CXL ENGINEERING CHANGE NOTICE

TITLE:	Metabits Storage Feature for HDM-H address region (UUID 3568da82-e69c-4518-95a2-446fe34ea865)
DATE:	Introduced: 05/22/2023 Updated: 07/10/2024
AFFECTED DOCUMENT:	CXL 3.1
SPONSOR:	Sanjay Goyal, Rambus Inc. Mahesh Natu, Intel Corporation Scott Lee, Microsoft Corporation Daniele Balluchi, Micron Technology Chet Douglas, Intel Corporation Lynn Jin, Montage Technology

Part I

1. **Summary of Functional Changes**
   This ECN defines a mechanism for the host to discover and control the support for storage of MetaValue bits for Meta0-State, As defined in “CXL 3.1 spec Table 3-36. Metadata Field Definition and TE State bit in the CXL device’s HDM-H address region. TE State bit stores TE State as defined in Section 11.5 “CXL Trusted Execution Environments Security Protocol (TSP)” of CXL Specification Revision 3.1, Version 1.0. This ECN is applicable only when TE State granularity is 64B. It is not applicable when TE State granularity is bigger than 64B. It is not applicable to CXL device’s HDM-DB address region. The intent of this ECN is about the capability related to storage of these bits. How these bits are updated is outside the scope of this ECN. This ECN does not cover Extended Meta-State (EMS) as “defined in “CXL 3.1 spec Table 3-36”

   Added an updated to TSP section to disable certain features if no TE State storage is configured using this feature.

2. **Benefits as a Result of the Changes**
   Device may be able to optimize its behavior, based on Metadata Value bits storage and TE State bit storage requirements from the host.

3. **Assessment of the Impact**
   This is an optional feature. It will have no impact on existing implementations.

4. **Analysis of the Hardware Implications**
   This is an optional normative feature. There will be no HW impact in the Host.
   For the device that supports this optional feature need to implement it as described in this document.

5. **Analysis of the Software Implications**
   This is an optional normative feature. To take advantage of this feature, software will be required to discover and configure the bits as described in this document.

6. **Analysis of the Compliance and Test Implications**
   This ECN does not repurpose any reserved bits and as such, does not impact existing C&I tests.

Part II

Detailed Description of the change

3.3 CXL.mem

3.3.1 Introduction

…

Support for memory with Metadata is optional but this needs to be negotiated with the Host in advance. If the device supports "Metabits Storage" Feature, this mechanism may be used to negotiate the Metadata configuration. Other~~The~~ negotiation mechanisms are beyond the scope of this specification.

8.2.9.6 Features

…

8.2.9.6.x.x Metabits Storage Feature Discovery and Configuration

The Feature Identifier of this feature is: 3568da82-e69c-4518-95a2-446fe34ea865.

This feature allows the host to discover and configure the support for storage of Metadata Value bits and TE State in the CXL device's HDM-H address region. It is not applicable to HDM-DB address region. This Feature is not applicable when TE State granularity is bigger than 64B.

Table 8-xx shows the information returned in the Get Supported Features output payload for the **Metabits Storage** Feature. Some feature attributes are changeable.

"Below Text added to the ECN on July 10, 2024"

Any changes to HDM-H Metabits Storage Configuration require a Conventional reset to take effect. Saved across Reset bit in Set Feature Input Payload shall be set to 1, otherwise the device shall return Invalid Input. Changes to HDM-H Metabits Storage Configuration may result in changes to the device capacity and CDAT.

An SH-MLD, MH-MLD or MH-SLD that support this feature shall report Set Feature Size=0 and Bit[0] of Attribute Flags Bit[0] = 0, over CCI exposed to individual hosts indicating that the Feature Data cannot be modified over these CCI.

An SH-MLD, MH-MLD or MH-SLD that support this feature shall report Set Feature Size=1 and Bit[0] of Attribute Flags Bit[0] = 1, over CCI exposed to the FM indicating that the Feature Data can be modified over these CCI.

After a successful CXL reset, a Conventional Reset or a successful Secure Erase operation, a subsequent read to any device cacheline (DPA) shall return Metafield=00b (Meta0-State abbreviation MS0) and MetaValue=00b, if the device is configured with non-zero Metadata bits via this Feature. As per section 12.2.3, a device must set the MetaField to No-Op in the CXL.cachemem return response when the Metadata is suspect.

Table 8-xx Supported Feature Entry for Metabits Storage Feature

Byte offset	Length in Bytes	Attribute	Description
00h	10h	Feature Identifier	3568da82-e69c-4518-95a2-446fe34ea865
10h	2	Feature Index	Device Specific
12h	2	Get Feature Size	3 Bytes
14h	2	Set Feature Size	1 Bytes
16h	4	Attribute Flags	• Bit[0]: Vendor-specific value (Changeable) • Bits[3:1]: 010b (Deepest Reset Persistence = Hot Reset). Conventional reset will restore the saved value. • Bit[4]: 1 (Persist across Firmware Update) • Bit[5]: 1 (Default Selection Supported) • Bit[6]: 1 (Saved Selection Supported) • Bits[31:7]: Reserved
1Ah	1	Get Feature Version	01h
1Bh	1	Set Feature Version	01h

Byte offset	Length in Bytes	Attribute	Description
1Ch	2	• Set Feature Effects	• Bit[0]: 1 (Configuration Change after Cold Reset) • Bit[1]: 0 (Immediate Configuration Change) • Bit[2]: 0 (Immediate Data Change) • Bit[3]: 0 (Immediate Policy Change) • Bit[4]: Vendor-specific value (Immediate Log Change) • Bit[5]: Vendor-specific value (Security State Change) • Bit[6]: 0 (Background Operation) • Bit[7]: Vendor-specific value (Secondary Mailbox Supported) • Bit[8]: 0 (Request Abort Background Operation Supported) • Bit[9]: 1 (CEL[11:10] Valid) • Bit[10]: 1 (Configuration Change after Conventional Reset) • Bit[11]: 0 (Configuration Change after CXL Reset) • Bits[15:12]: 0h
1Eh	18	Reserved	

Table 8-xx shows the output payload returned by a Get Feature command with Selection set to 0h (Current value), 1h (Default value) or 2h (Saved Value).

Table 8-xx Metabits Storage Feature Readable Attributes

Byte offset	Length in Bytes	Description
00h	2	**HDM-H Metabits Storage Capabilities:** • Bit[0]: **2 bits of Metadata** are supported, As defined for Meta0-State in CXL3.1 spec Table 3-36. Metadata Field Definition. 2 bits of storage supported. • Bit[1]: **No Metadata** is supported. No storage supported. • Bit[2]: **1-bit of Metadata is** supported. bit-0 of Meta0-State Value will be stored. One bit of storage supported. • Bit[3]: **1-bit of Metadata is** supported. bit-1 of Meta0-State Value will be stored. One bit of storage supported. • Bit[4]: **2 bits of Metadata + 1 TE State bit** are supported. Three bits of storage supported. • Bit[5]: **No Metadata + 1 TE State bit** is supported. One bit of storage supported. • Bit[6]: **1-bit of Metadata + 1 TE State bit** are supported. bit-0 of Meta0-State Value will be stored. Two bits of storage supported. • Bit[7]: **1-bit of Metadata + 1 TE State bit** are supported. bit-1 of Meta0-State Value will be stored. Two bits of storage supported. • Bits[15:8] - Reserved
02h	1	• Bits[2:0]: **HDM-H Metabits Storage Configuration**. Default value is Vendor Specific. - 0h: 2 bits of Metadata - 1h: No Metadata - 2h: 1 bit of Metadata, bit-0 of Meta0-State Value - 3h: 1 bit of Metadata, bit-1 of Meta0-State Value - 4h: 2 bits of Metadata + 1 TE State bit - 5h: No Metadata + 1 TE State bit - 6h: 1 bit of Metadata, bit-0 of Meta0-State Value + 1 TE State bit - 7h: 1 bit of Metadata, bit-1 of Meta0-State Value + 1 TE State bit • Bits[7:3]: Reserved

Table 8-xx shows the input payload for Set Feature command.

Table 8-xx Metabits Storage Feature Writable Attributes

Byte offset	Length in Bytes	Description
0h	1	• Bits[2:0]: Values are as defined for **HDM-H Metabits Storage Configuration** field in Table 8-xx Metabits Storage Feature Readable Attributes • Bits[7:3]: Reserved

8.2.9.9.4.3 Clear Poison (Opcode 4302h)

…

Change from,

This command must not modify the content of the Extended Metadata field associated

with this address.

To,

This command must not modify the content of the Extended Metadata field associated

with this address. If the device is configured with non-zero Metadata bits as defined by HDM-H Metabits Storage Configuration field in Table 8-xx Metabits Storage Feature Readable Attributes, for subsequent read to the DPA, the device shall return Metafield=00b (Meta0-State abbreviation MS0) and MetaValue=00b.

14.16.3.17 Inject Memory Device Poison

Table 14-81. Memory Device Media Poison Injection Request

Change from,

Clear Poison Write Data: When Protocol = 2 and Action = 1, the device

shall write this replacement data into the requested physical address,

atomically, while clearing poison.

To,

Clear Poison Write Data: When Protocol = 2 and Action = 1, the device

shall write this replacement data into the requested physical address,

atomically, while clearing poison. If the device is configured with non-zero Metadata bits as defined by HDM-H Metabits Storage Configuration field in Table 8-xx Metabits Storage Feature Readable Attributes, for subsequent read to the DPA, the device shall return Metafield=00b (Meta0-State abbreviation MS0) and MetaValue=00b.

9.18.1.6 CXL System Description Structure (CSDS)

Table 9-25. CSDS Structure

Field	Length in Bytes	Byte Offset	Description
Type	1	00h	4 = Indicates that this is a CSDS entry
Reserved	1	01h	Reserved
Record Length	2	02h	Length of this record = 08h
System Capabilities	2	04h	A bitmap that describes system-wide capabilities. More than one bit within this field is permitted to be set. • Bit[0]: Cmp-M: — 1 = System is configured for use with devices that return modified data using the Cmp-M response. • Bit[1]: No Clean Writeback: Specifies the clean writeback behavior of the host. — 0 = The host may or may not generate clean writebacks — 1 = The host guarantees to never generate clean writeback transactions at the host's cacheline granularity • Bits[5:3]: Metabits Storage Configuration. Upon hot-add, the OS may configure the device to match host metadata storage requirements - 0h: 2 bits of Metadata - 1h: No Metadata - 2h: 1 bit of Metadata, bit-0 of Meta0-State Value - 3h: 1 bit of Metadata, bit-1 of Meta0-State Value - 4h: 2 bits of Metadata + 1 TE State bit - 5h: No Metadata + 1 TE State bit - 6h: 1 bit of Metadata, bit-0 of Meta0-State Value + 1 TE State bit - 7h: 1 bit of Metadata, bit-1 of Meta0-State Value + 1 TE State bit • Bits[15:6~~2~~]: Reserved

Add the following changes to 11.5.4.5 TE State Changes and Access Control:

The granularity utilized for TE State changes shall be consistent with the interleave granularity being configured. For example, if the host utilizes a 4K TE State change granularity on each target that is part of a 16-way interleave set with a 256B interleave granularity, each target will utilize 256B of DPA space to change 4K of TE State.

If the target was configured with no TE State storage in the device, by utilizing Set Features with Metabits Storage, then it is assumed the target has no TE State tracking capabilities and the target shall disable the following in Get Target Capabilities Response:

- Implicit TE State Change

- Explicit In-band TE State Change when TE State Granularity is set to 64B
- Explicit Out-of-band TE State Change when TE State Granularity is set to 64B