Achieving DRM Robustness

securing the device from the silicon up to the application

A Certicom White Paper
November 2005

certicom™

Introduction

Digital Rights Management (DRM) is a critical business enabler for the digital content ecosystem. The ecosystem's success depends very much on robust end-to-end security for content protection. End-user device security is a particular concern because this is where content is rendered and DRM is most vulnerable to circumvention.

Providing DRM robustness in end-user devices is not a simple process. There are a number of issues to contend with, depending on the type of device, the level of integration with the application environment, and the encryption technologies underpinning the DRM agent. Security involves the integration of hardware and software to leverage hardware protection mechanisms, using safeguards wherever sensitive keying material could be exposed.

This paper will explore how providing a robust DRM implementation requires:

- securing all software components, including the DRM agent, with a comprehensive security architecture;
- securing the silicon; and
- securing the device manufacturing environment.

certicom™

DRM Schemes and the Environment They Operate In

Today, many DRM schemes exist, including standards-based as well as proprietary schemes. It remains to be seen which of these will become the defacto standard or if the difference schemes will actually co-exist on the same device.

Whether standards-based or proprietary, securing digital content with a software agent and mitigating security threats is a challenge, because of the nature of the environment in which they operate.

It is easiest to comment on the Open Mobile Association's version 2 DRM specification (OMAv2 DRM) because it is based on open standards. With its standards-based approach, OMAv2 DRM does not attempt to benefit from security by obscurity as proprietary DRM schemes do, because there is no secrecy of design.

Proprietary DRM schemes from Microsoft, Apple, RealNetworks and others do attempt to benefit from security by obscurity – but trying to keep the implementation details a secret doesn't make them any more robust. Any widespread DRM deployment presents a very high profile target for attack.

DRM agents in open application platform devices running Windows, Symbian or Linux are particularly vulnerable. These devices are more accessible to would-be attackers, and because of enhanced connectivity, successful agent compromise in one device might be propagated very quickly to other devices.

Mass-market multimedia handsets with closed application platforms, where content-based business models have real potential, are vulnerable too – simply from a different class of attacker with more elaborate tools. While the typical consumer does not have access to a JTAG debugger or sophisticated tools, professional hackers do. In fact, it wasn't too long ago that a conditional access vendor sued its biggest competitor for compromising its pay TV conditional access system.

Achieving DRM Robustness: Securing the Software

In order for a software DRM agent implementation to be robust, precautions must be used to protect the DRM agent's critical system code and data - especially private keys and digital certificates.

Security also needs to protect against DRM database rollback – thwarting users attempting to "extend" their license rights by manipulating data stores. DRM agents are responsible for protecting these rights objects and stateful files, but they rely on system services to do so. System services, therefore, should be DRM aware.

Applications that consume digital content also need to be trusted – and secured. They should be authenticated with a DRM agent before being granted its services. This includes components such as the media player and file browser. Only "trusted" applications should have access to sensitive DRM resources, and this trust must be based on verifiable security properties such as application authentication and integrity checking.

The need to tie together the security of multiple applications simultaneously (browsers, media players and various DRM agents), and the need to protect against other threats such as device cloning, malware, and viruses drives the requirements for a complete security architecture.

To frame that security architecture we take a page from the Trusted Computing Group's work for mobile devices. Security in complex mobile devices involves a transitive[1] trust process – where a trusted boot program establishes trust in a kernel image, and the kernel image establishes trust in application components.

In order to anchor device security, the transitive model requires a core root of trust – essentially an initial load of known, trusted code. This is achieved by storing boot software and trust proofs (cryptographic hashes or root certificates) in tamperproof memory.

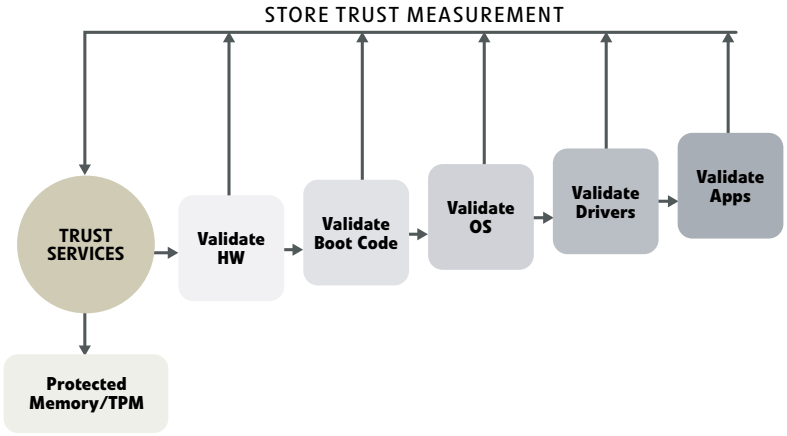

Figure 1: **The Secure Boot Process**

Root of trust boot code can be made immutable, but for a variety of business reasons – from over-the-air updates to factory rework, OEMs and operators need the flexibility to update system software. Digital signatures must be used to extend the chain of trust to updated software components. Using well-established public-key methodologies can help lower system deployment and maintenance costs while maintaining very strong device security.

The process of using digital signatures to verify components and configuration data can be as fine-grained as desired – from a single signed monolithic system image to signatures on individually signed drivers.

With digitally signed code, dynamically loadable modules and drivers can even be developed and deployed by trusted third parties, adding new flexibility to the value chain. A platform security architecture that supports public key can help nurture the growing digital content ecosystem.

[1] Tran-si-tive \'tran(t)s-*t-iv, 'tranz-; 'tran(t)s-tiv\ adj 2 : being or relating to a relation with the property that if the relation holds between a first element and a second and between the second element and a third, it holds between the first and third elements

certicom™

Achieving DRM Robustness: Securing the Silicon

Features like true random number generators (RNG) and large number operators or finite field multipliers are frequently used to improve system performance. When properly designed, they can improve security robustness as well.

Such hardware should include facilities for secure key operations — shrouding sensitive private key operations and allowing system software to address keys by reference rather than letting them cross the memory bus in plain text form. If shrouded memory is not available, system software should restrict access to crypto hardware via kernel access controls.

Key storage presents another problem. Specialized crypto hardware on system silicon with shrouded memory does not typically have enough RAM to store keys for multiple applications. Keys need to be fed into the crypto processor on demand - typically at the beginning of each new crypto operation. Long-lived keys must be in protected or encrypted storage when not in use.

The problem of key exposure can occur when multi-threading is supported on a crypto interface. Device drivers need to protect keys and contexts, for example, during a system context switch.

Another way to support key protection is for the hardware design to wrap keys leaving the crypto module. If that facility isn't available it's critical to at least store the key in an encrypted file or memory object. A well-thought out security architecture should handle that requirement automatically rather than leaving it to system programmers.

The benefits of hardware for crypto acceleration are clear. Designed into the system ASIC, crypto hardware can speed computationally intensive hashing or digital signature verification of system software. That is especially important for reducing secure startup time. Specialized crypto hardware can also reduce system power consumption. That's critical for multimedia handsets — they don't generate revenue if they're always in the charger!

While there is a definite trend towards hardware-based crypto, it's important to consider that not all hardware implementations are created equal. A significant amount of expertise is required to implement the complex arithmetic operations required to effectively accelerate public key operations, in either hardware or software. It's not at all uncommon to see optimized software beat hardware implementations like the point addition operations for elliptic curve cryptography if not implemented efficiently. It is also important to create hardware designs which do not leak, eliminating the opportunity for side-channel attacks.

Given the complexity of modern crypto algorithms, it's critical that low-level crypto building blocks are rigorously tested with high-quality test cases. Poor implementations can lead to very difficult to detect problems with the resulting DRM agent.

certicom™

A Practical Approach to the Security Architecture

A comprehensive security architecture should take advantage of the crypto acceleration and security features that chip vendors provide – yet fill in the gaps between the hardware implementation and optimized software algorithms.

A clean low-level abstraction layer is very important when integrating higher-level security protocols and functions because it paves the way for code re-use. Porting security applications like DRM from one chip to another should be easy – and not require weeks of re-working platform specific application code.

Likewise, the security architecture should offer a common trust model to enable faster development of robust DRM implementations. The trust model should allow applications like DRM and media players to interact without exposing developers to the complexities and potential pitfalls of sharing secret keys, doing key management and coding to hardware crypto interfaces.

Achieving DRM Robustness: Securing Device Manufacturing

The final step in creating a secure DRM ecosystem with robustly protected end-user devices is securing the device manufacturing process.

The manufacturing process is a hidden threat to securing the digital content environment – a place where breaches in security would have a very widespread impact. Unique keying or certificate material injected into a device at the time of manufacture acts as a unique identifier, linking it to premium services, content and billing systems. If a device is cloned, content owners lose the ability to control access to their content and service providers lose the ability to charge for services. Thus any unique keying or certificate material that is loaded into the device needs to be secured.

As more vendors outsource the manufacture of devices to contract manufacturers there is a risk of device cloning, where dishonest employees could create handsets with duplicate certificates. Selling compromised devices on the black market would endanger the digital content ecosystem. It is a risk that must be mitigated.

OEM liability is a growing issue – illustrated by the licensing authorities for digital media such as OMA DRM's Content Management License Administrator (CMLA) and the High-bandwidth Digital Content Protection's (HDCP) Digital Content Protection, LLC. The device vendor is liable for millions of dollars if the keys get leaked. Having the wrong or non-unique certificates is also a risk – making it critical to protect the quality of the manufacturing process too with a secure control and auditing process for DRM license administration.

Device vendors need a cryptographically wrapped key injection system that allows them to maintain control over keying material, even when they employ a contract manufacturer. The system should not only control access to the keying information but also provide metering and reports on key usage.

Certicom Security for DRM

To provide a strong foundation on which to achieve robust DRM on the device, and to ensure that the platform security requirements described in this paper are met, Certicom has partnered with key players such as DRM specialist Beep Science and major silicon vendors such as Freescale and Intel. The result is Certicom Security for DRM which provides a foundation for platform security and trust services, as well as securing the manufacturing environment. The solution includes the Certicom Security Architecture and Certicom KeyInject, with the following components:

- A pre-integrated OMA DRM agent, extensible to other DRM schemes
- Embedded Trust Services (ETS) for secure key storage and key management
- Hardware IP Cores to secure media co-processor designs
- Board Support Packages (BSPs) to expose hardware cryptographic providers in leading processors
- Trusted key injection to secure device manufacture

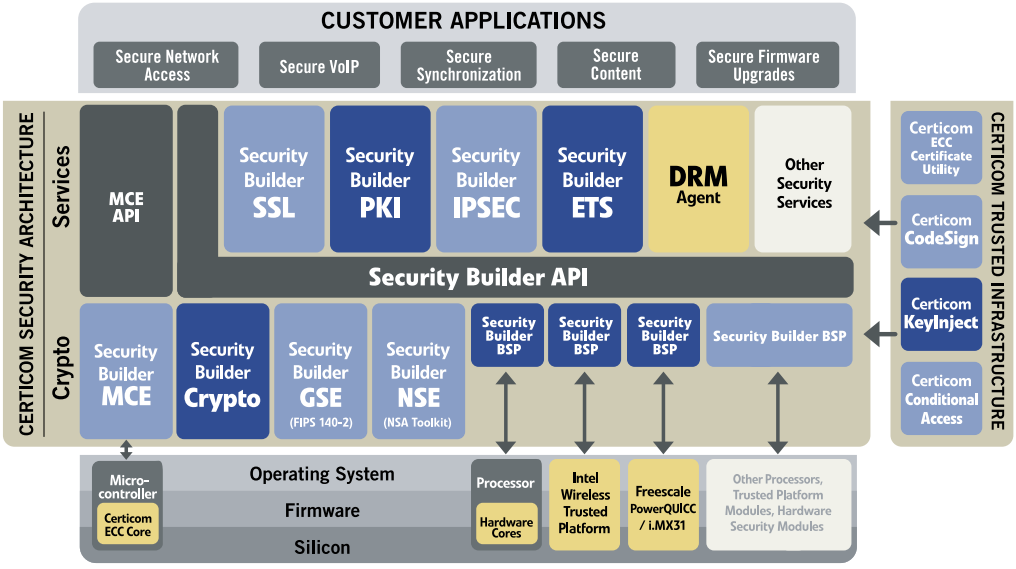

Figure 2: **Certicom Security for DRM**

Certicom Security for DRM helps to achieve DRM security and robustness in devices by:

Securing the Software Components: in addition to pre-integrated DRM agents, the Certicom Security Architecture offers a full range of security services including SSL, PKI and IPSec. As well it offers software cryptographic providers that support both ECC and RSA. All components of the solution are tightly integrated into a trusted platform that ensures protection from the boot process up to the application.

Securing the Silicon: Certicom Security Architecture offers board support packages (BSP) that expose the cryptographic functionality available in hardware, and provides Embedded Trust Services (ETS) to ensure keys are protected if they leave the hardware. Components of the Architecture are unified by a common application programming interface (API) that sits between the security services or applications and the cryptographic providers. By standardizing on a cryptographic API, manufacturers maximize portability and code re-use in products with different chipsets. In addition, Certicom offers a full range of hardware crypto cores that are pre-integrated into the solution, for customers that want to create their own media co-processor designs.

Securing Device Manufacturing: By combining software-based Certicom® KeyInject™ with hardware-based IP cores, security is extended right to the silicon. Additionally, Certicom uses government-rated, tamper resistant Hardware Security Modules (HSMs) to protect key and logging data every step of the way, so designs are protected even in the most hostile environments.

Certicom KeyInject protects unique keying information such as IMEI numbers and CMLA certificates from the place of issuance all the way to the assembly line. It even provides a tracking mechanism that helps accounting for parts consumed from a digital Bill of Materials (BOM).

At each step in the chip manufacturing process, Certicom KeyInject is used to inject key information into the OTP (One Time programmable memory). At power on reset, the KeyGen/ KeyExpander combines this keying information to derive a single key, which is then fed into the AES decryption block and used to decrypt data and instructions from encrypted ROM. If the device has followed the defined manufacturing process, it will function correctly, otherwise it will be crippled.

Conclusion

Certicom Security for DRM provides a comprehensive framework for robust DRM implementations. It has been designed from the ground up to leverage hardware-based security, filling in the critical pieces of trusted boot, secure key management and application authentication.

Tight integration with hardware components extending all the way into the manufacturing environment provides device manufacturers a solution they can rely on to mitigate risk in building end-user devices for the digital content ecosystem.

certicom™

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

Using Digital Signatures to Cut Down on Bank Fraud Loss

Preparing for Unlicensed Mobile Access

Building Trust for Embedded Systems

Making the Grade - meeting government security requirements – Suite B

Current Public Key Cryptographic Systems

Injecting Trust to Protect Revenue and Reputation

certicom™

About Certicom

Certicom protects the value of your content, software and devices with government-approved security. Adopted by the National Security Agency (NSA) for classified and sensitive but unclassified government communications, Elliptic Curve Cryptography (ECC) provides the most security per bit of any known public-key scheme. As the undisputed leader in ECC, Certicom security offerings are currently licensed to more than 300 customers including General Dynamics, Motorola, Oracle, Research In Motion and Unisys. Founded in 1985, Certicom's corporate offices are in Mississauga, ON, Canada with worldwide sales headquarters in Reston, VA and offices in the US, Canada and Europe. Visit www.certicom.com

Contact Certicom

Corporate Headquarters

5520 Explorer Drive, 4th Floor
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
Reston, Virginia 20190
Tel: 703-234-2357
Fax: 703-234-2356
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

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

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

U.S. Western Regional Office
393 Vintage Park Drive, Suite 260
Foster City, CA 94404
Tel: 650-655-3950
Fax: 650-655-3951
E-mail: sales@certicom.com

www.certicom.com