TDX status update

Isaku Yamahata
<isaku.yamahata@intel.com>

@ymhtq

Introduction: What's TDX

Trust Domain Extensions(TDX)

- Hardware protected VM
  - Memory and CPU state
  - Base for confidential computing
- Hardware + software
  - CPU ISA extensions
  - Memory encryption
  - Firwmare(a.k.a. the TDX module)

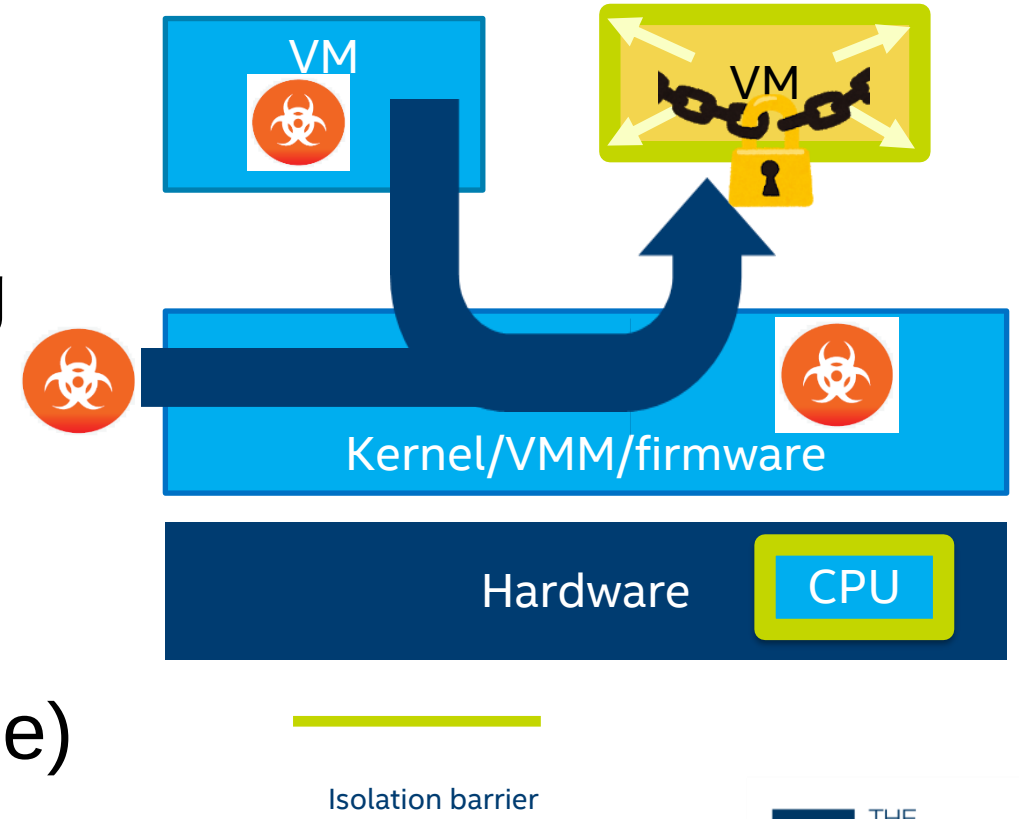

Touchpoints

- Linux Host
- Linux KVM
- Qemu
- Linux Guest

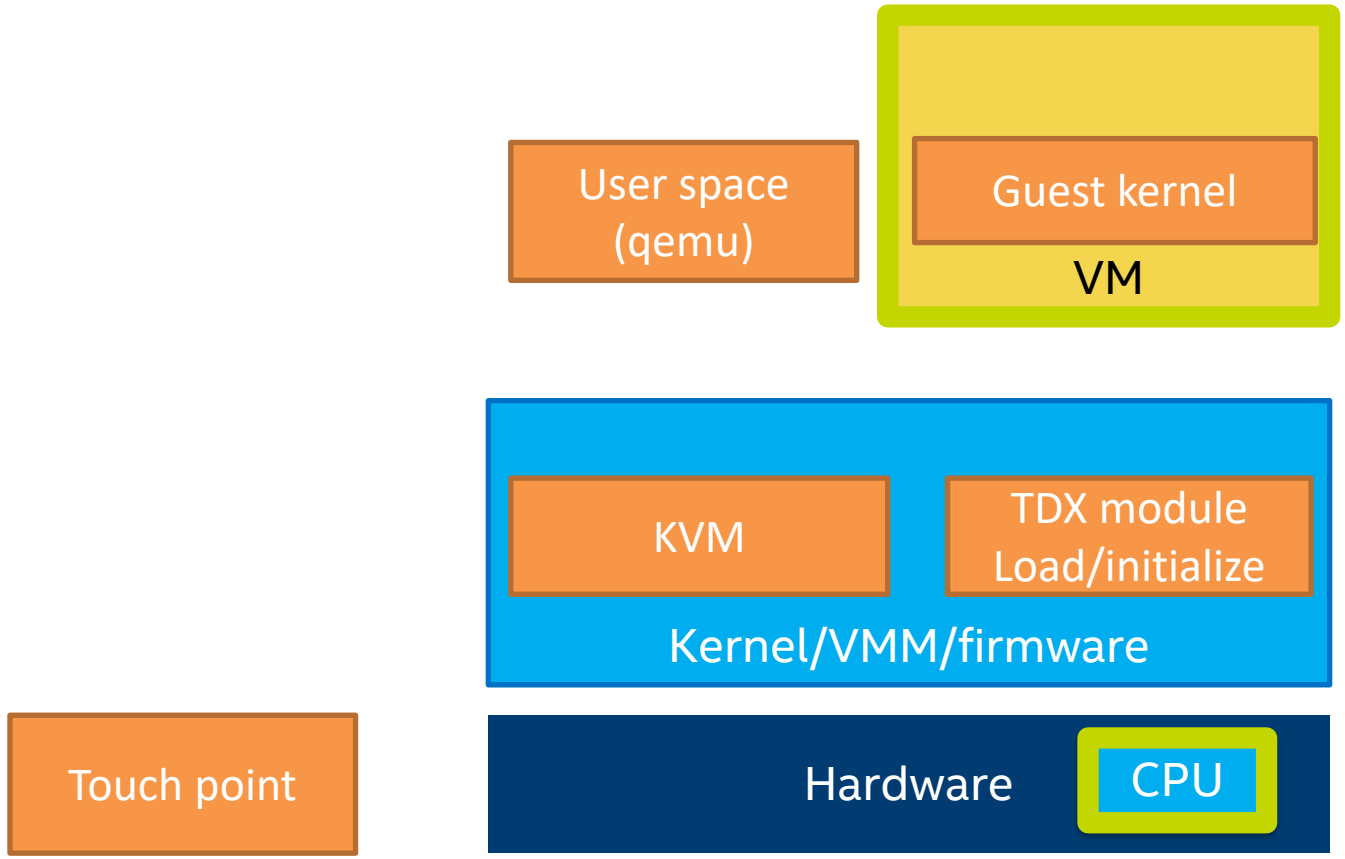

Linux host

- Load/initialize the firmware a.k.a. the TDX module
- Sysfs
- Runtime update of the firmware(after first merge)
- Kexec/kdump(after first merge)

KVM

- Initialization/teardown
  - Replace basic VMX operation with firmware call(a.k.a. SEAM call)
- KVM MMU: the next slide
- Debugger support(after first merge)
  - Qemu GDB support

KVM MMU

- MMU support
  - Add more operations for SEAM call
- Large(2M) page support
- TDP MMU support: WIP

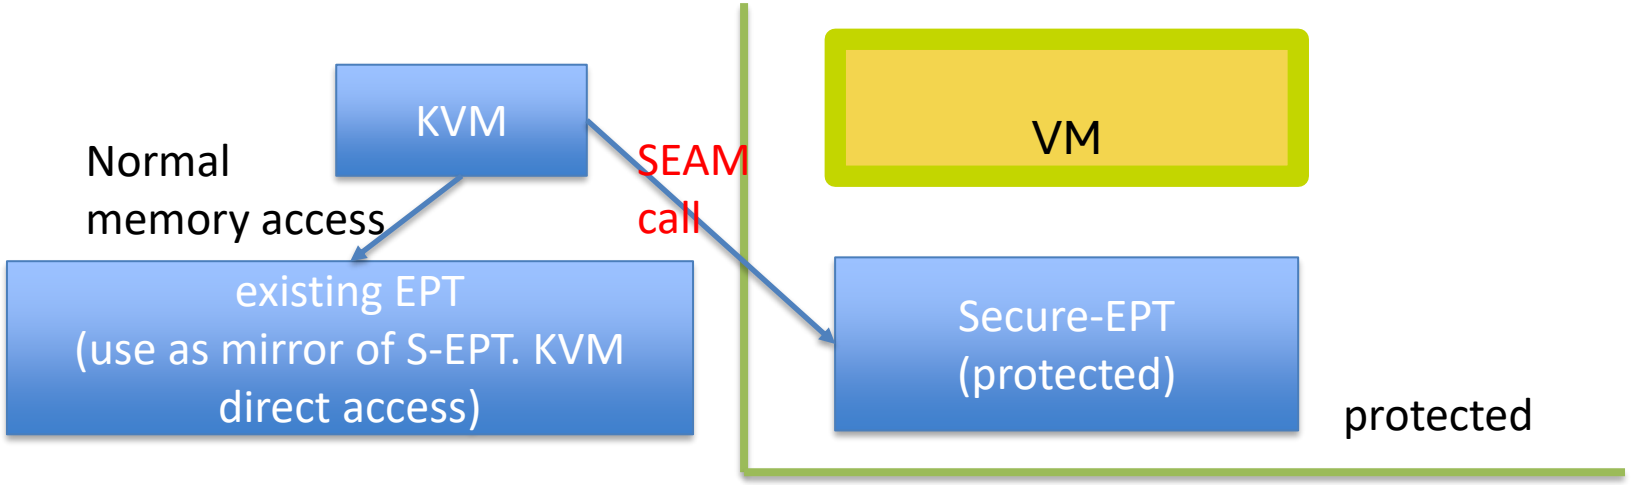

Unmapping private pages from user space

- There are two proposals
- Protection=NONE mapping with user space mapping
  - There is working PoC.
- Fd-based without any user mapping.
  - KVM MMU code needs change
  - Discussion in the community
- See other session.

Linux Guest

- #VE handling for hypercall
- Swiotlb for bounce buffer
- Device filtering
  - Don’t use untrusted devices

Qemu changes

- Loading guest bios
  - Special initialization is needed.
- Disabling features
  - Some features can’t be emulated. E.g. SMI, reboot
  - Feature advertise: CPUID, ACPI tables
- GDB support(after first upstream)

Related software upstream

- Guest UEFI bios(edk2)
- Grub2(boot loader)
- Libvirt and upper management system(openstack)

Status summary

Item	status
Linux firmware loader/initialization	Implementation WIP
KVM TD creatin/destruction	Under review
KVM MMU	Under review TDP: implementing WIP
KVM unmapping user space	Under discussion
Qemu guest BIOS loading	Under review
Qemu cpuid	WIP
Qemu guest attestation	WIP
Qemu unmapping user space	Under discussion

KVM FORUM