print | close

Standardizing Trust for Embedded Systems

Electronic Design

Stefan Thom, Steve Hanna, and Stacy Cannady

Tue, 2016-01-26 09:12

If you haven't been concerned about malicious players hacking into your products in the past, or haven't found success with previous efforts, it's time for renewed attention and action. Hacking efforts aren't slowing and, in fact, are on the rise. These days, hackers can accomplish far more than ever before—and the repercussions are far more costly.

Related

Understanding Trusted Computing From The Ground Up

The Many Faces of Embedded Security

End-to-End IoT Security Starts with the Infrastructure

While successful hacks to private industry and the government have been widely reported, perhaps the more personal ones indicate the extent of the attacks and provide more dramatic examples. Hacking a baby monitor is certainly a major concern to any parent, several instances of which were reported in 2015.

For many years, cybersecurity experts have scared motorists with the possibility that connectivity in the car and to the Internet could compromise a vehicle's control systems. In July 2015, an unmodified vehicle was successfully hacked from over 10 miles away. Taking control from the driver is no longer a theoretical issue.

Recognized Need for Trust in IoT

As the Internet of Things (IoT) expands, hackers will have even more gateways into unauthorized territory. For the IoT, 51% of respondents in published survey from last year (*Electronic Design*, October 2015) said that security in products is currently very important, and 54% said security will be even more important in future products.

The Trusted Computing Group (TCG) recognized the need for establishing trust in computing products over a decade ago. Experts from member companies developed the specification for a hardware component called the Trusted Platform Module (TPM). More recently, TCG has greatly simplified the ability to take advantage of TPM-based security and has turned its attention to embedded applications.

For those not familiar with TCG and the TPM, the TPM timeline provides some useful background.

Trusted Computing Group and the Trusted Platform Module

In 2003, the Trusted Computing Group (TCG) is created to enable extension of trusted computing beyond the personal computer (PC) into the enterprise. It then announces its first standard: the Trusted Platform Module (TPM). Figure 1 shows the timeline of establishing and increasing trust.

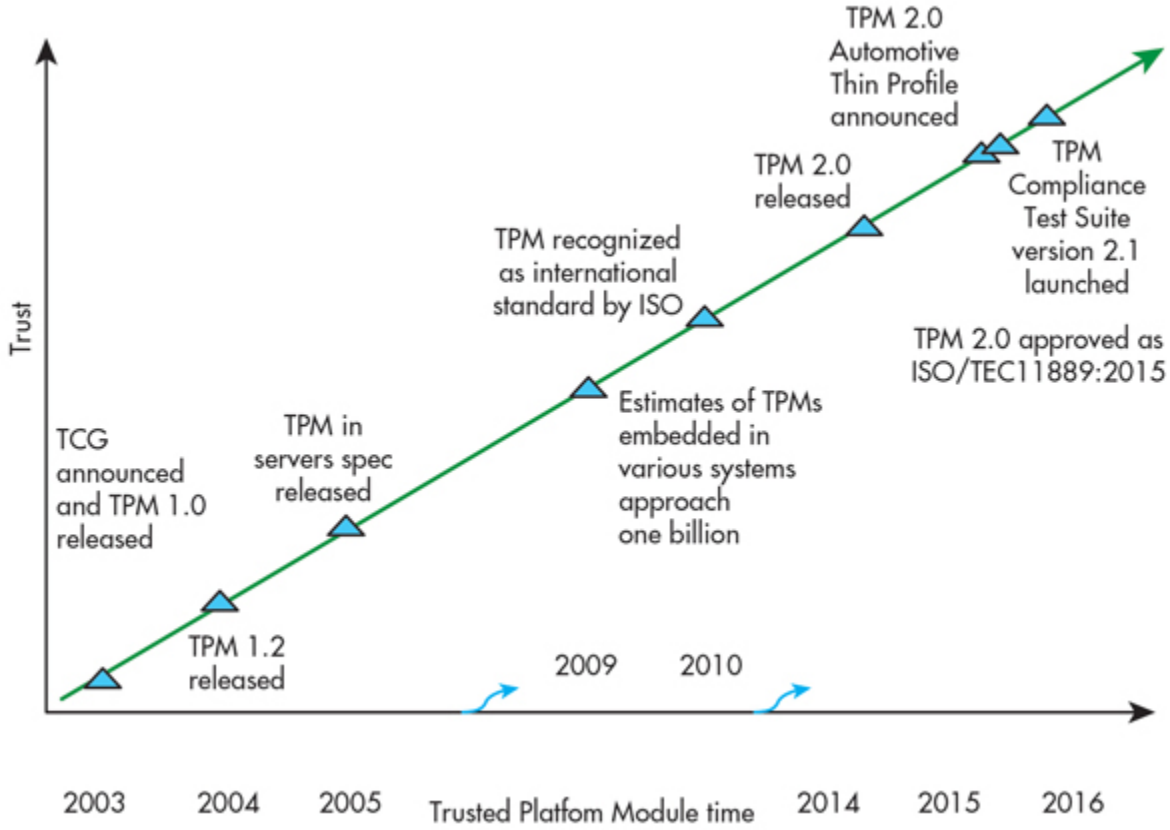

One of the major TPM milestones occurred in 2015, when JTC 1, a joint committee of the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC), approved the Trusted Platform Module Library Specification 2.0 as the International Standard ISO/IEC 11889:2015, Parts 1-4. The ISO seal of approval can create a compelling reason to investigate and implement the TPM in many products, including those in the embedded space.

The recent launch of the TCG TPM Compliance Test Suite Version 2.1 for implementations of the TPM 2.0 Library Specification will ensure TPMs based on the specification provide required security elements and consistent implementations.

For those aware of the initial TPM from over 10 years ago, it has come a long way. History buffs need to reset their understanding of its initial real and/or perceived shortcomings and weaknesses.

Trust Concepts

Trust and security are frequently used interchangeably, but they are different. A secure system with security measures that can be bypassed without detection cannot be trusted. A trusted system that has undiscovered software vulnerabilities is not secure.

To be trusted, a system will behave as expected and possess the ability to identify and communicate that something is wrong. A hardware-based root of trust (RoT) provides far greater protection than software-only protection techniques.

For Trusted Computing, the TPM is a hardware-based RoT that provides fundamental services such as strong identity, secure storage, integrity measurement, and more. The benefits of implementing a TPM RoT to generate random numbers, store and use long-term keys, and verify system integrity include reduced risk of compromise of long-term keys and undetected system compromise.

erview of TPM Architecture and TPM 2.0

In the PC world and server equipment area, the TPM has been very well received. Today, it's essentially an integral part of many companies' authentication, health checks, and more. This hasn't always been the case, though. When it was initially introduced, the TPM required activation by the user. Activation typically meant that someone in the information technology organization had to physically set up the TPM in each machine.

In the early years, the deployment software and tools didn't exist to expedite and simplify this process. This had both good and bad aspects. If problems occurred with password authentication, a user could be locked out of their own system.

However, the implementation, maintenance, and diagnostic aspects created a level of effort that was hard to justify. As a result, many companies were willing to take the chance that they would not be the ones being used as examples of successful hacking.

As more supplier companies recognized the implementation and usage challenges, software and other tools were developed and offered by TCG member companies to simplify the activation of the TPM, including remote and more automated activation. This greatly improved the situation, but it still required additional changes.

TPM 2.0 (officially, Trusted Platform Module Library Specification, Family "2.0," Level 00, Revision 01.16—October 2014) has made significant improvements. Recognizing its value for improved security, TPM 2.0 was accepted as an international standard, namely ISO/IEC 11889:2015, in 2015. With the international community behind the TPM, it creates a substantial reason and the right time to act for those who have been hesitant to implement higher security measures.

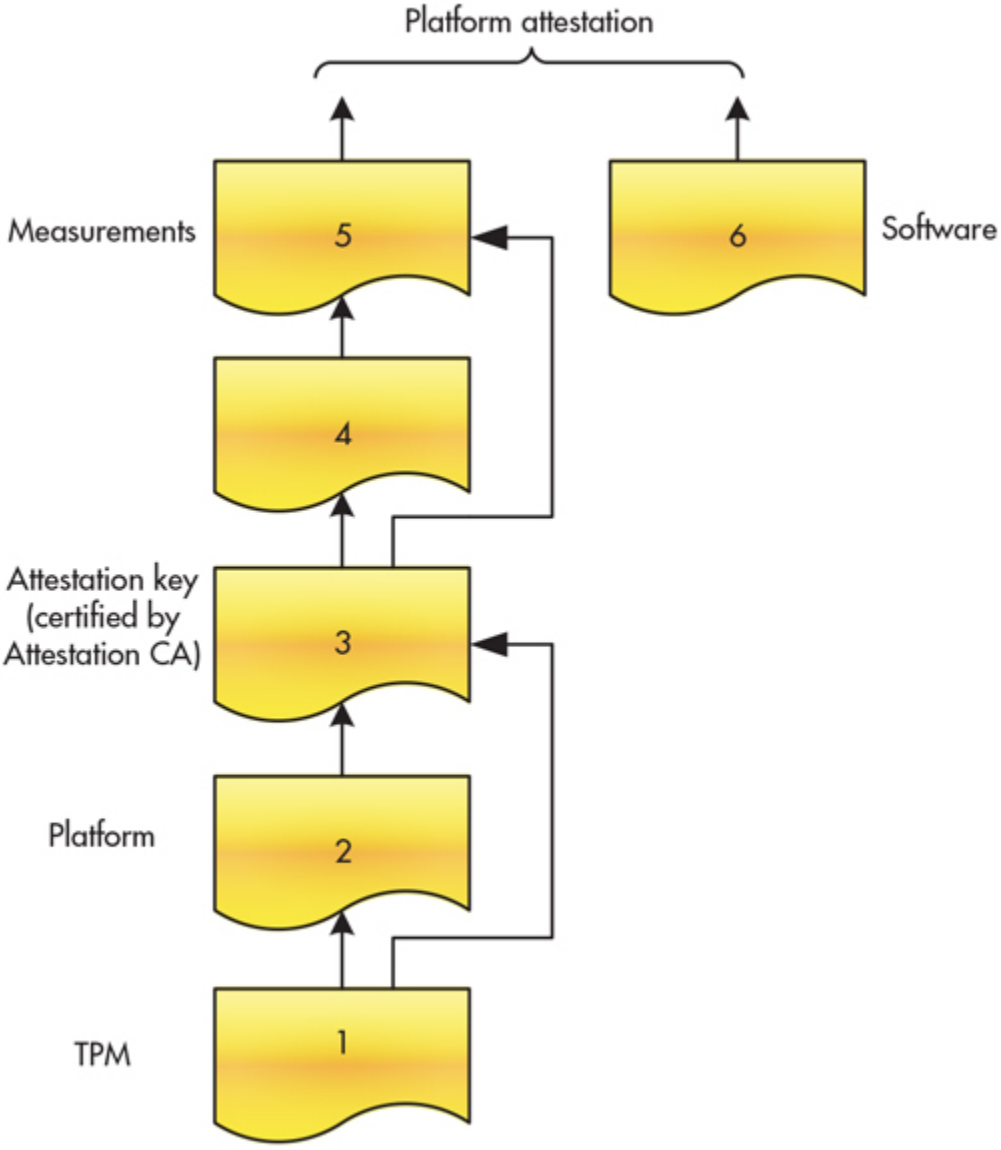

To establish and verify trust, the TPM employs a hierarchy of attestation levels. At the first level, an external entity attests to a TPM to verify that the TPM is genuine and complies with this TPM specification *(Fig. 2)*. This attestation uses an asymmetric key embedded in the TPM and a credential that vouches for the public key of that pair.

At step 2, an external entity attests to a platform to verify that the platform contains a Root of Trust for Measurement (RTM), a genuine TPM, plus a trusted path between the RTM and the TPM.

Attestation of types 3 and 4 use a key to sign the contents of shielded locations. At level 5, a trusted platform attests to a measurement to verify that a particular software/firmware state exists in a platform. At level 6, an external entity or third party attests to a software/firmware measurement to verify software/firmware.

Figure 3 shows the overall operation of the TPM and the functional units required for its operation.

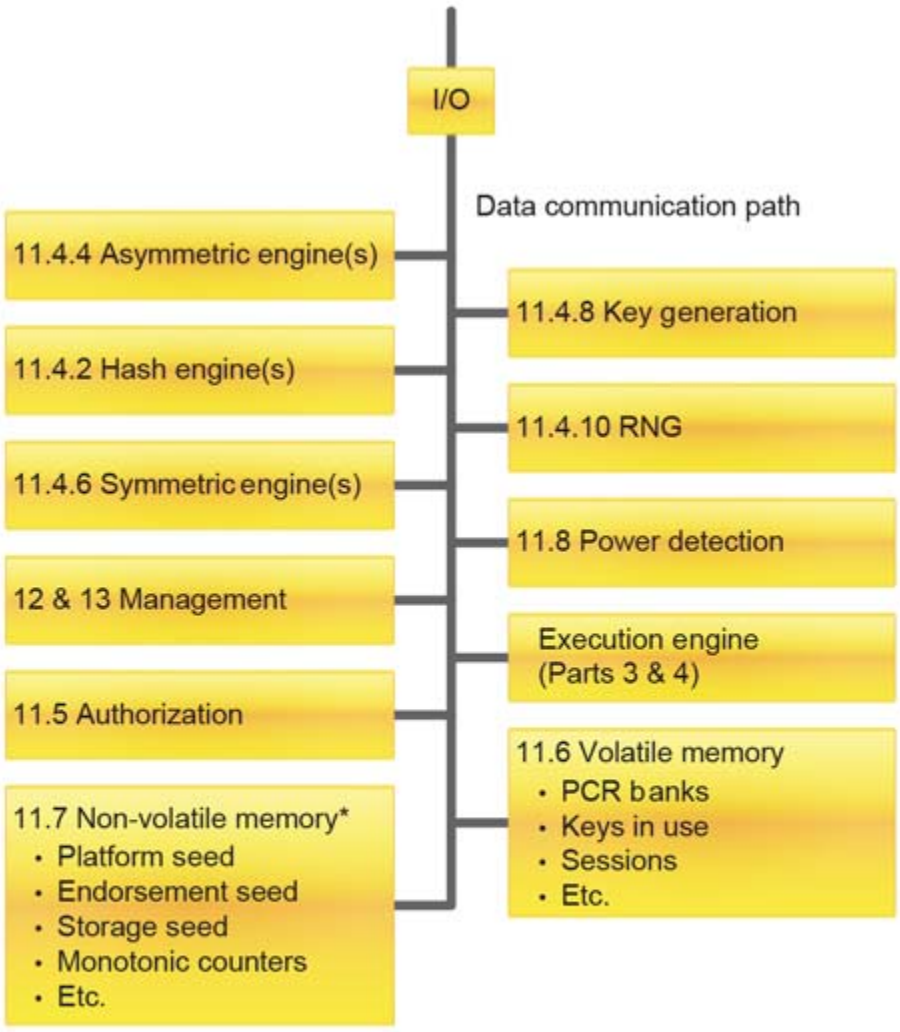

* NV memory may be provided by a system chip with the data going to/from NV in a protected form. What is kept in the "TPM" in that case is a cached copy of the NV contents.

Making the TPM Embedded-Friendly

While the TPM has been well received in certain industries, it's a rather large component for some applications. Recently, a TCG committee actively addressed this limitation. As a result, there's an industry effort to make it easier for system designers to establish trust at the microcontroller (MCU), system-on-chip (SoC), and other component levels.

Because of the cost associated with the TPM, TCG is working closely with chip manufacturers to obtain the most condensed version of TPM hardware-based security. At the same time, the team is making it easier to integrate the TPM into an embedded system.

To use the TPM for strong identity or measurement, the spec defines a common infrastructure with well-defined expectations of how to establish a shared secret or key with the device based on attestation that can be performed with the manufacturer in the back end.

To avoid the need for embedded-system designers to learn a different security scheme for each MCU, SoC, or other IC, the committee wants to establish a common way for secure boot, measured boot, and establishing identity that's a standard fabric of the controller. If this is accomplished the same way on several controllers, in

higher layers of the software, the TPM can rely on specific information being there. To do this, TCG is working on defining measured and secure boot in a higher-level language so that the resulting constructs are very similar.

In addition to Atmel, Infineon, and STMicroelectronics, which have been TCG members for many years, the committee is reaching out to other MCU suppliers. Therefore, once the new embedded TPM spec is published, other users of TCG standards can interface with as many TPMs as possible in a variety of systems.

Architecting an Embedded System that Includes a TPM

The minimalistic approach to establishing a root of trust can be applied to an MCU, SoC, ASIC, DSP, FPGA, or virtualized processor—anything with a strong identity and a variable part that gets loaded dynamically afterward.

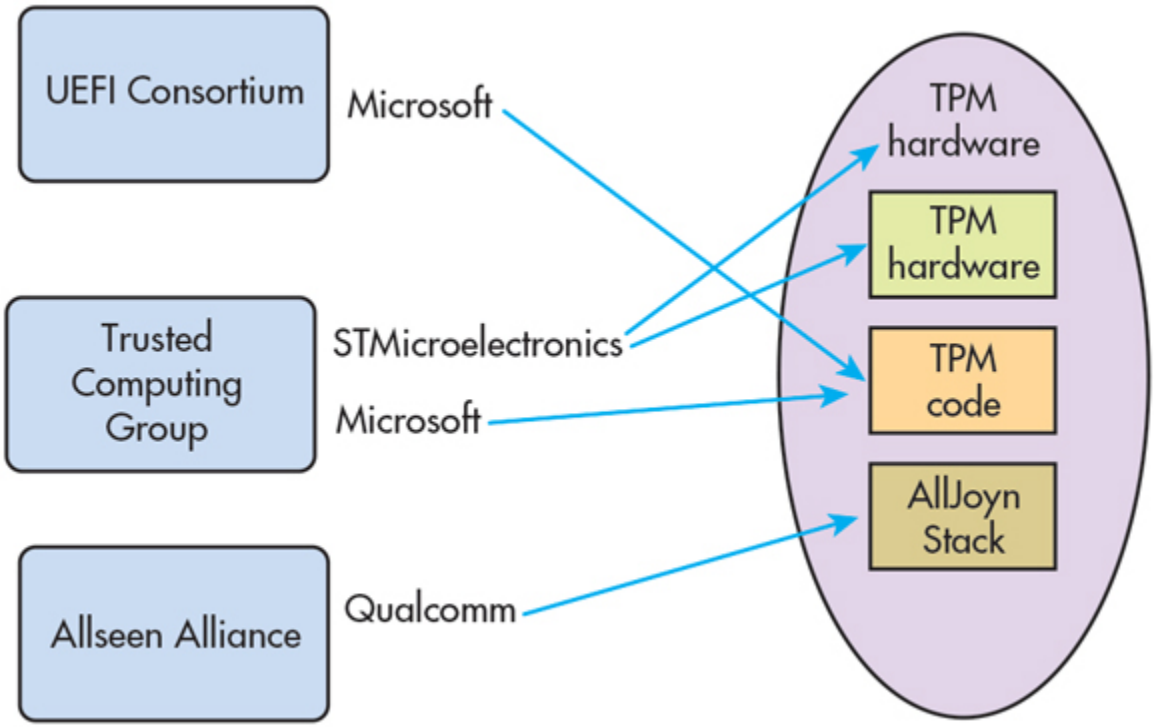

Figure 4 shows a prototype being built by Microsoft with help from STMicroelectronics to provide the chip hardware and the Allseen Alliance to provide a lightweight AllJoyn stack to discover nearby devices. With AllJoyn communication technology, devices within a proximity network can communicate with each other without explicitly knowing each other. The communication stack layers on top of the TPM. This enables it to discover a device in a network, ask for its identity, and through the back end of the cloud, determine if its identity is actually trusted through attestation. When the answer from the cloud verifies each device and the trustworthiness of the software on each, a key is sent to establish communication through a secure channel.

Examples of the TPM at Work

A smart building takes advantage of automation to save power and make it more efficient, increase convenience and make occupants comfortable, and collect data regarding usage and even information from video cameras/systems. For a more secure system, the camera system could have a TPM in it.

When the system is triggered, the TPMs on both ends of the transmitted data between the camera and monitoring console can communicate with each other to verify that there's been no tampering with the camera.

like a tampering or spoofing event that's been depicted in many movie (also known as the Ocean's 11 or a n-in-the-middle attack), attention is not diverted with the TPM linkage and the robbery can be averted.

Remotely controlling a heating appliance in an industrial setting or even a home requires policies and trust. For a furnace, certain policies need to be enforced between the gas valve and the ignitor, so that ignition doesn't occur if the gas valve is open longer than three seconds, thus avoiding an explosion *(Fig. 5)*. This policy needs to be an integral and unchangeable part of the gas valve and ignitor.

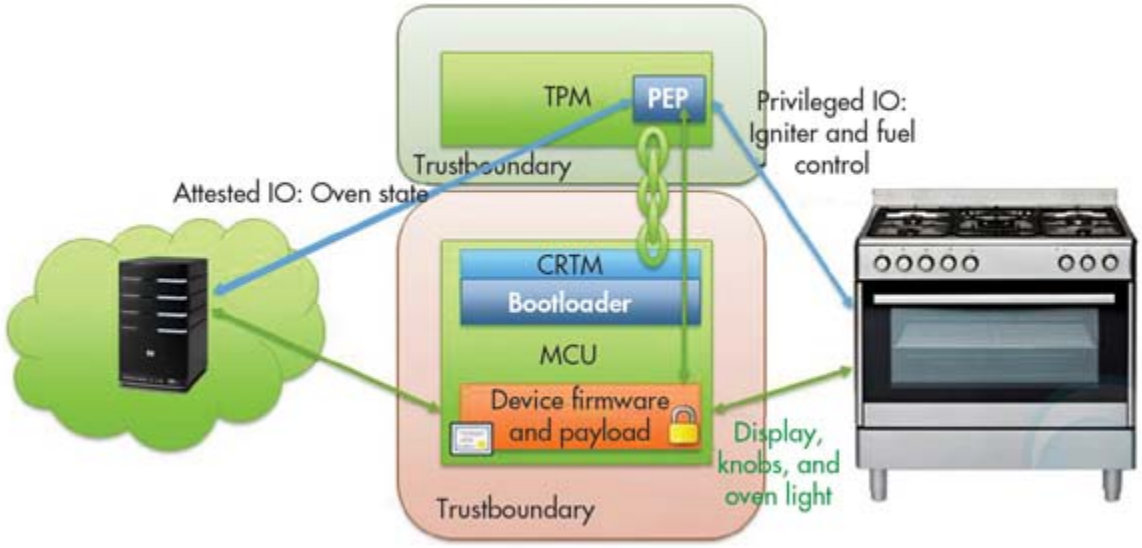

Windows software cannot be the go-between to prevent a problem because this would create a flexible application. However, if an MCU with a TPM controls the ignition process, it can have a policy firmly established; one that can't be changed. If someone flashes bad firmware into it, Windows can use attestation to determine that it should not use the ignition process because the measurement says it's not trusted. Thus, no further communications occur.

Implementing Improved Security

While not hacked to date, which doesn't mean it couldn't happen in the future or imply that it's failsafe, the additional effort to breach or penetrate the added defense provided by the TPM is well worth it. Developers of the specification, who represent the best experts at companies that pioneered software and hardware for the computer industry, have their collective reputations on the line.

It's one of the first steps that companies can take to prevent outsiders from taking control of a company's computers, networks, customer information, intellectual property, and control systems. With TPM 2.0-compliant products already available from four chip manufacturers, trusted communications between things is on the rise.

Reference:

Trusted Platform Module Library Specification, Family "2.0", Level 00, Revision 01.16—October 2014

About the authors:

Stefan Thom, Principal Software Development Engineer and Platform Security Architect, Microsoft Corp. Previously, he contributed to Xbox One Console Security and Attestation, Windows BitLocker Drive Encryption, Windows Virtual Smart Card, Windows Attestation, and Firmware TPM initiatives.

ve Hanna, Senior Principal at Infineon Technologies, is a frequent speaker at leading conferences such as A. He has authored numerous technical papers and standards including IETF RFCs 2730 and 5793 and Trusted Computing Group IF-IMC and IF-IMV. He is a member of IETF's Security Area Directorate and holds 43 U.S. patents.

Stacy Cannady, Engineer Technical Marketing, Cisco, has worked in the field of trusted computing for a number of years. As a subject-matter expert in trusted computing, his responsibilities require an in-depth understanding of the trusted-computing market, including advances in hardware and software security as well as vendor and customer market dynamics.

Source URL: http://electronicdesign.com/embedded/standardizing-trust-embedded-systems