Hardware-based Systems Security

A brief high-level view of developments and directions from an industry perspective

Ron Perez, Systems Security Fellow
Rambus Cryptography Research Division

Trusted Computing Base (TCB)

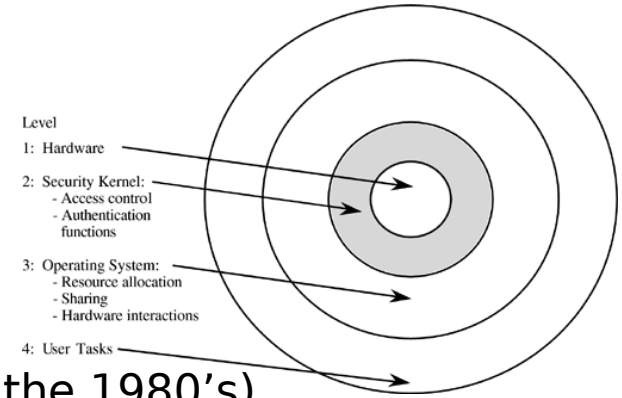

Developed over thirty (30) years ago
focused on assurance of computing systems
E.g., Department of Defense Trusted Computer System Evaluation Criteria (DoD 5200.28-STD in the 1980's)

Generally understood that TCB is everything (e.g., HW, SW, FW) in a "system" needed to support the security properties of that system
Where "system" can be a standalone HW and SW/FW computer – e.g., a server, a PC, a mobile device, or an embedded system

Or "system" can more generally be applied to distributed middleware or services platform infrastructures – e.g., an enterprise infrastructure/intranet, communications/messaging platforms, Software-as-a-Service platforms, or perhaps even the Internet itself (some TCBs are better/worse than others)

Hardware Security Module (HSM)

TCB assurance has proven to be difficult to achieve, at least for general purpose computing systems

- Leveraged HW: CPU-based support logical and temporal separation
- Chiefly in support of operating systems security for applications
- TCBs were/are generally large and complex

Is there an easier way? What about dedicated security subsystems?
HSMs are hardware-based self-contained subsystems developed largely for cryptographic key management and processing

- Leverage hardware to both accelerate and “harden”
- Physical and cryptographic separation from the host system
- Dedicated resources in a constrained environment higher assurance

Examples of HSMs in card form factor

also available in form factors from USB stick to rack-mount device

- Most often used in specialized systems, by specialized software
  - With some exceptions – e.g., AWS CloudHSM

Hardware Root of Trust (RoT)

Term developed in the late 1990's as part of a move to define "hardware-rooted" system security capabilities

- Hardware-based primitives used as a foundation on which to build TCBs
- Origins related to the Trusted Computing Group (TCG) and the Trusted Platform Module (TPM)

TPM: discrete chip, building on smart card and HSM technologies

100s millions (>1B?) TPM in PCs, servers, …

- Largely unused before Win8 due to complexity/usability and other challenges

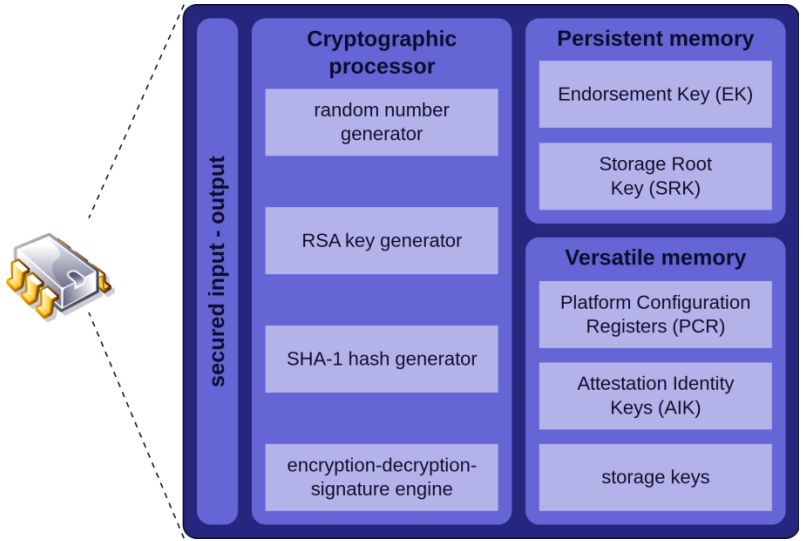

Example TPM concepts

Measured boot vs. secure boot
- Provide PCR “Quotes” to verifying party

Platform Configuration Registers (PCRs)
- Measuring system state
- PCRs are “extended”, not written:
  PCRn+1 = SHA(PCRn || Value)
- “Sealing” data to PCR value

Multiple Roots of Trust & Trust Chains
- **RTM** (measurement: static & dynamic)
- **RTS** (storage: TPM), **RTR** (reporting: TPM)

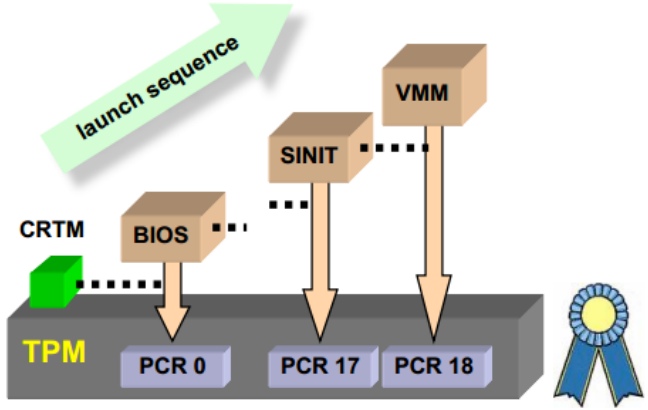

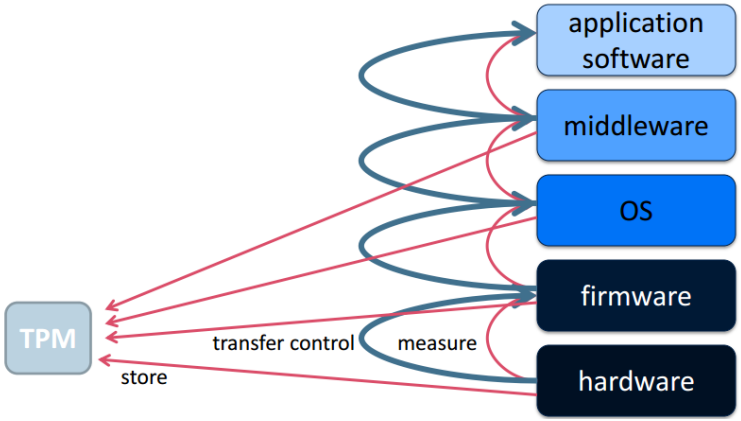

ARM TrustZone & Trusted Execution Environments

Migration to on-chip dedicated security capabilities through CPU mode, similar to virtualization
GlobalPlatform standardization

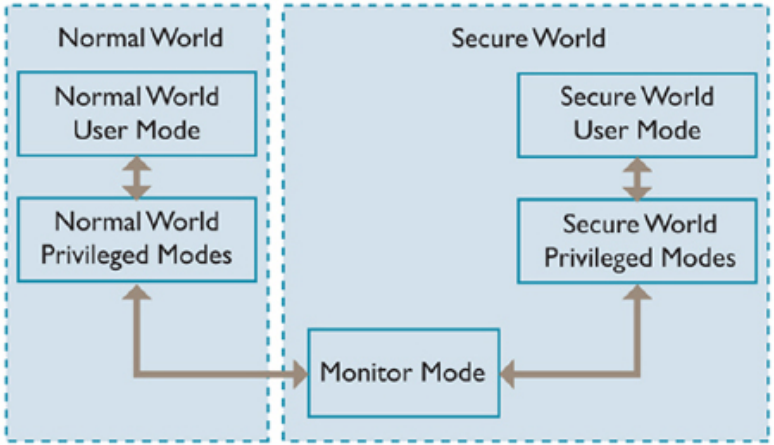

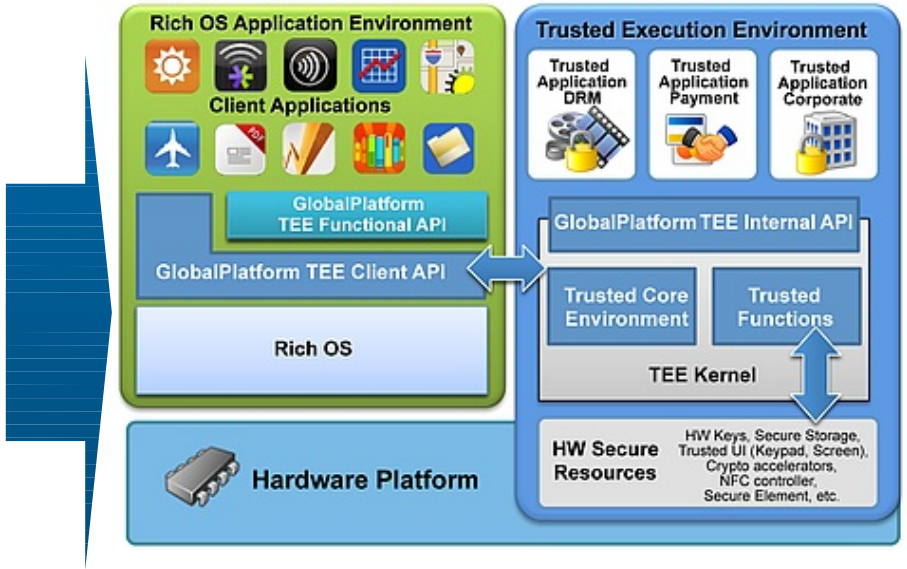

Additional On-Chip Secure Execution Environments

System-on-Chip (SoC) and security subsystems with dedicated resources and RoTs
- E.g., Intel Security Engine,
- AMD Platform Security Processor

Intel Software Guard Extensions (SGX)
- "inverse" sandbox (enclave) protecting applications from privileged malware

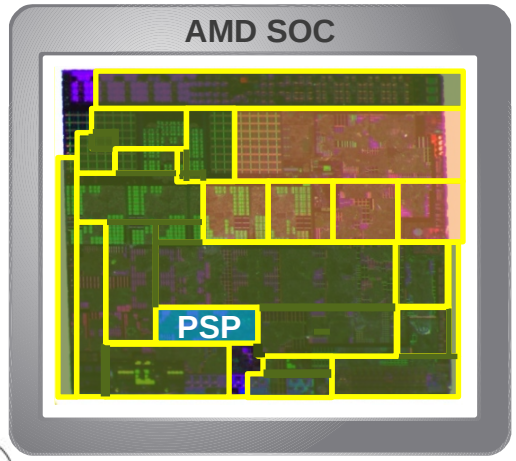

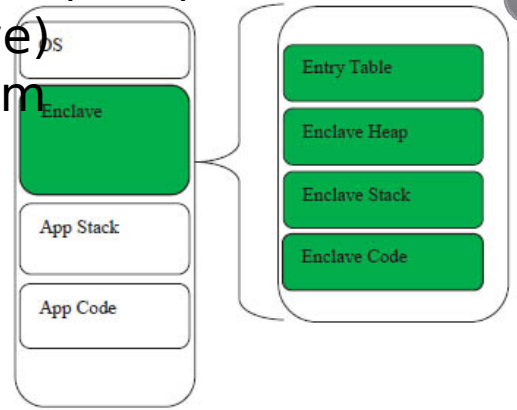

Multiple RoTs now being defined in context

Absolute RoT vs. “RoT for…”

A Root of Trust foundationally supports one or more conceptually higher-level security properties and/or security mechanisms, providing assurances to anyone/anything relying upon those properties/mechanisms – e.g.,

RoT **for** secure/measured/verified boot

RoT **for** attestation (measuring and/or reporting)

RoT **for** cryptographic key management (e.g., key provisioning, protected/access-controlled storage and usage)

RoT **for** supply chain integrity

RoT **for** …

Hardware-based system security challenges

SW/HW interfaces: usability by system and/or application SW?
Life-cycle: birth to death and rebirth again?
Does RoT == Reference Monitor?

RoT is very much analogous to if not exactly fitting the definition of an RoT

Assurance of policy-driven and highly-integrated hardware security mechanisms?
Composition: one vs. many RoTs?

Most likely, there are many RoTs! – e.g., RoT for secure execution may be composed of a RoT for key management, one for secure boot, and one for attestation (perhaps all three of which are implemented in the same piece of hardware/firmware/software)

- But is this a RoT for secure execution or a Trusted Computing Base?

Questions / Discussion

Rambus CRYPTOGRAPHY RESEARCH

HSM physical security layers

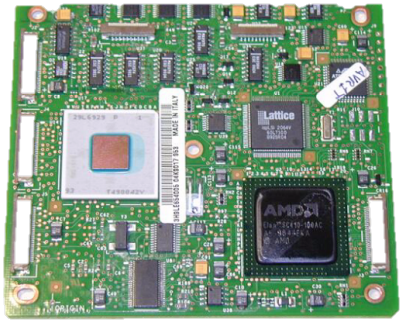

1. Daughter card

2. Inner enclosure

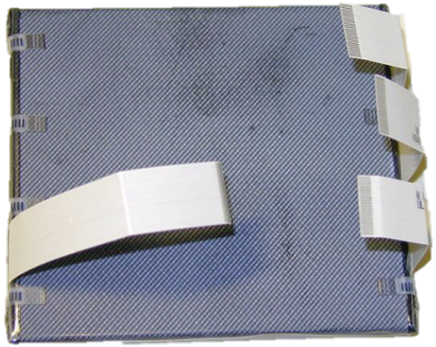

3. Tamper-sensing membrane

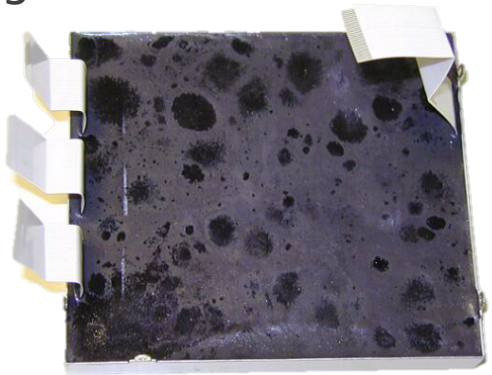

4. Outer shell & potting material

TPM concepts continued

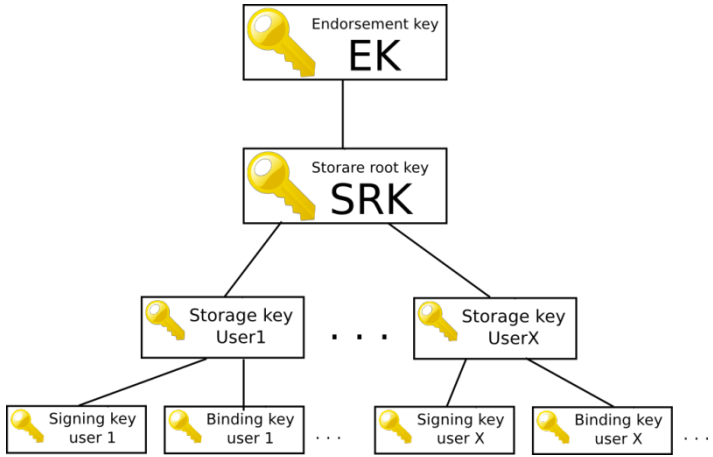

Unique, on-chip keys
- Endorsement Key, Storage Root Key
- Basis for [off-chip] hierarchies
- Generate new SRK on ownership change

Attestation Identity Keys
- Signed by EK
- Privacy issues and need for Privacy CA/TTP
- Addressed by Direct Anonymous Attestation (DAA) – not used

Monotonic counters
- Anti-replay, force sequencing, etc.

END