CXL ENGINEERING CHANGE NOTICE

TITLE:	Add support for HDM-DB targets with TSP
DATE:	Introduced date: 02/06/2024 Updated: 06/20/2024
AFFECTED DOCUMENT:	CXL 3.1
SPONSOR:	Chet Douglas, Intel Corporation

Part I

1. **Summary of Functional Changes**

   Add HDM-DB support for confidential computing using TSP

2. **Benefits as a Result of the Changes**

   HDM-DB targets can now make use of TSP with confidential computing

3. **Assessment of the Impact**

   This ECN will have an impact on target HW/SW, initiator HW/SW, RoT

4. **Analysis of the Hardware Implications**

   May effect HW

5. **Analysis of the Software Implications**

   May effect SW

6. **Analysis of the Compliance and Test Implications**

   May have test impacts

Part II

Detailed Description of the change

Add the following changes to 1.2 Terminology/Acronyms:

Term/Acronyms	Definition
DTRCS	Device Tracked Requestor Cache State – The requestor cache coherency state tracked by the device.
RCS	Requestor Cache State – The cache coherency state tracked by the host or initiator

Add the following changes to 3.3.5 M2S Request:

Table 3-35 M2S Req Memory Opcodes (Sheet 1 of 2)

Opcode	Description	Encoding
MemSpecRd	Memory Speculative Read is issued to start a memory access before the home agent has resolved coherence to reduce access latency. This command does not receive a completion message. The Tag, MetaField, MetaValue, and SnpType are reserved. See Section 3.5.3.1 for a description of the use case.	1000b
MemInvNT	This is similar to the MemInv command except that the NT is a hint that indicates the invalidation is non-temporal and the writeback is expected soon. However, this is a hint and not a guarantee. If the target is locked utilizing TSP, the target shall decode this opcode as MemInvP. If the target is not locked, the target shall decode this opcode as MemInvNT. See section 11.5 for TSP.	1001b
MemClnEvct	Memory Clean Evict is a message that is similar to MemInv, but intent to indicate host going to I-state and does not require Meta0-State return. This message is supported only to the HDM-DB address region.	1010b

Table 3-35 M2S Req Memory Opcodes (Sheet 2 of 2)

Opcode	Description	Encoding
MemSpecRdTEE	Same as MemSpecRd but with Trusted Execution Environment (TEE) attribute. See Section 11.5.4.5 for description of TEE attribute handling.	1100b
TEUpdate	Update of the TE state for the memory region. The memory region update is defined by the length-index field (passed in SnpType bits). The lower address bits in the message may be set to allow routing of the message to reach the correct interleave set target; however, the lower bits are masked to the natural alignment of the length when updating TE state. The MetaValue field defines the new TE state that supports 00b to clear and 01b to set. See details of the use of this message in Section 11.5.4.5.3.	1101b
MemInvTEE	Same as MemInv but with the Trusted Execution Environment (TEE) attribute. See Section 11.5.4.5 for description of TEE attribute handling.	0111b
MemInvP	Memory invalidation with precise TE State. If the target is locked utilizing TSP, the target shall decode this opcode as MemInvP. If the target is not locked, the target shall decode this opcode as MemInvNT. See section 11.5 for TSP.	1001b
MemInvPTEE	Same as MemInvP but with the Trusted Execution Environment (TEE) attribute. See Section 11.5.4.5 for description of TEE attribute handling.	1011b
MemClnEvctU	Same as MemClnEvct but TE State is not conveyed and assumed to be unknown.	1111b
MemClnEvctTEE	Same as MemClnEvct but with the Trusted Execution Environment (TEE) attribute. See Section 11.5.4.5 for description of TEE attribute handling.	1110b
~~Reserved~~	~~Reserved~~	~~<Others>~~

Add the following changes to 3.3.8 S2M Back-Invalidate Snoop (BISnp):

Table 3-47 S2M BISnp Opcodes

Opcode	Description	Encoding
BISnpCur	Device requesting Current copy of the line but not requiring caching state.	0000b
BISnpData	Device requesting Shared or Exclusive copy.	0001b
BISnpInv	Device requesting Exclusive Copy.	0010b
BISnpCurBlk	Same as BISnpCur except covering 2 or 4 cachelines that are naturally aligned and contiguous. The Block Enable encoding is in Address[7:6] and defined in Table 3-48. The host may give per cacheline response or a single block response applying to all cachelines in the block.More details are in Section 3.3.8.1.	0100b
BISnpDataBlk	Same as BISnpData except covering 2 or 4 cachelines that are naturally aligned and contiguous. The Block Enable encoding is in Address[7:6] and defined in Table 3-48. The host may give per cacheline response or a single block response applying to all cachelines in the block.More details are in Section 3.3.8.1.	0101b
BISnpInvBlk	Same as BISnpInv except covering 2 or 4 cachelines that are naturally aligned and contiguous. The Block Enable encoding is in Address[7:6] and defined in Table 3-48. The host may give per cacheline response or a single block response applying to all cachelines in the block.More details are in Section 3.3.8.1.	0110b
BISnpCurTEE	Same as BISnpCur but with the Trusted Execution Environment (TEE) attribute. See Section 11.5.4.5 for description of TEE attribute handling.	1000b
BISnpDataTEE	Same as BISnpData but with the Trusted Execution Environment (TEE) attribute. See Section 11.5.4.5 for description of TEE attribute handling.	1001b
BISnpInvTEE	Same as BISnpInv but with the Trusted Execution Environment (TEE) attribute. See Section 11.5.4.5 for description of TEE attribute handling.	1010b
BISnpCurBlkTEE	Same as BISnpCurBlk but with the Trusted Execution Environment (TEE) attribute. See Section 11.5.4.5 for description of TEE attribute handling.	1100b
BISnpDataBlkTEE	Same as BISnpDataBlk but with the Trusted Execution Environment (TEE) attribute. See Section 11.5.4.5 for description of TEE attribute handling.	1101b
BISnpInvBlkTEE	Same as BISnpInvBlk but with the Trusted Execution Environment (TEE) attribute. See Section 11.5.4.5 for description of TEE attribute handling.	1110b
Reserved	Reserved	<Others>

Add the following changes to 3.3.9 S2M No Data Response (NDR):

Table 3-50. S2M NDR Opcodes (Sheet 2 of 2)

BI-ConflictAck	Completion of the Back-Invalidate conflict handshake.	100b
CmpTEE	Completion for Writes (MemWr*) with TEE intent. Does not apply to any M2S Req.	101b
CmpTEE-S	Indication from the DCOH to the Host for Shared state with TEE intent.	110b
CmpTEE-E	Indication from the DCOH to the Host for Exclusive ownership with TEE intent.	111b
~~Reserved~~	~~Reserved~~	~~<Others>~~

Add the following changes to 11.5.2 Scope:

This CXL security content scope focuses on features that are needed for confidential computing utilizing CXL Type 3 memory expander devices, referred to as targets in the TSP, directly connected to CXL Root Ports owned by the host which is an initiator in TSP. TSP defines the security objectives, capabilities, and interfaces, and the host, initiator, and target behaviors that are required to create a secure CXL memory hierarchy that meets the needs of confidential computing. The scope does not include details on initiator or target security implementation.

- This scope includes support for the following:
    - SPDM 1.2 or newer for authentication and attestation
    - Directly connected LDs, SLDs, and MH-SLDs
    - Dynamic Capacity devices
    - HDM-H memory
    - ~~—~~HDM-DB memory
    - 256B and PBR flit formats
    - Memory pooling — Multiple initiators accessing the same physical memory on a device but not sharing access to it
    - Comprehensive Trust security model
    - Selective Trust security model

  - Implicit 64B cache line TE State Access Control
  - Explicit TE State Access Control
- This scope does not include the following:
  - CXL switches
    - Devices connected via a CXL switch, including MLDs, GFDs
    - Direct P2P using CXL.mem
    - Direct P2P using UIO over CXL.io
  - ~~UIO Direct Peer-to-peer (P2P) to HDM~~
  - Type 1 and Type 2 accesses to Type 3 HDM memory
  - HDM-D ~~or HDM-DB~~ memory
  - ~~PBR and~~ 68B flit format~~s~~
  - Memory sharing — Multiple initiators accessing the same physical memory on a device and simultaneously sharing access to it

Add the following changes to 11.5.3.1 Definitions:

The following additional terms are utilized in this threat model section:

…

- **Participant**: Initiator or target in a communication that utilizes a correct and error free implementation of the protocol.
- **Peer Device:** The peer device is an initiator that contains no CXL Root Ports.

Add the following changes to 11.5.3.2 Assumptions:

The threat model described below is based on the following assumptions:

…

- The target is directly connected to the host or peer device, either one acting as the initiator; thus, there are no attackers in the intermediaries in the ~~initial~~ TSP threat model. Targets connected via CXL switches have not been evaluated and the presence of switches are considered to be outside the threat model. Fabric-attached memory may require initiator-based memory encryption to keep the intermediaries out of the TCB and shall be addressed in a future version of the TSP.

Add the following changes to 11.5.4.1 Architectural Scope:

Figure 11-24 outlines the major components that the TSP considers to be inside the TCB or outside the TCB, the different connections between the TEE-capable initiator and TEE-capable target memory device, and those connections that are specified by the TSP. Hosts are the only initiators defined for ~~this~~ the original CXL 3.1 version of the TSP architecture for support of direct attached confidential computing in the TSP architecture. ~~Not all initiators are hosts because other CXL accelerators and PCIe P2P devices can function as initiators, the focus of future TSP additions.~~ With the addition of HDM-DB support to the TSP, CXL direct attached peer devices or accelerators are also considered initiators and may be utilized for confidential computing.

Replace the existing Figure 11-24 Reference Architecture with this one:

Figure 11-24. Reference Architecture

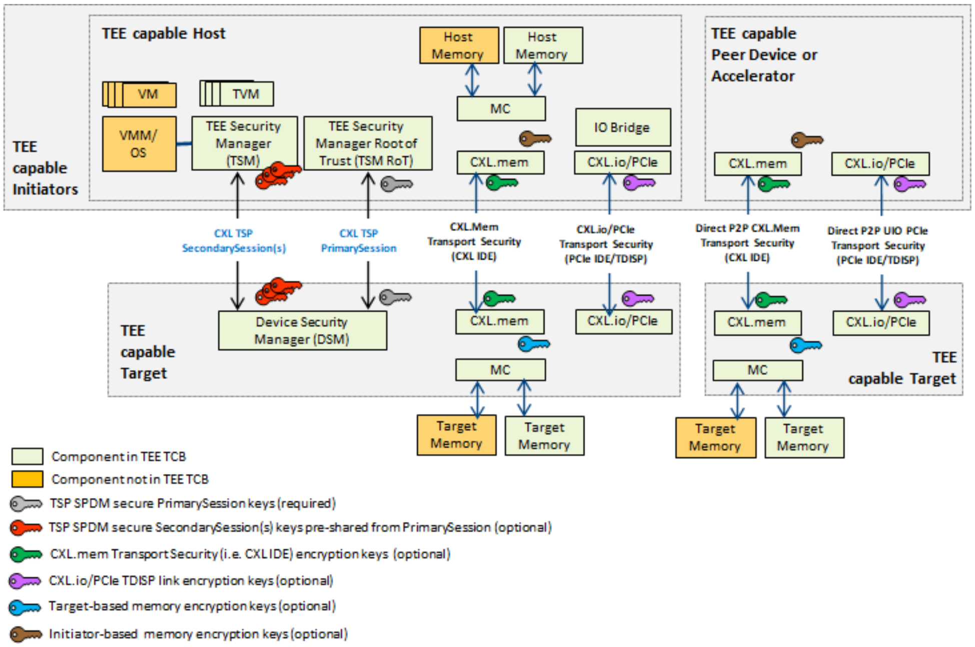
TEE capable Initiators
TEE capable Host
VM
TVM
VMM/ OS
TEE Security Manager (TSM)
TEE Security Manager Root of Trust (TSM RoT)
Host Memory
Host Memory
MC
IO Bridge
CXL.mem
CXL.io/PCIe
TEE capable Peer Device or Accelerator
CXL.mem
CXL.io/PCIe
CXL TSP SecondarySession(s)
CXL TSP PrimarySession
CXL.Mem Transport Security (CXL IDE)
CXL.io/PCIe Transport Security (PCIe IDE/TDISP)
Direct P2P CXL.Mem Transport Security (CXL IDE)
Direct P2P UIO PCIe Transport Security (PCIe IDE/TDISP)
TEE capable Target
Device Security Manager (DSM)
CXL.mem
CXL.io/PCIe
MC
Target Memory
Target Memory
CXL.mem
CXL.io/PCIe
MC
TEE capable Target
Target Memory
Target Memory
Component in TEE TCB
Component not in TEE TCB
TSP SPDM secure PrimarySession keys (required)
TSP SPDM secure SecondarySession(s) keys pre-shared from PrimarySession (optional)
CXL.mem Transport Security (i.e. CXL IDE) encryption keys (optional)
CXL.io/PCIe TDISP link encryption keys (optional)
Target-based memory encryption keys (optional)
Initiator-based memory encryption keys (optional)

Add the following new sections starting with 11.5.4.11 HDM-DB:

11.5.4.11 HDM-DB

The following sections describe the additional challenges with utilizing HDM-DB memory with the TSP and the resulting initiator and target requirements and behaviors needed for confidential computing with this type of memory.

With HDM-H memory the host is responsible for maintaining the cache coherency state of memory. With HDM-DB memory, the target owning the HDM-DB memory maintains the cache coherency state for the memory. The initiator and target utilize the BISnp and BIRsp channels to resolve coherency.

HDM-DB support in TSP enables the following with confidential computing:

- Target side compute

The target HDM Decoders shall be programmed before the target is locked through TSP. This allows the target to utilize the BI indicator in the programmed HDM decoders to determine if HDM-DB specific requirements and behaviors outlined here are to be utilized. It will also allow the host TEE architecture to ensure HDM-DB support is only enabled if it is capable of supporting such a configuration.

To correctly pass TEE Intent and TE State, additional request and response opcodes are required as outlined below. The new opcodes required for HMD-DB are defined in the M2S Req Memory Opcodes definitions, S2M BISnp Opcodes definitions, S2M NDR Opcodes definitions, and Appendix C.

For the TSP to operate correctly with the HDM-DB protocol, the following sub-sections outline additional requirements, initiator behaviors and target behaviors that define HDM-DB use with confidential computing. This includes:

- New initiator and target requirements for handling requestor cache state and TE State changes
- New M2S request opcodes to carry TEE Intent in support of HDM-DB.
    - MemInvTEE
    - MemInvP/MemInvPTEE – Memory invalidation requiring precise TE State
    - MemClnEvctU – Memory clean eviction with unknown TE State
    - MemClnEvctTEE
- New S2M BISnp opcodes to carry TEE Intent in support of HDM-DB.
    - BISnpCurTEE
    - BISnpDataTEE
    - BISnpInvTEE
    - BISnpCurBlkTEE
    - BISnpDataBlkTEE
    - BISnpInvBlkTEE
- New S2M NDR response opcodes to report TE State
    - MemInvP/MemInvPTEE returns Cmp, CmpTEE, CmpTEE-S, or CmpTEE-E

11.5.4.11.1 Determining TSP Support with HDM-DB

A target's support of HDM-DB memory is determined by looking at the BI bit in each HDM Decoder Control Register. HDM ranges with the BI flag set are enabled for HDM-DB. The target reports support for TSP with the TSP Capable bit in the DVSEC CXL Capability register.

Targets that report HDM-DB support and are TSP capable, shall support all of the request and response opcodes that are described here.

11.5.4.11.2 Requestor Coherency State (RCS)

The Requestor Coherency State (RCS) is the cache state maintained by the initiator. There are a variety of existing initiator implementations that handle RCS in fundamentally different ways. The following requirements take that into account and outline the expected initiator behaviors for maintaining RCS with HDM-DB and TSP, independent of implementation.

TSP behaviors for HDM-DB initiators:

- Initiators may update RCS without regard to the TE State. These initiators may utilize implicit and/or explicit TE State changes on the target.
- Initiators may update RCS by TE State. These initiators shall utilize explicit TE State changes on the target.
- When receiving a BISnp command, initiators may invalidate all RCS for a given address, regardless of whether the TE State specified in the BISnp command matches the TE State held in the RCS. However, initiators may safely retain RCS that does not match the TE State following a BISnp, as long as the initiator can guarantee that no internal or external entity can observe stale cache data.
- Initiators shall take additional actions (i.e. software-initiated cache flushes) to ensure RCS consistency on the target after a TE State mismatch when the target reports this requirement in the Additional Capabilities of Get Target Capabilities Response.

11.5.4.11.3 Device Tracked Requestor Coherency State (DTRCS)

The Device Tracked Requestor Coherency State (DTRCS) is the initiators cache state that is maintained by the target. There are a variety of existing target implementations that handle DTRCS in fundamentally different ways. The following requirements take that into account and outline the expected initiator and target behaviors for maintaining DTRCS with HDM-DB and TSP, independent of implementation.

TSP behaviors for HDM-DB initiators and targets:

- Targets that update DTRCS after a TE State mismatch shall require no special handling.
- Targets that do not update DTRCS after a TE State mismatch shall require one of the following target behaviors:
    - When the target receives a request on the M2S Req channel that results in a TE State mismatch, the target completes the request, then issues a BISnp with the current TE State being tracked for the address in the request, and blocks all new M2S Req requests to the address (including requests internal to the device) from the time the request causing the mismatch is processed until the BISnp completion is received, OR
    - The target requires the initiator to take additional actions after a TE State mismatch occurs.
        - The target indicates this dependency by setting Initiator Actions Required Following TE State Mismatch in Get Target Capabilities Response.
        - When this is indicated, the initiator is responsible for ensuring the coherency is maintained after the mismatch. Initiator responsibilities may include software-initiated target cache flushes or disallowing the mismatched line from allocating in the initiator's cache.
- Targets shall relax buried state rules to avoid unexpected state downgrade on MemRd, MemRdTEE, MemRdData, and MemRdDataTEE that result in a TE State mismatch as described in the buried state section that follows and the HDM-DB updates to Appendix C.

11.5.4.11.4 TE State Changes

TSP behaviors for HDM-DB targets:

- Targets shall support explicit and/or implicit TE State changes as specified in section 11.5.4.5.
- Targets shall snoop back all addresses affected by a TE State change using BISnp, before any memory contents or TE State is updated. While the snoop-back cycle is in progress:
    - The target shall block access to the affected memory (it is legal to block the request channel for short amounts of time without causing timeouts, but the RwD channel cannot be blocked without risk of deadlock), OR
    - The target shall handle the received transactions that address the same memory region that is undergoing the snoop back as a TE State mismatch and shall follow the mismatch behavior outlined in the TE State Changes and Access Control section 11.5.4.5 *and* the following subsections.

TSP behaviors for HDM-DB initiators:

- Initiators that retain the data following a BISnp that was requested with a TE State mismatch, shall utilize an explicit TE State change command.

11.5.4.11.5 BISnp S2M Requests with TE State

The BISnp requests are extended to encode a TE State. HDM-DB targets shall include TE State when sending BISnp. This is provided for initiators that may require accurate TE State to correctly resolve RCS for the target.

The TE State contained in the BISnp request shall match the current TE State tracked by the target for the address being snooped.

If the BISnp is occurring in response to an explicit TE update, then all the BISnp associated with the TE State update shall complete before the TE State is updated.

All HDM-DB capable targets utilizing TSP shall support reporting TE State with all BISnp request opcodes.

The following table outlines the required BISnp request opcodes the targets shall support:

S2M Request opcode	TEE State	Description
BISnpCur BISnpData BISnpInv BISnpCurBlk BISnpDataBlk BISnpInvBlk	0	Back invalidate the memory with current TE State 0.
BISnpCurTEE BISnpDataTEE BISnpInvTEE BISnpCurBlkTEE BISnpDataBlkTEE BISnpInvBlkTEE	1	Back invalidate the memory with current TE State 1.

11.5.4.11.6 MemRd M2S Requests with TEE Intent

MemRd requests shall include TEE Intent utilizing MemRd or MemRdTEE request opcodes. The intent is provided for targets that may require an accurate TE State to process the read request. The existing MemRd request is utilized for TE Intent = 0 and the new MemRdTEE request is utilized for TE Intent = 1.

All HDM-DB capable targets utilizing TSP shall support the MemRd/MemRdTEE request opcodes.

The following table outlines the required MemRd request opcodes the target shall support:

M2S Request opcode	TEE Intent	Target behavior
MemRd	0	Read memory with TEE Intent 0.
MemRdTEE	1	Read memory with TEE Intent 1.

11.5.4.11.7 MemRd S2M Responses with TE State

MemRd/MemRdTEE S2M DRS responses shall return TE State utilizing MemData or MemDataTEE. The target shall respond with the current TE State associated with the underlying data being read.

MemRd/MemRdTEE with MetaValue I is not supported when the target has been locked with TSP. If this request is received while TSP is enabled, the target shall respond with MemData with all 1's data, optionally return poison and no TE State shall be inferred by the initiator. This allows differentiation in behavior from a valid MemRd with MetaValue I that is received when TSP is not utilized.

All HDM-DB capable targets utilizing TSP shall support the MemData/MemDataTEE responses for MemRd/MemRdTEE request opcodes.

There are additional requirements for targets that maintain DTRCS that if a TE State mismatch is detected when executing the MemRd/MemRdTEE, the target shall not degrade the final DTRCS when handling the response. See Appendix C for special cases for not downgrading DTRCS on a TE State mismatch.

The following table outlines the valid MemRd S2M DRS response opcodes the target shall support when the current TE State matches the TEE Intent of the MemRd:

M2S Request opcode	Valid S2M DRS Response	Target behavior
MemRd	MemData	-Memory read with TE State = 0
MemRdTEE	MemDataTEE	-Memory read with TE State = 1

The following table outlines the valid MemRd S2M DRS response opcodes the target shall support when the current TE State does not match the TEE Intent of the MemRd:

M2S Request opcode	Valid S2M DRS Response	Target behavior
MemRd	MemDataTEE	-Memory read with TEE Intent = 0 resulted in a mismatch
MemRdTEE	MemData	-Memory read with TEE Intent =1 resulted in a mismatch

11.5.4.11.8 MemInv M2S Requests with TEE Intent

MemInv requests shall include TEE Intent utilizing MemInv or MemInvTEE request opcodes. TEE Intent is provided for targets that may require an accurate TE State in order to change the state of the cache line. The existing MemInv request is utilized for TEE Intent = 0 and the new MemInvTEE request is utilized for TEE Intent = 1. The TEE Intent shall indicate the intended TE State of memory following the DTRCS update.

The MemInv/MemInvTEE S2M NDR response does not convey TE State and shall not be utilized as an indicator of TE State. Initiators requiring precise TE State in the response shall utilize MemInvP/MemInvPTEE requests.

All HDM-DB capable targets utilizing TSP shall support the MemInv/MemInvTEE request opcodes.

The following table outlines the required MemInv request opcodes the target shall support:

M2S Request opcode	TEE Intent	Target behavior
MemInv	0	Invalidate the memory with TEE Intent 0. Initiator does not require TE State in the response as described below.
MemInvTEE	1	Invalidate the memory with TEE Intent 1. Initiator does not require TE State in the response as described below.

11.5.4.11.9 MemInvP M2S Requests with TEE Intent

MemInvP/MemInvPTEE are new MemInv request opcodes defined to indicate the TEE Intent of the invalidate and that the initiator requires a precise TE State to accompany the MemInvP/MemInvPTEE completion response. TEE Intent is provided for targets that may require an accurate TEE Intent in order to change the DTRCS of the cacheline. The TEE Intent shall indicate the intended TE State of memory following the DTRCS update.

Initiators that retain RCS following a BISnp shall utilize MemInvP/MemInvPTEE if knowledge of the TE State being invalidated is required for that initiators cache implementation.

Targets shall determine the current TE State of the memory being invalidated before responding to these requests.

All HDM-DB capable targets utilizing TSP shall support the MemInvP/MemInvPTEE request opcodes.

The following outlines the required MemInvP request opcodes the target shall support:

M2S Request opcode	TEE Intent	Target behavior
MemInvP	0	-Invalidate the memory with TEE Intent 0 -Report the precise TE State in the response.
MemInvPTEE	1	-Invalidate the memory with TEE Intent 1 -Report the precise TE State in the response

11.5.4.11.10 MemInv & MemInvP S2M Responses with TE State

MemInvP/MemInvPTEE S2M NDR responses shall return TE State utilizing Cmp or CmpTEE. The target shall respond with the current TE State associated with the underlying data being invalidated. This may require the responding target to look up the TE State prior to completing the MemInv request even though no data will be returned.

All HDM-DB capable targets utilizing TSP shall support the Cmp/CmpTEE responses for MemInvP/MemInvPTEE request opcodes.

The following table outlines the valid MemInv S2M NDR response opcodes the target shall support when the current TE State matches the TEE Intent of the MemInv:

M2S Request opcode	Valid S2M NDR Response	Target behavior
MemInv	Cmp Cmp-S Cmp-E	-Memory invalidated -Return no TE State in the response
MemInvTEE		
MemInvP	Cmp Cmp-S Cmp-E	-Memory invalidated -Return current TE State in the response
MemInvPTEE	CmpTEE CmpTEE-S CmpTEE-E	

The following table outlines the valid MemInv S2M NDR response opcodes the target shall support when the current TE State does not match the TEE Intent of the MemInv:

M2S Request opcode	Valid S2M NDR Response	Target behavior
MemInv	Cmp Cmp-S Cmp-E	-Memory invalidated (since precise TE State is not required there is no reason not to invalidate the memory for the mismatch case) -Return Cmp -Optionally log an event
MemInvTEE		
MemInvP	CmpTEE CmpTEE-S CmpTEE-E	-Do not invalidate the memory -Return current TE State -Optionally log an event
MemInvPTEE	Cmp Cmp-S Cmp-E	

11.5.4.11.11 MemRdData M2S Req Requests with TEE Intent

MemRdData requests shall include TEE Intent utilizing MemRdData or MemRdDataTEE request opcodes. The intent is provided for targets that may require an accurate TE State to process the read request. The existing MemRdData request is utilized for TE Intent = 0 and the new MemRdDataTEE request is utilized for TE Intent = 1.

All HDM-DB capable targets utilizing TSP shall support the MemRdData/MemRdDataTEE request opcodes.

The following table outlines the required MemRdData request opcodes the target shall support:

M2S Request opcode	TEE Intent	Target behavior
MemRdData	0	Read memory with TEE Intent 0.
MemRdDataTEE	1	Read memory with TEE Intent 1.

11.5.4.11.12 MemRdData S2M DRS Responses with TE State

MemRdData/MemRdDataTEE S2M DRS responses shall return TE State utilizing MemData or MemDataTEE. The target shall respond with the current TE State associated with the underlying data being read.

All HDM-DB capable targets utilizing TSP shall support the MemData/MemDataTEE responses for MemRdData/MemRdDataTEE request opcodes.

There are additional requirements for targets that maintain DTRCS that if a TE State mismatch is detected when executing the MemRdData/MemRdDataTEE, the target shall not degrade the final DTRCS when handling the response. See Appendix C for special cases for not downgrading DTRCS on a TE State mismatch.

The following table outlines the valid MemRdData S2M DRS response opcodes the target shall support when the current TE State matches the TEE Intent of the MemRdData:

M2S Request opcode	Valid S2M DRS Response	Target behavior
MemRdData	MemData	-Memory read with TE State = 0
MemRdDataTEE	MemDataTEE	-Memory read with TE State = 1

The following table outlines the valid MemRdData S2M DRS response opcodes the target shall support when the current TE State does not match the TEE Intent of the MemRdData:

M2S Request opcode	Valid S2M DRS Response	Target behavior
MemRdData	MemDataTEE	-Memory read with TEE Intent = 0 resulted in a mis-match
MemRdDataTEE	MemData	-Memory read with TEE Intent = 1 resulted in a mismatch

11.5.4.11.13 MemSpecRd M2S Req Requests with TEE Intent

MemSpecRd requests shall include TEE Intent utilizing MemSpecRd or MemSpecRdTEE request opcodes. The intent is provided for targets that may require an accurate TE State to process the speculative read request. The existing MemSpecRd request is utilized for TEE Intent = 0 and the new MemSpecRdTEE request is utilized for TEE Intent = 1.

All HDM-DB capable targets utilizing TSP shall support the MemSpecRd/MemSpecRdTEE request opcodes.

The following table outlines the required MemSpecRd request opcodes the target shall support:

M2S Request opcode	TEE Intent	Target behavior
MemSpecRd	0	Speculatively read memory with TEE Intent 0.
MemSpecRdTEE	1	Speculatively read memory with TEE Intent 1.

11.5.4.11.14 MemClnEvct M2S Req Requests without TEE Intent

MemClnEvctU is a new memory request opcode that may be utilized by initiators that don't know the TE State of the memory being clean evicted. The MemClnEvctU M2S req request does not convey TE State and shall not be utilized as an indicator of TE State.

Initiators should avoid MemClnEvctU and should utilize MemClnEvct or MemClnEvctTEE whenever possible for best performance. Initiators that utilize MemClnEvctU shall not track TE State when maintaining RCS.

HDM-DB targets that require an accurate TE State in order to process eviction requests and receive MemClnEvctU may evict utilizing current TE State, may evict both TE States, or may not evict anything. Initiators that require specific target behavior should utilize MemClnEvct or MemClnEvctTEE.

Targets should take extra measures to find and clean the DTRCS associated with the eviction request since failure to complete a clean eviction may result in extra BISnp requests, potentially impacting system performance.

All HDM-DB capable targets utilizing TSP shall support the MemClnEvctU.

11.5.4.11.15 MemClnEvct M2S Req Requests with TEE Intent

MemClnEvct requests shall include TE Intent utilizing MemClnEvct or MemClnEvctTEE request opcodes. TEE Intent is provided for targets that may require the TE State in order to process the eviction request and reset the state of the cacheline. The MemClnEvct request is utilized for TEE Intent 0 and the request MemClnEvctTEE is utilized for TEE Intent 1.

All HDM-DB capable targets utilizing TSP shall support the MemClnEvct and MemClnEvctTEE request opcodes.

The following table outlines the required MemClnEvct request opcodes that target shall support:

M2S Request opcode	TEE Intent	Target behavior
MemClnEvctU	N/A	Perform clean evict independent of TE State
MemClnEvct	0	Perform clean evict using TEE Intent 0
MemClnEvctTEE	1	Perform clean evict using TEE Intent 1

11.5.4.11.16 MemClnEvct S2M NDR Responses with TE State

Since MemClnEvctU/MemClnEvct/MemClnEvctTEE are provided for performance and not correctness, none of these requests require TE State to be reported in the response.

The following table outlines the valid MemClnEvct S2M NDR response opcodes the target shall support when the current TE State matches or mismatches the TEE Intent of the MemClnEvct:

M2S Request	Valid S2M NDR Response	Target behavior
MemClnEvctU	Cmp	The current state of the memory evicted is unknown
MemClnEvct		The current state of the memory evicted is TE State = 0
MemClnEvctTEE		The current state of the memory evicted is TE State = 1

11.5.4.11.17 Buried State Behavior

For targets that maintain DTRCS and support TE State tracking, if the target detects a TE State mismatch when the initiator is requesting S state, shall not downgrade the final DTRCS. MemRd, MemRdTEE and MemRdData, MemRdDataTEE shall not downgrade DTRCS for a TE State mismatch and is outlined in Appendix C.

Targets that don't update DTRCS after a TE State mismatch and rely on additional host actions to correct RCS may leave the final device cache and/or DTRCS unchanged after the mismatch occurs, relying on software actions to correct any coherency issues. See the "UCM" cases in the Device Cache and DTRCS columns of Appendix C.

Add the following changes to 11.5.5.1 TSP Request Overview:

Table 11-26 outlines the TSP Request payloads, defined in the sections that follow.

Table 11-26. TSP Request Overview

TSP Request Message		Message Support[1]	Payload Size	Legal TSP State	TSP Usage
Opcode	Name	HDM-H HDM-DB Devices			

Add the following changes to 11.5.5.2 TSP Response Overview:

Table 11-27 outlines the TSP Response payloads, defined in the sections that follow.

Table 11-27. TSP Response Overview

TSP Response Message		Message Support[1]	Payload Size
Opcode	Name	HDM-H HDM-DB Devices	

Add the following changes to 11.5.5.5.2 Get Target Capabilities Response:

Table 11-32. Get Target Capabilities Response (Sheet 1 of 2)

0Ch	2	**TE State Change and Access Control Features Supported**: The TE State change and access control features that the target supports. 0 or more bits may be set. 1 indicates supported, 0 indicates not supported. • Bit[0]: **Write Access Control**: When set indicates that the target supports dropping writes that fail the verification of TEE Intent to stored TE State. When set, explicit state changes shall be supported and one or more of Bits[4:3] shall also be set. • Bit[1]: **Read Access Control**: When set indicates that the target supports returning all 1's for read data in response to reads that fail the verification of TEE Intent to stored TE State. When set, one or more of Bits[4:2] shall also be set. • Bit[2]: **Implicit TE State Change**: When set indicates that the target supports implicit TE State changes using a 64B granularity, Explicit In-band TE State Change shall be set, and Explicit In-band TE State Granularity support for 64B shall be set. • Bit[3]: **Explicit Out-of-band TE State Change**: When set, indicates that the target supports the CMA/SPDM out-of-band explicit Set Target TE State change message and the Supported Explicit Out-of-band TE State Granularity field shall be valid. Support is optional for targets that support implicit TE State changes or explicit in-band TE State changes. • Bit[4]: **Explicit In-band TE State Change**: When set, indicates that the target supports explicit TE State changes utilizing the TEUpdate memory transaction and the Supported Explicit In-band TE State Granularity field shall be valid. Support is required for targets that support implicit TE State changes and optional for targets that support explicit out-of-band TE State changes. • Bit[5]: **Explicit TE State Change Sanitize**: When set, indicates that the target supports overwriting data that is affected by the explicit state change with 0s when the explicit request is received and before the change is considered complete by the target. When set, one or more of Bits[4:3] shall also be set. • Bits[15:6]: **Reserved**.
0Eh	1	**Additional Capabilities:** Other security related features and capabilities of the target. • Bit[0]: **Initiator Actions Required Following TE State Mismatch**: When set, indicates that the HDM-DB capable target will require initiator actions (i.e. software-initiated cache flushes) to ensure correct DTRCS is maintained on the target following a TE State mismatch. When clear, the target does not require additional initiator actions to maintain DTRCS following a TE State mismatch. This bit is only valid if the target reports Device Coherent for Supported Coherency Models in the HDM Decoder Capability Register and BI is supported in the HDM Decoder Control Register. • Bits[7:1]: **Reserved**.
~~0Eh~~0Fh	~~2~~1	**Reserved**

C.1.1 Appendix C Updates for HDM-DB w TSP

Add the following changes to Appendix C Memory Protocol Tables:

(Update 'Host State' column name to 'Device Tracked Requestor Coherency State' or 'DTRCS' in these tables)

C.1.1 ~~HDM-D and~~ HDM-DB Requests with TEE Support

Table C-3 defines messages on the request channel of CXL.mem protocol. Table C-5 additionally defines the Forward flows that apply only for HDM-D memory regions for Type 2 devices when the devices are accessing device-attached memory. Table C-6 defines the BISnp channel method of managing device-attached memory coherence for the HDM-DB memory region of Type 2 devices or Type 3 devices.

New footnote added to disallow downgrade of "host state" from current state if there is a TE mismatch detected at the target for read access. Example is if current state is A and normal resulting state is S-state then you may not downgrade to S-state if TE mis-match.

A new term is added of Un-Changed Mismatch ("UCM") to reflect cases where the device cache or requester state in the HDM-DB target is not changed in the case of a TEE mis-match. Targets that take advantage of this allowance will rely on software clean-up of coherence violation that may result. For additional details related to requirements for this behavior see Section X.X.X.

Table C-4. ~~HDM-D/~~HDM-DB Memory Requests with TE state (Sheet 1 of 6)

Legal	Host Request				Device Response					Final Device State			Description
	M2S Req	MetaField	MetaValue	SnpType	S2M NDR	S2M DRS	MetaField	MetaValue	DRS Trailer	Device Cache	Host State	EMD	
O-1	MemRd/MemRdTEE <Add description of target allowance for check or not for these opcodes as part of coherence resolution>	MS0	A	SnpInv	Cmp-M[1]	MemData/MemDataTEE	MS0-NO	<any>	N/A	I or UCM	A or UCM	UC	Supported for host that can accept M-state.
O-1 & O-2		EMS					EMS-NO		EMV-NA				
Y(1)		MS0			Cmp-E		MS0-NO		N/A				The host wants an exclusive copy of the cacheline.
O-2		EMS					EMS-NO		EMV-NA				
N				SnpData									
N				SnpCur									
N				No-Op									
N			S	SnpInv									
Y(1)		MS0		SnpData	Cmp-S		MS0-NO	<any>	N/A	S or UCM	S[2] or UCM	UC	The host is requesting a shared copy of the cacheline, but Rsp types allow the device to return S-state or E-state to the host. Cmp-E response is not recommended because the host did not request this state.
O-2		EMS					EMS-NO		EMV-NA				
Y(1)		MS0			Cmp-E		MS0-NO		N/A	I or UCM	A or UCM		
O-2		EMS					EMS-NO		EMV-NA				
N				SnpCur									
N				No-Op									

Table C-4. ~~HDM-D/~~HDM-DB Memory Requests with TE state (Sheet 2 of 6)

Legal	Host Request				Device Response					Final Device State			Description
	M2S Req	MetaField	MetaValue	SnpType	S2M NDR	S2M DRS	MetaField	MetaValue	DRS Trailer	Device Cache	Host State	EMD	
Y	MemRd <Only for Non-TSP targets>	MS0	I	SnpInv	Cmp	MemData	MS0-NO	<any>	N/A	I	I	UC	The host is requesting a non-cacheable but current value of the cacheline and is forcing the device to flush its cache.
O-2		EMS					EMS-NO		EMV-NA				
N				SnpData									
Y		MS0		SnpCur	Cmp		MS0-NO	<any>	N/A	<any>	I	UC	The host is requesting a non-cacheable but current value of the cacheline and is thereby leaving data in the device's cache.
O-2		EMS					EMS-NO		EMV-NA				
N				No-Op									
Y(1)	MemRd/MemRdTEE	No-Op	N/A	SnpInv	Cmp	MemData/MemDataTEE	MS0-NO	<any>	N/A	I or UCM	UC	UC	The host wants to read the cacheline without changing the state expected in the host cache, and the device should invalidate the cacheline from its cache.
O-2		EMS	No-Op				EMS-NO		EMV-NA				
N		EMS-NO	N/A	SnpData									
N		EMS	E-No-Op										
Y(1)		No-Op	N/A	SnpCur	Cmp		MS0-NO	<any>	N/A	<any>	UC	UC	The host wants a current value of the cacheline without changing the state expected in the host cache.
O-2		EMS	E-No-Op				EMS-NO		EMV-NA				
Y		No-Op	N/A	No-Op	Cmp		MS0-NO	<any>	N/A	<any>	UC		The host wants the value of the memory location without snooping the device cache and without changing the cache state expected in the host cache. A use case for this would be if the host includes E-state or S-state without data so that the host is requesting data only and doesn't want to change the cache state, and because the host has E-state or S-state, the host can know that the device cache does not need to be snooped.
O-2		EMS	E-No-Op				EMS-NO		EMV-NA				
Y		<all>	<all>	<all>	<none>	MemData-NXM	No-Op	N/A	N/A	N/A	N/A	N/A	The special case MemData-NXM response is used if the fabric or the device is unable to positively decode the Address. This is a common response type for both Type 2 devices and Type 3 devices to avoid an ambiguous case in which the host is unsure of whether the host should expect an NDR message.

Table C-4. ~~HDM-D/~~HDM-DB Memory Requests with TE state (Sheet 3 of 6)

Legal	Host Request				Device Response					Final Device State			Description
	M2S Req	MetaField	MetaValue	SnpType	S2M NDR	S2M DRS	MetaField	MetaValue	DRS Trailer	Device Cache	Host State	EMD	
Y(1)	MemInv~~/~~ ~~MemInvNT~~ MemInvTEE <Add description of target allowance for check or not for these opcodes>	MS0	A	SnpInv	Cmp-E	<none>	<any>	<any>	N/A	I or UCM	A or UCM	UC	The host wants ownership of the cacheline but does not require the data.
N				SnpData									
N				SnpCur									
N				No-Op									
N			S	SnpInv									
Y				SnpData	Cmp-S		<any>	<any>		S or I or UCM	S or UCM	UC	The host wants the device to degrade to S-state in its caches, and wants the shared state for the cacheline (but does not require the data).
N				SnpCur									
N				No-Op									
Y(1)			I	SnpInv	Cmp		<any>	<any>		I or UCM	I	UC	The host wants the device to invalidate the cacheline from its caches and does not require the data.
N				SnpData									
N				SnpCur									
N				No-Op									
Y		No-Op	N/A	SnpInv	Cmp		<any>	<any>		I or UCM	UC	UC	The host wants the device to invalidate the cacheline from its caches and does not require the data.
N				SnpData									
N				SnpCur									
N				No-Op									
N		EMS	N/A										

Table C-4. ~~HDM-D/~~HDM-DB Memory Requests with TE state (Sheet 4 of 6)

Legal	Host Request				Device Response					Final Device State			Description
	M2S Req	MetaField	MetaValue	SnpType	S2M NDR	S2M DRS	MetaField	MetaValue	DRS Trailer	Device Cache	Host State	EMD	
Y(1)	MemInvP/ MemInvPTEE~~/~~ ~~MemInvNT~~ <Target Required to lookup TE state>	MS0	A	SnpInv	Cmp-E/ CmpTEE-E	<none>	<any>	<any>	N/A	I or UCM	A or UCM	UC	The host wants ownership of the cacheline but does not require the data.
N				SnpData									
N				SnpCur									
N				No-Op									
N			S	SnpInv									
Y				SnpData	Cmp-S/ CmpTEE-S		<any>	<any>		S or I or UCM	S or UCM	UC	The host wants the device to degrade to S-state in its caches, and wants the shared state for the cacheline (but does not require the data).
N				SnpCur									
N				No-Op									
Y(1)			I	SnpInv	Cmp/ CmpTEE		<any>	<any>		I or UCM	I	UC	The host wants the device to invalidate the cacheline from its caches and does not require the data.
N				SnpData									
N				SnpCur									
N				No-Op									
Y		No-Op	N/A	SnpInv	Cmp/ CmpTEE		<any>	<any>		I or UCM	UC	UC	The host wants the device to invalidate the cacheline from its caches and does not require the data.
N				SnpData									
N				SnpCur									
N				No-Op									
N		EMS	N/A										

Table C-4. ~~HDM-D/~~HDM-DB Memory Requests with TE state (Sheet 5 of 6)

Legal	Host Request				Device Response					Final Device State			Description
	M2S Req	MetaField	MetaValue	SnpType	S2M NDR	S2M DRS	MetaField	MetaValue	DRS Trailer	Device Cache	Host State	EMD	
N	MemRdData/ MemRdDataTEE	<all>	N/A	SnpInv		MemData/ MemDataTEE							
Y(1)		MS0-NO		SnpData	Cmp-E		MS0	I or A	N/A	I or UCM	A or UCM	UC	The host wants a cacheable copy in either E-state or S-state.
O-2		EMS					EMS		EMV				
Y(1)		MS0-NO			Cmp-S		MS0	S	N/A	I or S or UCM	S[2] or UCM		
O-2		EMS					EMS		EMV				
Y		<all>			Cmp-E		No-Op	N/A	N/A	I or UCM	A or UCM		
Y					Cmp-S					I or S or UCM	S[2] or UCM		
N				SnpCur									
N				No-Op									
Y				<all>	<none>	MemData-NXM	No-Op	N/A	N/A	N/A	UC	UC	The special case MemData-NXM response is used if the fabric or the device is unable to positively decode the Address. This is a common response type for both Type 2 devices and Type 3 devices to avoid an ambiguous case in which the host is unsure of whether the host should expect an NDR message.
N	MemSpecRd/ MemSpecRdTEE	MS0-EMD	<all>	<all>									
N		No-Op	N/A	Snp*									
O-3				No-Op	<none>	<none>	<none>	<none>	N/A	UC	UC	UC	Speculative memory read. A Demand read following this with the same address will be merged in the device. No completion is expected for this transaction. Completion is returned with demand read.
N	MemClnEvctU <Requester does not know TE state>	MS0	A or S	<all>									
Y			I	No-Op	Cmp	<none>	No-Op	N/A	N/A	UC	I	UC	The host will only issue from E-state or from S-state. Target should make best estimate of TE state if required for coherence resolution.
N				Snp*									
N		EMD-NO	N/A	<all>									

Table C-4. ~~HDM-D/~~HDM-DB Memory Requests with TE state (Sheet 6 of 6)

Legal	Host Request				Device Response					Final Device State			Description
	M2S Req	MetaField	MetaValue	SnpType	S2M NDR	S2M DRS	MetaField	MetaValue	DRS Trailer	Device Cache	Host State	EMD	
N	MemClnEvct/ MemClnEvctTEE	MS0	A or S	<all>									
Y			I	No-Op	Cmp	<none>	No-Op	N/A	N/A	UC	I	UC	The host will only issue from E-state or from S-state. Requester should use these commands if TE state is known.
N				Snp*									
N		EMD-NO	N/A	<all>									
N	~~MemRdTEE/~~ ~~MemRdDataTEE/~~ ~~MemSpecRdTEE~~												~~TEE is not supported for HDM-DB/HDM-D in this revision of the specification.~~
~~Sub-Table~~[3]N	MemRdFwd	~~See Table C-5.~~											Not Supported (only used with HDM-D).
	MemWrFwd												

1. Cmp-M response is dis-allowed from target if TE mismatch is detected.
2. Target must not downgrade to S-state on TE-mismatch if current state tracked is A-state.
3. ~~Applicable only to HDM-D memory regions.~~