

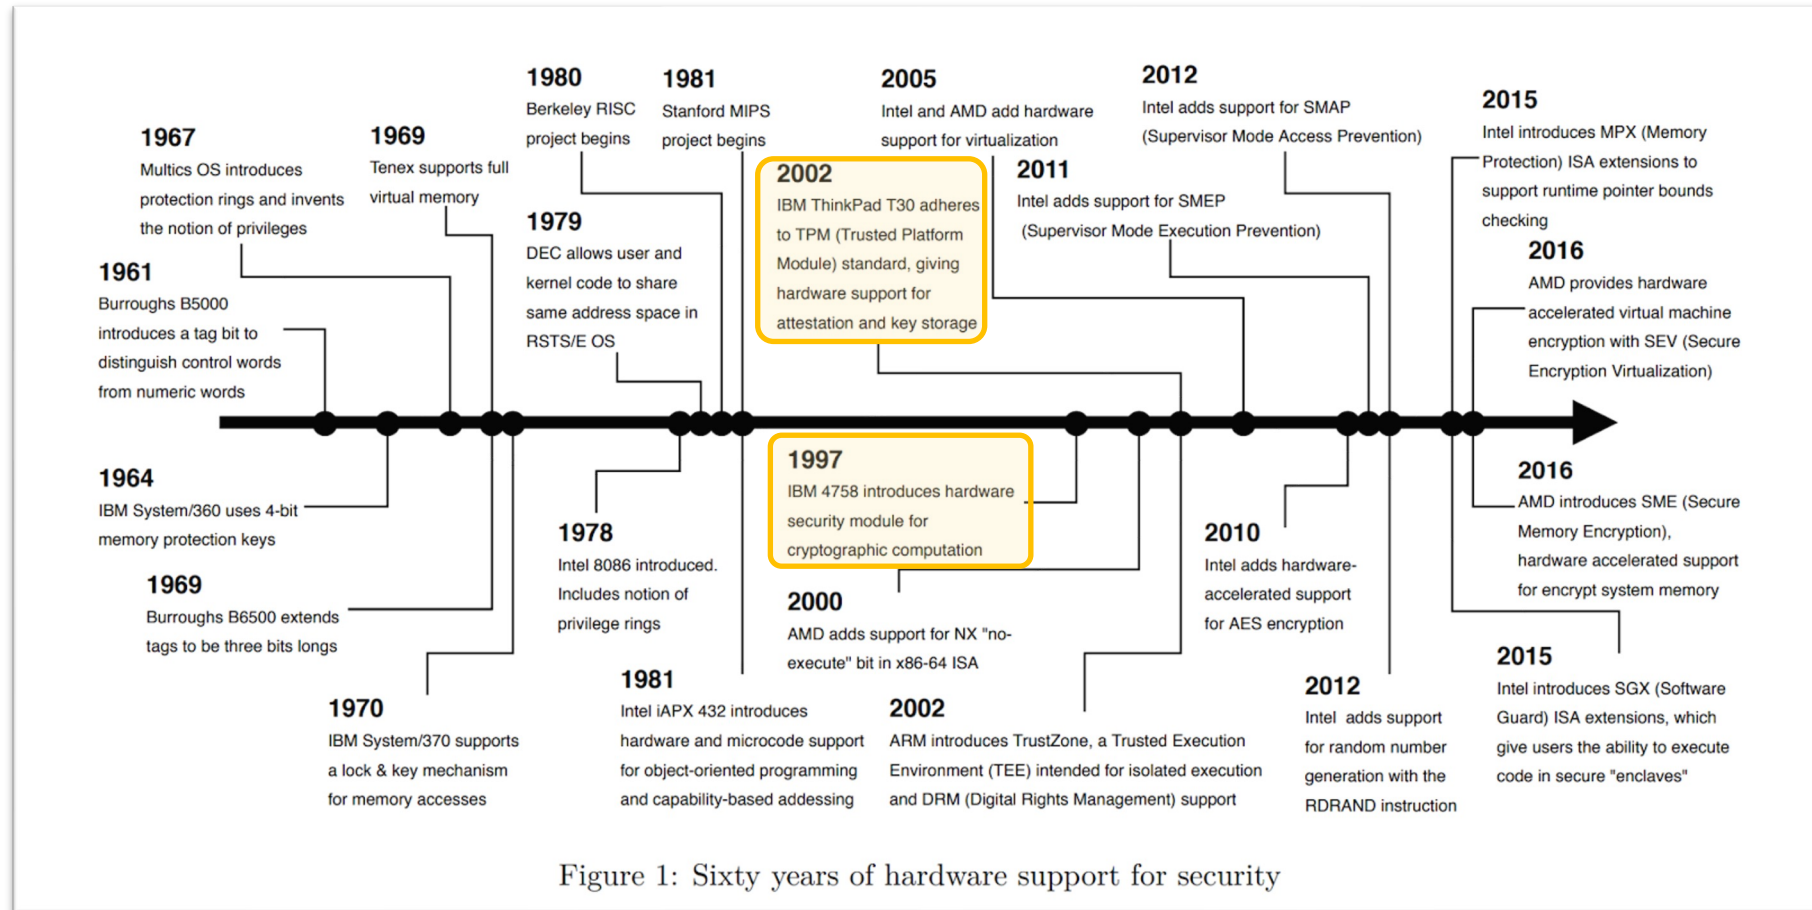
# Hardware Security Module

Mengjia Yan

Spring 2024

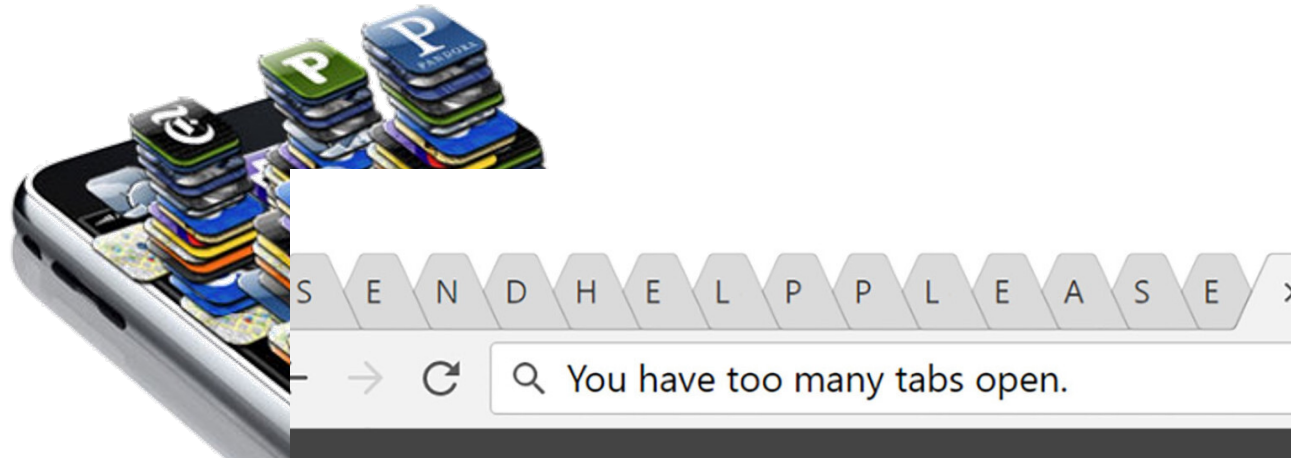


# Secure Processors/HSM



Apple  
Secure  
Enclave

# Security Contexts #1




```
CPU [ 0.0%] Tasks: 103, 181 thr; 2 running
Mem [ 611M/993M] Load average: 0.00 0.03 0.15
Swp [ 11.0M/2.00G] Uptime: 00:43:30

PID USER PRI NI VIRT RES SHR S CPU% MEM% TIME+ Command
9942 root 20 0 117M 2148 1432 R 0.0 0.2 0:00.08 htop
2671 mysql 20 0 885M 83308 1056 S 0.0 8.2 0:00.35 /usr/libexec/mysqld --basedir=/usr --
2589 root 20 0 540M 14188 3612 S 0.0 1.4 0:00.31 /usr/bin/python -Es /usr/sbin/tuned -
1 root 20 0 120M 3040 1752 S 0.0 0.3 0:01.32 /usr/lib/systemd/systemd --switched-r
461 root 20 0 35096 2536 2352 S 0.0 0.2 0:00.18 /usr/lib/systemd/systemd-journald
487 root 20 0 123M 684 684 S 0.0 0.1 0:00.00 /usr/sbin/lvmtd -f
498 root 20 0 43700 1156 940 S 0.0 0.1 0:00.17 /usr/lib/systemd/systemd-udev
609 root 16 -4 51208 1148 1024 S 0.0 0.1 0:00.00 /sbin/auditd -n
599 root 16 -4 51208 1148 1024 S 0.0 0.1 0:00.01 /sbin/auditd -n
613 root 12 -8 80220 784 680 S 0.0 0.1 0:00.00 /sbin/auditd -n
610 root 12 -8 80220 784 680 S 0.0 0.1 0:00.01 /sbin/auditd -n
612 root 16 -4 26200 704 656 S 0.0 0.1 0:00.00 /usr/sbin/sedispd
624 root 39 19 16752 820 788 S 0.0 0.1 0:00.00 /usr/sbin/alsactl -s -n 19 -c -E ALSA
653 root 20 0 395M 2804 1984 S 0.0 0.3 0:00.04 /usr/libexec/accounts-daemon
686 root 20 0 395M 2804 1984 S 0.0 0.3 0:00.00 /usr/libexec/accounts-daemon
625 root 20 0 395M 2804 1984 S 0.0 0.3 0:00.17 /usr/libexec/accounts-daemon
682 root 20 0 280M 2800 2428 S 0.0 0.3 0:00.03 /usr/sbin/rsyslogd -n
683 root 20 0 280M 2800 2428 S 0.0 0.3 0:00.01 /usr/sbin/rsyslogd -n
628 root 20 0 280M 2800 2428 S 0.0 0.3 0:00.07 /usr/sbin/rsyslogd -n
629 root 20 0 4372 508 488 S 0.0 0.0 0:05.68 /sbin/rngd -f
634 dbus 20 0 30316 2896 1224 S 0.0 0.3 0:00.55 /bin/dbus-daemon --system --address=s
F1Help F2Setup F3Search F4Filter F5Tree F6SortBy F7Nice F8Nice F9Kill F10Quit
```

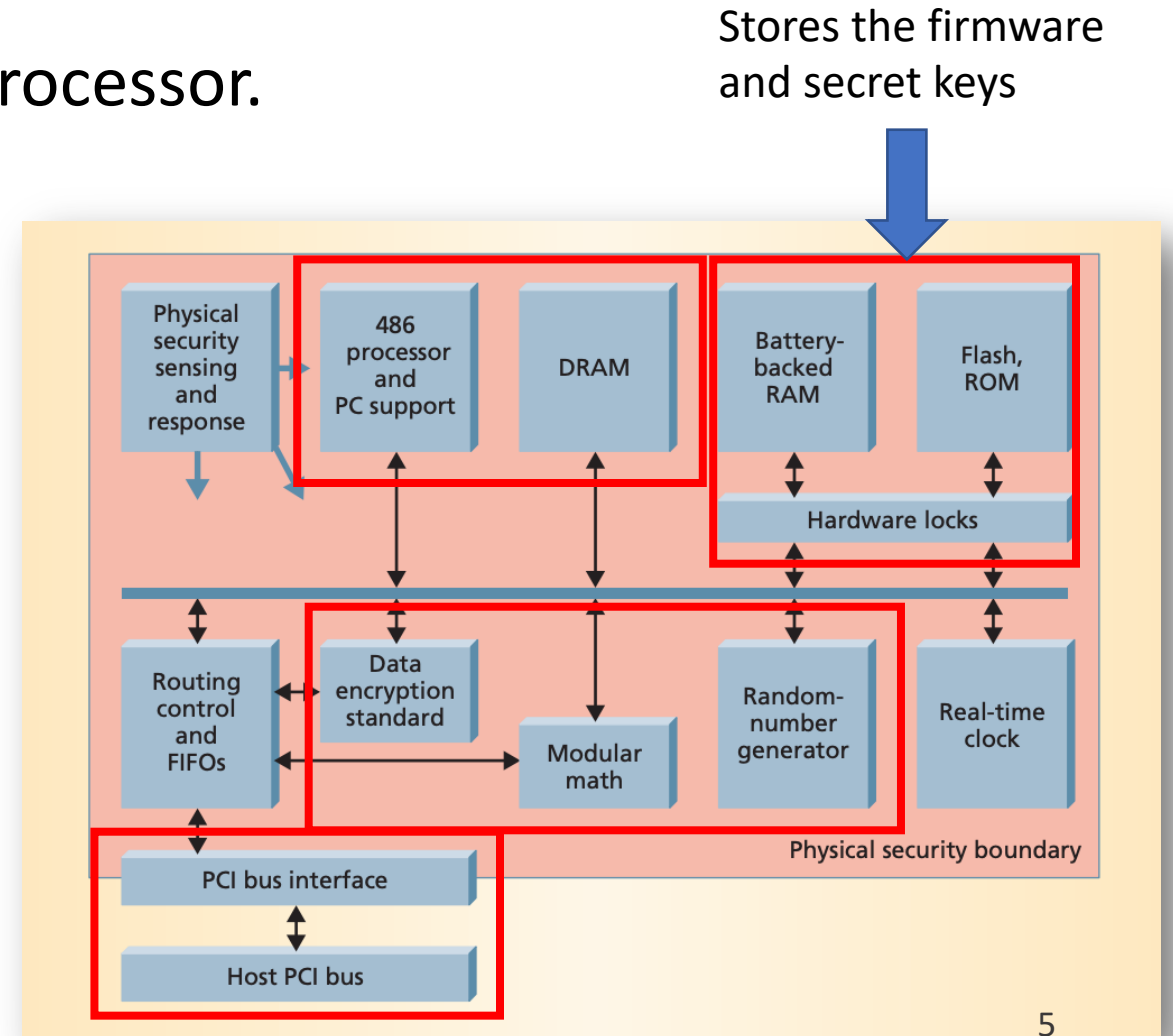
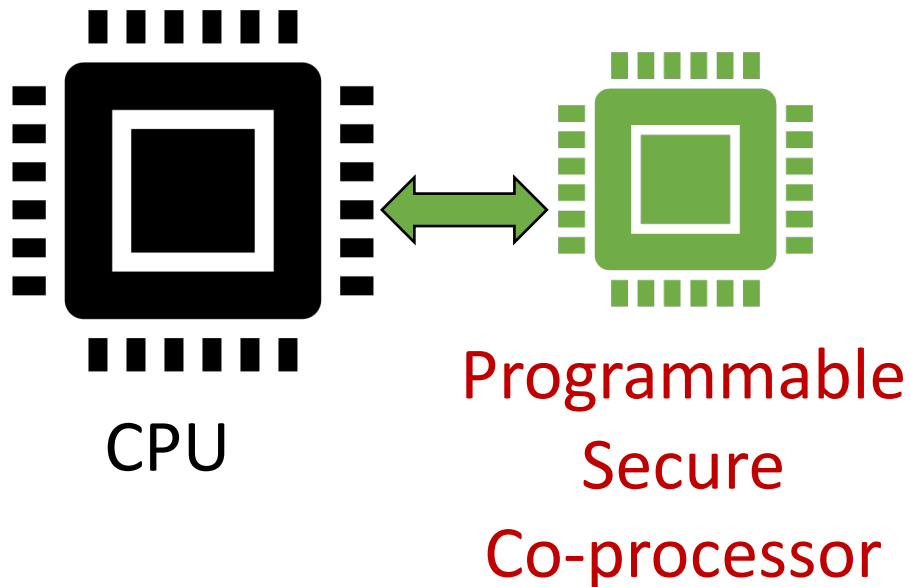
- Software can be buggy (or sometimes malicious)
- Running daily applications together with security-sensitive applications
- Can we do better than software-based isolation?

# Before IBM 4758 (1999): Crypto Accelerators



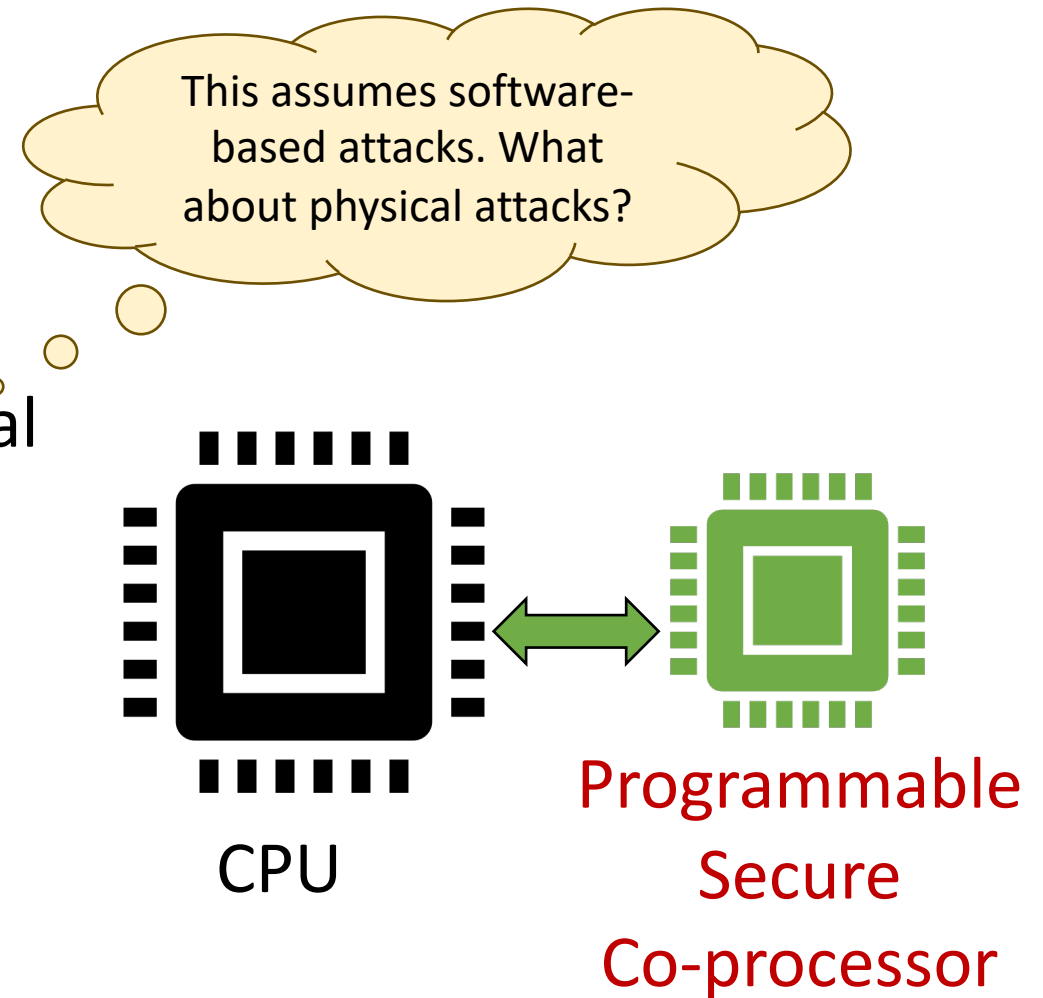
# IBM 4758 (1999) -- 4765 (2012)

- Goal: a programmable, secure co-processor.
- High level idea: virtual locker room



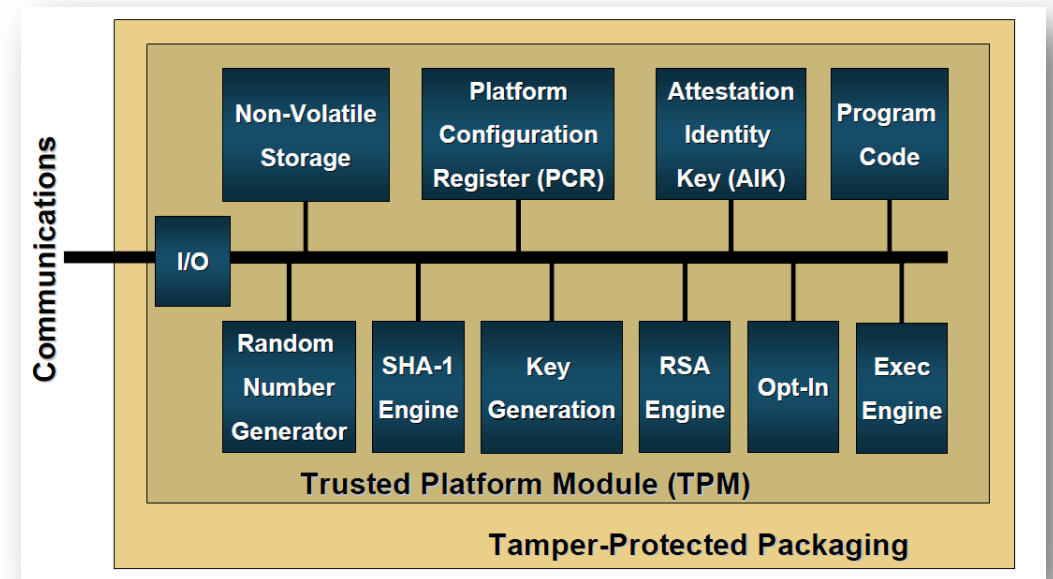


# Why this is more secure?

- Physical isolation (Not share physical memory)
- Narrow interface, only interact with external worlds via APIs (keys do not leave the co-processor)
- Simpler software on co-processor, so fewer bugs (maybe can be formally verified)
- Problem?
  - **The SWOFTWARE!** Bad programmability.



# Trusted Platform Module (TPM)

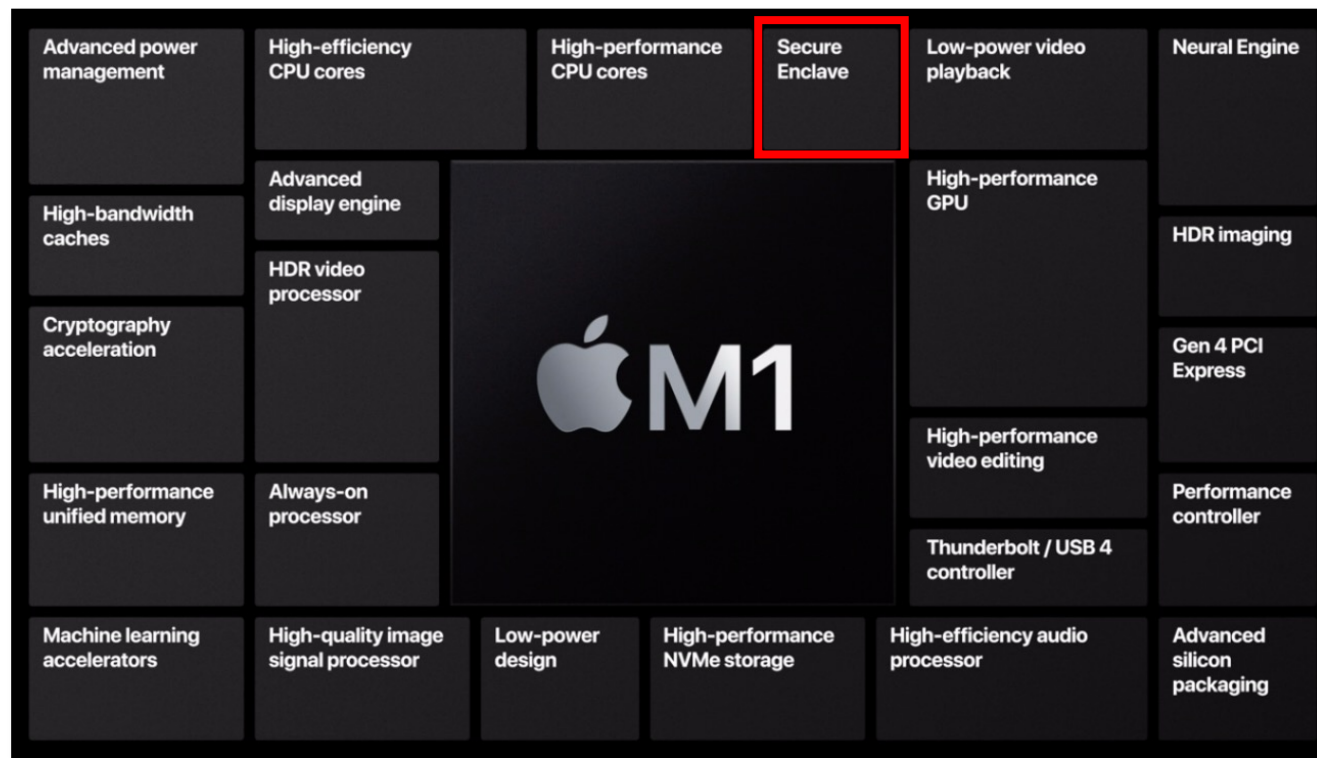
- “*Commoditized* IBM 4758”: Standard LPC interface attaches to commodity motherboards



<https://scotthelme.co.uk/upgrading-my-pc-with-a-tpm/>

# Apple Secure Enclave

- Advantage: one company controls both the hardware and the software
- Apple secure enclave runs a customized formally verified micro-kernel OS

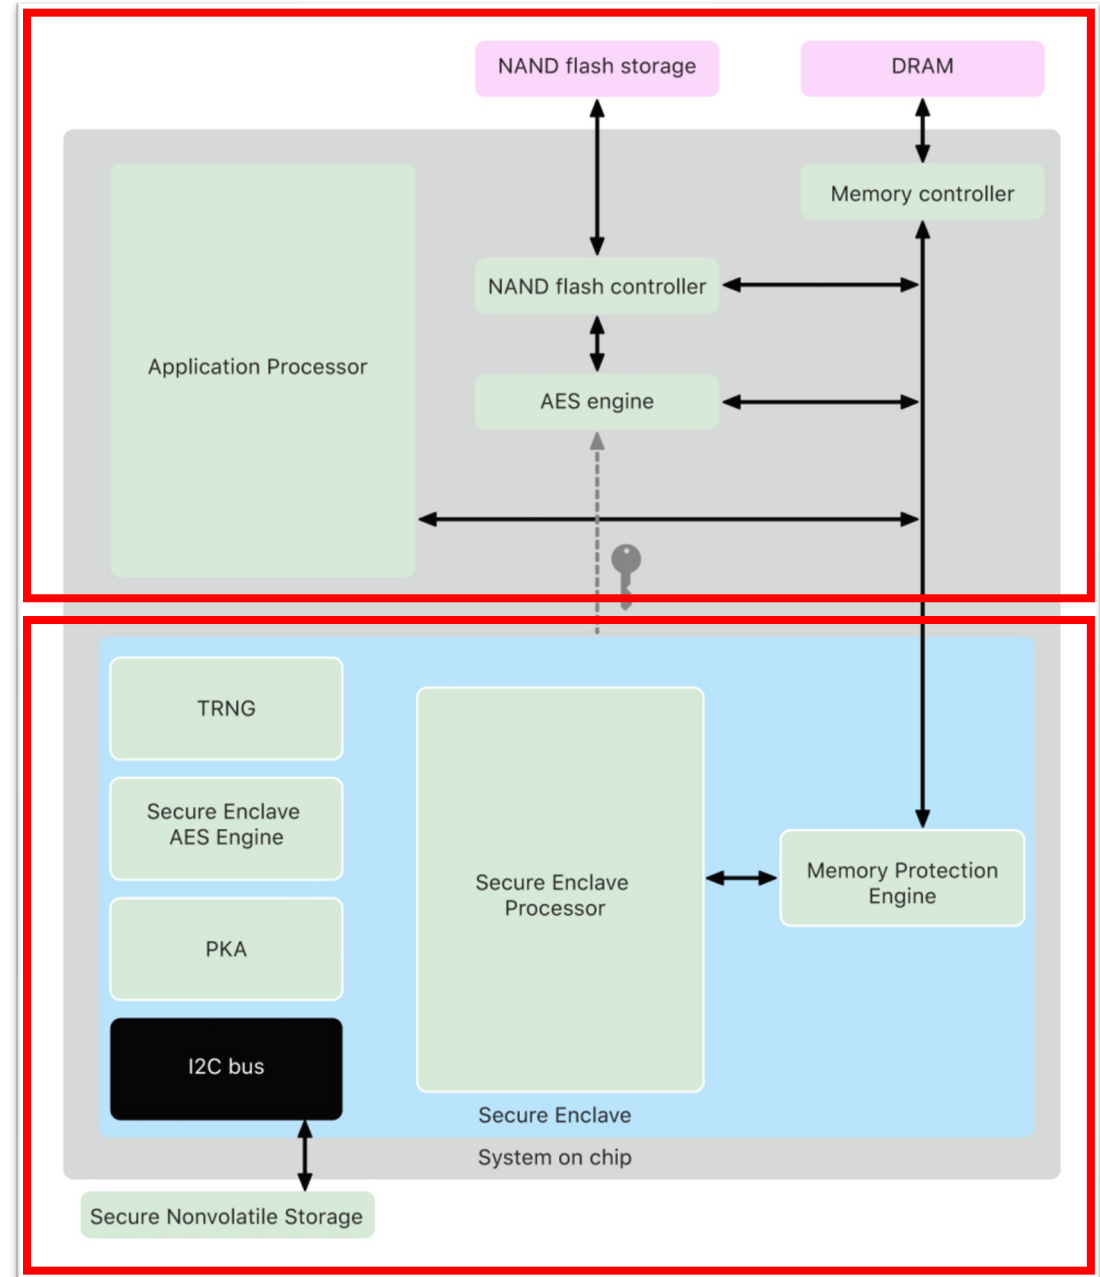




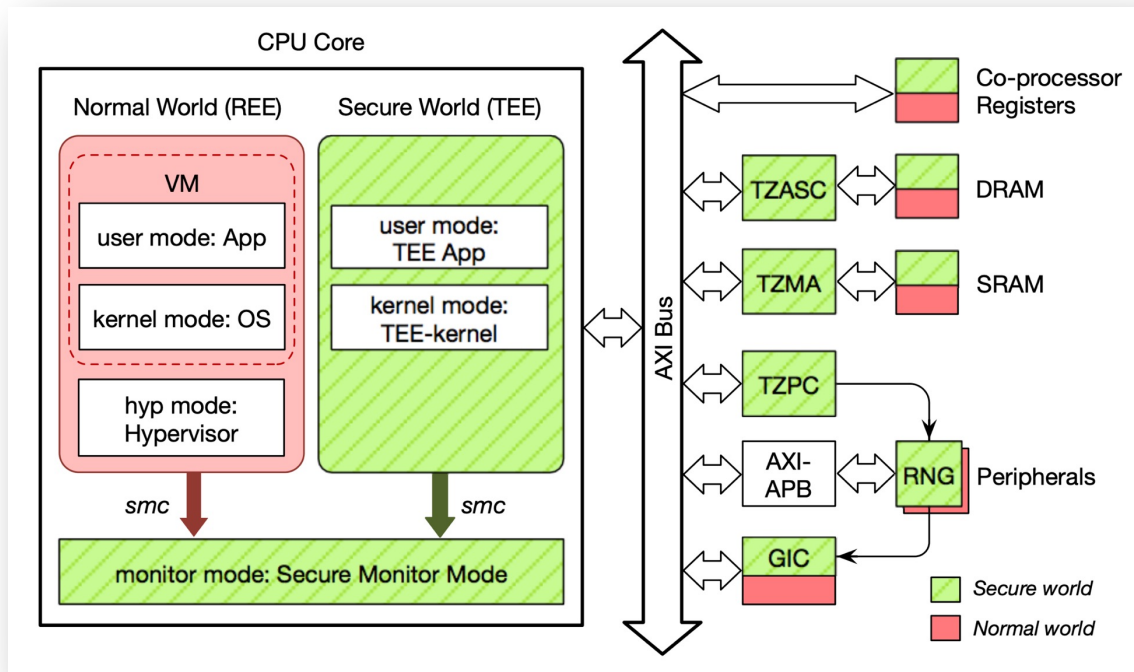
# Separate Cores

The secure enclave:


- Not general-purpose, only run secure enclave functionality, no user code
- Block vulnerabilities due to software bugs (running L4 microkernel) and side channels



# The Trends (isolation with some sharing?)



ARM TrustZone



Intel SGX model

Security?

Usability?



Fixed Design (Static)

Flexible Design (Dynamic)

# Security Contexts #2




- Disk lost or removed, leading to confidentiality leakage.
- Data encryption with weak passwords, such as, 6-digit passcode.

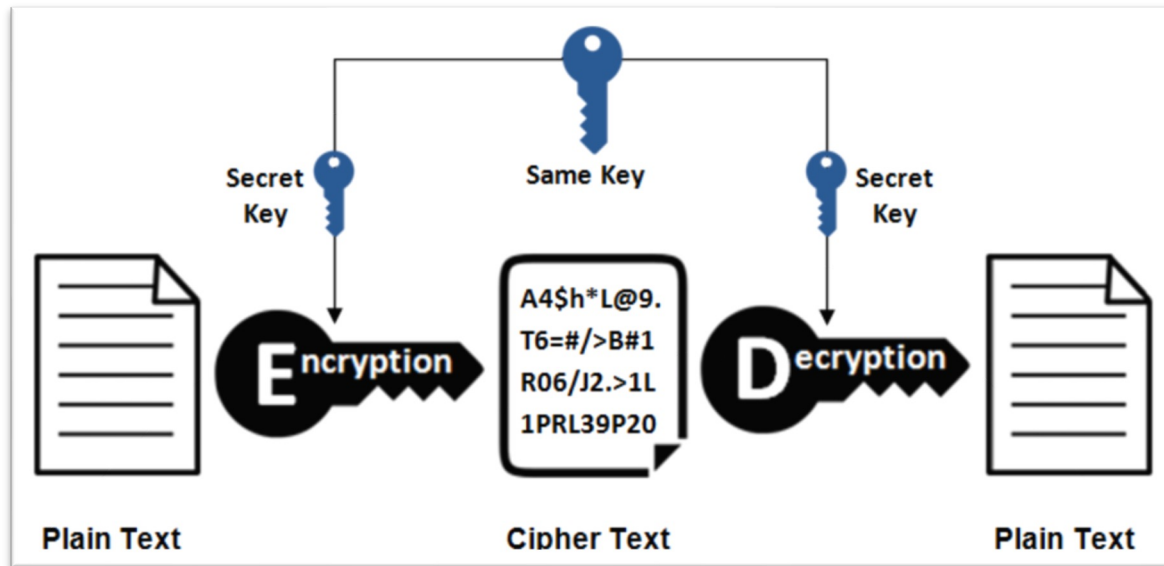
Bind data/application with hardware using crypto.

# Security Property and Crypto Primitives

- Confidentiality
  - Symmetric
  - Asymmetric
- Integrity
- Freshness



# Symmetric Cryptography



- One-time-pad (OTP)

Encryption:  
 $\text{ciphertext} = \text{key} \oplus \text{plaintext}$


Decryption:  
 $\text{plaintext} = \text{key} \oplus \text{ciphertext}$

How about encrypting arbitrary length message? Any problems?


# Block ciphers (e.g., DES, AES)

- Divide data in blocks and encrypt/decrypt each block
- AES block size can be 128, 192, 256 bits

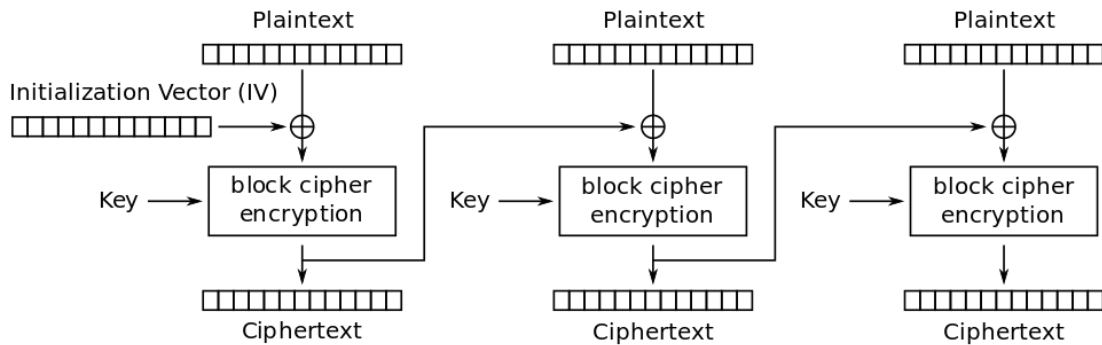
**ECB IS NOT  
RECOMMENDED**



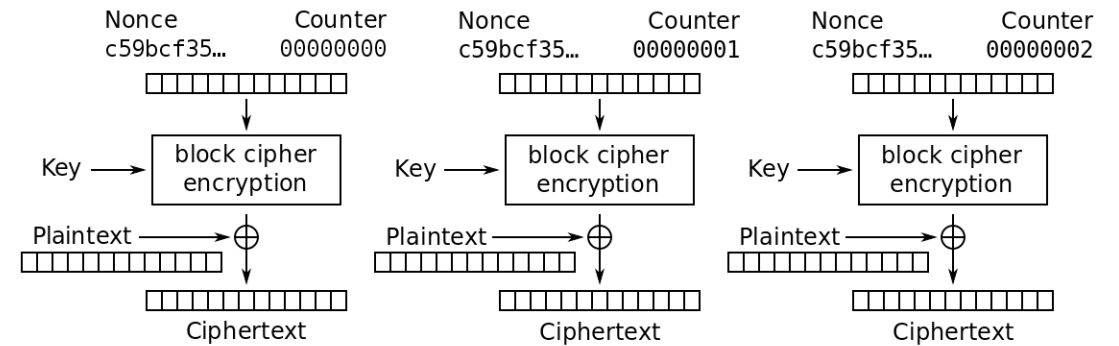
Electronic Codebook (ECB) mode encryption



# Other Block cipher Modes



Cipher Block Chaining (CBC) mode encryption



Counter (CTR) mode encryption

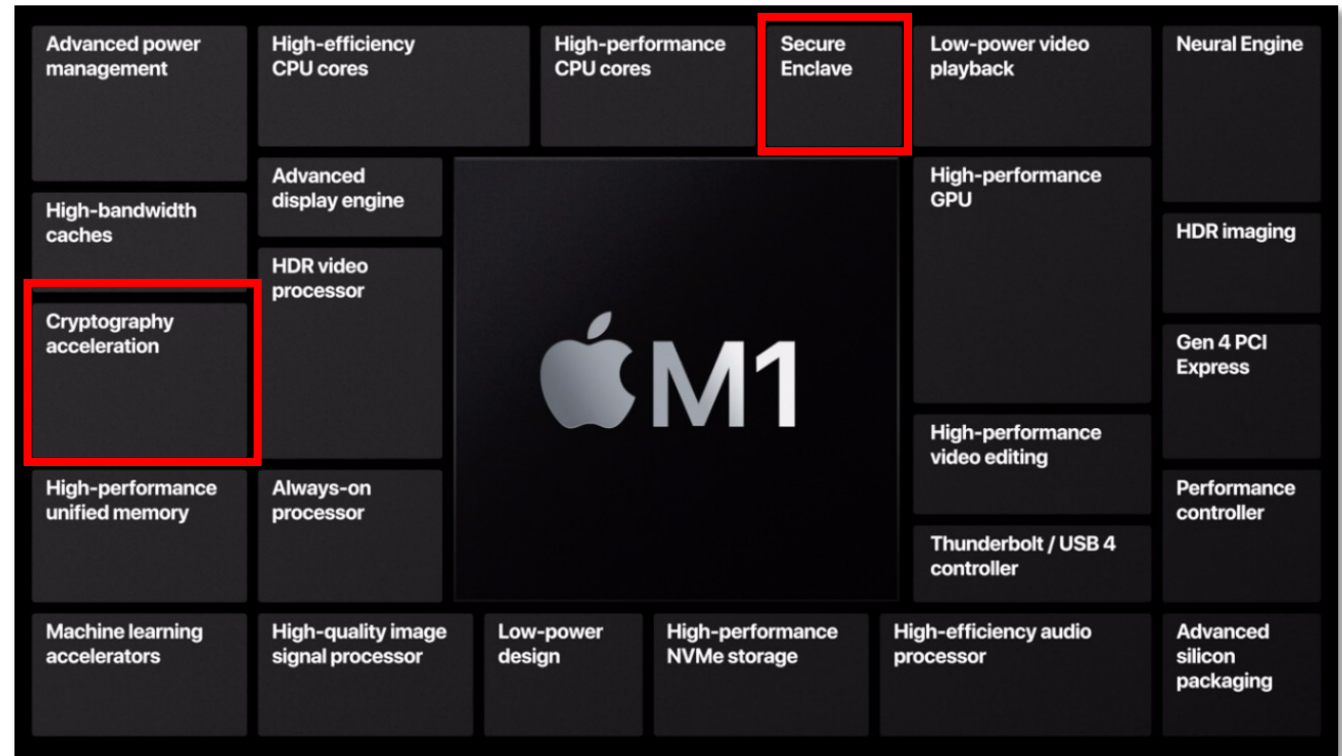
IV can be public, but need to ensure to not reuse IV for the same key.

- Real-world application: file/disk encryption and memory encryption.
- How to exchange the shared key between two parties?

Diffie-Hellman  
Key Exchange



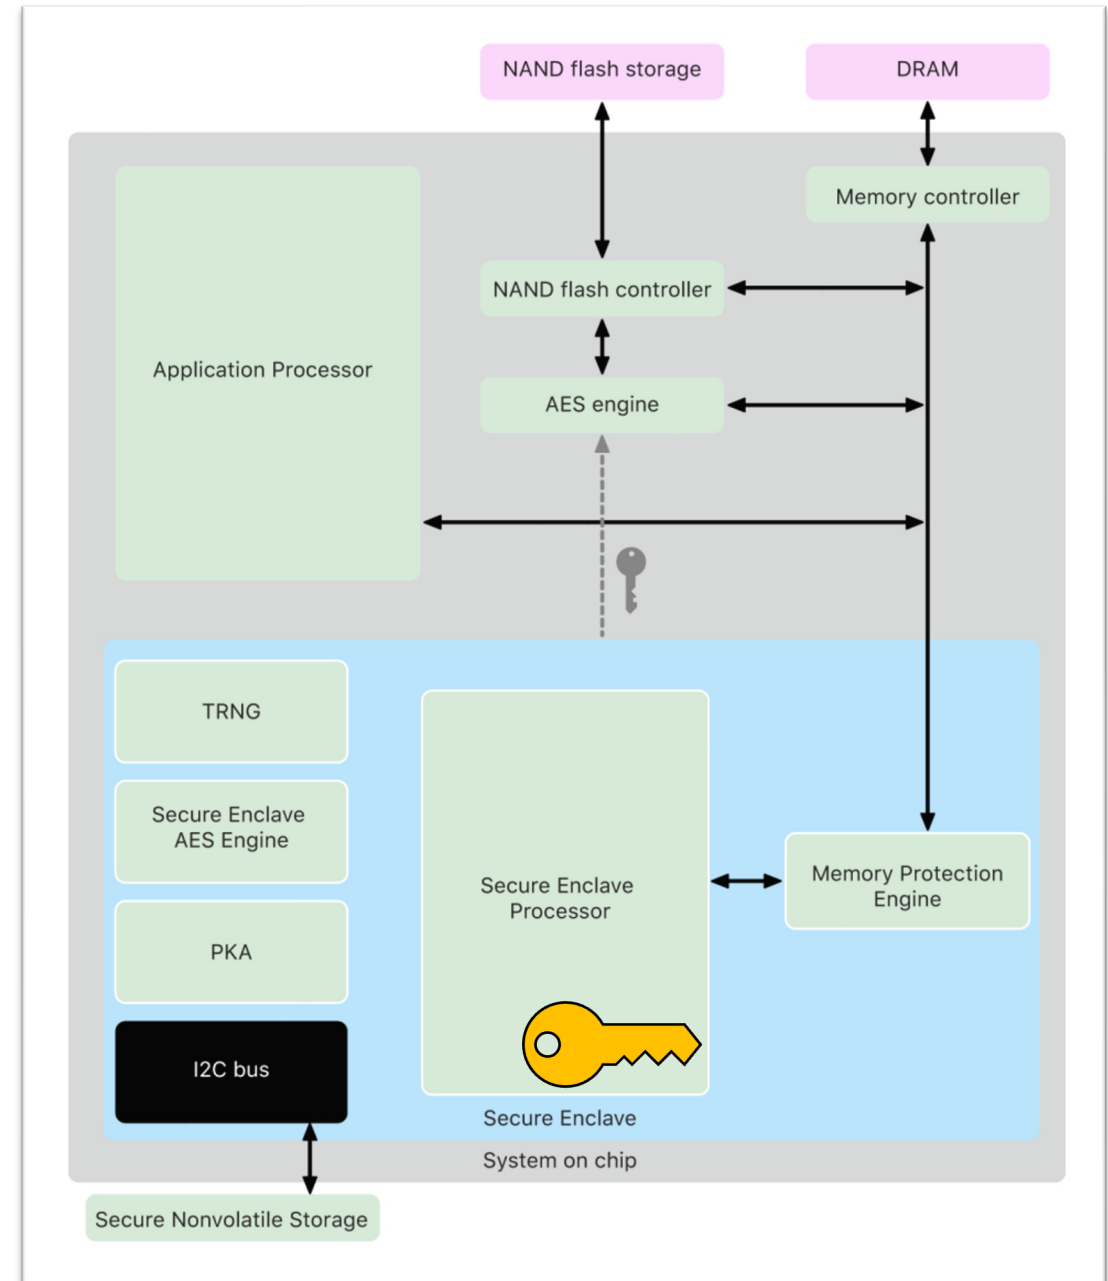
# Apple Secure Enclave



# Crypto Keys

The Secure Enclave includes a unique ID (**UID**) root cryptographic key.

- Unique to each device
- Randomly generated
- Fused into the SoC at manufacturing time
- Not visible outside the device



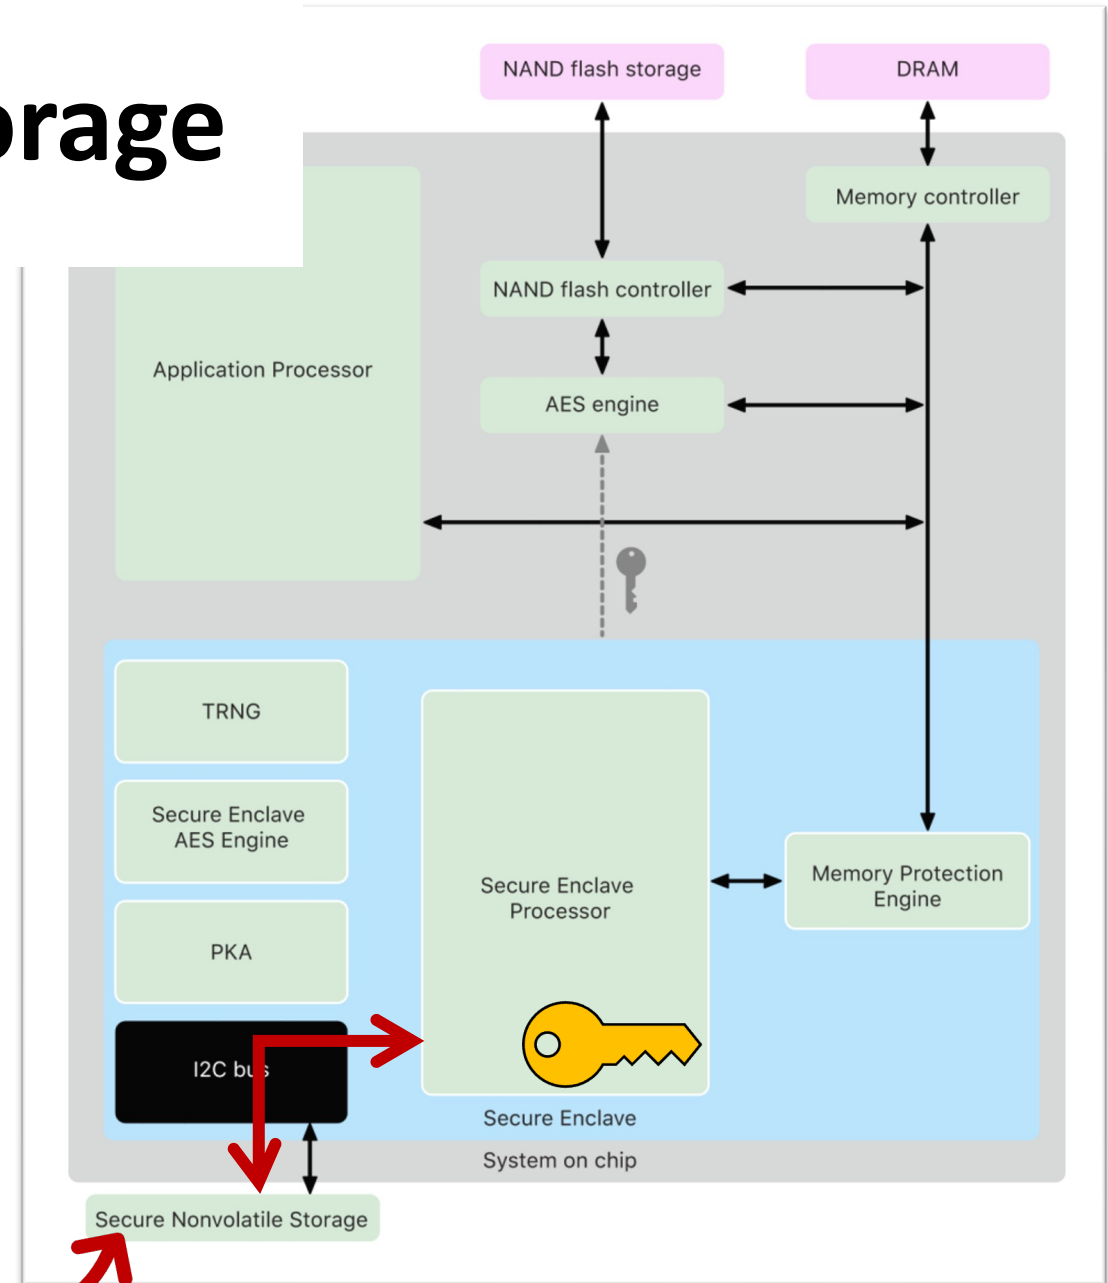
# Secure Non-volatile Storage

For easy to use: short passcode. But weaker security?

Passcode + **UID** -> passcode entropy

Brute-force has to be performed on the **device under attack** (not create a copy of the software and brute-force in parallel)


- Escalating time delays
- Erase data when exceeding attempt count



All user data encryption keys

# Security Contexts #3


Hardware establishes  
root of trust.



- a) An end-user wants to trust a remote server, e.g., bank server.
- b) A remote server wants to trust an end-user, e.g., when joining a company's highly-secure network.
- c) Lost device, rootkits? Are you sure you are running your trusted OS?

# Asymmetric Cryptography (e.g., RSA)

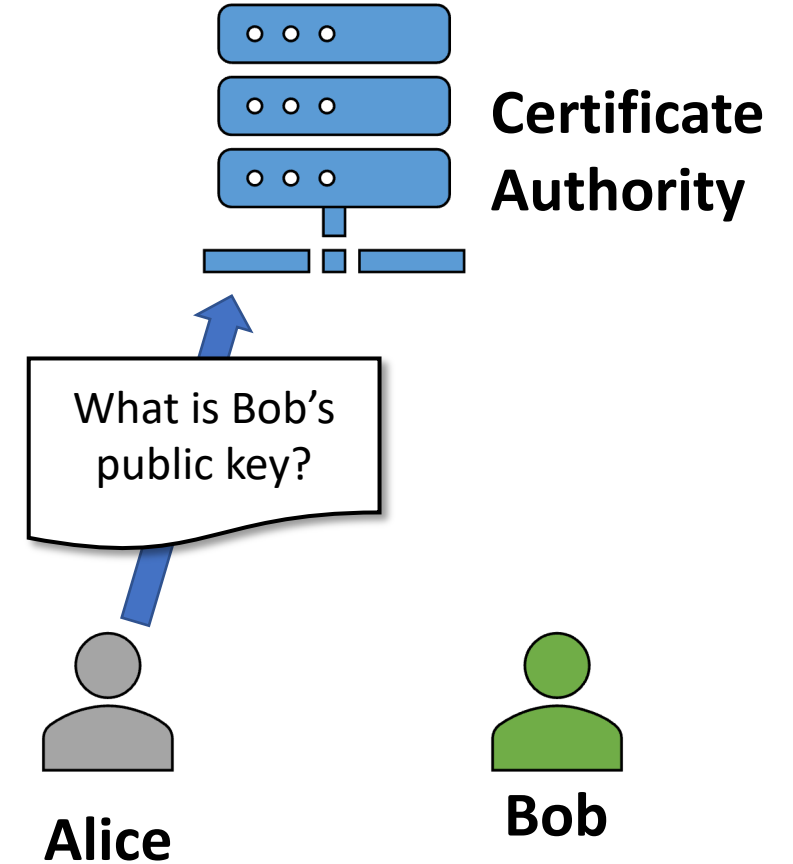
- A pair of keys:
  - Private key ( $K_{\text{private}}$  – kept as secret)
  - Public key ( $K_{\text{public}}$  – safe to release publicly)
- Computation:
  - $\text{Encrypt}(\text{plaintext}, K_{\text{public}}) = \text{ciphertext}$
  - $\text{Decrypt}(\text{ciphertext}, K_{\text{private}}) = \text{plaintext}$
- Computationally more expensive, so usually use asymmetric cryptography to negotiate a shared key (e.g., DKE key exchange), then use symmetric cryptography
- How to announce and obtain the public key?



Mail box is public;  
Box key is private

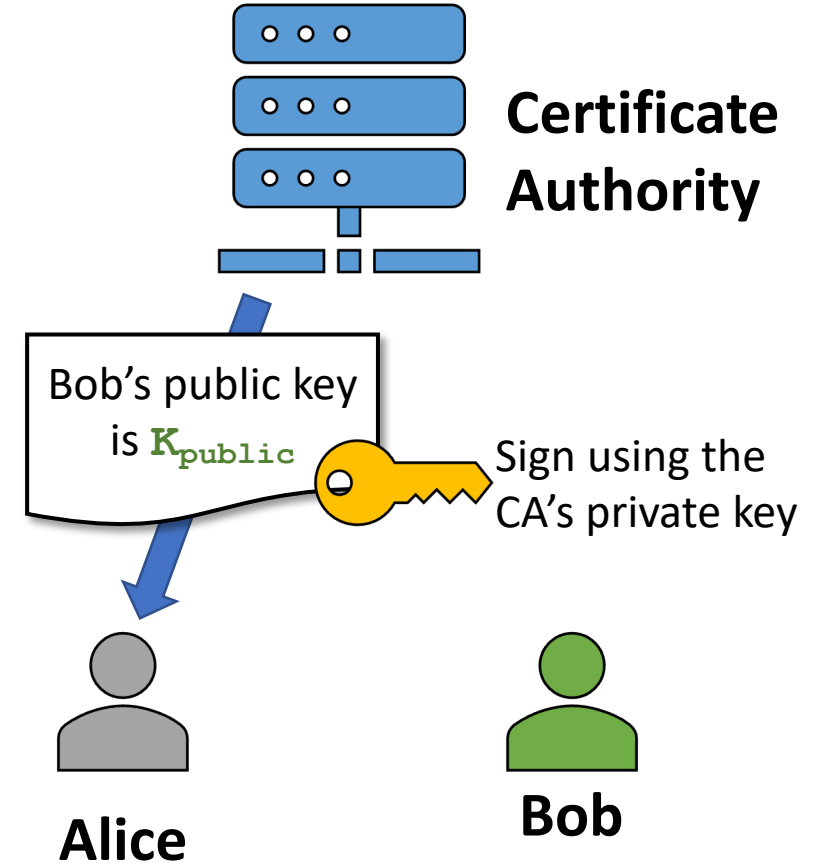
# Public Key Infrastructures (PKIs)

- Analogy: public key is like a government-issued ID, need to be validated by an authority.
- Bob has a private key  $K_{\text{private}}$  and wants to claim he corresponds to a public key  $K_{\text{public}}$




# Public Key Infrastructures (PKIs)

- Analogy: public key is like a government-issued ID, need to be validated by an authority.
- Bob has a private key  $K_{\text{private}}$  and wants to claim he corresponds to a public key  $K_{\text{public}}$
- Establish a chain of trust
- **Real-world use cases:** identify website, identify hardware chips/processors



# Integrity (MAC/Signature)

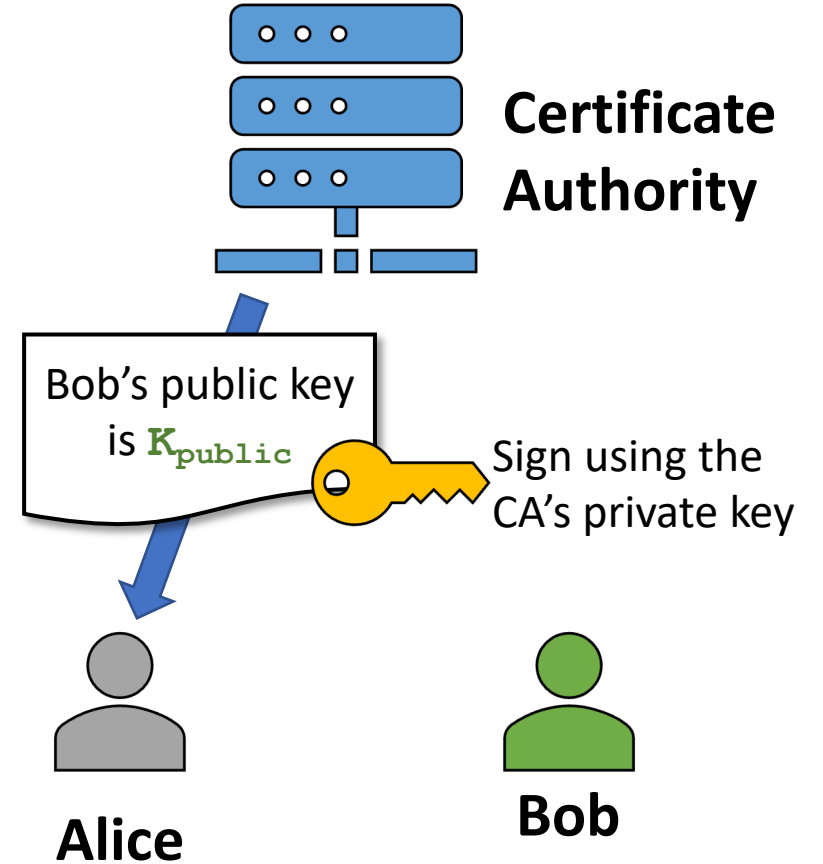


- One-way hash
  - Practically infeasible to invert, and difficult to find collision
- Avalanche effect
  - “Bob Smith got an A+ in ELE386 in Spring 2005” → 01eace851b72386c46d
  - “Bob Smith got an B+ in ELE386 in Spring 2005” → 936f8991c111f2cefaw
- When message is long
  - Divide message into blocks, and keep extending the hash by adding previous hash




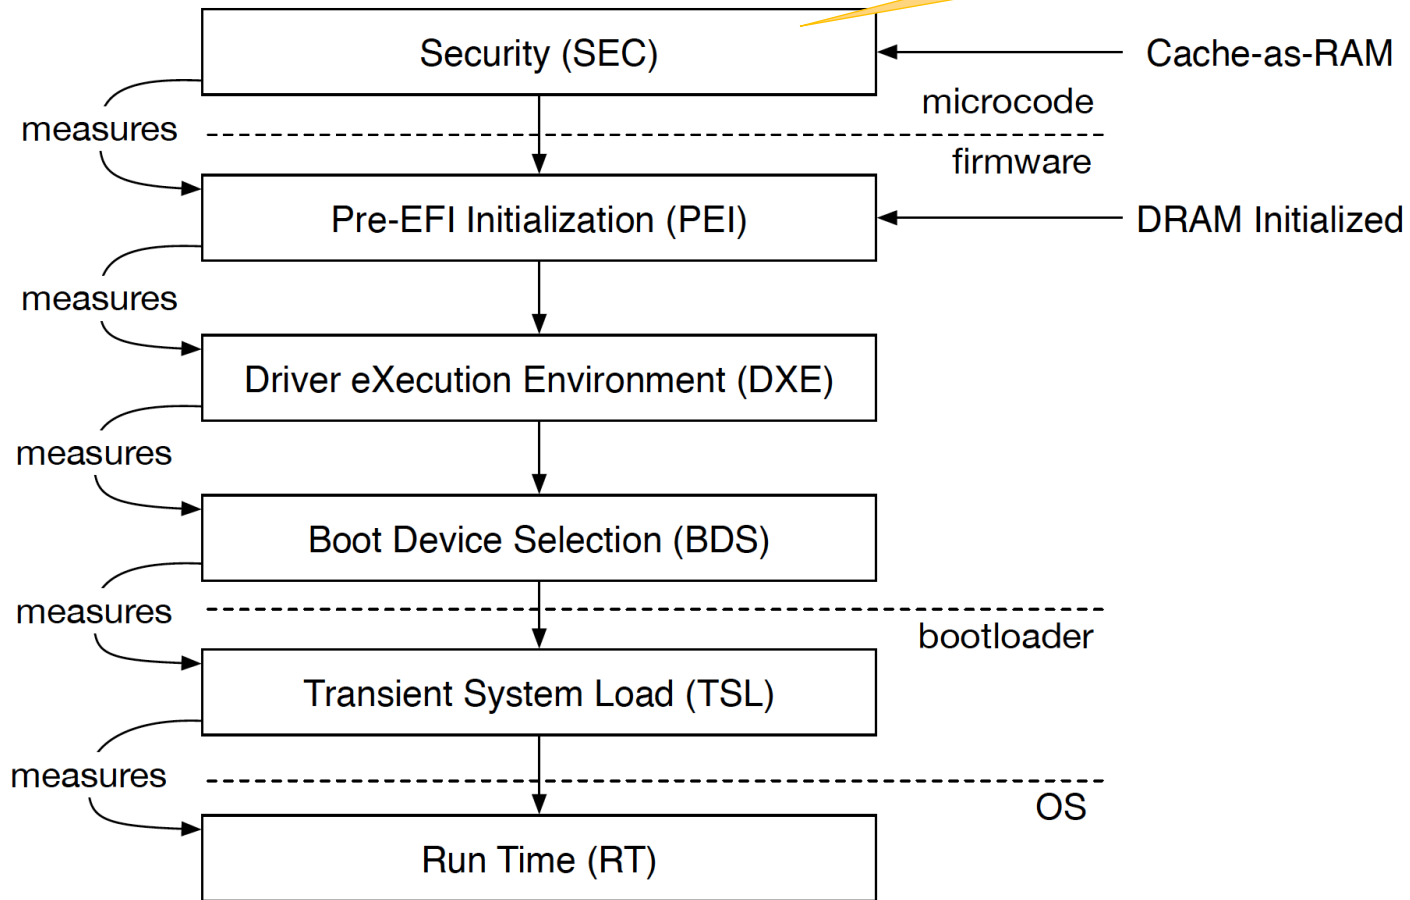
# Integrity + Crypto

- Using symmetric crypto:
  - $\text{hash} = \text{SHA}(\text{message})$
  - $\text{HMAC} = \text{enc}(\text{hash}, \text{key})$
- Using asymmetric crypto:
  - Sign:  $\text{sig} = \text{enc}(\text{hash}, K_{\text{private}})$
  - Verify:
    - $\text{hash}' = \text{SHA}(\text{message})$
    - $\text{sig} = \text{enc}(\text{hash}', K_{\text{public}})$




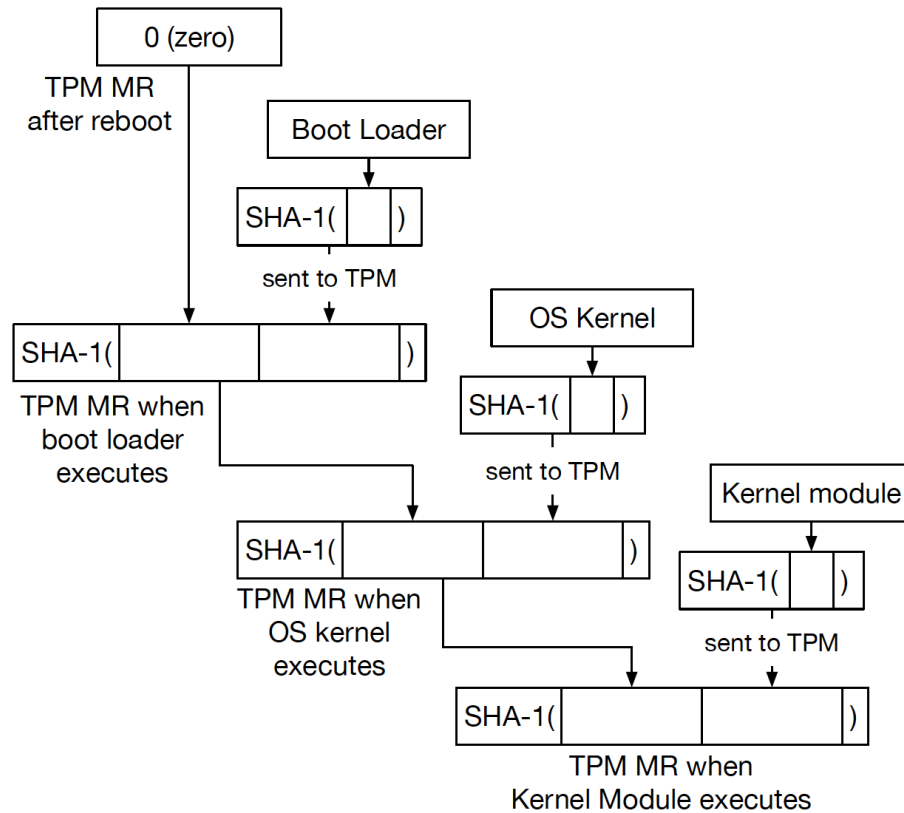

# Boot Process (UEFI)

Root of trust



How to perform the measurement?

# Secure Boot using TPM


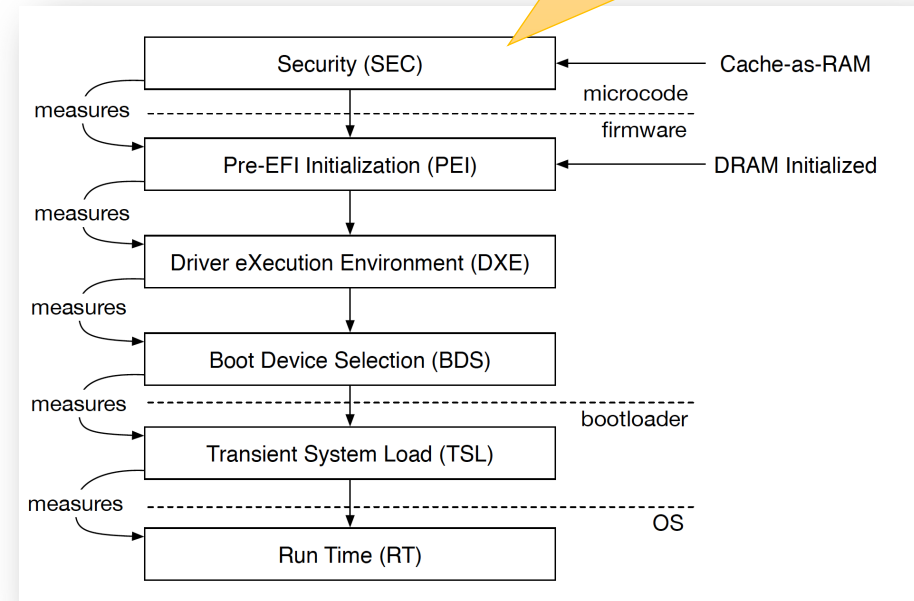


Each step, TPM compares to expected values locally or submitted to a remote attester.


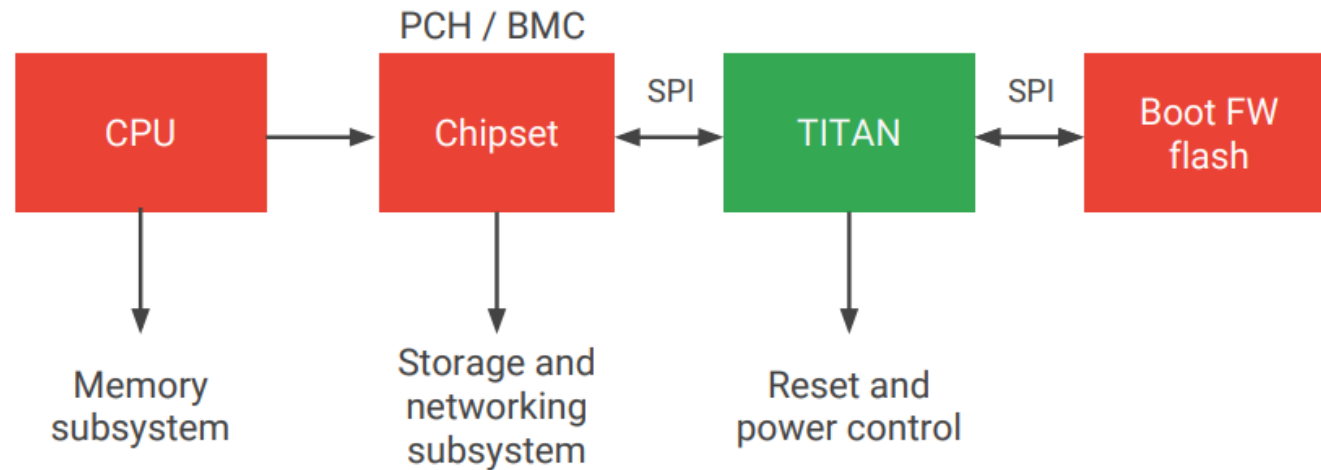
# Security Problems of Using TPM

- Assume the first-stage bootloader is securely embedded in motherboard
- Not easy to use with frequent software/kernel update
- Time to check, time to use
- TPM Reset attacks
  - exploiting software vulnerabilities and using software to report false hash values

Root of trust

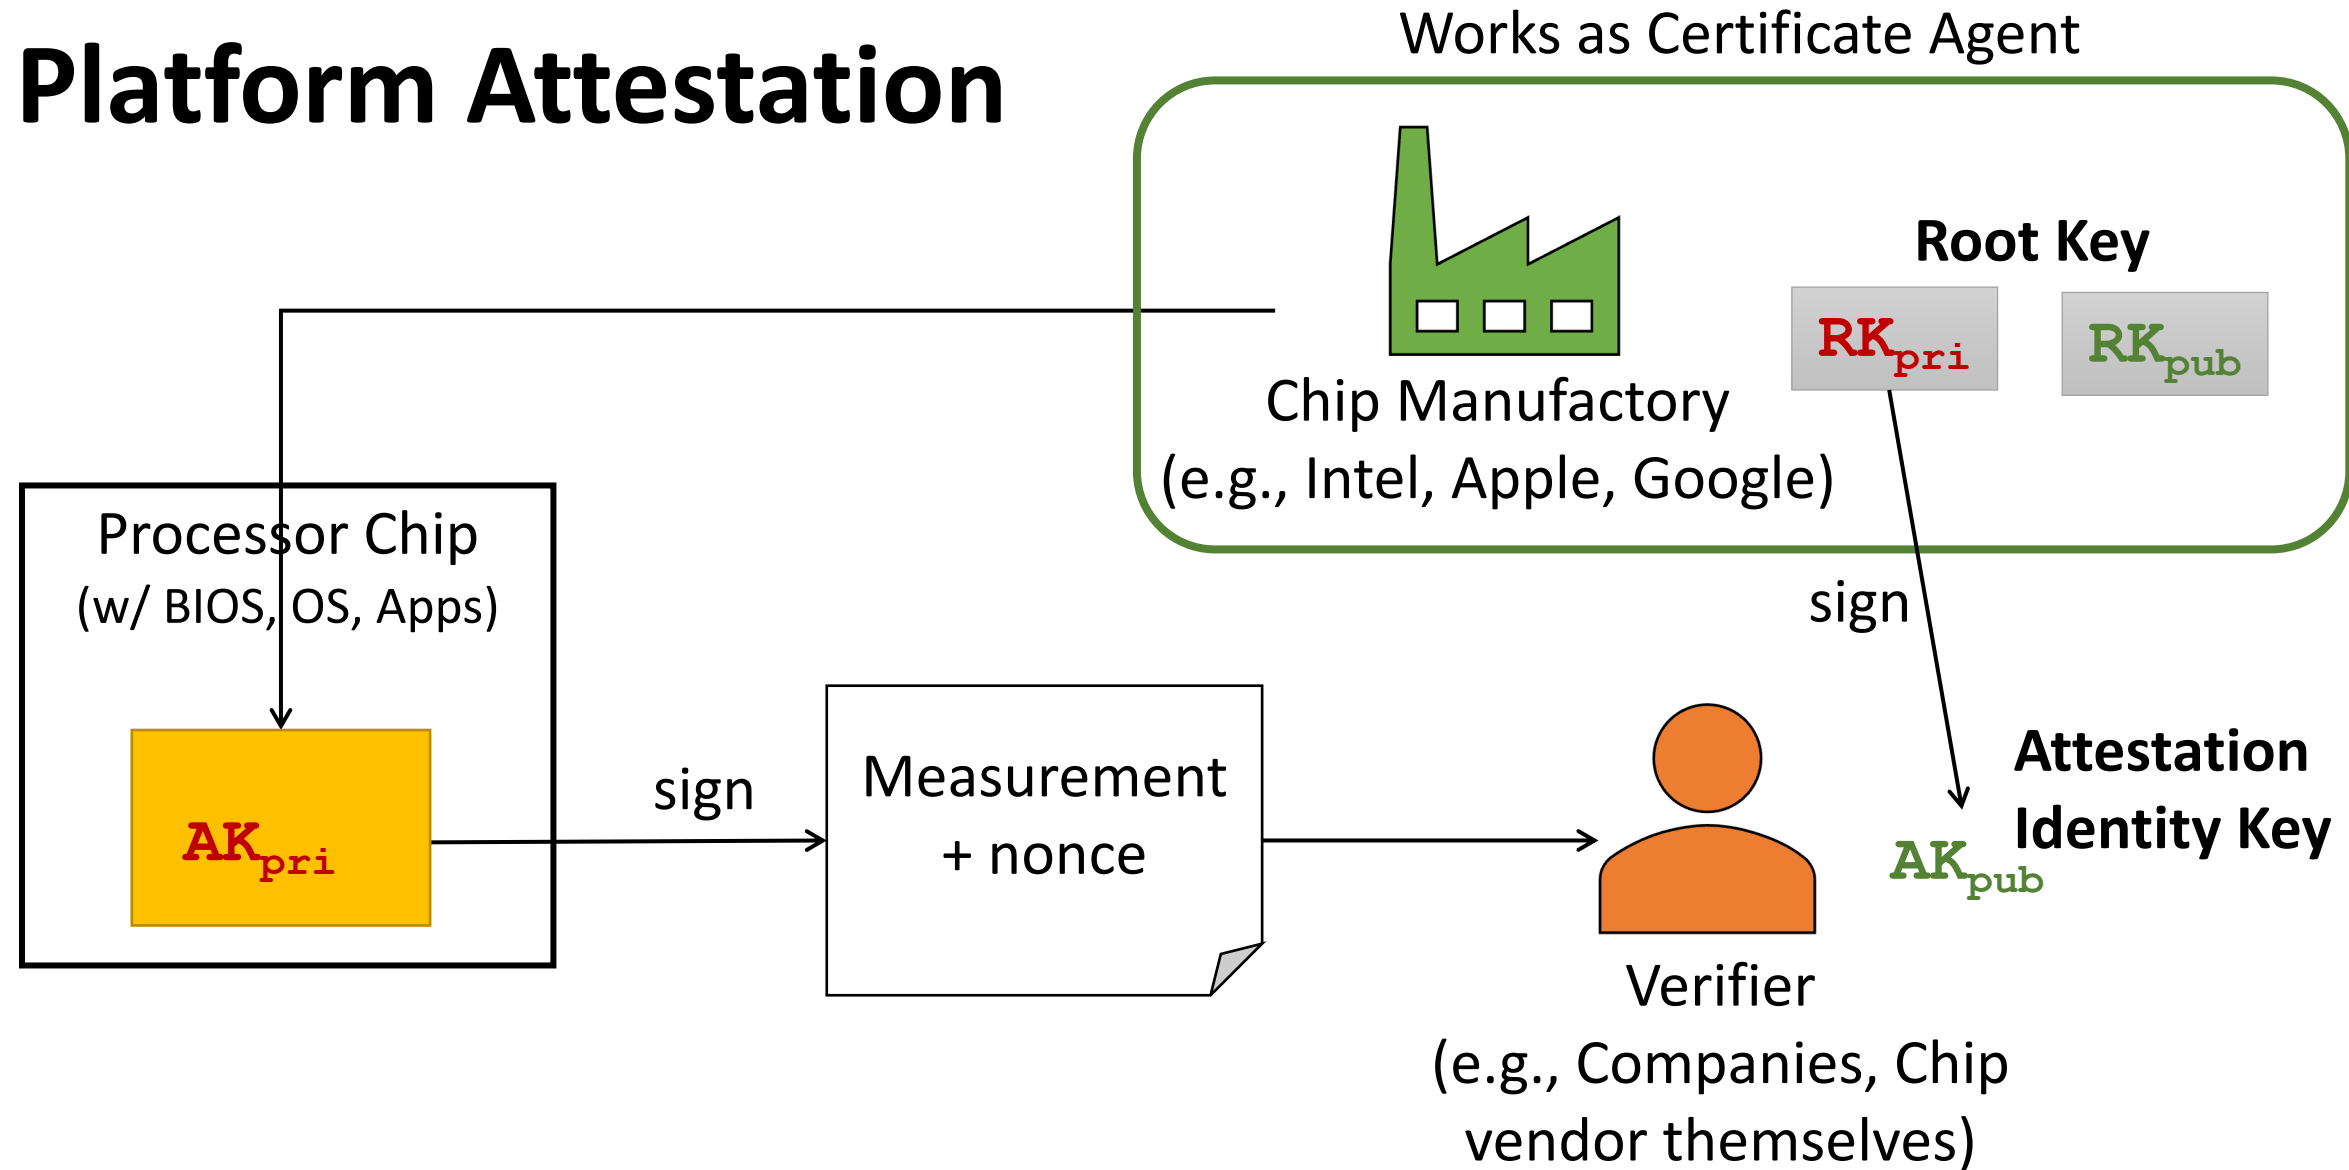


# Open-source Choice: Google Titan

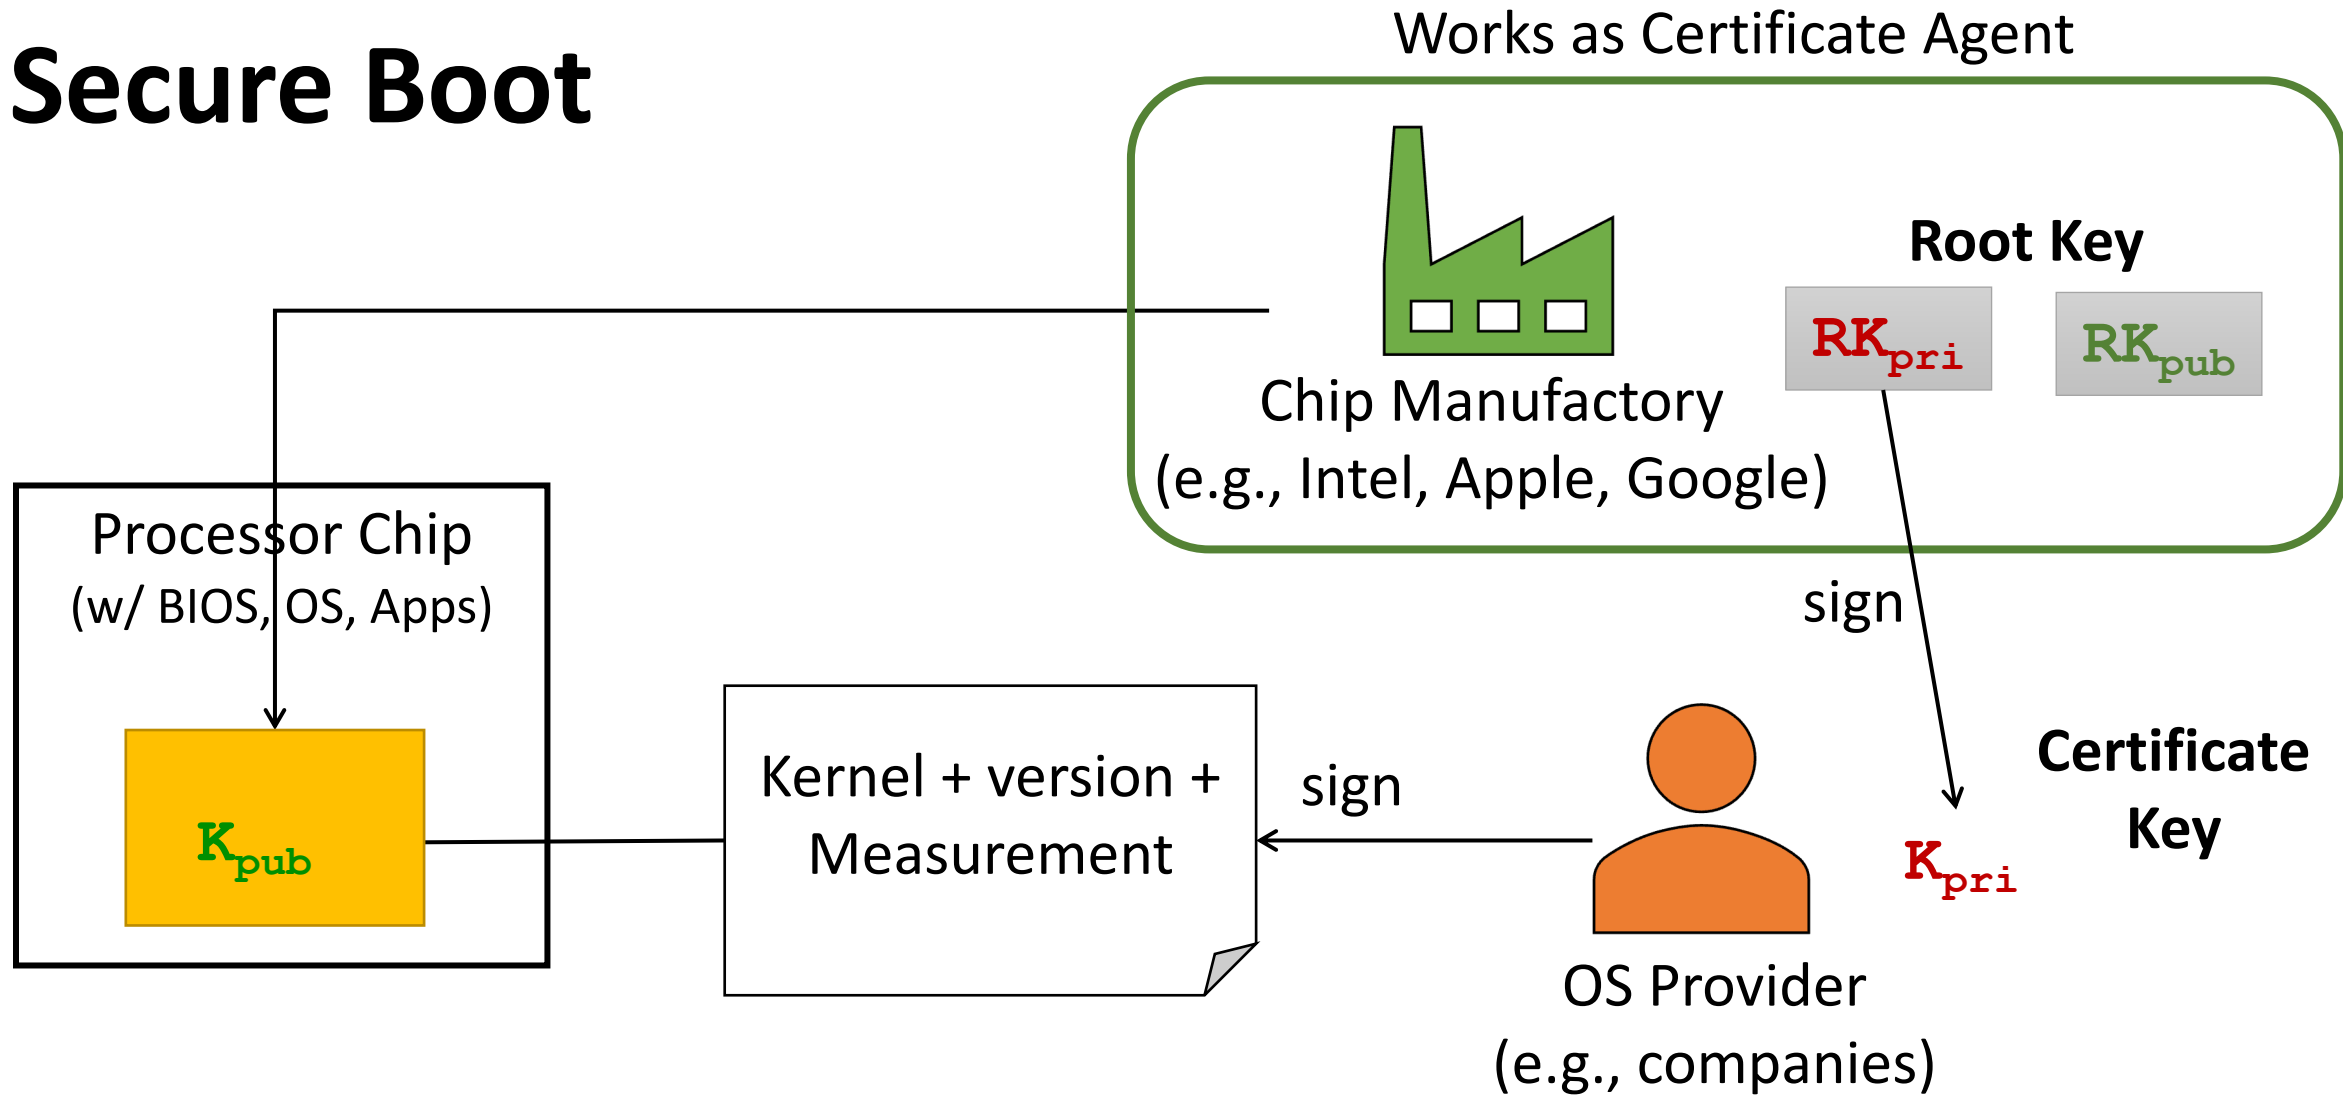


from [https://www.hotchips.org/hc30/1conf/1.14\\_Google\\_Titan\\_GoogleFinalTitanHotChips2018.pdf](https://www.hotchips.org/hc30/1conf/1.14_Google_Titan_GoogleFinalTitanHotChips2018.pdf)

# Platform Attestation



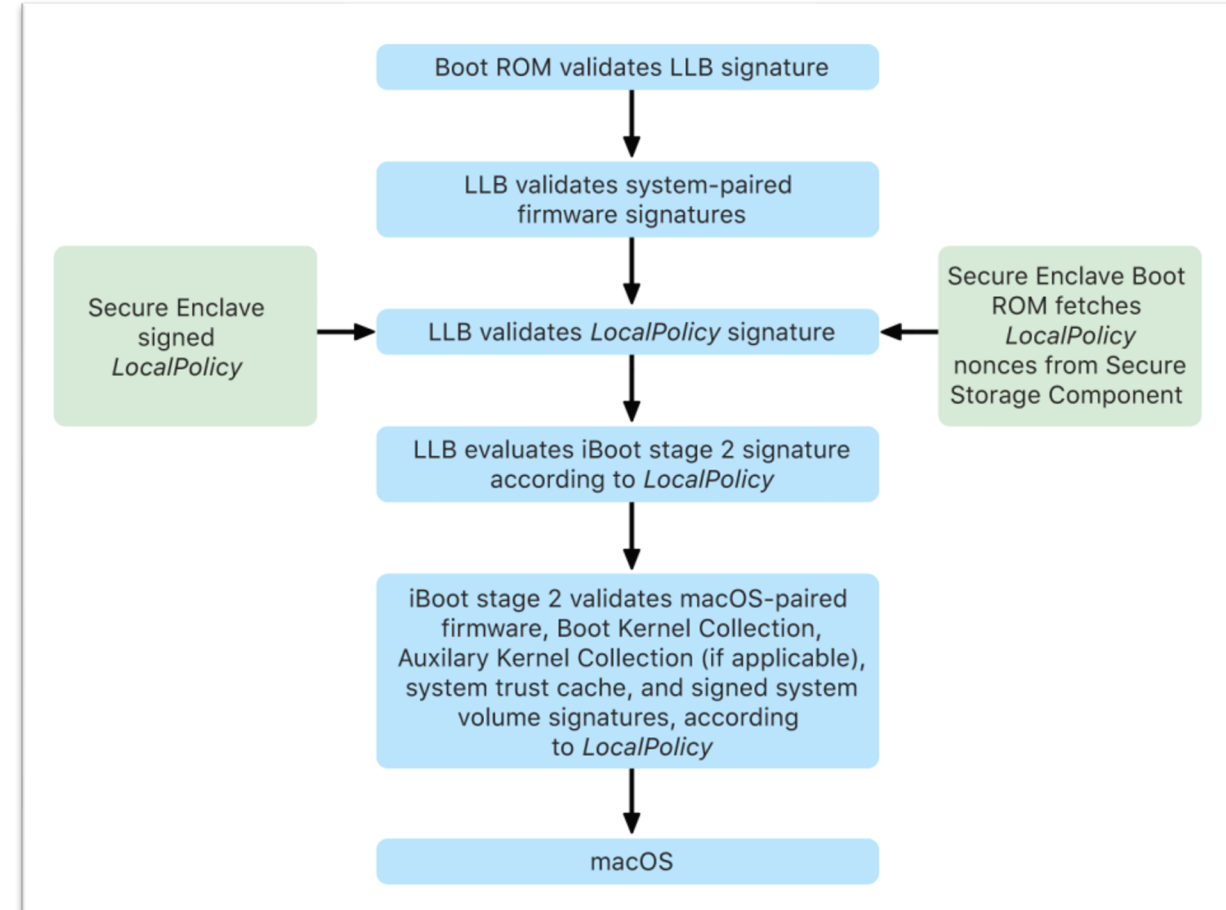
# Secure Boot



# Secure Boot

## Similar to TPM but with more constraints

- Each step is signed by Apple to prevent loading non-Apple systems
- Verify more components, including operating system, kernel extensions, etc.
- Keep track of version number to prevent rolling back to older/vulnerable versions





# Summary

What Can Hardware Security Modules Offer?

- Physical isolation
- Bind data and applications with the hardware device
- Establish root of trust
- More efficient

Challenges: software support. Programmability.

# Next: Physical Attacks

(with many demos 🛠️ 🎩 😄 ✨ )