Continental Automated Buildings Association

Information Series

IS 2005-24

Building Trust for Embedded Systems Starting at the Platform Foundation Layer

CABA
www.caba.org

certicom™

Building Trust for Embedded Systems

Starting at the Platform Foundation Layer

March 2005

The Need for Secure, Smart Devices

Security is becoming an essential requirement for all computing devices as we begin to trust and rely upon them to control our environments and protect our information, identity and ultimately our lives. In fact, the Trusted Computing Group (TCG) initiative was established in recognition of that dependence and the need to develop a common secure computing framework.

As security vulnerabilities and risks grow, the concept of trusted computing has spread quickly from personal computers (PCs) to the mobile (handheld/wireless) environment. And now awareness of security issues is impacting the wider embedded device market.

TCG has developed a comprehensive model of how trusted devices should function. It has produced specifications for an architecture, a set of functions and interfaces that provide a baseline trust model for a wide range of computing platforms.

New working groups are adapting these specifications for a wider range of devices such as PDAs, mobile handsets and network devices as well as platform peripherals. What's required in these adaptations is a solution that scales the TCG model to one appropriate to a wider range of embedded devices that are constrained in memory, processing power and bandwidth available. This pragmatic approach will give device manufacturers the ability to offer the device security that their customers are increasingly demanding – scaled to meet their needs.

> **"Trust is the expectation that a device will behave in a particular manner for a specific purpose. A trusted platform should provide at least three basic features: protected capabilities, integrity measurement and integrity reporting."**
>
> – Trusted Computing Group

What is a Trusted Device?

A “trusted device” in this context is one which is known to be running authenticated, authorized, known-origin software in a secure computing environment. New mobile handsets and PDAs are just the tip of the iceberg in terms of devices where “trust” is relevant.

Equally important to consider are “closed” devices like a voice over IP (VoIP) terminal adapter, an Internet appliance that can download games, a commercially available Linux-based home gateway with open source software, or a telecom server blade running an application on Carrier Grade Linux.

For example, take the case of an FDA-approved heart defibrillator – a sophisticated device with complex electronics and software. Properly operated and maintained, it provides a clear benefit to people in cardiac arrest. While the firmware in a defibrillator is rigorously tested, bugs are inevitable – and they can be deadly. After equipment is deployed, it’s important that checks are done to ensure that a device hasn’t been recalled and that its firmware is still valid. Each device that needs to be recalled costs the manufacturer money at a minimum. In this case, failing to recall a device could cost someone dearly. A trusted framework for medical devices reduces risks, saves money, and could save lives.

While the risk to the user of a broadband gateway running Linux is not life threatening, the popularity of “patches” for these devices from the open source community can create vulnerabilities. What happens, for instance, if a software patch that is supposed to provide Quality of Service (QoS) support actually creates a hidden hole in the firewall – putting sensitive information on your network at risk 24x7.

In a trusted computing scenario, devices are designed to protect against vulnerabilities. Checks are done automatically, with steps taken to ensure that the firmware hasn’t been recalled, patches are authentic, devices are in good working order, and able to receive new certified software updates as soon as they are available.

While device functionality varies widely, embedded platforms need only some basic mechanisms to establish trustworthiness:

1. **The ability to securely store cryptographic keys and sensitive information.**

2. **The ability to use authentic code, configuration data and user identities.**

Cryptography, and public-key systems in particular, provide the foundation and enable the required functionality for a trusted platform. This overview will look at how you can bring trust to embedded systems and outlines a solution to address these requirements. But first, let’s look at some of these embedded markets that require security and the challenges that exist when adding security to a device.

The World of Embedded Security

Many real world examples exist, across a number of market segments, that establish security threats in embedded devices. Although devices differ widely, every market segment has two things in common:

1. **Devices are increasingly networked; and**
2. **Network connectivity and flash technology are used extensively to enable in-the-field firmware and configuration updates.**

Consumer Electronics Devices

Consumer products, especially those associated with service providers, such as mobile handsets, are well-known targets of hackers. In the mobile phone industry the impact has driven new requirements for secure boot and secure software update. Now, with the advent of mobile digital rights management (DRM), security is the center of attention.

Other segments, such as broadband customer premise equipment (CPE) like cable modems and IPTV set-top boxes, are exposed to similar risks from fraud, theft of service, denial of service, spyware or other vulnerabilities. As these devices become platforms for new value-added services, security threats create a liability for both consumers (loss of privacy) and service providers (loss of revenue) alike. Everyone in the value chain is impacted by device security.

The common requirement among these devices, especially ones subsidized by service providers, is the need for secure boot, a service lockdown mechanism, remote provisioning, service activation and firmware updates. The challenge original equipment manufacturers (OEMs) and service providers face is that they already are preoccupied with their core business.

How can they deliver robust, reliable security on a multitude of platforms and still innovate in their core technology?

Medical Equipment

As noted above in the defibrillator example, medical electronics also have a number of security, privacy and safety issues. Complicating those requirements is that these devices are increasingly networked – creating a device security vulnerability.

The medical market's regulatory and safety requirements create a compelling case for secure software download, device management and configuration control. Given the cost and safety issues of recalling bad firmware, using digital certificates to track and control critical care device software is an ideal way to leverage trusted computing technology.

Beyond medical devices, the US government's 1996 HIPAA (Health Insurance Portability and Accountability Act) drives security requirements down to printers and data storage devices – adding encryption requirements in a wide range of products.

Industrial Automation

Although some equipment can and does access the Internet, industrial automation devices generally operate on closed networks. Of special concern are critical plant control and instrumentation systems (e.g. chemical processing plants or refineries, power grid, water works). They very frequently use remote firmware upgrade technologies, and often include remote terminal facilities, which make them vulnerable to embedded firmware sabotage.

Depending upon the risk assessment, retrofitting the software in these devices to be more secure could be a very good idea. The challenge is to do so in a cost-effective manner without redeploying expensive new hardware. A secure software approach using trusted devices concepts could make these systems secure from tampering without that cost.

Telecoms and Network Infrastructure

A variety of telecom control-plane and network management systems (NMS) and billing applications are run with an embedded computing infrastructure in secure network operation centers (NOCs) and co-location facilities throughout the world.

Network security is not taken lightly. London's Telehouse sites, purpose-built co-location facilities, are said to be among the most secure buildings in Europe. The telecom industry's embedded computing software platforms need to be equally secure.

While management and provisioning is generally done on a close network, remote software update is a security risk that requires authenticated software.

Military and Aerospace

The military/aerospace market covers a wide range of products characterized by specialized design requirements for ruggedness and security. In fact, there are specific security requirements for products used by government agencies, including Federal Information Processing Standard (FIPS) 140-2 Validation and the use of elliptic curve cryptography (ECC) for mission critical security information and sensitive but unclassified data as specified by the National Security Agency (NSA).

The challenge for military/aerospace designers is how to build mechanisms that can survive the long service lifetime of military and aerospace equipment. The military's Joint Tactical Radio System (JTRS) program, for example, started several years ago but radios designed for this system must remain secure and interoperable for over 25 years. Recently, the NSA selected ECC as the public-key technology for government communications, support for elliptic curve digital signature algorithm (ECDSA) for digital signatures and elliptic curve Menezes-Qu-Vanstone (ECMQV) for key exchange are required cryptographic schemes for this market.

The Business Challenge of Security

Pervasive computing is driving the need for secure, smart devices. As reprogrammable embedded devices become networked to each other, hackers are exposing and exploiting vulnerabilities to corrupt system software. To protect against these potential attacks, device manufacturers need to address platform security.

Therefore, to maintain or grow market share, vendors should embed security to differentiate their offerings. However, vendors face an extremely competitive market where they must balance innovation costs against profit margins. Any additional functionality, including security, must add value and sell more devices or it will be cut in favour of some feature that will drive more revenue and profits.

The bottom line is that innovation adds to development costs in an environment where time-to-market is exceedingly important. Embedded device security is moving from an optional feature to an established requirement. It demands specialized skill sets that can lie outside the core expertise of vendors and integration which must be repeated with each new design. Embedded security must therefore be reusable and portable across product families and generations to avoid hindering release schedules.

Differentiating through embedded security also allows device manufacturers to help customers address the growing problem of theft and fraud in the industry. The US Secret Service estimates losses from telecommunication fraud at more than a billion dollars each year with cell phone cloning representing a large part of this figure.[1] Device manufacturers that help their customers reduce losses through security will be well positioned to become the supplier of choice.

Security is also an absolute requirement for lucrative markets such as the US Government. According to a US Department of Defense directive issued in March 2003, US government departments are required to use solutions that are FIPS 140-2 Validated for sensitive data communications. FIPS 140-2 Validated solutions are also increasingly recognized by other industries that require guaranteed security: including financial and healthcare. While these markets are attractive, the validation process represents a significant barrier to entry.

Embedding a previously validated cryptographic module can further differentiate a device and provide quick access to these markets. In fact, using off-the-shelf security addresses key requirements for a specific market quickly and easily. This brings us to the point of finding a solution that addresses the technical challenges of integrating security and building a trusted platform.

[1] http://www.secretservice.gov/financial_crimes.shtml#Telecommunications

The Technical Challenge of Security

One of the key principles for secure system design is that security should be built into the platform foundation. Otherwise, developers add security patches to address specific vulnerabilities within individual applications. These additions may or may not be interoperable and lead to uncertainty regarding the overall security of the platform.

Without a system-wide approach to security, latent vulnerabilities may only be recognized after a successful attack. If the original patch was added on, these weaknesses can be difficult to address. This is also true when trying to extend patchwork security to cover the new features of an evolving design.

Embedding security from the ground up adds measurable value. In addition to limiting vulnerabilities by design, new threats can be addressed as they appear. A ground-up approach also allows developers to facilitate the migration of security information by users from one platform to another when they upgrade.

However, this is only part of the story. The device itself needs to be secured so that all applications and protocols run on a trusted platform. Within the TCG you can find the types of functionality that will be required in embedded systems. But this functionality must be streamlined and modular to work in a variety of embedded systems.

Beyond traditional security integration challenges, there is a significant range of processing power, memory and bandwidth limitations in all resource-constrained devices. As a result, a scalable security system is required and must meet the available memory and battery power requirements, without impacting performance. The problem is compounded by processor and bandwidth-intensive cryptography calculations that negatively impact areas such as performance and battery life.

The level of effort required to optimize the size and performance of a security architecture creates another hurdle in the form of higher switching and time-to-market costs when a new chipset is required. These costs can be incurred repeatedly across the entire product line. Rather than duplicating effort with each device, you need a solution that is portable across multiple chipsets and offers a future-proof security strategy that allows manufacturers to use their existing code when migrating to a new chipset.

Before we look at a solution that addresses these technical and business requirements, let’s take a closer look at the types of security functionality that need to be embedded in all devices.

A Trusted Platform for Embedded Systems

The Root of Trust and Secure Boot

A trusted computing environment must be designed from the boot process upwards to be secure. The device must have mechanisms to validate the hardware environment and code signatures for any modifiable software or firmware, including boot code and configuration files.

Boot code can be stored in on-chip one-time programmable (OTP) memory or reprogrammable memory, typically flash. It is when the code can be reprogrammed that security risks occur. To ensure system integrity, code signatures must be authenticated, and sensitive information such as the keys stored in non-volatile memory must be encrypted. If the code signature cannot be authenticated by the key stored in memory, the code is deemed untrusted.

This functionality provides a “root of trust” for the device. In the PC world, secure storage for encryption keys is done using secure, tamperproof non-volatile memory or a trusted platform module (TPM). The TPM holds keys injected by the manufacturer to allow the device manufacturer, device owner and users to be authenticated before accessing sensitive information. In the embedded device the same facility could work using a TPM, similar on-chip hardware, or even software − to authenticate device software and control system access.

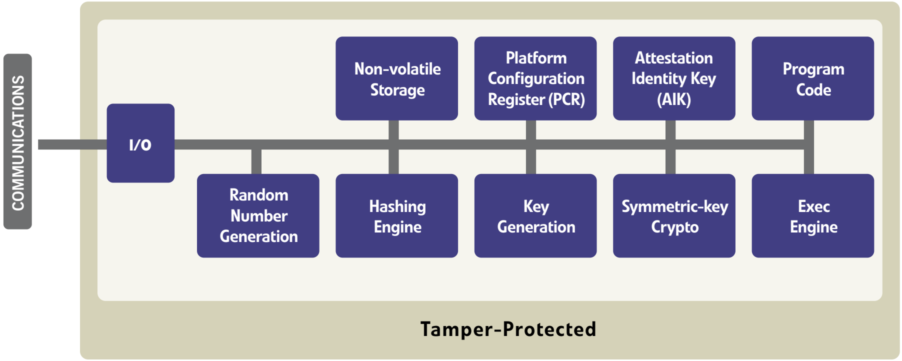

Figure 1: Trusted Platform Module (TPM) showing functionality required for secure storage and secure boot.

Extending the Chain of Trust

In the simplest case of a closed system it is only the main code image that must be authenticated. Many embedded devices today, however, have several modules or code images and configuration files that must be loaded separately. Critical to extending this root of trust beyond the boot process is a chain of hardware and software authentication for each component.

This measurement or code authentication is required to validate a secure runtime environment. The technology for establishing integrity in firmware and the runtime code it loads is relatively straightforward.

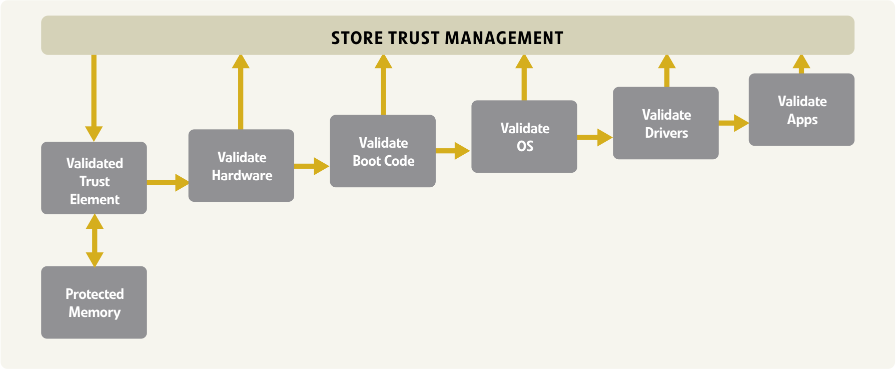

Figure 2: Extending the root of trust beyond the TPM through the secure boot process.

Authentication can be as simple as executing a series of routines which perform a cryptographic hash of the code image and then compare it to a known value - such as one placed in secure storage. Authentication complexity increases when dynamic loading is permitted but the same principles apply.

In a complex (but not fully open) embedded device, an authentication module would manage the loading process, digitally authenticating and exchanging information amongst various management and control entities in the device. In an open application environment, such as a Windows Mobile, Symbian or an embedded Linux environment, the authenticator would need to be part of the kernel, where it could help manage restricted access to kernel resources which are typically controlled by user privilege levels.

The device itself should load only properly authenticated code unless there are specific hardware provisions such as MMU and kernel level security which allow untrusted applications to have safe access to non-critical resources. Additional tools are required to provide code and configuration file signing facilities. In more complex devices once boot has been established there will also need to be facilities for secure download. There may also need to be facilities to manage lists of revoked certificates for code and device configuration.

In the end, security requirements are driven by market need – whether to protect sensitive information, safety-critical equipment, prevent theft-of-service or intellectual property. Whatever the reason for establishing a trusted environment, the core requirements are very similar. What is different is the hardware and software value chain that must be leveraged to pursue these markets.

The Future

Embedded device vendors will soon face an unforgiving marketplace in which they either hold a competitive advantage through products differentiated by comprehensive security or they lose to those competitors who do. Finding a secure architecture to build in to their platform foundation layer will allow these vendors to seize a low-cost market opportunity to lock in continuously evolving security without increasing the time-to-market of their rapid release products.

About Certicom

Certicom Corp. (TSX: CIC) is the authority for strong, efficient cryptography required by software vendors and device manufacturers to embed security in their products. Adopted by the US government's National Security Agency (NSA), Certicom technologies for Elliptic Curve Cryptography (ECC) provide the most security per bit of any known public key scheme, making it ideal for constrained environments. Certicom products and services are currently licensed to more than 300 customers including Motorola, Oracle, Research In Motion, Terayon, Texas Instruments and Unisys. Founded in 1985, Certicom is headquartered in Mississauga, ON, Canada, with offices in Ottawa, ON; Herndon, VA; San Mateo, CA; Saltsjo-Boo, Sweden and London, England. Visit www.certicom.com.

Certicom White Papers

To read other Certicom white papers, visit www.certicom.com/whitepapers.

The Inside Story

Many Happy Returns: The ROI of Embedded Security

Wireless Security Inside Out (authored by Texas Instruments and Certicom)

Welcome to the Real World: Embedded Security in Action

Sum Total: Determining the True Cost of Security

The Elliptic Curve Cryptosystem for Smart Cards

Elliptic Curve DSA (ECDSA): An Enhanced DSA

Formal Security Proofs for a Signature Scheme with Partial Message Recovery

Postal Revenue Collection in the Digital Age

An Elliptic Curve Cryptography Primer

ECC in Action: Real World Applications of Elliptic Curve Cryptography

Using ECC for Enhanced Embedded Security: Financial Advantages of ECC over RSA or Diffie-Hellman

Good Things Come in Small Packages: Certicom Security Architecture for Mobility

Meeting Government Security Requirements: An Overview of the Certicom Security Architecture for Government

The Benefits of Digital Signatures for Reducing Bank Fraud Losses: An Overview of the Certicom Security Architecture for Check 21

Contact Certicom

Corporate Headquarters

5520 Explorer Drive
Mississauga, Ontario
L4W 5L1
Tel: +1-905-507-4220
Fax: +1-905-507-4230
E-mail: info@certicom.com

Sales Offices

Worldwide Sales Headquarters

1800 Alexander Bell Dr., Suite 400
Herndon, Virginia 20190
Tel: 703-234-2357
Fax: 703-234-2356
E-mail: sales@certicom.com

Canada

5520 Explorer Drive
Mississauga, Ontario
L4W 5L1
Tel: 905-507-4220
Fax: 905-507-4230
E-mail: info@certicom.com

Ottawa

84 Hines Road
Suite 210
Ottawa, Ontario
K2K 3G3
Tel: 613-254-9270
Fax: 613-254-9275

U.S. Western Regional Office

1810 Gateway Drive, Suite 220
San Mateo, CA 94404
Tel: 650-655-3950
Fax: 650-655-3951
E-mail: sales@certicom.com

Europe

Golden Cross House
8 Duncannon Street
London WC2N 4JF UK
Tel: +44 20 7484 5025
Fax: +44 (0)870 7606778

Engelska Huset
Trappv 9
13242 Saltsjo-Boo
SWEDEN
Tel: +46 8 747 17 41
Mobile: +46 70 712 41 61
Fax: +46 708 74 41 61

www.certicom.com