slack space will continue to exist. Short/long file names and deleted data will

[5] This assumes there is no backdoor.

[6] Carrier makes a distinction between essential (trusted) and nonessential (untrusted) data. For example, he considers file system information such as content addressing to be essential, otherwise the system would be unable to read the file; whereas data and time stamps are nonessential because they can be easily manipulated by the user [3:12-13]. Non-essential data that can be easily manipulated is more likely to be challenged in court.

continue to be recovered in the same manner. Data will continue to be written to the hard drive using the same allocation algorithms. Now, for the bad news. There is no reason to expect that Microsoft will follow in the same footsteps.[7] In fact, given Microsoft's track record, there is every reason to believe otherwise. Most likely—which may account for all of the delays—Microsoft is poised to come out with an entirely new file system that is not backward compatible, retains no structures in common with NTFS and cannot be reverse-engineered (without running afoul of the DMCA). All of which does not bode well for cyberinvestigators.

6. UNINTENDED CONSEQUENCES

Under the current guidelines, trusted computing based on hardware encryption uses a key generated internally (which is a function of the computer identity, the software encryption identity and possibly other system parameters). What happens if LEAs do not have access to the decryption key or worse still, there is a hardware malfunction? Does this mean that all data on the hard drive is lost, in the sense that it is encrypted and the system cannot compute the required decryption key so the information that is stored is lost forever.

In fact it is possible to get the key, provided cyberinvestigators have access to the computer ID and the software encryption ID. Consider two possibilities:

A: The hardware is designed so that it is impossible to get the computer ID (note that it must be easy to get the software ID, otherwise the computer will not be able to generate a key for encrypting/decrypting). In this case, it will be impossible to compute the decryption key and therefore to decrypt stored data, even by the pc owner. If trusted computing is implemented this way, it is doomed, because any hardware failure will result in all stored data being lost forever—and that does not make good business sense, so it is unlikely to prevail.

B: It is possible to extract the ID from the hardware so the owner can recover the data. For the same reason, the agent can recover the data, as indeed anybody else who has physical access to the pc. For example, even a thief. The only way around this that we can see (so that the agent can, but the thief can't) is to protect the computer ID. It must not be in the clear, and the manufacturer must not know it (i.e., a malicious manufacturer may sell these Ids to hackers who can then compute the keys).

The solution to this dilemma would be to hardwire the pc with an "encryption" ID which is printed internally and stored in a way that it can't be easily recovered. To access the encryption key, the hardware would have to be destroyed and the TPM could never be used again to assert trust. But any lost or incriminating data would be recoverable. This will result in the pc getting a new protected hardware ID, while making it possible to access the encrypted data with the exposed key.[8] Lastly, we should point out that the practice of using encryption keys that are not stored on the computer (or cannot be internally generated by the computer) is the most serious threat to digital forensics. By the same token, whoever uses this practice, the one being recommended by the TCG, is also at great risk of losing all data stored on the hard drive if he/she loses the encryption keys. So these keys must be kept safely. This is where law enforcement must insist that the TCG rework their design to incorporate some type of key recovery mechanism even though we recognize that this solution is unlikely to be popular with pc owners. However, the alternative—losing valuable data--is even less appealing.

7. CONCLUDING REMARKS

The release of computer hardware devices based on TC is heralding a paradigm shift that will have profound implications for digital forensics. TC-enabled machines are expected to thwart everything from denial of service attacks, unauthorized access, phishing scams, to illegal downloads. What is often overlooked in this brave new world—where every bit is locked down—is the downside risks. Conducting a cybercrime investigation in an environment dominated by secure I/O, curtained memory, sealed storage and attestation technologies will present some unique challenges for law enforcement. Any increase in actionable evidence may be offset by encrypted data that cannot be recovered. Just as the Internet spawned spammers and hackers; no doubt trusted computing will create a new breed of cybercriminal who uses encryption and darknets to avoid detection. In conclusion, we ignore at our peril, McLuhan's admonition:

We shape our tools, and thereafter, our tools shape us.

8. REFERENCES

[1] **Against TCPA.** URL: http://www.againsttcpa.com/what-is-tcpa.html .

[2] **Anderson. R.** 'Trusted Computing' Frequently Asked Questions - TC / TCG / LaGrande / NGSCB / Longhorn / Palladium / TCPA. Version 1.1 August 2003. URL: http://www.ftp.cl.cam.ac.uk/ ftp/users/rja14/tcpa.pdf .

[3] **Carrier, B.** *File Systems and Forensics Analysis*. Addison-Wesley Professional: March 17, 2005

[4] **Chaum.** Security Without Identification: Transaction Systems to Make Big Brother Obsolete. *Commun. ACM*. 28(10): 1030-1044 (1985).

[5] **Lemos, R.** Hardware security sneaks into PCs, *CNET News.com*. 3/16/2005. URL: http://news.com.com/ Hardware+security+sneaks+into+PCs/2100-7355_3-5619035.html .

[6] **Safford, D.** Clarifying Misinformation on TCPA/Palladium/ DRM. October, 2002. URL: http://www.research.ibm.com/ gsal/tcpa/tcpa_rebuttal.pdf .

[7] **Schoen, Seth.** Trusted Computing: Promise and Risk. URL: http:// www.eff.org/Infrastructure/trusted_computing/2003 1001_tc.php .

[8] **Trusted Computing Group.** TCG Specification Architecture Overview. Revision 1.2. 28 April 2004. URL: https://www.trustedcomputinggroup.org/downloads/TCG_1_0_Architecture_Overview.pdf .

[7] An anonymous reviewer points out: "WinFS is not going to be an entirely new file system (as was originally hinted at). Instead it is adding relational components to the existing NTFS structure." From what we can ascertain, it seems that Vista will incorporate two files systems: WinFS and NTFS—the details of how they will interoperate are still unknown.

[8] In fact, a TC-enabled machine will need several ids, because some may have to be published for attestation purposes.