



# Exploring I/O Support for Virtualization-Based Trusted Execution Environment

Hao Wu hao.wu@intel.com

Sep. 2022 / KVM Forum 2022

# Disclaimers

Intel Corporation (“Intel”) provides these materials as-is, with no express or implied warranties.

All products, dates, and figures specified are preliminary, based on current expectations, and are subject to change without notice. Intel does not guarantee the availability of these interfaces in any future product. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

The products described might contain design defects or errors known as errata, which might cause the product to deviate from published specifications. Current, characterized errata are available on request.

Intel technologies might require enabled hardware, software, or service activation. Some results have been estimated or simulated. Your costs and results might vary.

No product or component can be absolutely secure.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted that includes the subject matter disclosed herein.

No license (express, implied, by estoppel, or otherwise) to any intellectual-property rights is granted by this document.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice.



# Agenda

- Background
- Current direct I/O support for TEE VM (TVM)
  - Working model and challenges
- TDISP support for TVM
  - PCIe* TEE Device Interface Security Protocol (TDISP) Overview
  - Intel® Trust Domain Extensions (TDX) with TEE-IO (TDISP) support
- Summary

# Background

- Virtualization techniques are used to provide an increased security guarantee for Trusted Execution Environment (TEE) such as TEE Virtual Machine (TVM)
- Confidential computing inside TVM requires I/O support
  - Assistances or accelerations provided by external devices
- Focus on direct IO support discussion in this presentation

# Current direct I/O solution for TVMs – Overview

- Devices are not allowed to read / write the TVM's confidential memory
- No protection for data inside shared memory which can be accessed by VMM
- Data path between Host and Device is not trusted
  - IOMMU is not in the TCB
  - Physical Link is not protected

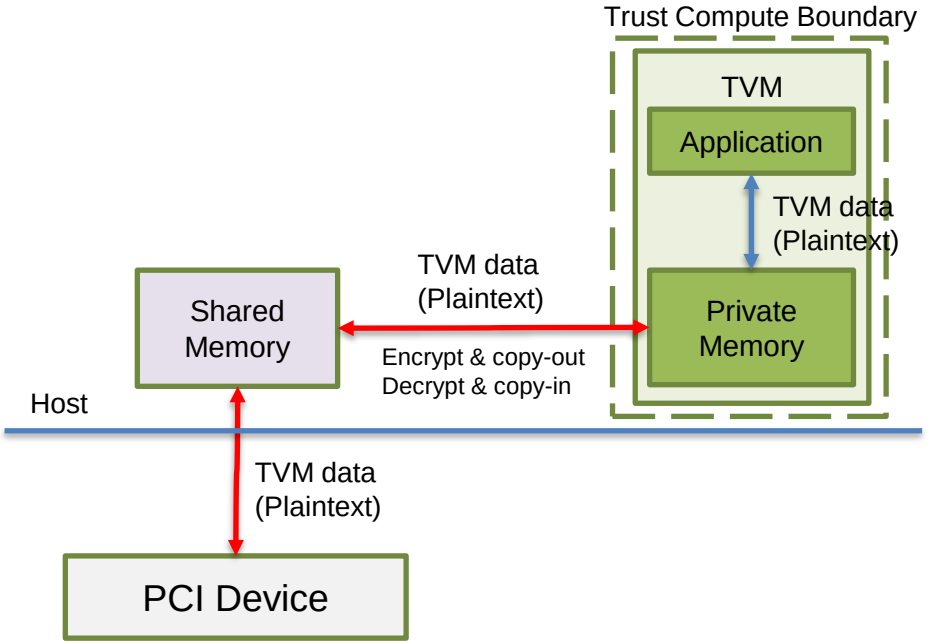

# Current direct I/O solution for TVMs – Working Model

- TVM Data can be consumed by either device (case 1) or peer (case 2)
- Secured data channel must be established to improve data confidentiality and integrity

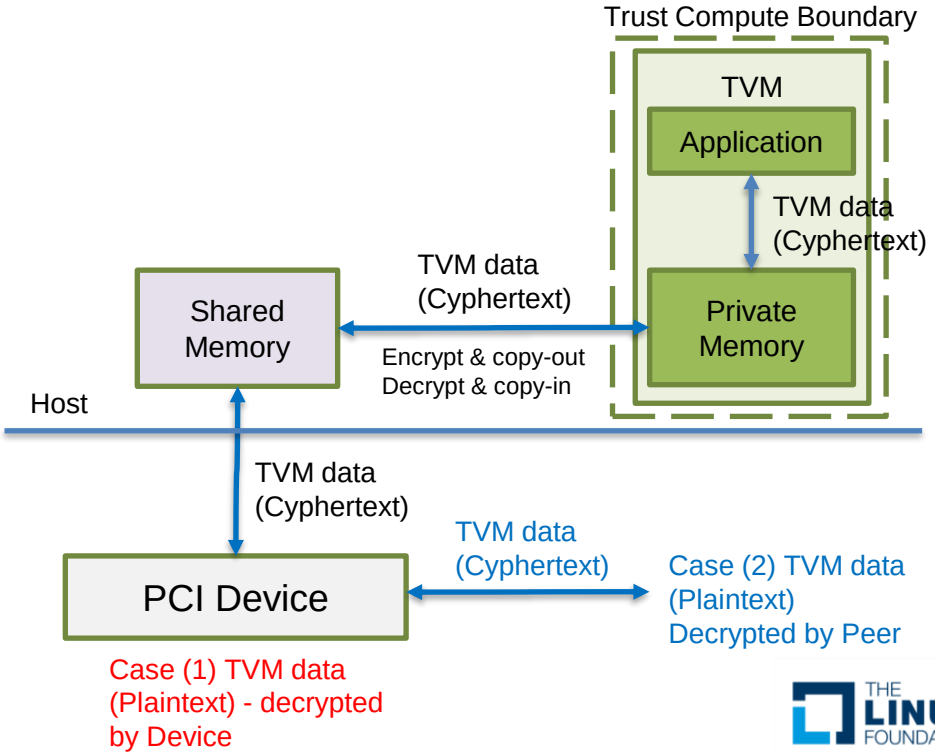

# Current direct I/O solution for TVMs – Challenges

- Additional cryptographic protections required for TVM data
- Performance overhead as extra steps needed for encrypt copy-out & decrypt copy-in to/from shared memory

# Include device into TVM's TCB?

- Is there any mechanism to allow TVM to include the target device into the TVM's TCB?

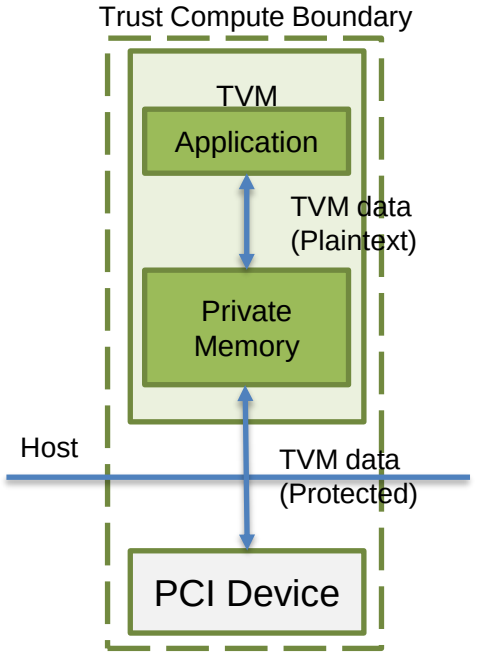

# TDISP Overview

- PCIe* TEE Device Interface Security Protocol (TDISP) defines an architecture of trusted I/O virtualization (TEE-IO)
  - Establishment a trust relationship between a TVM and a TDISP-compliant device
  - Help secure the data-path interconnect between the host and device
  - Support TDISP-compliant device assignment and removal life cycle in a trusted manner

# TDISP Overview

- TDISP builds upon the foundation provided by:
  - DMTF* Security Protocol and Data Model (SPDM)
  - PCIe* Component Measurement and Authentication (CMA)
  - PCIe* Integrity and Data Encryption (IDE)
  - PCIe* Data Object Exchange (DOE)

# TDISP – Architecture Overview

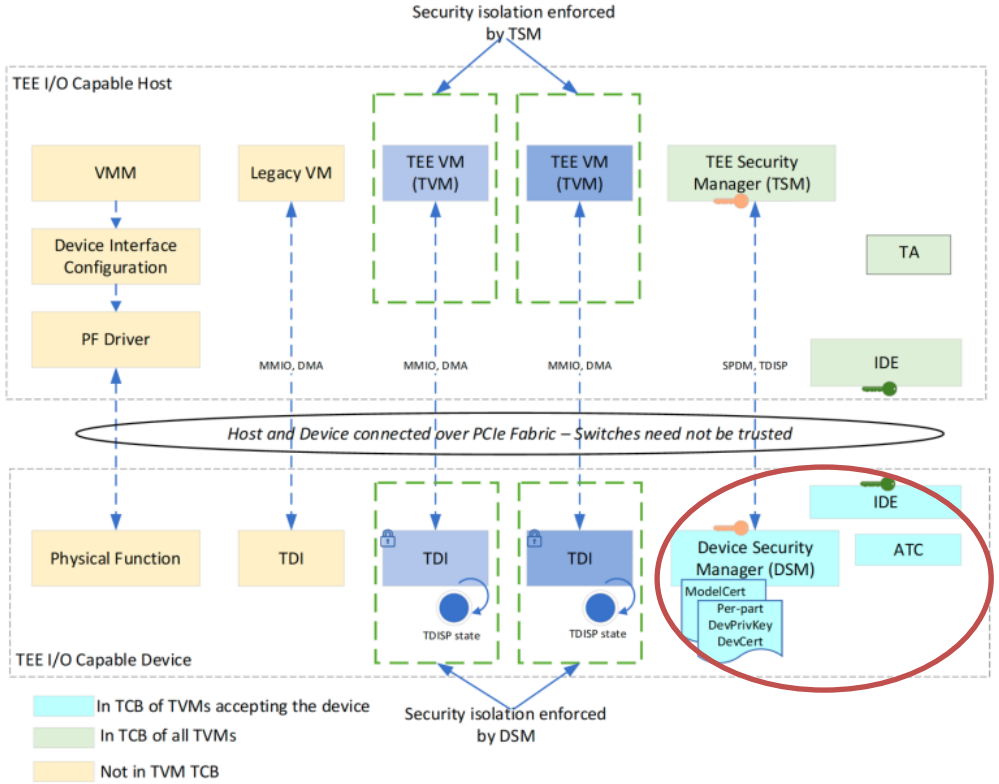


TDISP Host/Device Reference Architecture [From PCIe* TDISP spec]

- TDI: TEE Device Interface
  - Unit of device assignment
  - Can be an entire device, a non-IOV function, a PF or a VF
- DSM: Device Security Manager
  - Enforce security isolation for TDIs
  - Authentication of device identities and measurement reporting
  - IDE encryption keys configuration
  - TDI management (UNLOCKED, LOCKED, RUN and ERROR)
  - Access control
  - Security mechanisms to isolate TVM data

# TDISP – Architecture Overview

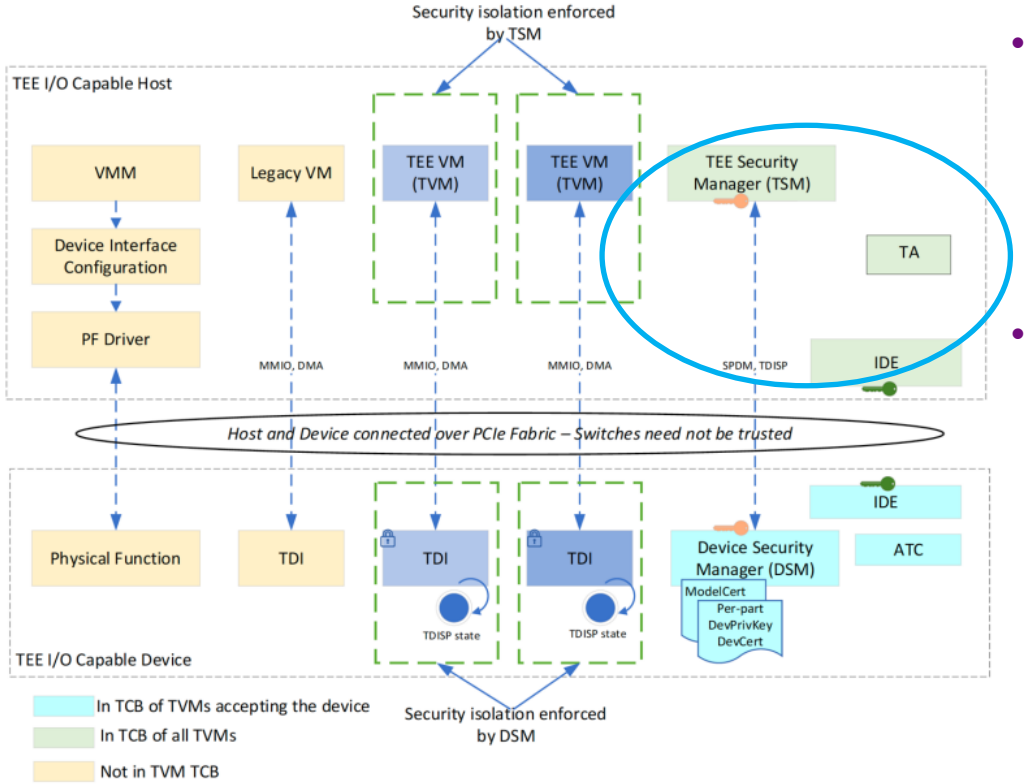


TDISP Host/Device Reference Architecture [From PCIe* TDISP spec]

- TSM: TEE Security Manager
  - Enforce security isolation for TVMs
  - Manage security states of TDIs
  - Security mechanisms and access controls
  - Establish and manage IDE Keys for the host
- Trusted MMIO/DMA
  - T bit in TLP IDE prefix
  - Used by device and host translation agent to provide access control

# Intel® TDX with TEE-IO support – Overview

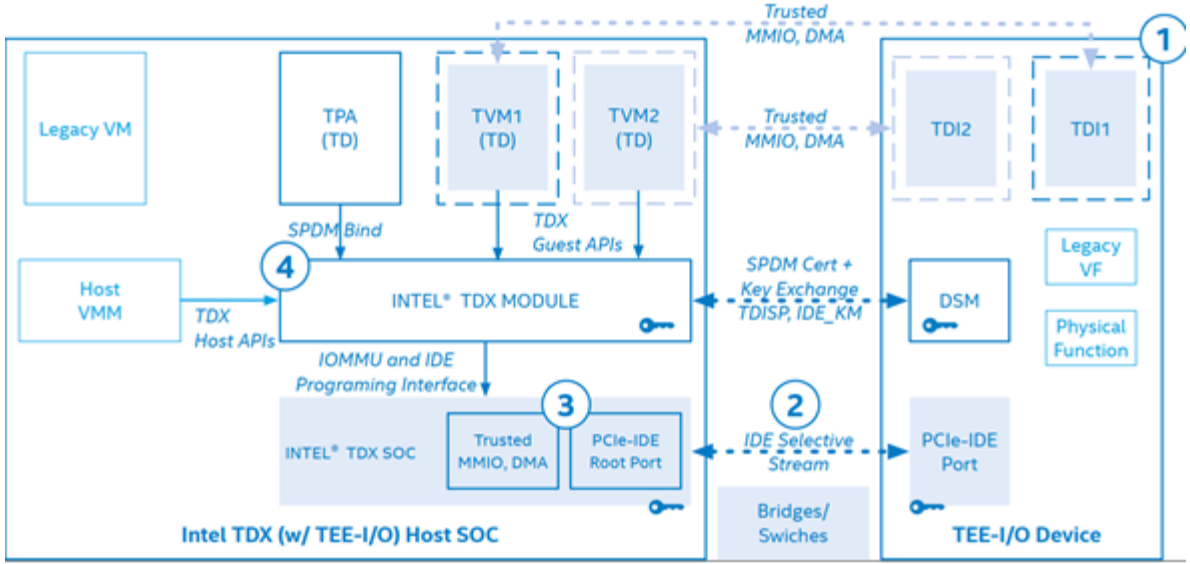


Intel® TDX with TEE-IO architecture [From: Software enabling for Intel® TDX in support of TEE-I/O]

1) TDISP-complaint device
   - Implement TDIs and DSM
2) IDE support
3) Trusted MMIO / DMA
   - Access control based on T bit
4) Intel® TDX Module + TPA
   - Function as TSM
   - TPA (TDX provision agent) is an architectural TD, helps TDX module with SPDM support

# Intel® TDX with TEE-IO – Software touchpoints

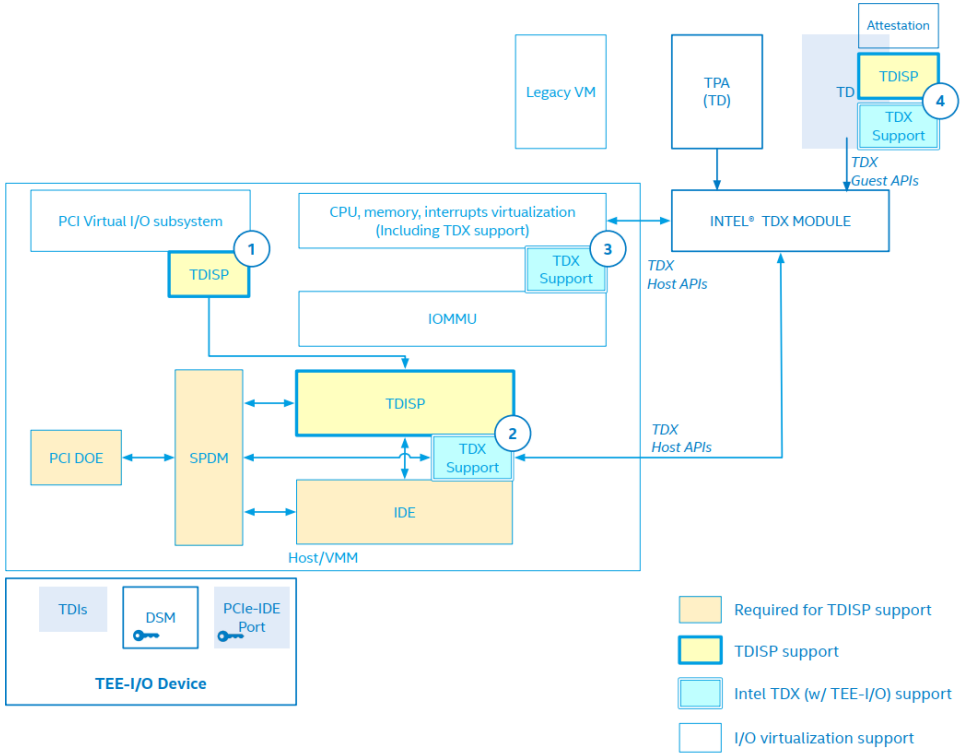


SW touchpoints [From: Software enabling for Intel® TDX in support of TEE-I/O]

1) VFIO: Expose TDI to TD
   - Identify TEE-IO capability of the device
   - Additional TDISP initialization / cleanup
2) PCI TDISP support
   - TDI state management
   - Request SPDM and IDE support for TDISP use case
   - Bind a TDI(s) to the target TVM
3) Trusted MMIO/DMA support
   - KVM: manage trusted MMIO via Secure EPT
   - IOMMU: New TDX mode for trusted address translation, and reuse Secure EPT as IO page table

# Intel® TDX with TEE-IO – Software touchpoints

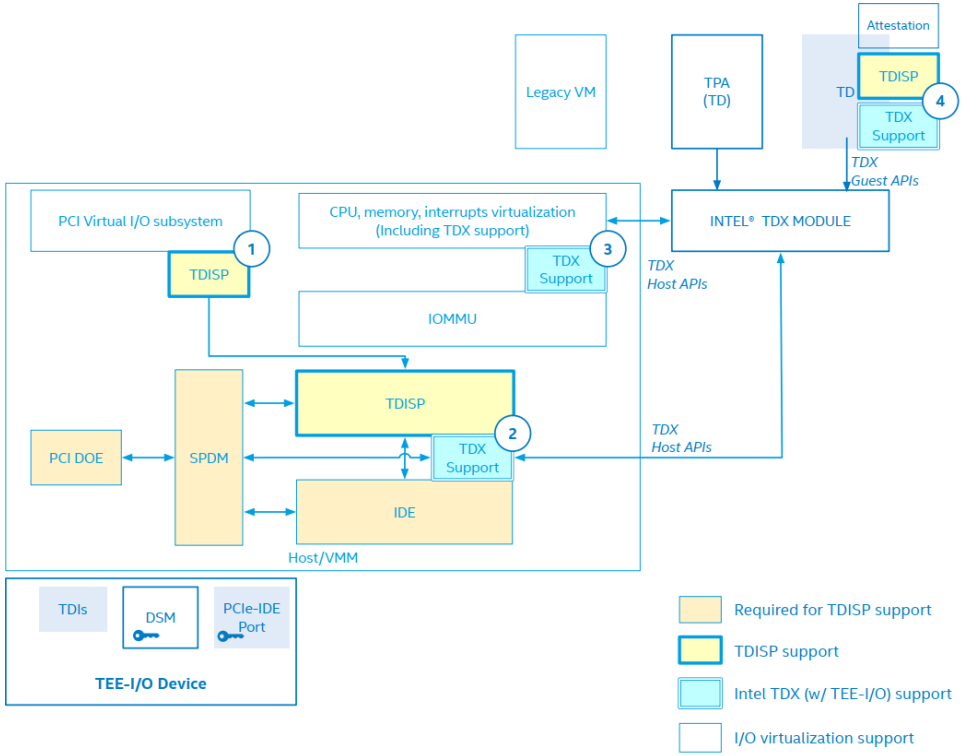


SW touchpoints [From: Software enabling for Intel® TDX in support of TEE-I/O]

4) PCI TDISP support (inside TD)

- TDI enumeration
- TDI attestation
- TDI acceptance
- Extensions to kernel APIs to support trust MMIO and DMA use case

# Summary

- IO support is important for confidential computing inside TVM. Current direct I/O solution has limitations and performance overhead as device can’t access TVM’s private memory.

- TDISP defines an architecture of trusted I/O virtualization. New architecture allows TDI to be accepted into the TVM’s TCB.

- Intel® TDX with TEE-IO is designed to implement the TDISP architecture. Besides platform and Intel® TDX module extensions, software changes to Linux / KVM are required, including common support for TDISP and specific implementation for Intel® TDX.

* Intel® TDX with TEE-IO is trying to reduce the performance overhead but actual performance results may vary.

# Reference

- PCIe* TDISP: https://members.pcisig.com/wg/PCI-SIG/document/18268
- PCIe* CMA: https://members.pcisig.com/wg/PCI-SIG/document/14236
- DMTF* SPDM: https://www.dmtf.org/dsp/DSP0274 v1.2+ & https://www.dmtf.org/dsp/DSP0277
- PCIe* IDE: https://members.pcisig.com/wg/PCI-SIG/document/16599
- PCIe* DOE: https://members.pcisig.com/wg/PCI-SIG/document/14143
- Intel® TDX: https://www.intel.com/content/www/us/en/developer/articles/technical/intel-trust-domain-extensions.html

# KVM FORUM