



CXL ENGINEERING CHANGE NOTICE

TITLE:	Common Event Record ECN
DATE:	Introduced: 2/22/2024 Updated date :5/17/2024
AFFECTED DOCUMENT:	CXL 3.1 V1.0, Aug7, 2022
SPONSORS:	Ankit Patel – Intel Corporation, Ariel Sibley - Microchip Technology Inc.

Part I

1. Summary of Functional Changes

This ECN focuses on changes needed in Common Event Record format to support more meaningful event transmission from CXL MLD devices to FM.

This ECN focuses on changes needed in Common Event Record format to support more meaningful event transmission from CXL Multi Headed Devices to FM.

Changes proposed in this ECN.

1. Update API definition for Common Event Record Format, Table 8-43

2. Benefits as a Result of the Changes

This ECN assures that MLD devices can report CXL events to FM with specific LD association. This ECN assures that Multi Headed CXL devices can report CXL events to FM with specific head association.

3. Analysis of the Hardware Implications

This ECN may have an impact on CXL MLD component hardware depending on the component micro-architecture.

4. Analysis of the Software Implications

Changes to MLD component firmware and CXL FM are needed to take advantage of capabilities brought by this ECN.

5. Analysis of the Compliance and Test Implications

To be updated based on CWG feedback.

Part II

Detailed Description of the change

In an MLD device, any LD can generate a CXL event. While decoding CXL event from an MLD device, FM has no way to figure out the exact source of the event within the CXL Device. In case of General Media Event or DRAM Event, FM must know LD-ID (the exact source) of the event, as DPA range are not unique among LDs of the MLD device. In cases of other events (for example Memory Module Event) information about the exact source could be optional or not required as the event may be global affecting all LDs.

Furthermore, for CXL event originating from of Multi-Headed MLDs (MH-MLDs), FM must know LD-ID and Head-ID(the exact source) of the event, as DPA range are not unique among LDs and not unique among heads of the MH-MLD device. In cases of other events (for example Memory Module Event) information about the exact source could be optional or not required as the event may be global affecting all LDs across all heads.

To fix above gaps, existing Common Event Record Format needs to be improved with addition of following information.

- Valid Flags
 - LD ID field valid flag
 - Head ID field valid flag
- LD ID field
- Head ID field

Red content shown in the following sections are changes originating from this ECN.

- Add following definition to table 1-1

Term/Acronym	Definition
SH-SLD	Single Headed Single Logical Device CXL component that contains a single CXL port, presenting an SLD.
SH-MLD	Single headed MLD. CXL component that contains a single CXL port, presenting an MLD.

- Updates in API definition for Common Event Record Format, Table 8-43

Table 8-43. Common Event Record Format

Byte Offset	Length in Bytes	Description
00h	10h	Event Record Identifier: UUID representing the specific Event Record format. The following UUIDs are defined in this spec: <ul style="list-style-type: none"> • fbcd0a77-c260-417f-85a9-088b1621eba6 – General Media Event Record(see Table 8-45)

		<ul style="list-style-type: none"> • 601dccb3-9c06-4eab-b8af-4e9bfb5c9624 – DRAM Event Record (see Table 8-46) • fe927475-dd59-4339-a586-79bab113b774 – Memory Module Event Record(see Table 8-47) • e71f3a40-2d29-4092-8a39-4d1c966c7c65 - Memory Sparring Event Record(see Table 8-48) • 77cf9271-9c02-470b-9fe4-bc7b75f2da97 – Physical Switch Event Record(see Table 7-77) • 40d26425-3396-4c4d-a5da-3d47263af425 – Virtual Switch Event Record(see Table 7-78) • 8dc44363-0c96-4710-b7bf-04bb99534c3f – MLD Port Event Record (see Table 7-79) • ca95afa7-f183-4018-8c2f-95268e101a2a - Dynamic Capacity Event Record(see Table 8-50)
10h	1	<p>Event Record Length: Number of valid bytes that are in the event record, including all fields.</p>
11h	3	<p>Event Record Flags: Multiple bits may be set.</p> <ul style="list-style-type: none"> • Bits[1:0]: Event Record Severity: The severity of the event.— 00b = Informational Event — 01b = Warning Event — 10b = Failure Event — 11b = Fatal Event • Bit[2]: Permanent Condition: The event reported represents a permanent condition for the device. This shall not be set when reporting Event RecordSeverity of Informational. • Bit[3]: Maintenance Needed: The device requires maintenance. This shall not be set when reporting Event Record Severity of Informational. • Bit[4]: Performance Degraded: The device is no longer operating at optimal performance. This shall not be set when reporting Event RecordSeverity of Informational. • Bit[5]: Hardware Replacement Needed: The device should immediately be replaced. This shall not be set when reporting Event Record Severity of Informational. If this bit is set and the Component Identifier field in the event record is valid, the hardware to be replaced is the hardware identified by the Component Identifier field. If this bit is set and the Component Identifier field in the event record is invalid or not part of the event record, the hardware to be replaced is the entire device. • Bit[6]: Maintenance Operation Subclass Valid Flag: If set, the Maintenance Operation Subclass is valid. This bit applies only to CXL devices. • Bits[23:7]: Reserved • Bits[7]: LD-ID Valid Flag. If set to 1, the LD-ID field is valid and this event is applicable to only the specified LD. If set to 0, and the Head ID Valid Flag is also 0, the event is applicable to the entire device. If set to 0, and the Head ID Valid Flag is 1, the event is applicable the entire device head.

		<ul style="list-style-type: none"> • Bits[8]: Head ID Valid Flag. If set to 1, the Head ID field is valid. MH-SLDs and MH-MLDs shall set this value to 1 if setting the LD-ID Valid Flag to 1. • Bits[23:9]: Reserved
14h	2	Event Record Handle : The event log unique handle for this event record. This is the value that the host shall use when requesting the device to clear events using the Clear Event Records command. This value shall be nonzero.
16h	2	Related Event Record Handle : Optional event record handle to another related event in the same event log. If there are no related events, this field shall be cleared to 0000h.
18h	8	Event Record Timestamp : The time the device recorded the event. The number of unsigned nanoseconds that have elapsed since midnight, 01-Jan- 1970, UTC. If the device does not have a valid timestamp, return all 0s.
20h	1	Maintenance Operation Class : This field indicates the maintenance operation the device requests to initiate. If the device does not have any requests, this field shall be cleared to 00h. This field applies only to CXL devices. See Table 8-110 .
21h	1	Maintenance Operation Subclass : This field indicates the maintenance operation subclass that the device recommends to initiate. If the device does not have a specific recommendation for the subclass, the Maintenance Operation Subclass Valid flag shall be 0. This field applies only to CXL devices. See Table 8-110 .
22h	0Eh	Reserved
22h	2	LD-ID : LD ID of LD from where the event originated. Only valid if LD-ID valid flag is set to 1.
24h	1	Head ID : ID of the Device head, from where the event originated. Only valid if head valid flag is set to 1.
25h	0Bh	Reserved
30h	50h	Event Record Data : Format depends on the Event Record Identifier

When reporting events using per-LD level Primary Mailbox or Secondary Mailbox CCI instances, the event consumer considers that LD as an independent device regardless of whether the LD is an SH-SLD, or one of multiple LDs in an SH-MLD, or LD is from an MH-SLD, or LD is from one of multiple LDs in an MH-MLD. As such, it is recommended that the device clear LD-ID Valid Flag to 0 and Head ID Valid Flag to 0, thereby not reporting any LD-ID or Head ID details.

When reporting events using CCI instances which are potentially common to multiple LDs (e.g. MCTP based CCI instances), the event consumer will be aware of all LDs discoverable through that CCI instance. As such, it is recommended that the device follows the below guidelines for

usage of the LD-ID Valid Flag, Head ID Valid Flag, LD-ID, and Head ID fields to inform the event consumer about the origin and scope of the reported events.

For SH-SLDs:

The device should choose to clear the LD-ID Valid Flag and Head ID Valid Flag to '0' indicating the LD-ID and Head ID fields are invalid. The event consumer would have already discovered that the device is an SH-SLD as part of the device discovery, initialization, and configuration process; hence, the event consumer already has head and LD association.

For SH-MLDs:

The device should choose to keep the Head ID Valid Flag set to '0' indicating the Head ID field is invalid, as the event consumer would have already discovered that the device is a Single Headed Device as part of the device discovery, initialization, and configuration process; hence, the event consumer already has head association.

If the given event impacts only one LD out of all LDs then the LD-ID Valid Flag should be set to '1' and LD-ID field should contain the LD ID of the impacted LD. If the given event impacts more than one LDs, the device should send multiple events (one for every impacted LD) with the LD-ID Valid Flag set to '1' and the LD-ID field indicating the impacted LD; additionally using "Related Event Record Handle" and "Event Record Timestamp" fields to link all related events together. Separate event records should be generated if the event applies to a specific physical memory location that is shared across multiple LDs, and that physical memory location is represented using different DPAs on those LDs. If the device is impacted by an event which impacts globally to all LDs, the device should choose to send the event with the LD-ID Valid Flag set to '0' indicating the LD-ID field is invalid and the event impacts the entire device.

For MH-SLDs and MH-MLDs:

If the given event impacts only one head out of all heads the Head-ID Valid Flag should be set to '1' and the Head ID field should contain the Head ID of the impacted head. If the given event impacts more than one heads, the device should send multiple events (one for every impacted head) with the Head ID Valid Flag set to '1' and the Head ID field indicating the impacted head; additionally using "Related Event Record Handle" and "Event Record Timestamp" fields to link all related events together. Separate event records should be generated if the event applies to a specific physical memory location that is shared across multiple LDs across multiple Heads, and that physical memory location is represented using different DPAs on those LDs across multiple Heads. If the device is impacted by an event which impacts globally to all heads, the device should choose to send the event with the Head ID Valid Flag set to '0' indicating the Head ID field is invalid and the event impacts the entire device.

Additionally, if the given event impacts only one LD out of all LDs on a given head, the LD-ID Valid Flag should be set to '1' and the LD-ID field should contain the LD-ID of the impacted LD. If the given event impacts more than one LDs on a given head, the device should send multiple events (one for every impacted LD) with the LD-ID Valid Flag set to '1' and the LD-ID field indicating the impacted LD; additionally using "Related Event Record Handle" and "Event Record Timestamp" fields to link all related events together. Separate event records should be generated if the event applies to a specific physical memory location that is shared across multiple LDs on given head, and that physical memory location is represented using different DPAs on those LDs on given head. If the device is impacted by an event which impacts globally

to all LDs on a given head, the device should choose to send the event with LD-ID Valid Flag set to '0' indicating the LD-ID field is invalid and the event impacts all LDs on the given head.

Please note that in case of MH-SLDs and MH-MLDs, if the device wants to report an event while setting the LD-ID Valid Flag to 1, the device must set the Head ID Valid Flag to '1' and provide a valid Head ID.