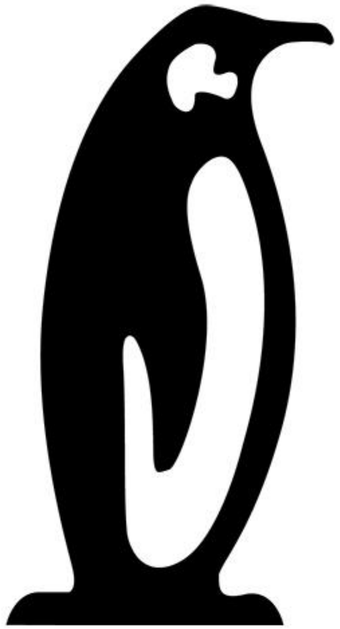

Linux Plumbers Conference

Dublin, Ireland September 12-14, 2022

Secure Bootloader for Confidential Computing

Jiewen Yao, Intel
Ken Lu, Intel

Linux Plumbers Conference | Dublin, Ireland Sept. 12-14, 2022

Notices & Disclaimers

- Intel technologies may require enabled hardware, software or service activation.
- No product or component can be absolutely secure.
- All product plans and roadmaps are subject to change without notice.
- The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.
- Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Agenda

- Measurement and Trust Chain ←
- Attestation and Disk Decryption

Confidential Computing in Cloud Usage

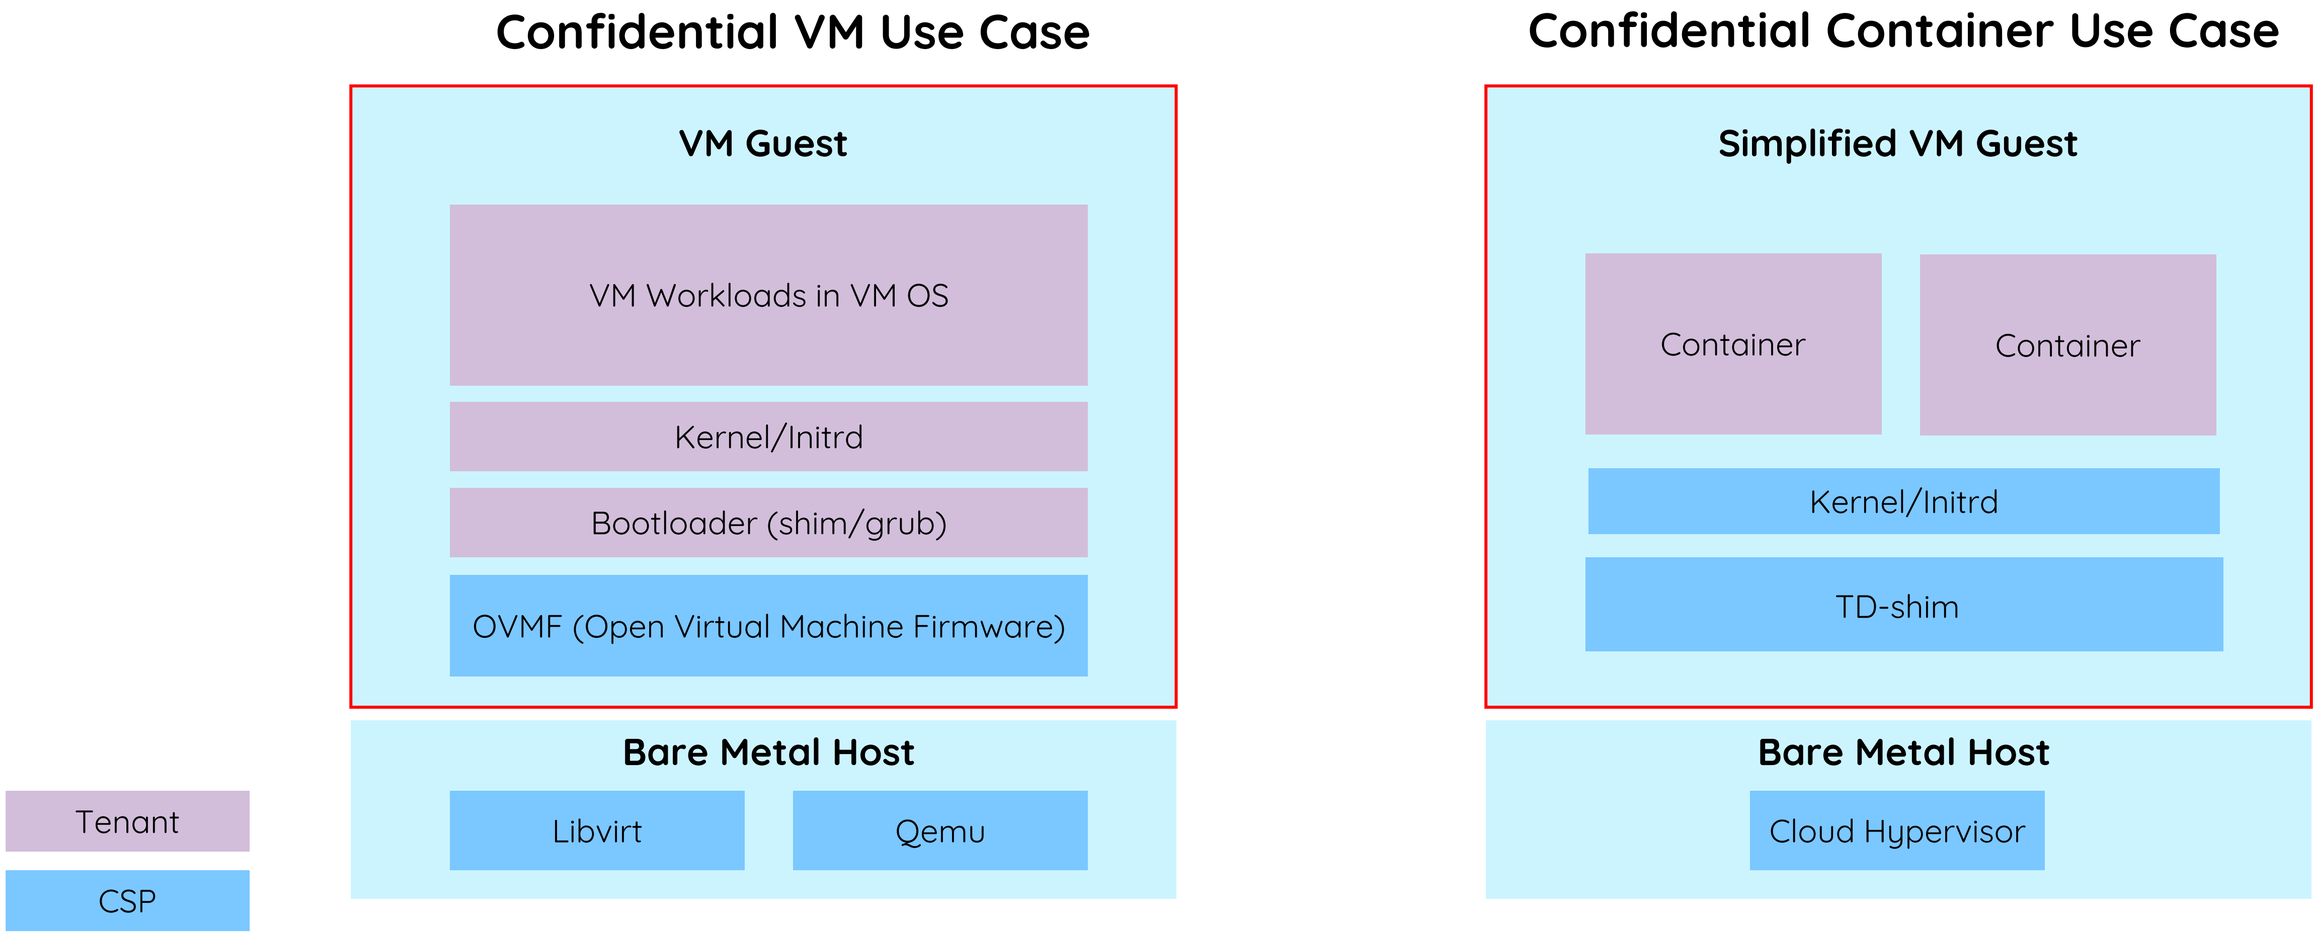

Measurement for Attestation

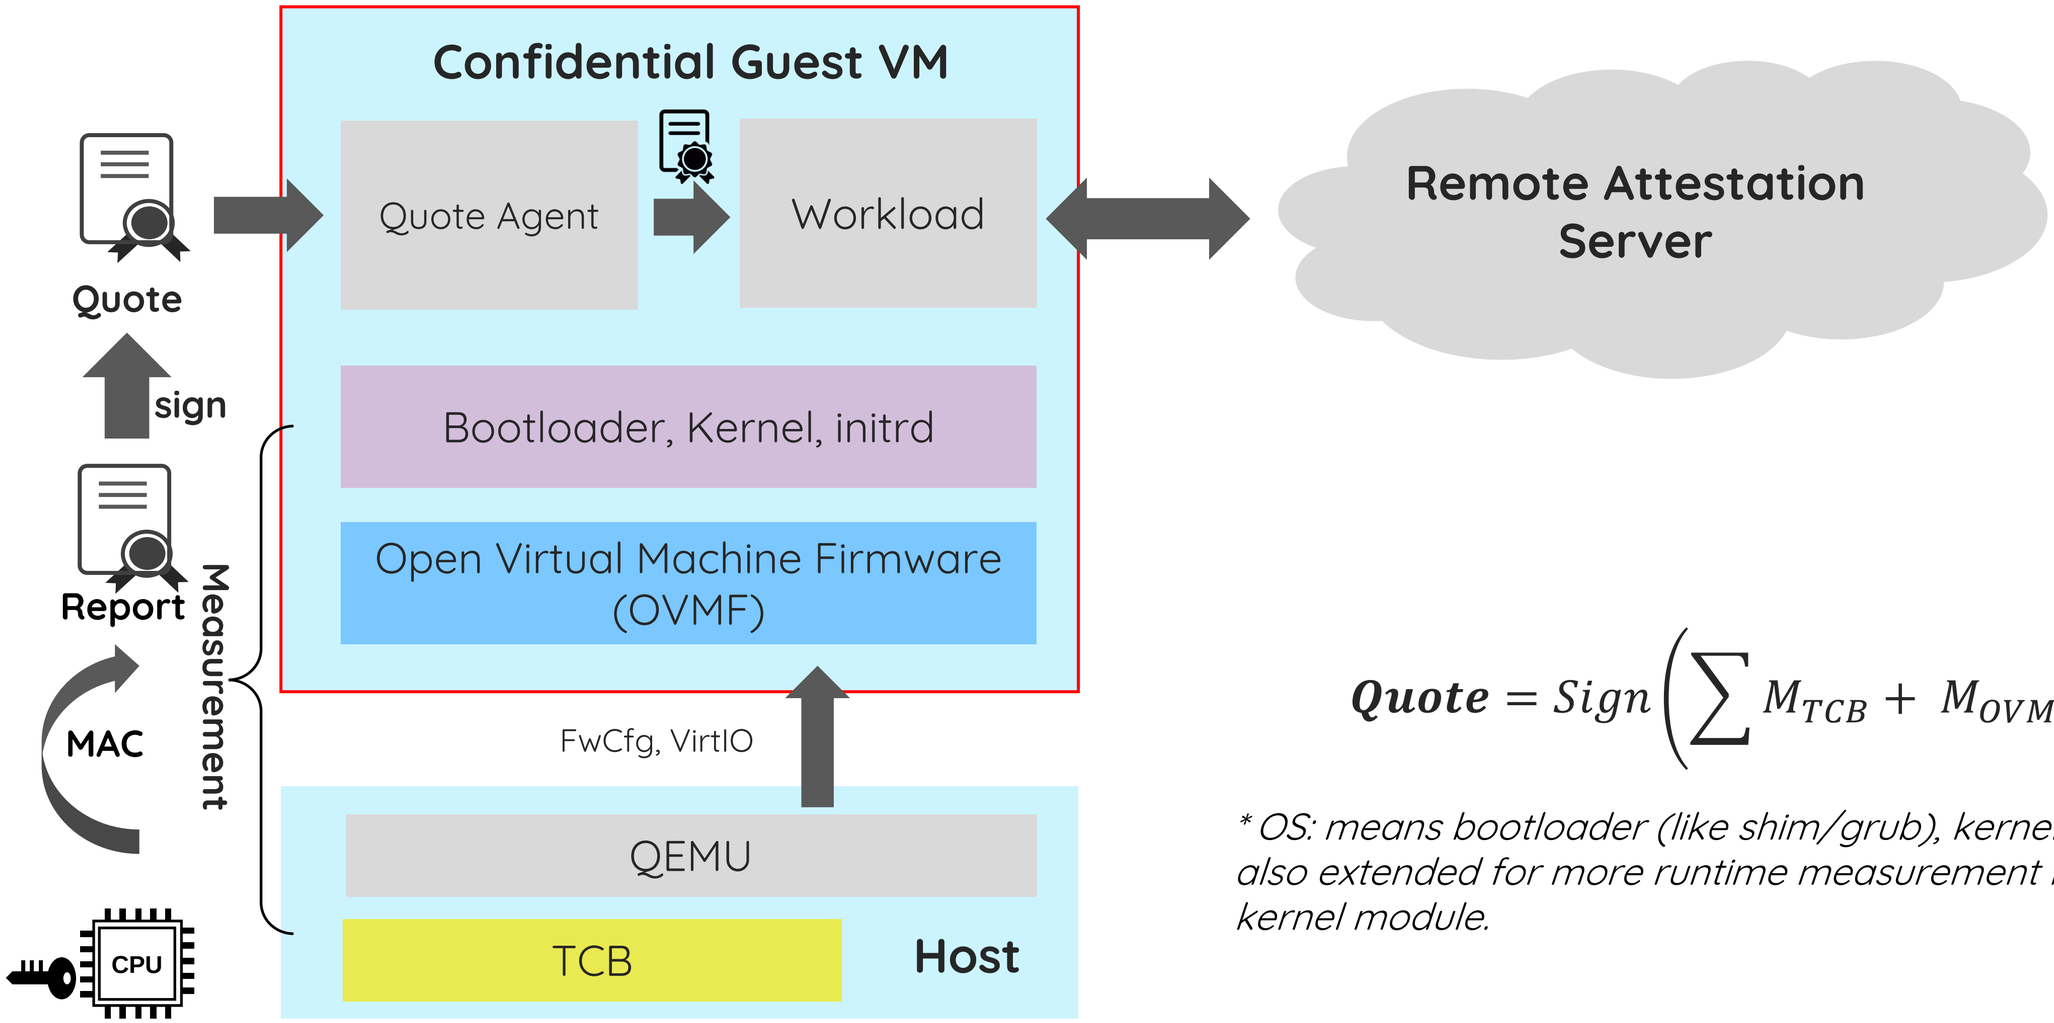

$$\boldsymbol{Quote} = Sign\left(\sum M_{TCB} + M_{OVMF} + \mathrm{M}_{OS}\right)$$

** OS: means bootloader (like shim/grub), kernel, initrd. It can be also extended for more runtime measurement like dynamic kernel module.*

TDREPORT and Measurement

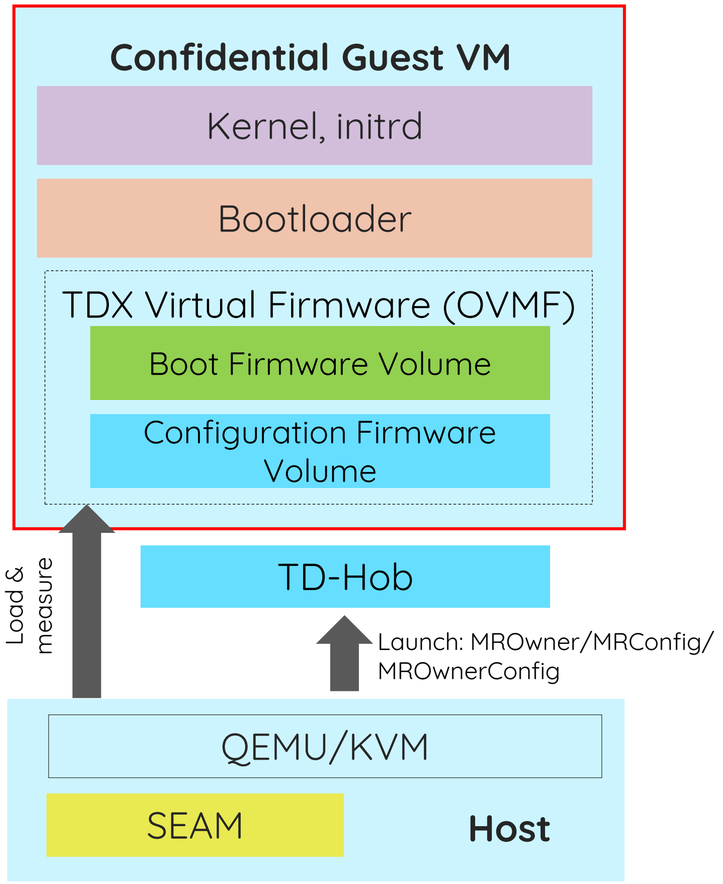

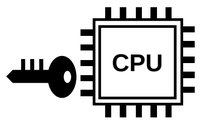

CPUSVN

0x000	REPORTMACSTRUCT	REPORTTYPE	RESERVER
0x010		CPUSVN	
0x020		TEE_TCB_INFO_HASH	
0x030			
0x040			
0x050			
0x060		TEE_INFO_HASH	
0x070			
0x080		REPORTDATA	
0x090			
0x0A0			
0x0B0			
0x0C0			
0x0D0		RESERVE	
0x0E0		MAC	
0x0F0			
0x100	TEE_TCB_INFO	VALID	TEE_TCB_SVN
0x110		TEE_TCB_SVN	MRSEAM
0x120		MRSEAM	
0x130			
0x140		MRSEAM	MRSIGNERSEAM
0x150		MRSIGNERSEAM	
0x160			
0x170		MRSIGNERSEAM	ATTRIBUTES
0x180		RESERVED	
0x190			
0x1A0			
0x1B0			
0x1C0			
0x1D0			
0x1E0			
0x1F0	RESERVED	RESERVED	

0x200	TDINFO	ATTRIBUTES	XFAM
0x210		MRTD	
0x220			
0x230			
0x240			
0x250		MRCONFIGID	
0x260			
0x270		MROWNER	
0x280			
0x290			
0x2A0		MROWNERCONFIG	
0x2B0			
0x2C0			
0x2D0		RTMR[0]	
0x2E0			
0x2F0			
0x300		RTMR[1]	
0x310			
0x320			
0x330		RTMR[2]	
0x340			
0x350			
0x360		RTMR[3]	
0x370			
0x380			
0x390		RESERVED	
0x3A0			
0x3B0			
0x3C0			
0x3D0			
0x3E0			
0x3F0			

Refer: https://github.com/tianocore/edk2/blob/master/MdePkg/Include/IndustryStandard/Tdx.h

UEFI Confidential Computing Interface

- **Interface**
  - EFI_CC_MEASUREMENT_PROTOCOL
    - Abstract the measurement for virtual firmware in confidential computing environment.
      - Similar to TPM2: EFI_TCG2_PROTOCOL
  - CCEL ACPI Table
    - Provide the CC event log to the operating system.
      - Similar to TPM2: TPM2 ACPI table

- **Code**
  - EDKII - CcMeasurement
    - https://github.com/tianocore/edk2/blob/master/MdePkg/Include/Protocol/CcMeasurement.h
  - Shim Enabling (Thanks Robbie Harwood)
    - https://github.com/rhboot/shim/commit/4fd484e4c29364b4fdf4d043556fa0a210c5fdfc
  - Grub Enabling (Thanks Daniel Kiper)
    - https://git.savannah.gnu.org/cgit/grub.git/commit/?id=4c76565b6cb885b7e144dc27f3612066844e2d19

- **Specification**
  - Intel GHCI specification (published)
    - https://cdrdv2.intel.com/v1/dl/getContent/726790
  - UEFI Confidential Computing Extension
    - UEFI 2.10 - https://uefi.org/specs/UEFI/2.10/38_Confidential_Computing.html#confidential-computing
    - ACPI 6.5 - https://uefi.org/specs/ACPI/6.5/05_ACPI_Software_Programming_Model.html#cc-event-log-acpi-table

CC Measurement Flow in Pre-Boot

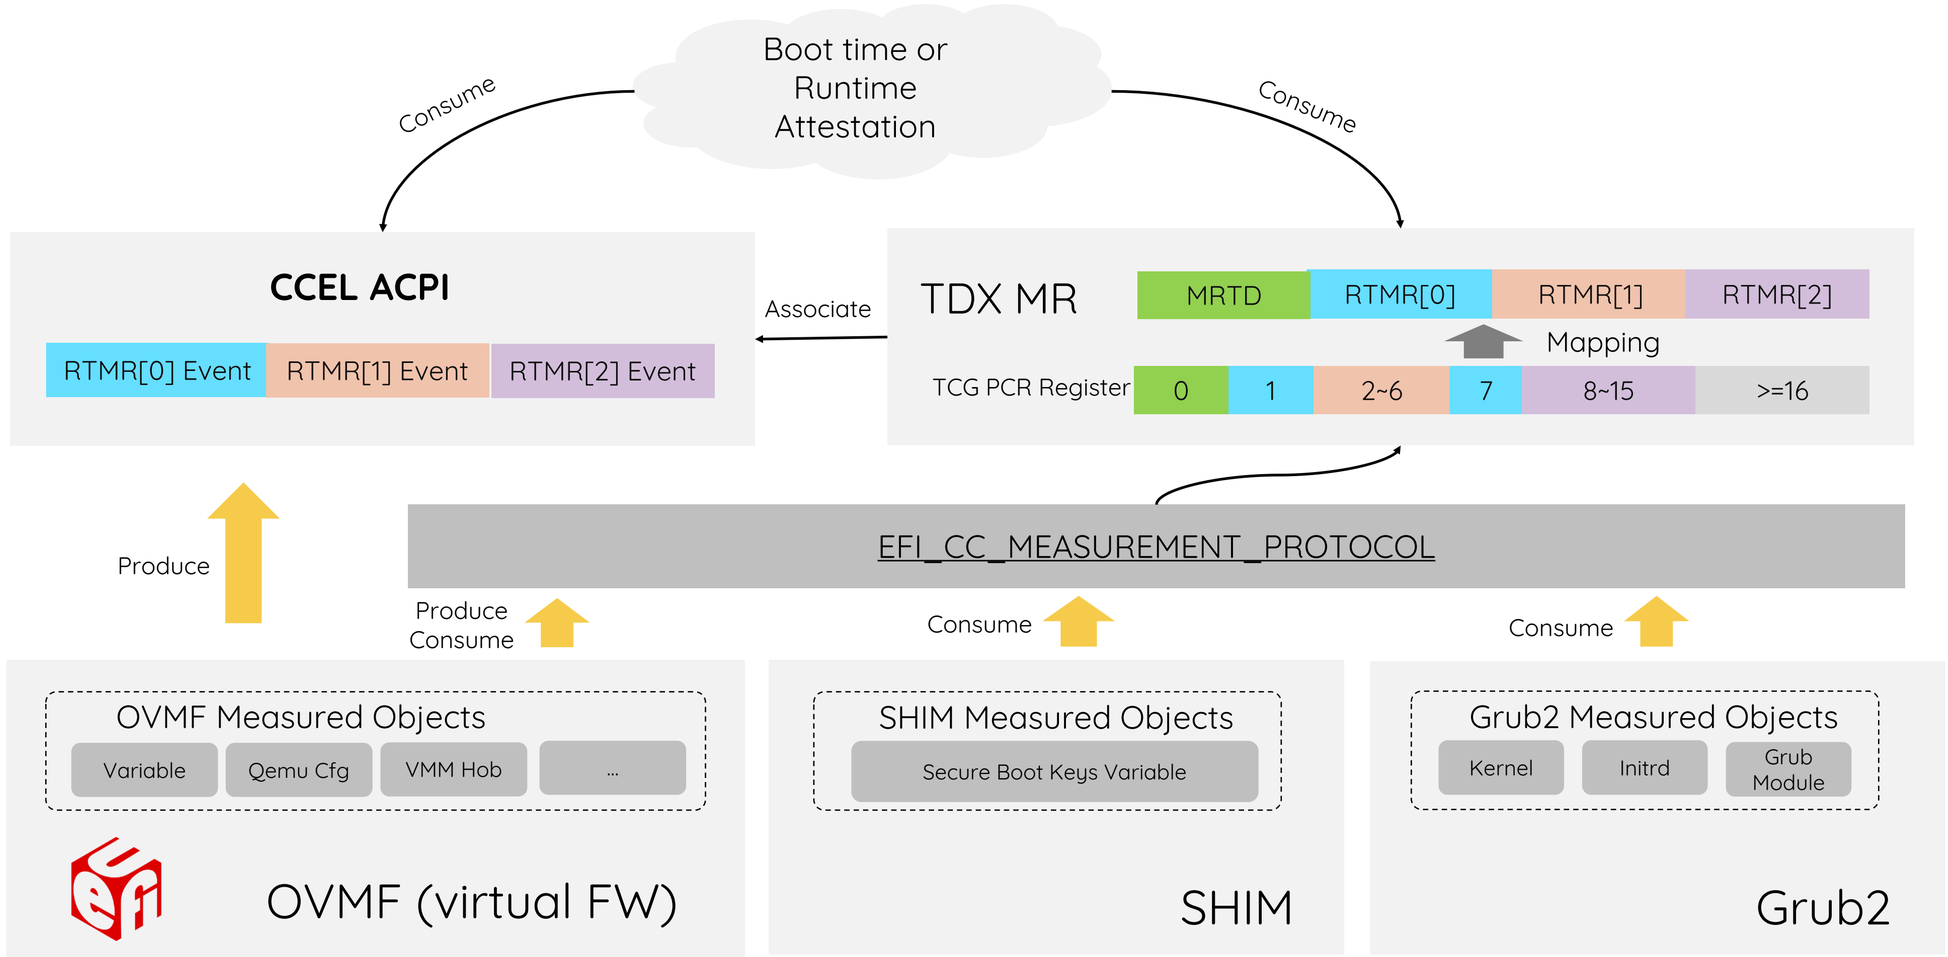

Linux Bootloader Phase

Dump CC Event Log

```
==== TDX Event Log Entry - 71 [0x1C331F50] ====
RTMR              : 2
Type              : 0xD (EV_IPL)
Length            : 161
Algorithms ID     : 12 (TPM_ALG_SHA384)
Digest[0] :
00000000  E8 09 21 B2 EE 7B D3 C9 24 2D 7D 2B 60 CA D2 8A  ..!..{..$-}+`...
00000010  7F E6 8E 4C 8E A3 5A AE F3 E6 21 81 3B 60 E8 FF  ...L..Z...!.;`..
00000020  BE BE 2A AF ED 50 03 23 BE D6 4E 99 B9 71 6A 70  ..*..P.#..N..qjp
RAW DATA: ---------------------------------------------
1C331F50  03 00 00 00 0D 00 00 00 01 00 00 00 0C 00 E8 09  ................
1C331F60  21 B2 EE 7B D3 C9 24 2D 7D 2B 60 CA D2 8A 7F E6  !..{..$-}+`.....
1C331F70  8E 4C 8E A3 5A AE F3 E6 21 81 3B 60 E8 FF BE BE  .L..Z...!.;`....
1C331F80  2A AF ED 50 03 23 BE D6 4E 99 B9 71 6A 70 5F 00  *..P.#..N..qjp_.
1C331F90  00 00 67 72 75 62 5F 63 6D 64 20 69 6E 69 74 72  ..grub_cmd initr
1C331FA0  64 20 28 68 64 30 2C 6D 73 64 6F 73 33 29 2F 62  d (hd0,msdos3)/b
1C331FB0  6F 6F 74 2F 69 6E 69 74 72 61 6D 66 73 2D 35 2E  oot/initramfs-5.
1C331FC0  31 35 2E 30 2D 53 50 52 2E 4D 56 50 2E 50 43 2E  15.0-SPR.MVP.PC.
1C331FD0  76 31 30 2E 34 2E 6D 76 70 34 30 2E 65 6C 38 2E  v10.4.mvp40.el8.
1C331FE0  78 38 36 5F 36 34 2B 67 75 65 73 74 2E 69 6D 67  x86_64+guest.img
1C331FF0  00                                               .
RAW DATA: ---------------------------------------------
```

- Pytdxmeasure (https://github.com/intel/tdx-tools/tree/main/utils/pytdxmeasure) retrieve/print the TDREPORT and CC Event log within TDX guest.
  - Event log example for Direct Boot
  - Event log example for Grub2 Boot

Measurement in Td-Shim

- Td-Shim: a lightweight TDX shim firmware
  - Support confidential container or service TD
  - Support kernel direct boot

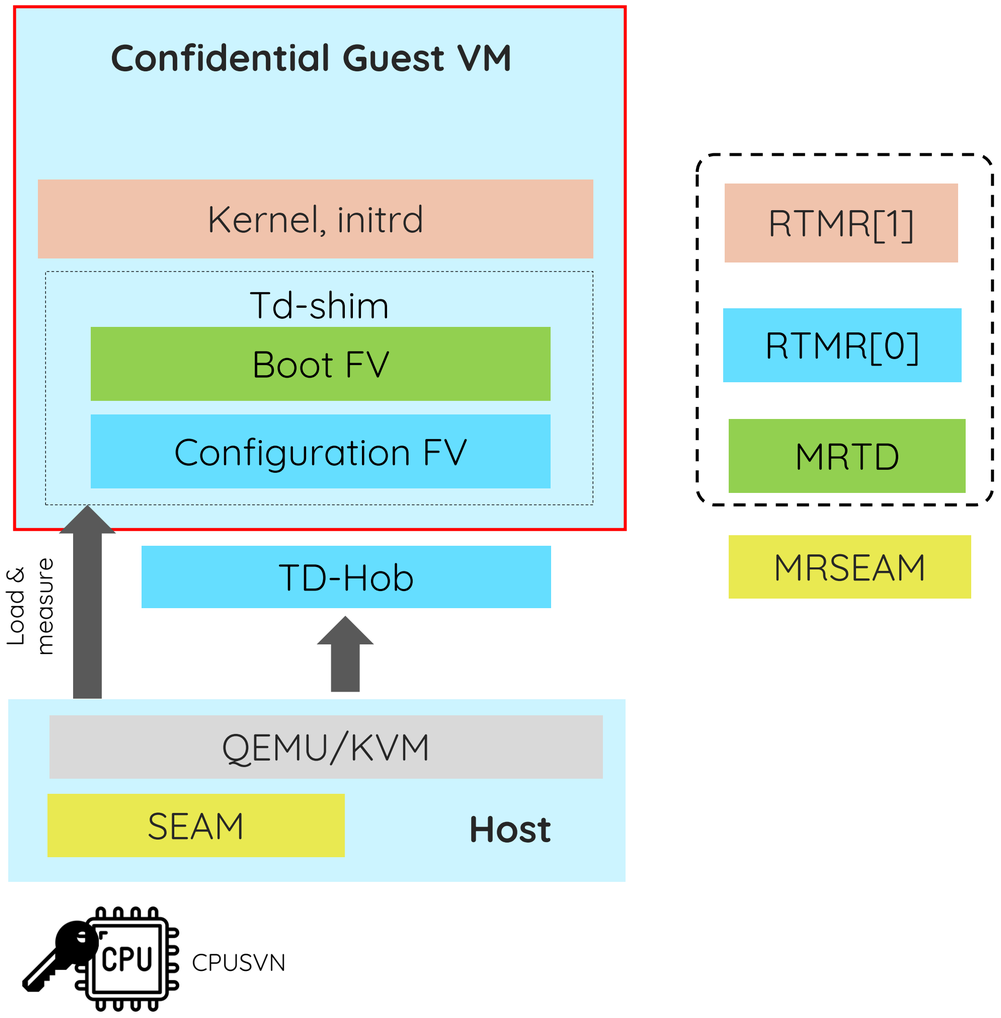

Agenda

- Measurement and Trust Chain
- Attestation and Disk Decryption ⬅

Attestation for Confidential Computing

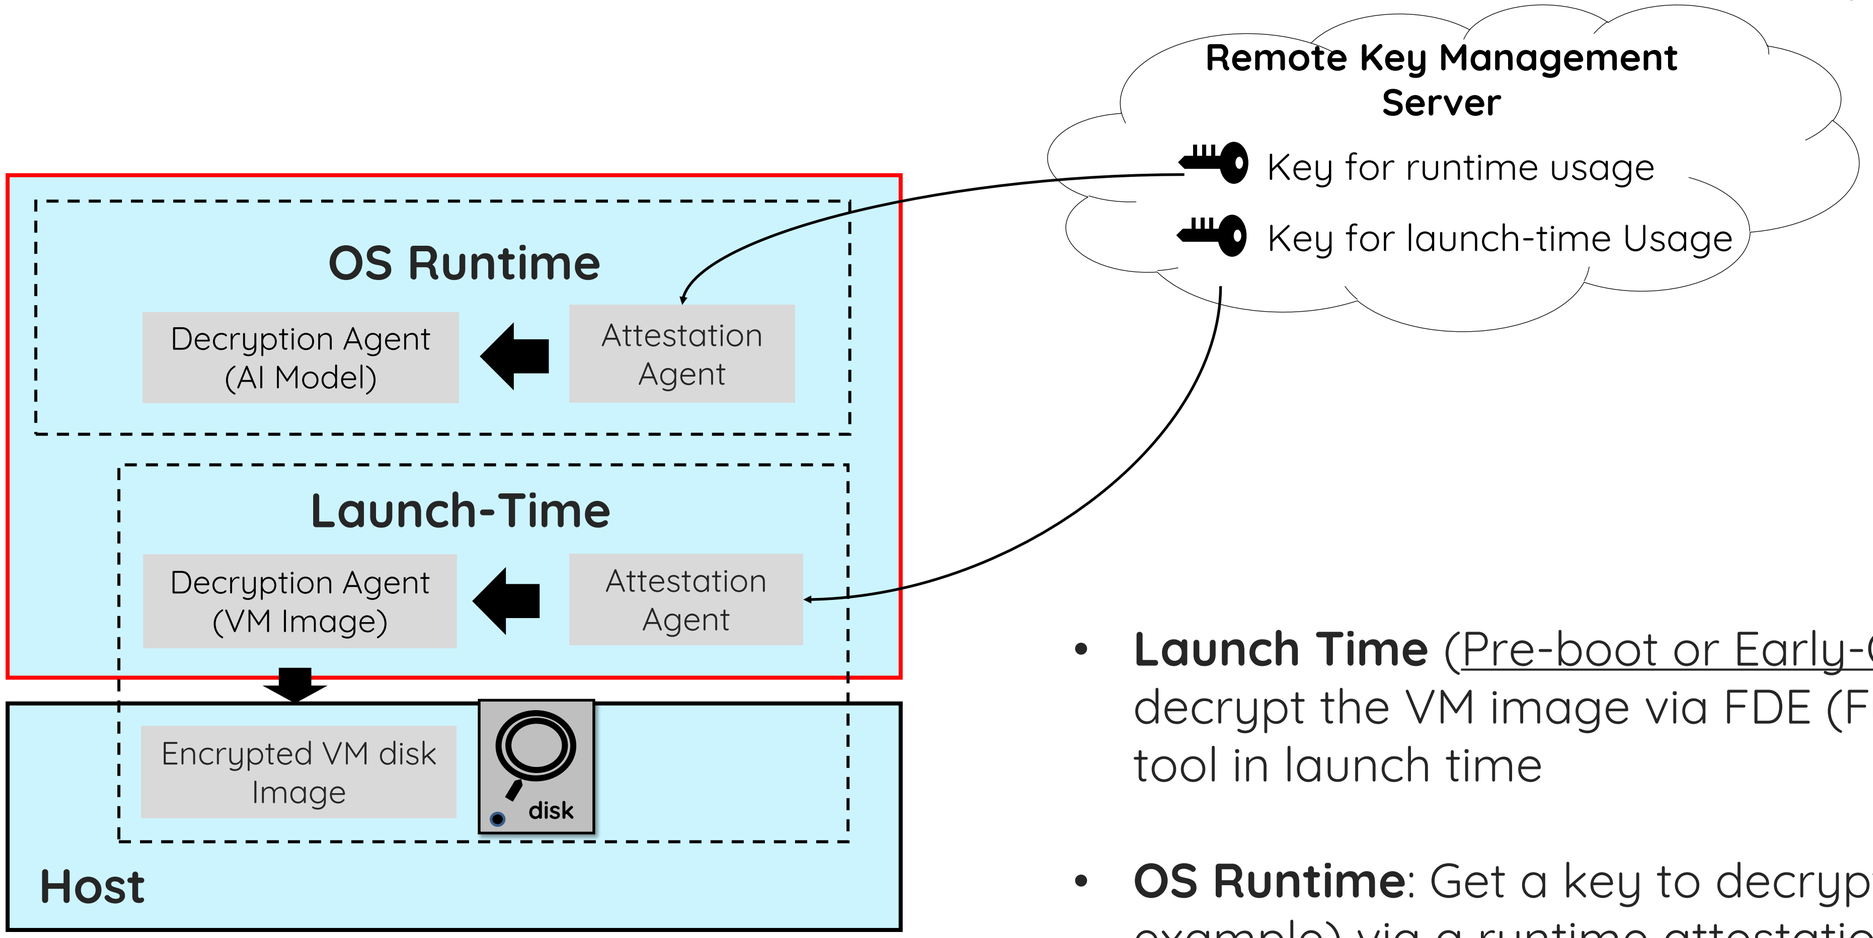

- **Launch Time** (Pre-boot or Early-OS boot): Get a key to decrypt the VM image via FDE (Full Disk Encryption) tool in launch time
- **OS Runtime**: Get a key to decrypt AI model (as example) via a runtime attestation agent

Considerations for Disk Decryption in Launch Time

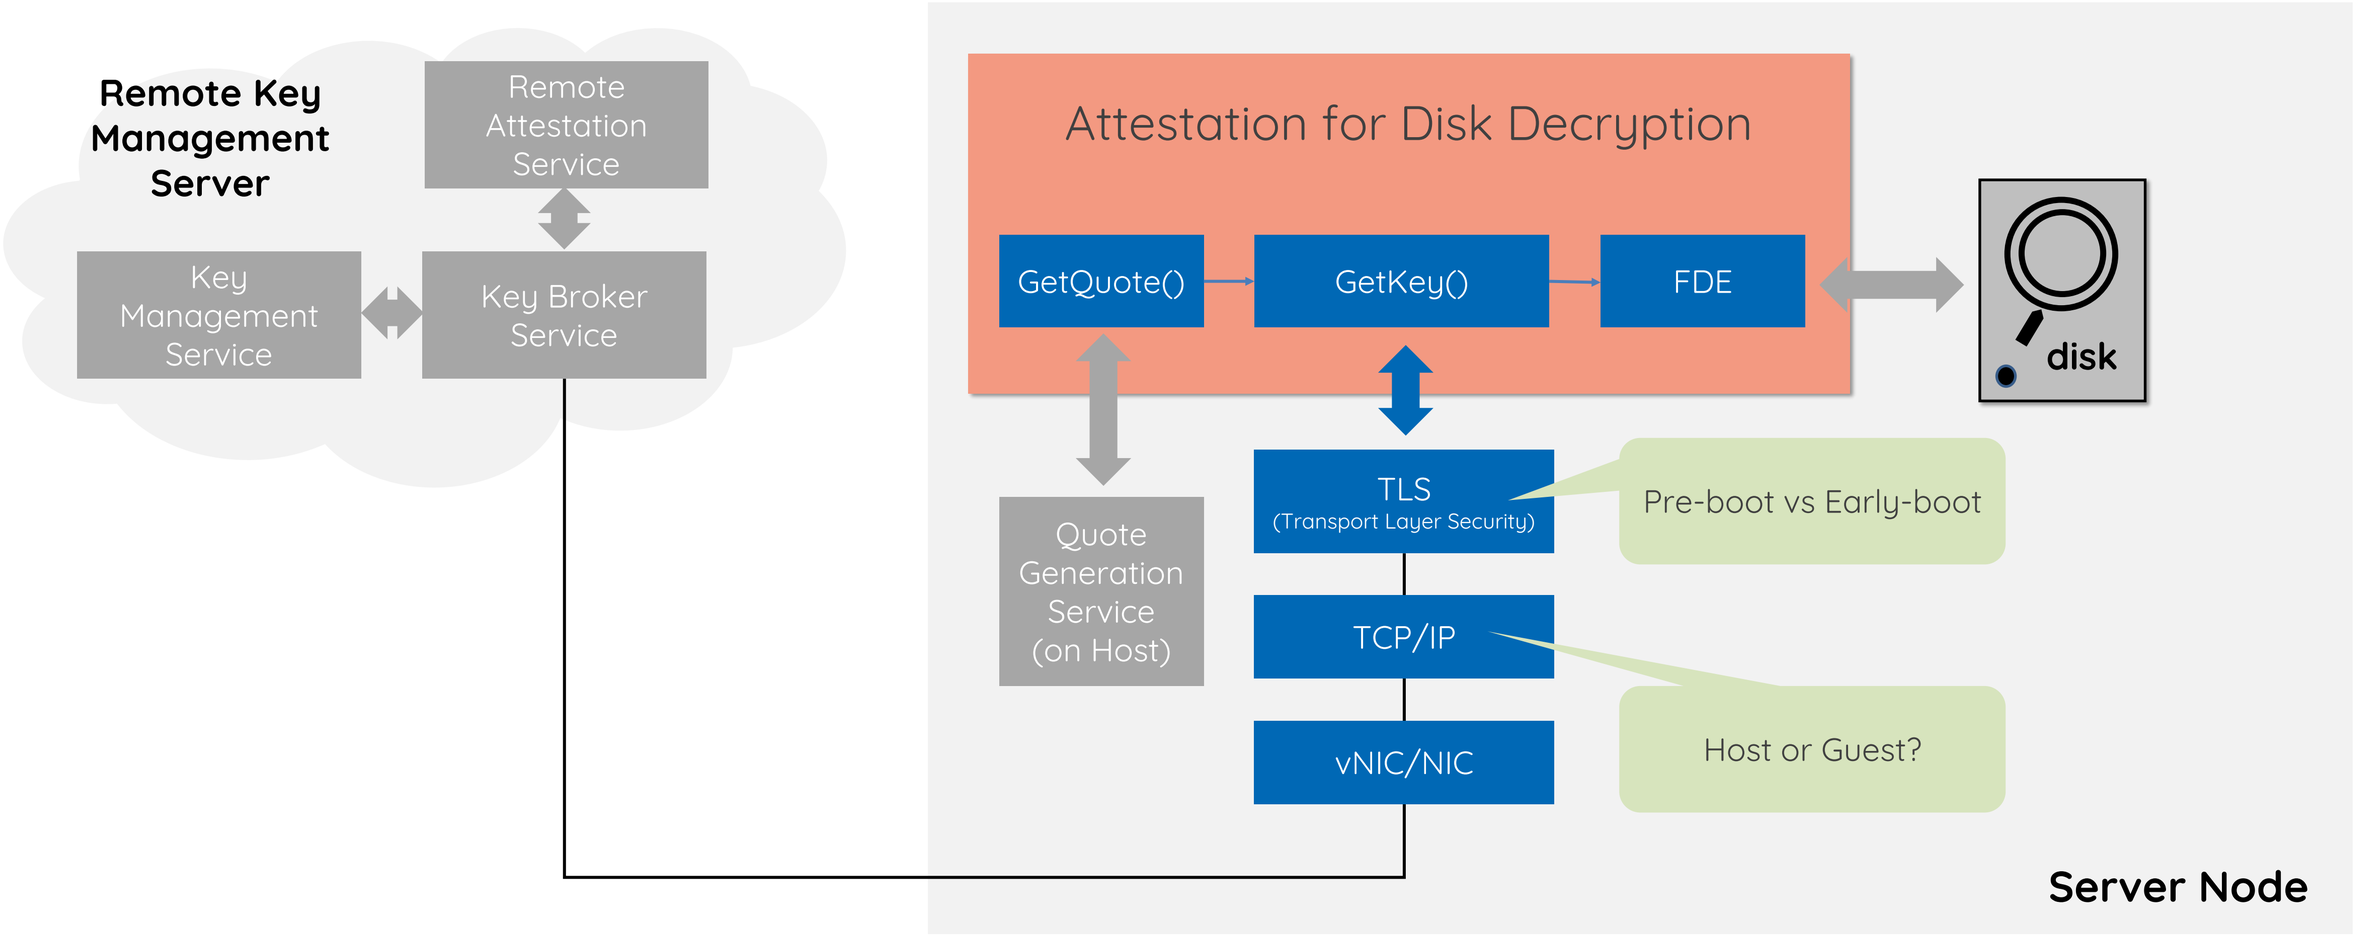

Pre-Boot Disk Decryption

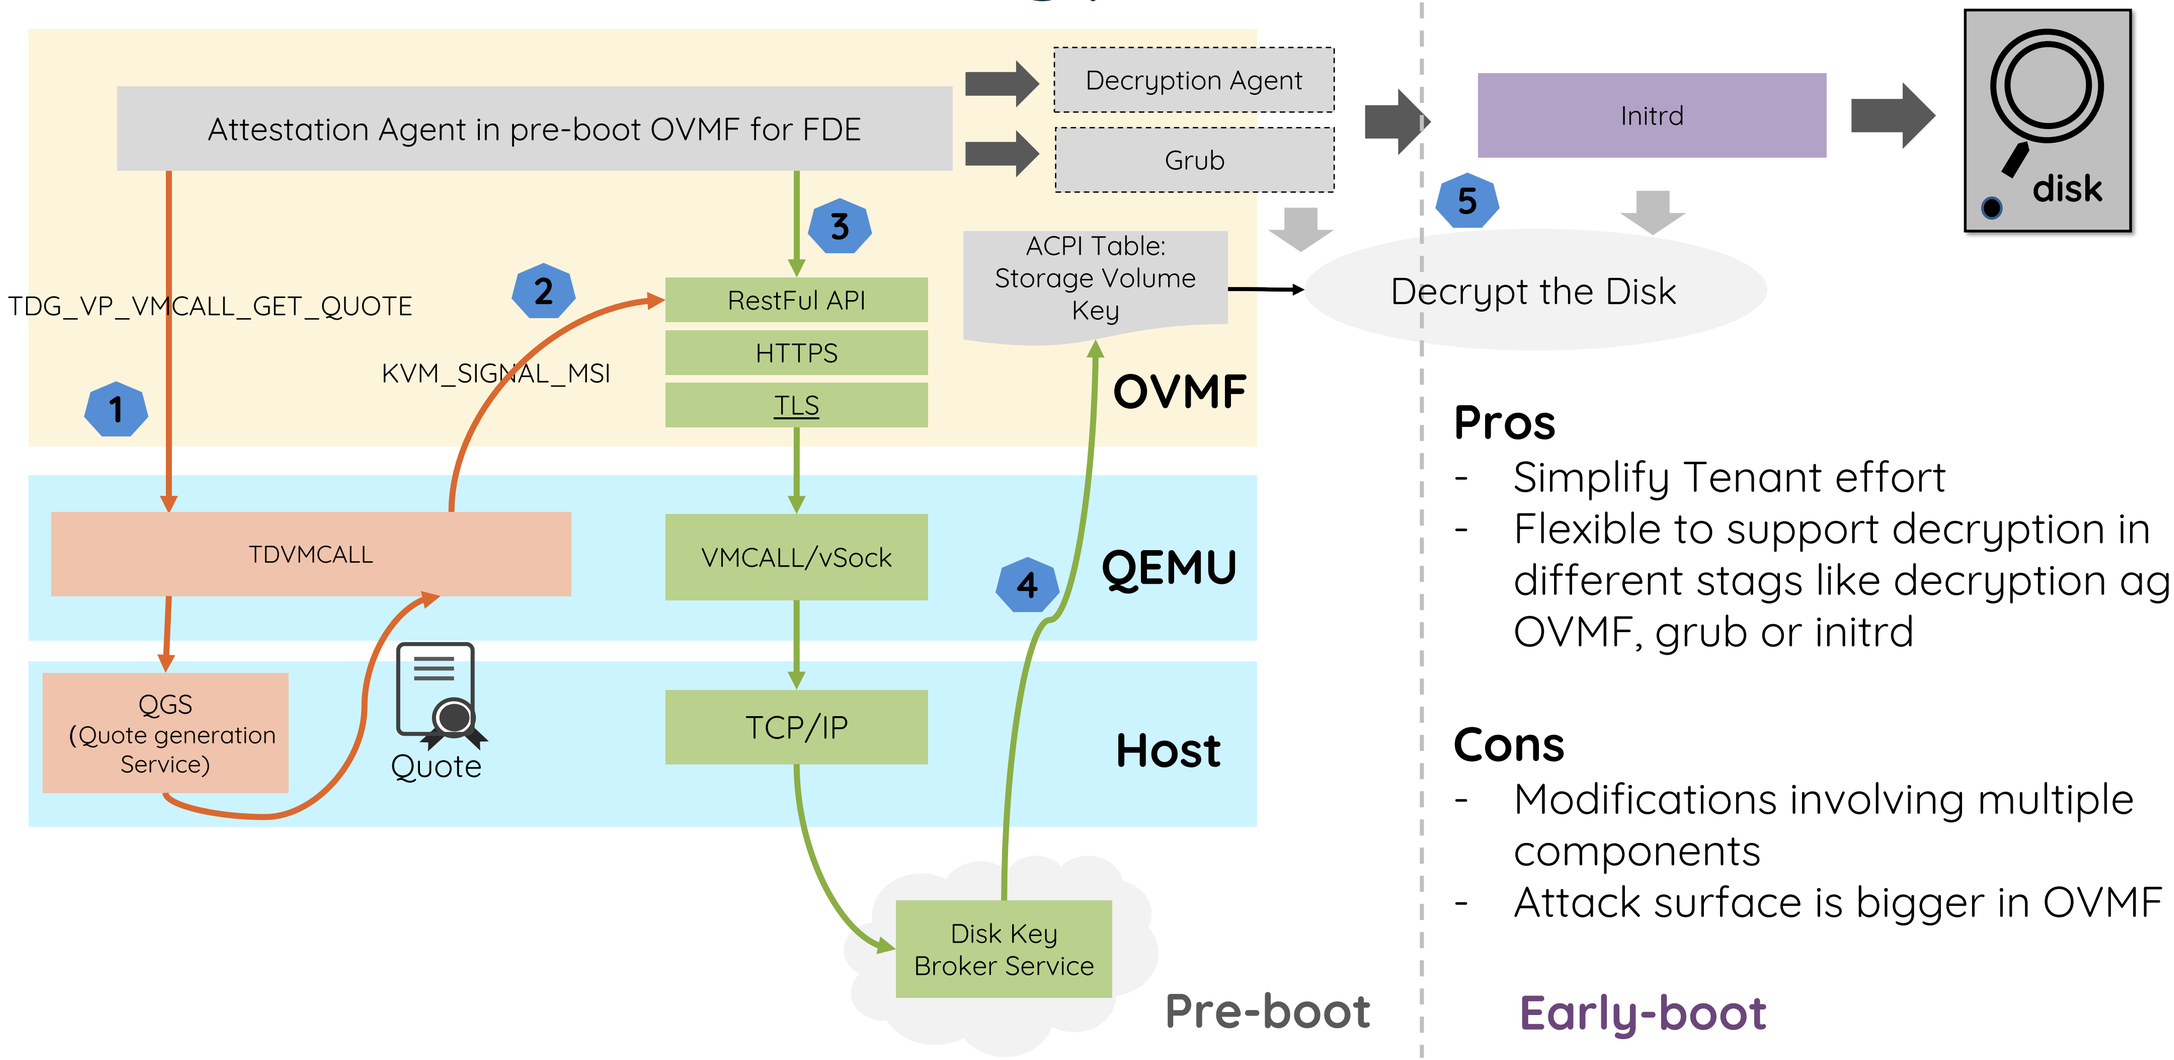

Pros

- Simplify Tenant effort
- Flexible to support decryption in different stags like decryption agent in OVMF, grub or initrd

Cons

- Modifications involving multiple components
- Attack surface is bigger in OVMF

Refer: Storage-Volume-Key is defined in GHCI specification chapter 4.4 or ACPI 6.5

Early-Boot Disk Decryption

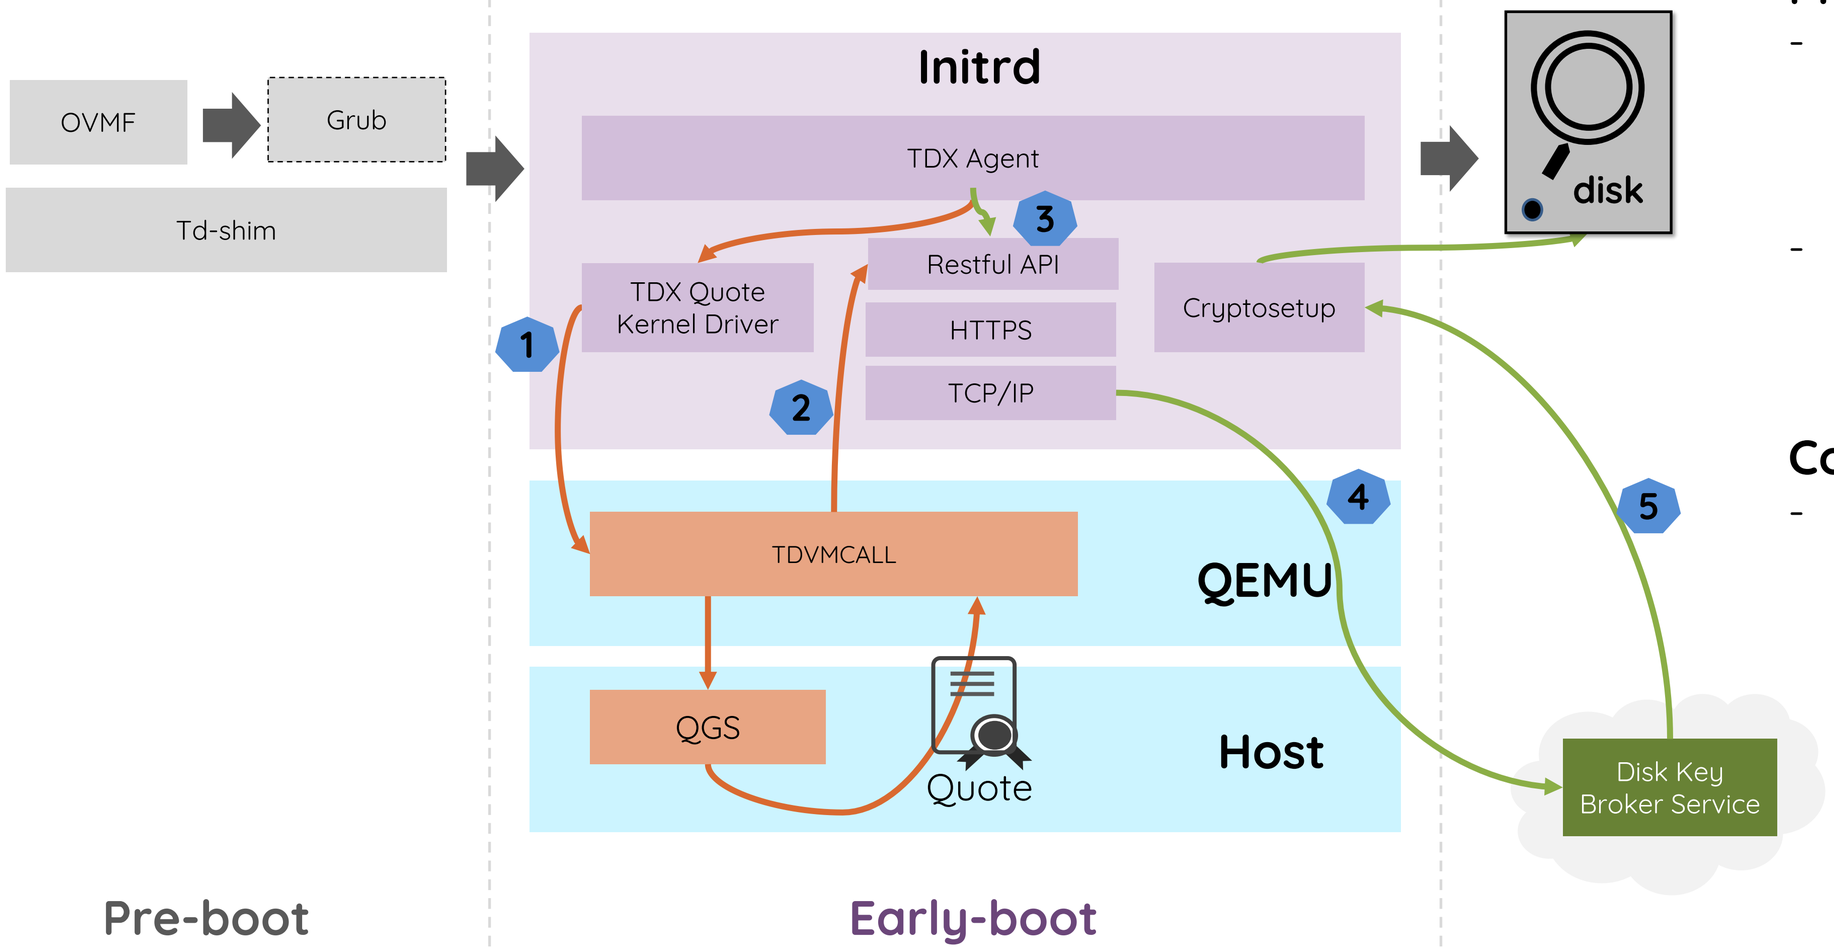

Pros

- The attestation process of disk decryption is running in user space like a normal runtime-attestation
- Allow different type vBIOS/bootloader like OVMF, simplified td-shim, etc.

Cons

- Need modify customer's VM image if put initrd in /boot partition instead of "-initrd" via launch command

Linux TDX MVP Stack

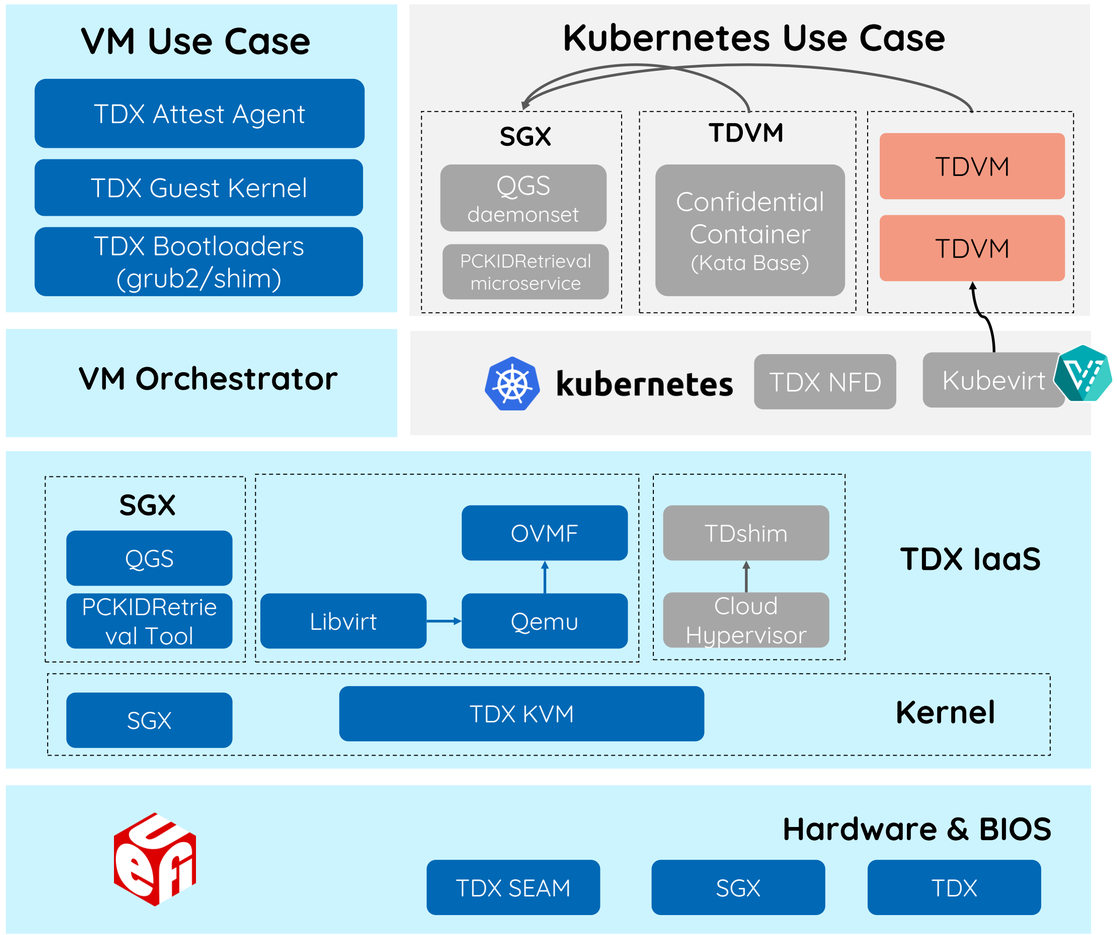

- TDX MVP Stack: github.com/intel/tdx-tools
- VM Orchestrator:
  - Launch TD VM via qemu & libvirt
  - Direct boot & Grub boot
  - Secure boot & measured boot
  - OS runtime attestation
  - Launch time attestation (WIP)
- Kubernetes Orchestrator (WIP):
  - Confidential Container (CNCF): Launch/Manage confidential container in TDVM transparently
  - Kubevirt: Launch/Manage the TDVM explicitly

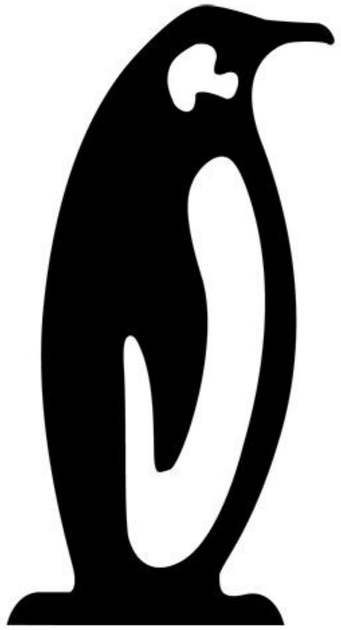

Linux Plumbers Conference

Dublin, Ireland September 12-14, 2022