JEDEC STANDARD

JEDEC® Memory Device Management Standard – for Compute Express Link® (CXL®)

JESD325

September 2024

JEDEC SOLID STATE TECHNOLOGY ASSOCIATION

NOTICE

JEDEC standards and publications contain material that has been prepared, reviewed, and approved through the JEDEC Board of Directors level and subsequently reviewed and approved by the JEDEC legal counsel.

JEDEC standards and publications are designed to serve the public interest through eliminating misunderstandings between manufacturers and purchasers, facilitating interchangeability and improvement of products, and assisting the purchaser in selecting and obtaining with minimum delay the proper product for use by those other than JEDEC members, whether the standard is to be used either domestically or internationally.

JEDEC standards and publications are adopted without regard to whether or not their adoption may involve patents or articles, materials, or processes. By such action JEDEC does not assume any liability to any patent owner, nor does it assume any obligation whatever to parties adopting the JEDEC standards or publications.

The information included in JEDEC standards and publications represents a sound approach to product specification and application, principally from the solid state device manufacturer viewpoint. Within the JEDEC organization there are procedures whereby a JEDEC standard or publication may be further processed and ultimately become an ANSI standard.

No claims to be in conformance with this standard may be made unless all requirements stated in the standard are met.

Inquiries, comments, and suggestions relative to the content of this JEDEC standard or publication should be addressed to JEDEC at the address below, or refer to www.jedec.org under Standards and Documents for alternative contact information.

Published by
©JEDEC Solid State Technology Association 2024
3103 10th Street North
Suite 240S
Arlington, VA 22201

PRICE: Contact JEDEC

Printed in the U.S.A.

DO NOT VIOLATE

THE

LAW!

For information, contact:

JEDEC Solid State Technology Association
3103 10th Street North
Suite 240S
Arlington, VA 22201
https://www.jedec.org/contact

This page intentionally left blank

JEDEC MEMORY DEVICE MANAGEMENT STANDARD – FOR COMPUTE EXPRESS LINK (CXL)

Contents

JEDEC MEMORY DEVICE MANAGEMENT STANDARD – FOR COMPUTE EXPRESS LINK (CXL)

Contents (cont'd)

JEDEC MEMORY DEVICE MANAGEMENT STANDARD – FOR COMPUTE EXPRESS LINK (CXL)

Contents (cont'd)

Figures

Page

JEDEC MEMORY DEVICE MANAGEMENT STANDARD – FOR COMPUTE EXPRESS LINK (CXL)

Contents (cont'd)

Tables

Page

JEDEC® Memory Device Management Standard - for Compute Express Link® (CXL®)

(From JEDEC Board Ballot JCB-24-32, formulated under the cognizance of the JC-40.7 subcommittee on Memory Support Logic for CXL (item number 701.12A) and JC-45.7 subcommittee on Memory Modules for CXL (item number 2286.98B)).

1 Scope

This standard provides a reference specification for systems and device management capabilities found in CXL memory devices. It is intended to target, but may not be limited to, CXL memory FRUs that are based on PCIe Gen 5 and compliant to the CXL 2.0 Specification or later.

2 Normative References

NOTE: Information on how these references apply is found in the place where they are cited and not in the Normative References clause.

- "Compute Express Link® 2.0 Specification." CXL Consortium, Nov 2020
- "PCI Express Base Specification Revision 5.0, Version 1.0." PCI-SIG, May 2019
- SNIA SFF-TA-1009 Enterprise and Datacenter Standard Pin and Signal Specification, Revision 3.1
- DMTF Management Component Transport Protocol (MCTP) Base Specification, Revision 1.3.1 (DSP0236)
- DMTF MCTP over PCIe VDM – Management Component Transport Protocol (MCTP) PCIe VDM Transport Binding Specification, Revision 1.2.0 (DSP0238)
- DMTF MCTP over SMBus/I2C – DMTF Management Component Transport Protocol (MCTP) SMBus/I2C VDM Transport Binding Specification, Revision 1.2.0 (DSP0237)
- DMTF Management Component Transport Protocol (MCTP) I3C Transport Binding Specification, Revision 1.0.0 (DSP0233)
- DMTF Management Component Transport Protocol (MCTP) IDs and Codes, Revision 1.9.0 (DSP0239)
- PCI-SIG Data Object Exchange (DOE) ECN, published March 26, 2020
- DMTF Platform Level Data Model (PLDM) Base Specification, Revision 1.1.0 (DSP0240)
- DMTF Platform Level Data Model (PLDM) over MCTP Binding Specification, Revision 1.0.0 (DSP0241)
- DMTF Platform Level Data Model (PLDM) for Platform Monitoring and Control Specification, Revision 1.2.2 (DSP0248)
- DMTF Platform Level Data Model (PLDM) for Firmware Update Specification, Revision 1.2 (DSP0267)
- DMTF Platform Level Data Model (PLDM) for Redfish Device Enablement Specification, Revision 1.1.2 (DSP0218)
- DMTF CXL™ Type 3 Device Component Command Interface over MCTP Binding Specification, Revision 1.0.0. (DSP0281)
- "Hardware Secure Boot." Open Compute Project, 2018, Revision 1.0
- "SP 800-193, Platform Firmware Resiliency Guidelines." National Institute of Standards and Technology, May 2018

2 Normative References (cont'd)

- DMTF Security Protocol and Data Model (SPDM) Specification, Revision 1.2.0 (DSP0274)
- DMTF Security Protocol and Data Model (SPDM) over MCTP Binding Specification, Revision 1.0.1 (DSP0275)
- DMTF Secured Messages using SPDM over MCTP Binding Specification, Revision 1.1.0 (DSP0276)
- DMTF Secured Messages using SPDM Specification, Revision 1.1.0 (DSP0277)
- "Data Object Exchange (DOE) Engineering Change Notice", PCI-SIG, March 2020
- "Component Measurement and Authentication (CMA) Engineering Change Notice", PCI-SIG, 2020
- "Chapter 12 Architectural Out of Band Management Engineering Change Notice", PCI-SIG, December 2023
- DMTF Platform Level Data Model (PLDM) Type 2 CXL Memory Device Modeling, Revision 1.0 (DSP2067)

3 Terms and Definitions

- **AIC:** Add-in card. Generally accepted as a PCIe CEM adapter.
- **BIOS:** Basic Input/Out System, modern usage synonymous with 'UEFI'
- **BMC:** Baseboard Management Controller
- **CAMM:** Compression Attached Memory Module
- **CCI:** Component Command Interface (CXL)
- **CMA:** Component Measurement and Authentication (PCI-SIG)
- **CMC:** JEDEC Memory Controller for Compute Express Link®
- **CMF:** JEDEC Memory FRU for Compute Express Link®. Generally accepted as a field replaceable memory device supporting the CXL protocol. Encompasses CMMs, AICs, and modularized trays form factors. See Figure 1.
- **CMM:** JEDEC Memory Module for Compute Express Link®. Generally accepted as a modularized memory device supporting the CXL protocol in an enclosed form factor (such as U.2, EDSFF, etc.) with embedded memory media components. See Figure 1.
- **Config:** Configuration
- **CXL:** Compute Express Link®
- **CXL Memory Device (CMD):** Any CXL device advertising itself to the host as a 'CXL Memory Device' on the PCIe bus (as described in the 'Memory Device Configuration Space Layout' section of the CXL Specification). Generally accepted as any CXL Type 3 device (independent of form factor) and may include the controller, media, and other components. See Figure 1.
- **DDR:** Double Data Rate
- **DIMM:** Dual Inline Memory Module
- **DMTF:** Distributed Management Task Force
- **DRAM:** Dynamic Random Access Memory
- **DOE:** Data Object Exchange (PCI-SIG)
- **E1:** A family of slender form factors within EDSFF (SNIA)
- **E3:** A family of form factors within EDSFF (SNIA)
- **ECDSA:** Elliptic Curve Digital Signature Algorithm
- **ECN:** Engineering Change Notice
- **EDSFF:** Enterprise and Datacenter Standard Form Factor (SNIA)
- **FRU:** Field Replaceable Unit
- **FRU Information Device:** Device which stores vital product data
- **FW:** Firmware
- **IDE:** Integrity and Data Encryption
- **IPMI:** Intelligent Platform Management Interface
- **MCTP:** Management Component Transport Protocol (DMTF)
- **Memory Media FRU:** Field replaceable memory module without a CXL memory controller (e.g., DIMM or CAMM)
- **MMIO:** Memory-mapped Input/Output
- **OOB:** Out of Band
- **OCP:** Open Compute Project
- **OS:** Operating System
- **PCIe:** PCI Express (PCI-SIG)
- **PLDM:** Platform Level Data Model (DMTF)

3 Terms and Definitions (cont'd)

- **RSA:** Rivest–Shamir–Adleman public-key cryptosystem
- **SHA:** Security Hash Algorithm
- **SNIA:** Storage Networking Industry Association
- **SPD:** Serial Presence Detect, modern usage referring the device on DIMM which stores DIMM vital product data and may perform other functions such as temperature sensing (JEDEC)
- **Temp:** Temperature
- **TS:** Temperature Sensor
- **SPDM:** Security Protocol and Data Model (DMTF)
- **UEFI:** Unified Extensible Firmware Interface, modern usage synonymous with 'BIOS' (UEFI Forum)
- **VDM:** Vendor Defined Message (PCI-SIG)
- **VPD:** Vital Product Data

Figure 1 represents some example CXL memory solutions and their relationships to the terms used throughout this document:

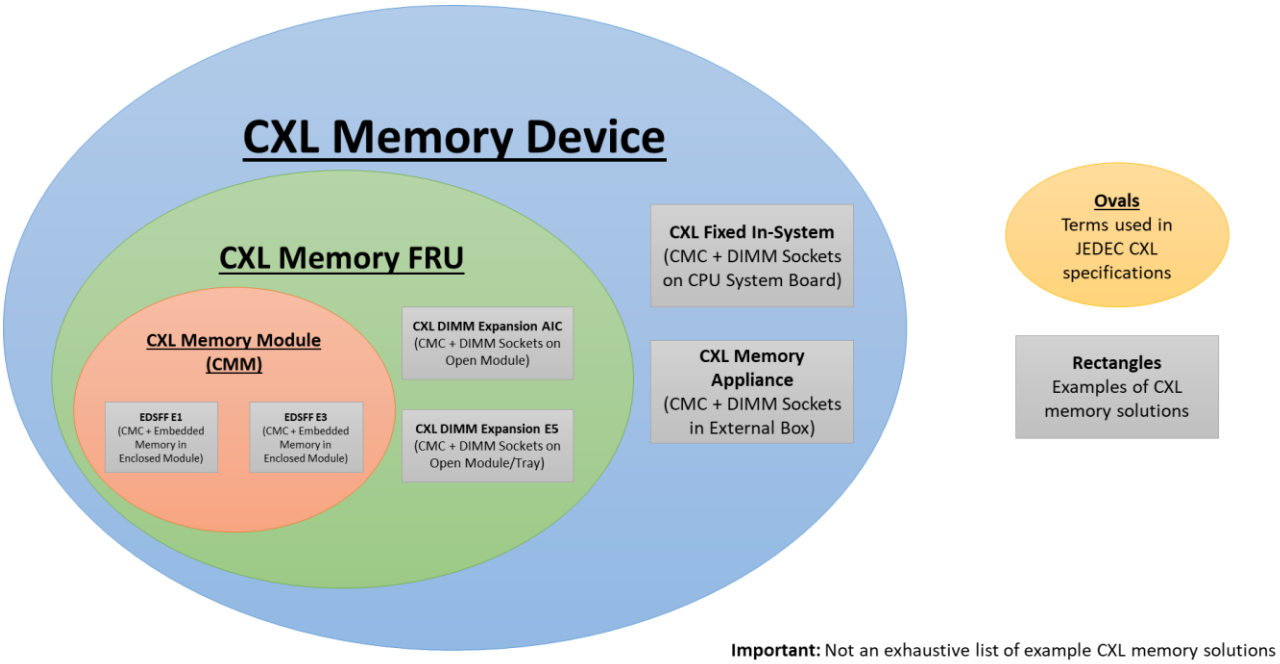

Figure 1 — Example Usage of Terms

4 Management Interface Requirements

The CXL memory device management interface is divided into three layers: Physical, Transport, and Protocol Layer. Supported interface layer combinations are shown in the table below, with expected usage models provided for reference.

Table 1 — Supported Management Interface Layer Combinations

Reference Usage Models	Device and Event Management and In-band FW Update	FW First Error/Event Management	Device Initialization during Boot	Device and Event Management and OOB FW Update	Platform Root of Trust based Attestation
Reference Managing Entity	Operating System	BIOS	BIOS	BMC	BMC
Protocol Layer	CCI	CCI	CCI	CCI, PLDM	SPDM
Transport Layer	Primary Mailbox	Secondary Mailbox	Primary Mailbox	MCTP	MCTP
Physical Layer	PCIe (Config Space and MMIO)	PCIe (Config Space and MMIO)	PCIe (Config Space and MMIO)	SMBus/I2C, PCIe (VDM)	SMBus/I2C, PCIe (VDM)

Optionally supported interface layer combinations are shown in Table 2, with expected usage models provided for reference.

Table 2 — Optionally Supported Management Interface Layer Combinations

Reference Usage Models	Host Device Security (Optional)
Reference Managing Entity	Operating System
Protocol Layer	SPDM over CMA
Transport Layer	DOE
Physical Layer	PCIe (Config Space and Extended Cap)

4.1 Management Physical Layer Requirements

CXL memory devices support two physical layers for device management: a two-wire serial bus and the host bus interface (PCIe).

4.1.1 Two-Wire Bus Interface Requirements

CXL memory devices support SMBus/I2C as described in the following sections of SNIA SFF-TA-1009 Enterprise and Datacenter Standard Pin and Signal Specification, Revision 3.1:

- Section 5.3.2 – SMBus Interface
- Section 6.2 – Timings
- Section 6.3 – 3.3 V Logical Signal Requirements

CXL memory devices optionally support I3C Basic as described in the following sections of SNIA SFF-TA-1009 Enterprise and Datacenter Standard Pin and Signal Specification, Revision 3.1:

- Section 5.3.3 – I3C Basic Interface
- Section 6.2 – Timings
- Section 6.4 – I3C Basic Signal Requirements
- Section 10 – I3C Basic Implementation

4.1.1.1 CXL Memory FRU SMBus/I2C Endpoint Requirements

CXL memory devices implemented as an FRU (e.g., CMM or AIC) supports the following SMBus/I2C endpoints:

Table 3 — CXL Memory FRU SMBus/I2C Endpoint Requirements

SMBus/I2C Endpoints	Default SMBus/I2C 8-Bit Address	Endpoint Presence	Management Power Only (Vmain = Off, Vmgmt = On)
SMBus/I2C Management (MCTP) Endpoint	3Ah	Required	Accessibility Not Required
FRU Information Endpoint	A6h	Required	Accessibility Required
SMBus/I2C Channel Multiplexer	E8h	Optional	Accessibility Required (if used)

Separate addresses are used if any of the above elements are combined/virtualized into one or more component devices.

The FRU Information Endpoint adheres to the FRU Information Device Requirements documented in the PCI-SIG Architectural Out of Band Management ECN (section 12.6.1). The FRU Information Endpoint contents (VPD) is described in Section 8.

CXL memory devices that support I3C will use the static I2C address listed above until it is assigned a dynamic address.

4.1.2 Host Bus Interface

PCI Express is used by CXL memory devices as a physical layer for both out-of-band mechanisms and in-band mechanisms.

For the out-of-band mechanism, CXL memory devices implement a management endpoint via PCIe Vendor Defined Messages (VDMs) and MCTP over PCIe VDM as described by the MCTP PCIe VDM Transport Binding Specification.

For the in-band mechanism, CXL memory devices implement CXL mailboxes via PCIe configuration space and MMIO as described by the CXL Specification. CXL memory devices may optionally implement a Data Object Exchange (DOE) mailbox via PCIe configuration space and extended capability registers.

4.2 Management Message Transport Requirements

CXL memory devices support two message transports for device management: MCTP and the CXL mailbox.

CXL memory devices may also optionally support a Data Object Exchange mailbox (DOE) for SPDM over Component Measurement and Authentication (CMA).

4.2.1 Management Component Transport Protocol (MCTP)

CXL memory devices support the Management Component Transport Protocol as described in the following DMTF specifications:

- Base MCTP – Management Component Transport Protocol (MCTP) Base Specification, Revision 1.3.1 (DSP0236)
- MCTP over PCIe VDM – Management Component Transport Protocol (MCTP) PCIe VDM Transport Binding Specification, Revision 1.2.0 (DSP0238)
- MCTP over SMBus/I2C – DMTF Management Component Transport Protocol (MCTP) SMBus/I2C VDM Transport Binding Specification, Revision 1.2.0 (DSP0237)

Unless otherwise defined in this document, CXL memory devices support required features and functions (including control messages) as required by the above DMTF specifications. Likewise, all features and functions described as optional in the DMTF specifications are also optionally supported by CXL memory devices.

CXL memory devices optionally support the Management Component Transport Protocol as described in the following DMTF specification(s):

- MCTP over I3C – Management Component Transport Protocol (MCTP) I3C Transport Binding Specification, Revision 1.0.0 (DSP0233)

4.2.2 CXL Mailboxes

CXL memory devices have both primary and secondary CXL mailboxes as described by the CXL Specification.

Note: The CXL Specification only requires a primary mailbox.

4.2.3 Data Object Exchange (DOE) Mailboxes

CXL memory devices optionally support a Data Object Exchange mailbox (DOE) for SPDM over Component Measurement and Authentication (CMA) as described in the PCI-SIG Data Object Exchange (DOE) ECN, published March 26, 2020.

4.3 Management Protocol Layer Requirements

CXL memory devices support three message data protocols for device management: PLDM, CCI, and SPDM.

4.3.1 Platform Level Data Model (PLDM)

CXL memory devices support the PLDM management protocol (MCTP message type 1) as described in the following DMTF specifications:

- Platform Level Data Model (PLDM) Base Specification, Revision 1.1.0 (DSP0240)
- Platform Level Data Model (PLDM) over MCTP Binding Specification, Revision 1.0.0 (DSP0241)

CXL memory devices support the following PLDM message types:

- Type 2 – PLDM for Platform Monitoring and Control, as described in DMTF Platform Level Data Model (PLDM) for Platform Monitoring and Control Specification, Revision 1.2.2 (DSP0248)
- Type 5 – PLDM for Firmware Update, as described in DMTF Platform Level Data Model (PLDM) for Firmware Update Specification, Revision 1.2 (DSP0267)

CXL memory devices optionally support the following PLDM message type(s):

- Type 6 – PLDM for Redfish Device Enablement, as described in DMTF Platform Level Data Model (PLDM) for Redfish Device Enablement Specification, Revision 1.1.2 (DSP0218)

4.3.2 CXL Component Command Interface (CCI)

CXL memory devices support the CXL Component Command Interface (CCI) over the following transports:

- Primary and Secondary CXL Mailboxes as described in the CXL Specification
- MCTP as described in the DMTF CXL™ Type 3 Device Component Command Interface over MCTP Binding Specification, Revision 1.0.0. (DSP0281)

4.3.3 Security Protocol and Data Model (SPDM)

CXL memory devices support SPDM as described in Security section 5.5.

4.4 Firmware Update Requirements

4.4.1 Firmware Update Interface Requirements

Based on the management interface content specified in sections 4.1-4.3, CXL memory devices support firmware update through the following methods and bindings:

4.4.1 Firmware Update Interface Requirements (cont'd)

- CCI Firmware Update over Primary CXL Mailbox
- CCI Firmware Update over MCTP over PCIe VDM
- CCI Firmware Update over MCTP over 2-Wire Interfaces
- PLDM Type 5 Firmware Update over MCTP over PCIe VDM
- PLDM Type 5 Firmware Update over MCTP over 2-Wire Interfaces

4.4.2 Firmware Update Security Requirements

CXL memory devices support secure firmware update as described in Security section 5.3.

5 Security

5.1 Signature and Hashing Algorithms

CXL memory devices support the following signature and hashing algorithms for Secure Boot, Secure FW update, and SPDM implementations:

CXL memory devices support ECDSA P384 or other asymmetric key signature algorithms that meet at least 192-bits of security strength as documented in NIST SP 800-57 Part 1 Revision 5 (Table 2). CXL memory devices may also support RSA 3072 instead. Supporting multiple signature algorithms is not required.

CXL memory devices support SHA2-384, SHA3-384, or other hashing algorithms that meet at least 192 bits of security strength as documented in NIST SP 800-57 Part 1 Revision 5 (Table 3). Supporting multiple hashing algorithms is not required.

5.2 Secure Boot

CXL memory devices implement a secure boot protocol as defined by the OCP “Hardware Secure Boot” specification revision 1.0, published 2018:

- https://www.opencompute.org/documents/secure-boot-2-pdf

This requirement does not require implementation of the separate OCP “Recovery” and “Attestation of System Components” specifications that are mentioned in the Secure Boot spec.

5.3 Secure Firmware Update

CXL memory device firmware update complies with the security guidelines of NIST SP800-193.

5.4 CXL Integrity and Data Encryption (IDE)

CXL memory devices optionally implement CXL IDE as described in section 11.1 of the CXL 2.0 specification.

5.5 Security Protocol and Data Model (SPDM)

CXL memory devices support SPDM messages and object exchanges (MCTP message type 5 and optionally type 6) as described in the following DMTF specifications:

- Security Protocol and Data Model (SPDM) Specification, Revision 1.2.0 (DSP0274)
- Security Protocol and Data Model (SPDM) over MCTP Binding Specification, Revision 1.0.1 (DSP0275)
- Secured Messages using SPDM over MCTP Binding Specification, Revision 1.1.0 (DSP0276)
- Secured Messages using SPDM Specification, Revision 1.1.0 (DSP0277)

5.5.1 SPDM Interface Requirements

Based on the management interface content specified in sections 4.1-4.3, CXL memory devices support SPDM through the following methods and bindings:

- SPDM/CMA over MCTP over 2-wire interfaces (I2C/SMBus or optionally I3C)
- SPDM/CMA over MCTP over PCIe VDM

5.5.1 SPDM Interface Requirements (cont'd)

CXL memory devices also optionally support SPDM via DOE/CMA over PCIe. SPDM via DOE/CMA is required if the device supports CXL IDE. In-band SPDM sessions are required to support IDE key exchange.

5.5.2 SPDM Features

5.5.2.1 Secure Sessions

CXL memory devices do not require Secure Sessions (unless CXL IDE or certificate provisioning outside of a trusted manufacturing environment are supported).

5.5.2.2 Mutual Authentication

CXL memory devices do not require mutual authentication.

5.5.2.3 Certificate Provisioning Support

CXL memory devices do not require certificate provisioning support and are outside the scope of this document.

5.5.3 SPDM Device Response Messages

Table 4 — SPDM Device Response Messages

Response	Code	CMM Requirement
DIGESTS	0x01	Required
CERTIFICATE	0x02	Required
CHALLENGE_AUTH	0x03	Required
VERSION	0x04	Required
CHUNK_SEND_ACK	0x05	Optional
CHUNK_RESPONSE	0x06	Optional
MEASUREMENTS	0x60	Required
CAPABILITIES	0x61	Required
ALGORITHMS	0x63	Required
KEY_EXCHANGE_RSP	0x64	Optional
FINISH_RSP	0x65	Optional
PSK_EXCHANGE_RSP	0x66	Optional
PSK_FINISH_RSP	0x67	Optional
HEARTBEAT_ACK	0x68	Optional
KEY_UPDATE_ACK	0x69	Optional
ENCAPSULATED_REQUEST	0x6A	Optional
ENCAPSULATED_REQUEST_ACK	0x6B	Optional
END_SESSION_ACK	0x6C	Optional
CSR	0x6D	Optional
SET_CERTIFICATE_RSP	0x6E	Optional
VENDOR_DEFINED_RESPONSE	0x7E	Optional
ERROR	0x7F	Required

5.5.4 SPDM Device Certificate Slots

DeviceCert Model support: Required

AliasCert Model support: Optional. If supported, the Alias certificate chain must be in Slot 0. Also, the Alias certificate chain must contain an immutable hardware identity certificate with hardware identity OID. All mutable certificates must include the mutable certificate OID format as specified in SPDM 1.2 specification.

5.5.4.1 Certificate Slots for CXL Memory FRUs

CXL memory FRUs must have at least one certificate slot provisioned. Two certificate slot configuration options are supported:

- Option 1
  - Slot 0 – Contains the certificate chain attesting the authenticity of the FRU.
  - Slots 1-7 – Not specified.

- Option 2
  - Slot 0 – Contains the certificate chain attesting the authenticity of the component executing the SPDM protocol.
  - Slot 1 – Contains the certificate chain attesting the authenticity of the FRU.
  - Slots 2-7 – Not specified.

Implementation Note: Per SPDM 1.2 requirements, if more than one certificate slot is being used (e.g., Option 2) then the SPDM protocol responder needs to have access to the private key corresponding to the common public key in the leaf certificate for the provisioned slots. SPDM challenge is supported for all provisioned slots by the responder. If the key pairs are different between provisioned certificate slots, then the SPDM protocol responder must have multi-key support per SPDM 1.3 or later.

5.5.5 SPDM Certificate Requirements

Unless otherwise defined in this document, CXL memory devices support the SPDM certificate fields as described as required in the DMTF SPDM specifications noted in section 10.4. Likewise, all SPDM certificate fields described as optional in the DMTF SPDM specifications are also optionally supported by CXL memory devices.

5.5.6 SPDM Challenge-Response

Device requirements for CHALLENGE_AUTH response:

Table 5 — SPDM Challenge-Response Requirements

CHALLENGE_AUTH Response	Implementation Requirement
0x0 = No Measurement Summary Hash	Required
0x1 = TCB Component Measurement Hash	Required
0xFF = All measurements Hash	Required

5.5.7 SPDM Firmware Measurements

CXL memory devices generate signed measurements and support the following measurement block types:

Table 6 — SPDM Firmware Measurement Requirements

DMTF Spec Measurement Value Type	Implementation Requirement
[7] - 0b: Hash, 1b: Raw Bit Stream	Hash - Required, Raw Bit Stream - Optional
[6:0] = 00h: immutable ROM	Required
[6:0] = 01h: mutable firmware	Required
[6:0] = 02h: hardware configuration, such as straps, debug modes	Required
[6:0] = 03h: firmware configuration, e.g., configurable firmware policy	Required
[6:0] = 04h: Measurement manifest	Required
[6:0] = 05h: Structured representation of debug and device mode.	Required
[6:0] = 06h: Mutable firmware's version number.	Optional but recommended
[6:0] = 07h: Mutable firmware's security version number, which should be formatted as an 8-byte unsigned integer.	Optional but recommended

CXL memory devices do not require support for recomputing all measurements without requiring a Reset. It’s acceptable that its capability flag MEAS_FRESH_CAP = 0.

6 Management Sensors and Effectors

CXL memory devices support and expose sensors and effectors for device management using PLDM Type 2 messages and optionally Redfish Device Enablement (RDE, PLDM Type 6). The minimum set of sensors and effectors provided through these interfaces are described in this section. Devices may have more sensors and effectors than mentioned in this reference specification. PLDM sensors and effectors are management abstractions that do not necessitate separate physical sensors or effectors.

To support PLDM Type 2 Sensors and Effectors as specified in this section, CXL memory devices support:

- PLDM State Set IDs as defined in DMTF Platform Level Data Model (PLDM) State Set Specification, Revision 1.1.0 (DSP0249)
- PLDM Sensor and Effecter Units as defined in section 27.4 (‘Sensor and effector units’) of DMTF Platform Level Data Model (PLDM) for Platform Monitoring and Control Specification, Revision 1.2.2 (DSP0248)

Guidance on how to model a CXL memory device in PLDM Type 2 semantics is provided by the following DMTF document(s):

- Platform Level Data Model (PLDM) Type 2 CXL Memory Device Modeling, Revision 1.0 (DSP2067)

Note: The DMTF PLDM Type 2 CXL Memory Device Model referenced above is not intended to be a document on required sensors or effectors. The model is a guidance document on how PLDM information should be constructed and conveyed by CXL memory devices. This document (JESD325) provides the minimum required sensors and effectors by JEDEC standard CXL memory devices.

6.1 CXL Memory Device Sensors and Effectors

CXL memory devices, regardless of form-factor, expose the following device-level sensors and effectors:

Sensor/Effector Requirement	Description	PLDM Type 2 Model Reference (DSP2067)
Aggregate Device Thermal Sensor	One sensor to report the device’s aggregate temperature reading – see Section 7 (Thermal Management) for further details on aggregate temperature reporting.	Sections 7.2.5, 8.3.4
Device Thermal Throttle Effector	One effector to enable/disable thermal throttling that attempts to keep temperatures on the device below the programmed warning, critical, or fatal limit for each temperature sensor – see section 7 (Thermal Management) for further details on thermal throttling.	Sections 7.2.6, 8.3.5
Device Power Sensor	One sensor to report the device’s current total power consumption, inclusive of all components associated with the device.	Sections 7.2.3, 8.3.2
Device Power Throttle Effector	One effector to enable/disable power throttling that attempts to keep CXL Memory Device’s total power consumption below the programmed warning, critical, or fatal limit for the module power sensor.	Sections 7.2.7, 8.3.6

6.2 CXL Memory Controller Sensors and Effectors

CXL memory controllers expose the following sensors and effectors:

Sensor/Effector Requirement	Description	PLDM Type 2 Model Reference (DSP2067)
Controller Temperature Sensor	One or more sensors to report the CXL memory controller's current temperature – see section 7 (Thermal Management) for further details on temperature reporting.	Sections 7.2.8, 8.4.1
Power Sensor	One or more sensors to report the CXL memory controller's current power consumption.	Sections 7.2.9, 8.4.2

6.3 Local Memory Media Sensors and Effectors

CXL memory devices that implement embedded memory media expose the following device-level sensors and effectors:

Sensor/Effector Requirement	Description	PLDM Type 2 Model Reference (DSP2067)
Memory Media Temperature Sensor	One or more sensors to report the local memory media component's current temperature – see section 7 (Thermal Management) for further details on temperature reporting. The CXL memory device exposes a PLDM temperature sensor for every physical temperature sensor implemented on the device intended to measure the memory media temperature.	Not Available in DSP2067 revision 1.0.

6.4 Local PMIC Sensors and Effectors

CXL memory devices that implement local power management ICs (PMICs) expose the following sensors and effectors for each PMIC:

Sensor/Effector Requirement	Description	PLDM Type 2 Model Reference (DSP2067)
Fault Sensor	One or more sensors to report the PMIC's current fault status.	Sections 7.2.13, 8.5.4
Power Sensor	One or more sensors to report the PMIC's current power reading.	Sections 7.2.11, 8.5.2
Current Sensor	One or more sensors to report the PMIC's current amperage reading.	Sections 7.2.12, 8.5.3

6.5 Memory Media FRU Sensors and Effectors

CXL memory devices that implement connectable memory module FRUs (such as DIMMs, CAMMs, etc.) expose the following sensors and effectors:

Sensor/Effector Requirement	Description	PLDM Model Reference (DSP2067)
Memory Media Sensors and Effectors	For each memory media FRU, the CXL memory device shall expose the same memory media sensors and effectors documented in section 6.3.	N/A
PMIC Sensors and Effectors	For each memory media FRU, the CXL memory device shall expose the same PMIC sensors and effectors documented in section 6.4.	N/A
Presence Sensor	For each memory media FRU connector, the CXL memory device shall expose a presence sensor to indicate whether a module has been installed in the connector.	Sections 7.2.14, 8.6.1

7 Thermal Management

Thermal management for CXL memory devices adhering to this management specification can be interfaced over two different protocols: the CXL Component Command Interface (CCI) and DMTF Platform Level Data Model (PLDM) – see section 4.3.

The thermal management functions built into CCI are basic and in some cases ambiguous. This section of the JEDEC Management Reference provides further details on how compliant CXL memory devices implement CCI based thermal management.

Conversely, thermal management capabilities via PLDM are powerful and scalable. This section of the specification provides details on a standard implementation of thermal management in PLDM semantics.

7.1 Temperature Sensors

Section 6 provides details on many of the CXL memory device's required temperature sensors as exposed in PLDM Type 2 or Type 6. However, also noted in section 6, these sensors are abstractions that do not necessarily correlate to specific physical or component-integrated temperature sensors. This section directly addresses the implementation of physical or component-integrated temperature sensors on a CXL memory device. As noted in section 7, compliant CXL memory devices expose a PLDM temperature sensor for every physical/integrated temperature sensor.

The CXL memory FRUs are divided into two common topologies: One is an assembled module with local memory media (e.g., JEDEC CMM) and the other is a memory media FRU installation module (e.g., AIC with DIMM sockets). Both topologies have integrated temperature sensors in the CXL memory controller and common PMICs.

Temperature sensors on a memory media FRU are defined by those module's standards. For example, the DDR5 temperature sensor on DIMM (TSOD) is defined by JEDEC with 3 temperature sensors on each DIMM. The number of temperature sensors and their placements on assembled modules with local memory media (e.g., CMM) are vendor defined.

DRAM ambient temperature sensor calibration should be owned by CMM or DIMM vendor to provide the offset values. System is capable of overriding initial offset values dependent on implementation specific. There are temperature sensors on CXL memory controller and PMIC as well for die temperature measurement. PMIC temperature sensors are not required to be exposed via PLDM. This is because the 3-bit value (as specified in register 0x33 of PMIC50x0 Power Management IC specification) does not translate well into a full numeric reading.

7.1 Temperature Sensors (cont'd)

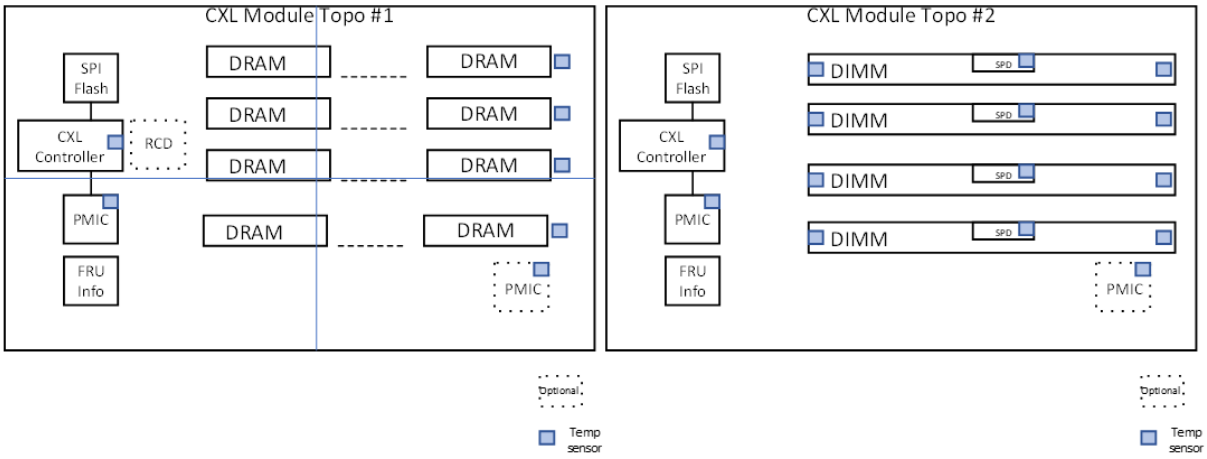

Figure 2 — Example CXL Memory Devices Topologies

7.2 CXL Memory Temperature Reporting Interface

There are two interfaces to report CXL memory temperature to system software/host. One is CXL CCI (Component Command Interface) reporting an aggregate device temperature and the other is PLDM (Platform Level Data Model) with aggregate and granular reporting.

7.2.1 CXL CCI

Temperature reported via Byte 4 ('Device Temperature') of the "Get Health Info" CCI command utilizes an aggregate temperature ($T_{aggregate}$) as defined in this specification. The aggregate temperature is also used when reporting Warning and Critical device temperature status (e.g., Bits[3:2] of in Byte Offset 2 of "Get Health Info" output payload). A CXL memory device will have multiple temperature sensors and is used in determining the aggregate temperature. In general, all PLDM temperature sensors measuring for a critical component temperature limit (e.g., DRAM, CMC, PMIC, etc.) should be calculated as part of the CCI aggregate temperature reading. The CXL memory controller continuously polls all temperature sensors at a default of frequency of 64 ms and is configurable according to controller vendor design. Calculation of aggregate temperature is reported as:

$$T_{aggregate} = 85 - \min((T_{sensor1\_thresh} - T_{sensor1\_Curr}), (T_{sensor2\_thresh} - T_{sensor2\_Curr}) ..)$$

Note that the Aggregate Temperature value is calculated based on multiple temperature inputs therefore it is not necessary to represent the actual temperature of any physical point within the CXL memory device. The PLDM method should be used for specific temperature component measurement of the device. Table 7 is an example for aggregate temperature calculation:

Table 7 — Example of Aggregate Temperature Calculation

PLDM Sensor	Tsensor_thresh	Tsensor_Curr
DRAM	85	80
PMIC	100	70
CXL Memory Controller	80	60
SPD	90	50

7.2.1 CXL CCI (cont'd)

$T_{aggregate}$ = 85 – min[(85-80), (100-70), (80-60), (90-50)] = 85 – min[(5), (30), (20), (40)] = 80

The controller register records current values for temperature for all sensors and aggregates the value (sticky until new value is read-shadow register in case of error event). The controller always presents the current aggregate temperature value to host via "Get Health Info" CCI command.

7.2.2 PLDM

The CXL memory controller polls all temperature sensors and keeps it in registers for management entities like BMCs to read individual temperature sensors value via PLDM. The same CCI aggregate temperature value discussed in section 7.2.1 is available via PLDM as well.

Each individual PLDM sensor on CMDs has its own PLDM Warning/Critical/Fatal threshold temperature which is configurable, and its default value depends on CMD implementation. Table 8 is provided as a reference value for each key device sensor:

Table 8 — Reference PLDM Sensor Threshold Values

PLDM Sensor	Warning	Critical	Fatal	Note
DRAM	85	95	100	Operating Case Temperature
PMIC	100	125	130	Operating Junction Temperature
CXL Memory Controller	80	90	100	Operating Junction Temperature
SPD	90	105	120	Operating Junction Temperature

Note: Please refer to the detail in the spec, JESD402-1 for temperature range and measurements of components and modules. Operating Junction temperature is the temperature reported by the internal temperature sensor of the component.

7.3 CXL Memory Thresholds and Alert Configuration

Two examples of CCI-based thresholding policies are provided in Figure 3. Implementation of these examples is not required but provided as a reference point for how event reporting can be used to facilitate system and local device actions.

7.3 CXL Memory Thresholds and Alert Configuration (cont'd)

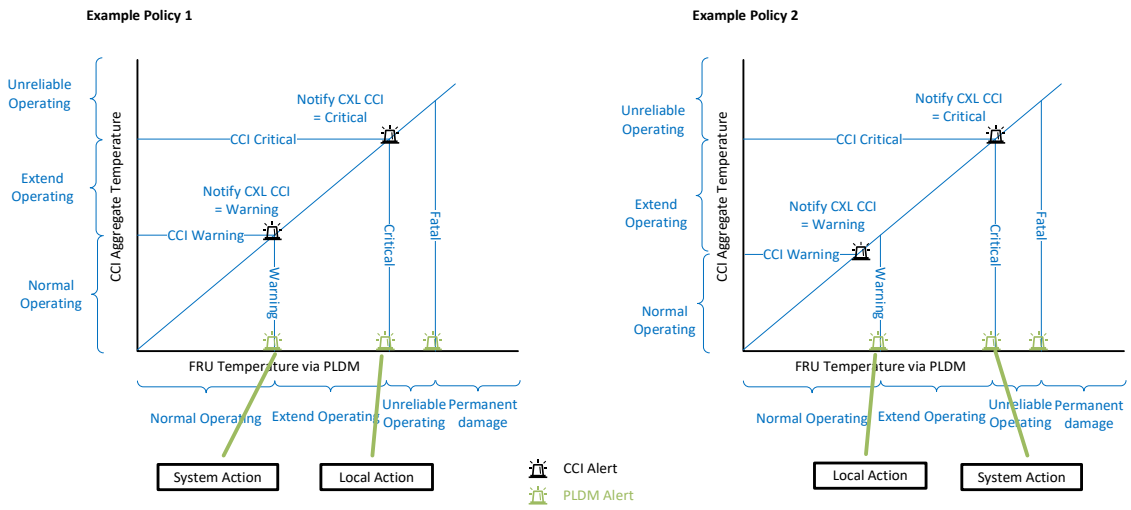

Figure 3 — Thermal Event Handling Examples

In Example Policy 1, the system is the first responder to increasing temperatures as it approaches the warning threshold. The device ultimately takes a local action at the device's critical thresholds.

In this example policy, 85 °C is recommended as the default value for CCI aggregate temperature Warning threshold. It may be configurable to a lower value by users in order to cool down the CMD in advance and prevent it from triggering 2x refresh (which in some circumstances cause worse thermal conditions.) Performance degradation could be observed once the temperature exceeds Warning threshold so system will be recommended to speed up fans, trigger thermal throttling correspondingly.

In this example policy, 95 °C is recommended as vendor defined CCI aggregate temperature Critical threshold. It is not recommended for the user to configure it since DRAM will not be reliable and the device may take a local action such as self-shutdown to prevent device damage.

In Example Policy 2, the response actions are reversed. Here, the device is the first responder to increasing temperatures as it approaches the warning threshold and then the system ultimately takes some action at the device's critical thresholds. Both 85 °C and 95 °C are still recommended for Warning and Critical thresholds respectively. As the device approaches the Warning threshold, it may take such actions as self-throttling. Beyond 85 °C it will trigger 2x refresh operation. At the Critical threshold, the system would like to take actions such as forced device shut down or CXL link disable to prevent temperature to keep rising and potential damage to components.

A hybrid policy of both local device and system actions at the Warning and Critical may also be taken but is not provided.

So, it is recommended CCI Warning and Critical aggregate temperature should be configured as below and showed some examples for the correlation between CCI and PLDM alerts.

- CCI Warning threshold: 85 °C (vendor defined but user configurable)
- CCI Critical threshold: 95 °C (vendor defined and do not recommend user to change it)

7.3 CXL Memory Thresholds and Alert Configuration (cont'd)

Table 9 — Examples of Correlation between PLDM and CCI Threshold Values

Sensor	FRU	Threshold Values	Example 1	Example 2	Example 3	Example 4	Example 5
PLDM Temperature Sensor	DRAM	Refer to Table 8	80	85	95	90	70
	PMIC		75	75	100	105	125
	Controller		70	70	80	85	60
	SPD		50	50	80	70	30
Aggregate Temperature	CMM	Warning > 85 Critical > 95	80	85	95	90	110

7.3.1 CXL CCI

System SW uses CXL CCI "Set Alert Config" to set warning threshold for aggregate temperature. The CXL memory device vendor defines the Critical Aggregate Temperature, considering the margin for system action delay. As mentioned in the previous section, 85 °C and 95 °C are recommended for Warning and Critical thresholds respectively. Table 10 shows examples for Local/System actions when hit threshold settings. It is not a requirement but for reference only.

Table 10 — Examples of Local/System Actions

Local Action	System Action
CXL controller polls all temperature sensors	Threshold value override
CXL controller defined aggregate temperature	Polls health status and PLDM FRU
Trigger DRAM 2x refresh	Fan speed up
Trigger Memory throttling	CXL link disable
Shut down PMIC	Event reporting

7.3.2 PLDM

Section 7.2.2 provides example default thresholds (warning, critical, fatal) for representative temperature sensor topology. It is implementation specific. The management entities like BMC may override the default threshold. The threshold severity levels are defined in PLDM for ***Platform Monitoring and Control Specification (DSP0248)*.**

Table 11 — The Reference Policy of Local/System Actions

Policy	Local Action at Warning Threshold	System Action on Warning Threshold	Local Action at Critical Threshold	System Action on Critical Threshold
Policy	Enable/disable	Enable/disable	Enable/disable	Enable/disable
DRAM temp > 84 ºC	1	0	0	0
PMIC temp > 90 ºC	0	1	0	0
DRAM temp > 94 ºC	0	0	1	1
PMIC temp > 98 ºC	0	0	1	1
DRAM temp > 98 ºC	0	0	1	1

7.4 Thermal Event Handling Examples

7.4.1 CXL CCI Based

The flow chart in Figure 4 shows the run time behavior. White boxes shown in flow chart are local actions. Blue boxes are the management entities actions as an implementation example.

Management entities set FRU temperature thresholds

- Controller polls all the sensors, calculates aggregate temperature, reports aggregate temperature via CCI Get Health Info and also keep register value of all the temperature values
- Management entities may poll via CCI Get health status and/or PLDM FRU temperature

When Threshold exceeds for a single temperature sensor

- Controller continues to report aggregate temperature, reports aggregate temperature via CCI Get Health Info and reports PLDM event for the temperature sensor that exceeded
- At this point, depending on the policy of handling, Management entities may start any system action (fan speed) and/or start polling the particular sensor

When Threshold exceeds for an aggregate temperature value

- Controller reports CCI event for aggregate temperature, and reports PLDM event(s) for the temperature sensor that exceeded
- At this point, depending on the policy of handling, Management entities may start any system action (fan speed) and/or start polling the particular sensor or CCI values

7.4.1 CXL CCI Based (cont'd)

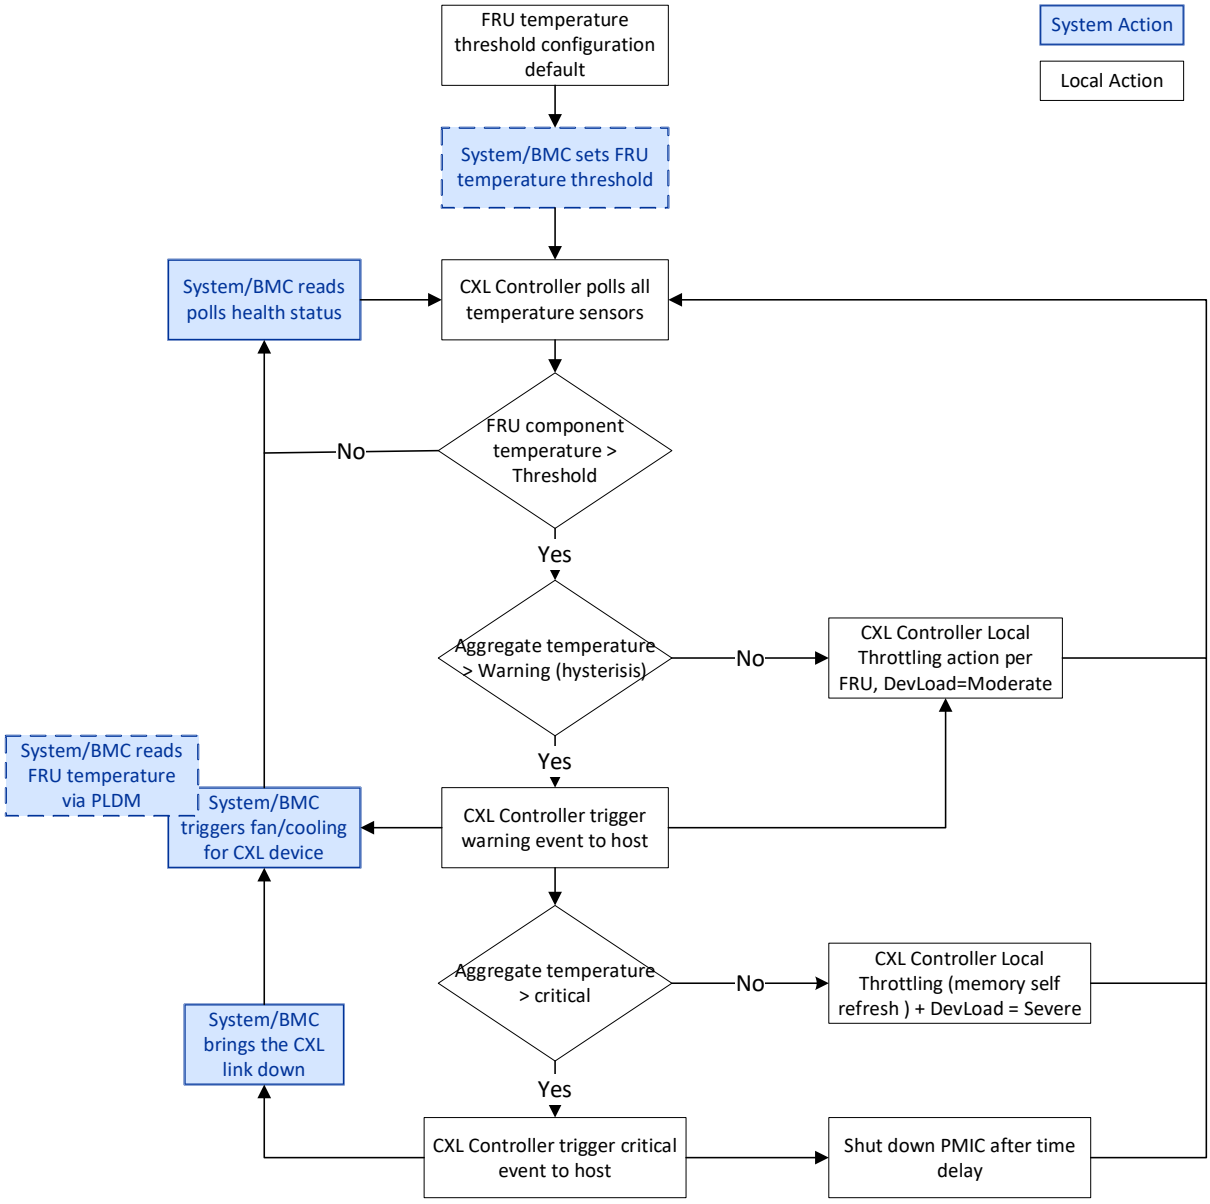

Figure 4 — Temperature Event Handing (CCI Based)

7.4.2 PLDM Based

Similar with CCI based temperature handling. The trigger condition of PLDM based control is when Threshold exceeds for a single temperature sensor. The system may optionally poll the PLDM sensor value after warning/critical/fatal events to tune actions further. Management entities can read individual temperature reading of each temperature sensor via PLDM either as polling or post a particular warning event being generated via PLDM path. In all, Management entities do not need to go through controller only.

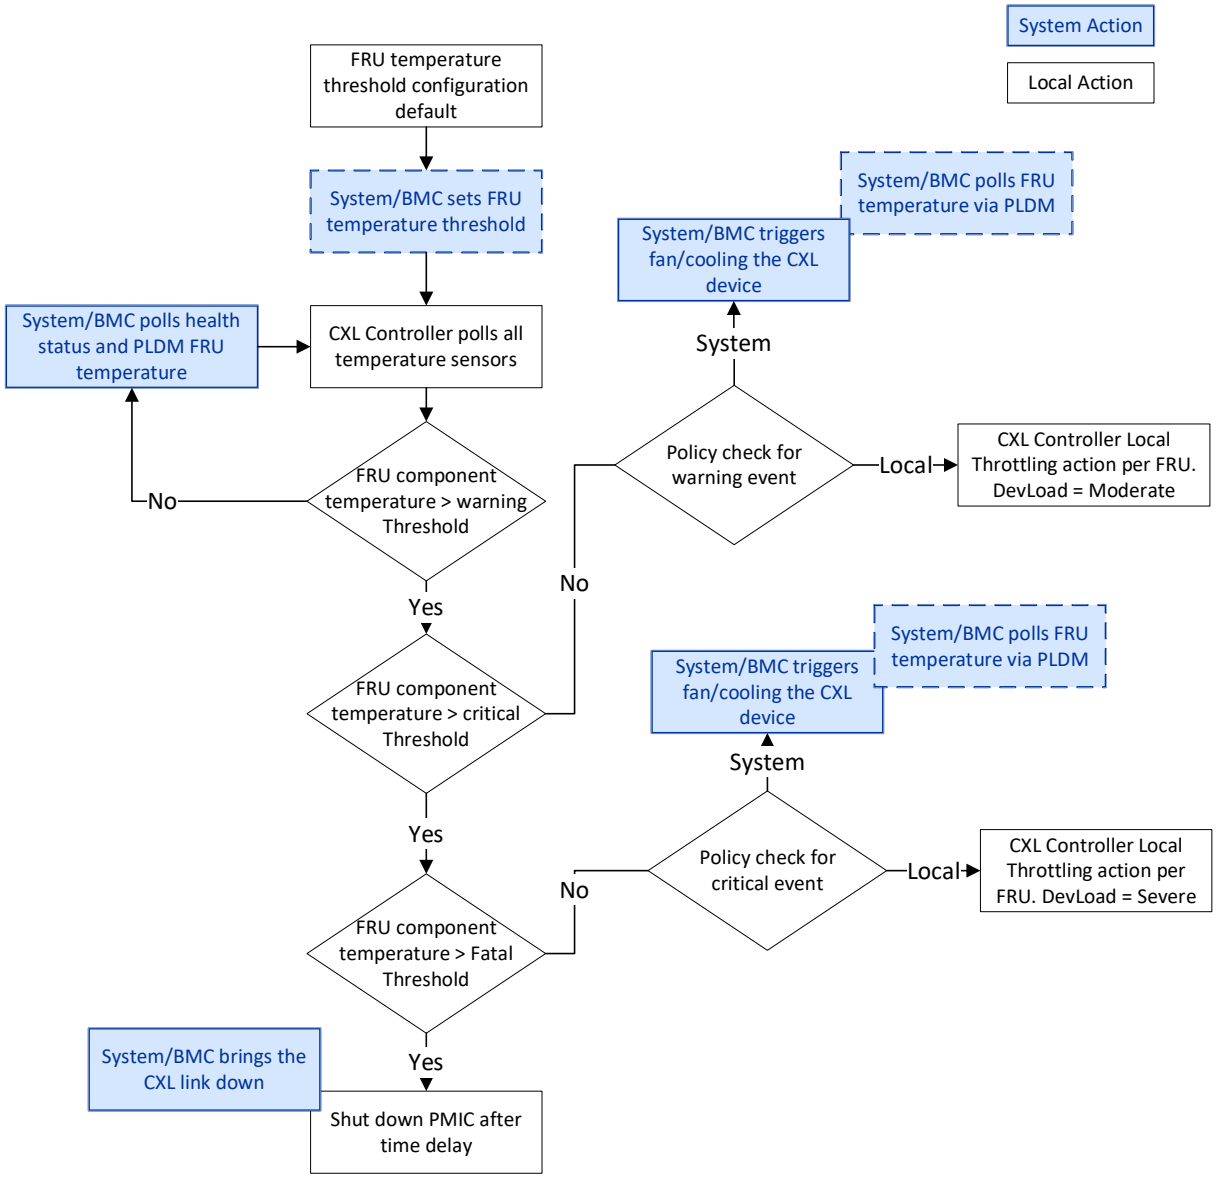

Figure 5 — Temperature Event Handing (PLDM Based)

8 FRU Information (VPD) Content

Vital Product Data (VPD) is data that describes a CXL memory device Field Replaceable Unit (FRU) to a host system to enable discovery and configuration. Each CXL memory device must have (at minimum) an FRU information device (i.e., EEPROM) which contains the defined VPD content in the required format. FRU may be virtual or physical construct. CXL memory devices may also include sub-FRU information coming in from any/all connected DIMM modules. This data is also referred as Serial Presence Detect (SPD) data.

8.1 VPD Formats

The VPD format uses a nested structure to define and describe the VPD content information. The structure describes the inclusion of a Multi-Record Area within the IPMI FRU Format definition containing a JEDEC Record for CXL Memory Devices, JEDEC Thermal Info Record, and JEDEC Power Info Record. The JEDEC Record for CXL Memory Devices definition includes Record header information as well as Element Descriptors for each FRU record element. A structure is defined for each FRU record element. Finally, a structure for fields within each Record Element is provided. A visual representation of the nested VPD structure is shown in Figure 6. Please follow the color coding to map the published content in next sections of this document:

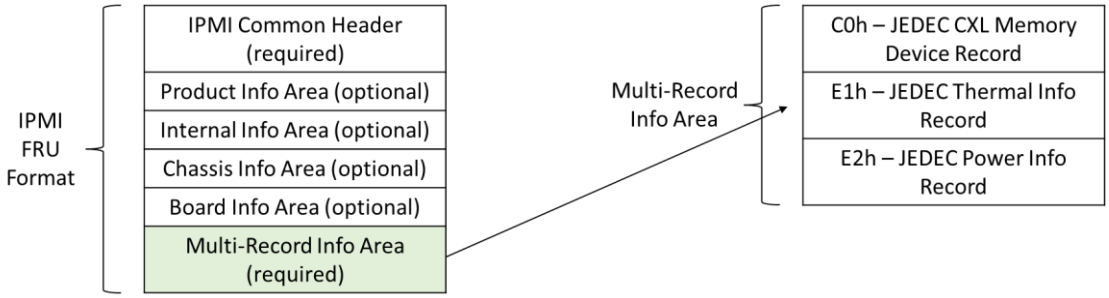

Figure 6 — High Level VPD Layout

8.1.1 IPMI VPD Structure

The following tables list the supported VPD element areas and IPMI VPD format length (Table 12), and detailed IPMI-based VPD definitions (Table 13).

Table 12 — JEDEC VPD Areas for CXL Memory Devices

FRU Element Areas	IPMI Format	
	Used	Length (Bytes)
Common Header	YES	8
Product Info Area	OPTIONAL	Implementation Specific
Multi-Record Info Area	YES	Implementation Specific
Internal Info Area	OPTIONAL	Implementation Specific
Chassis Info Area	OPTIONAL	Implementation Specific
Board Info Area	OPTIONAL	Implementation Specific

8.1.1 IPMI VPD Structure (cont'd)

Table 13 — CXL Memory Device VPD Areas – Detailed Definition

VPD Areas	Byte Offset	Factory Default	Description	Notes
Common Header	00h	01h	IPMI Format Version Number (IPMIVER)	
	01h	00h	Internal Use Area Starting Offset (IUAOFF)	Optional – 00h not present
	02h	00h	Chassis Info Area Starting Offset (CIAOFF)	Optional – 00h not present
	03h	00h	Board Info Area Starting Offset (BIAOFF)	
	04h	00h	Product Info Area Starting Offset (PIAOFF)	
	05h	08h	Multi-Record Info Area Starting Offset (MRIOFF)	The Multi-Record Info Area contains CXL Device Records, and Sub-FRU Records
	06h	00h	Reserved	
	07h	SKU Specific	Common Header Checksum (CHCHK)	Zero checksum
Multi-Record Area	SKU Specific			C0h – JEDEC Record for CXL Memory Devices E1h – JEDEC Thermal Info Record E2h – JEDEC Power Info Record

8.2 JEDEC Record Type 0xC0 – JEDEC Record for CXL Memory Devices

Table 14 — Layout of JEDEC Record for CXL Memory Devices

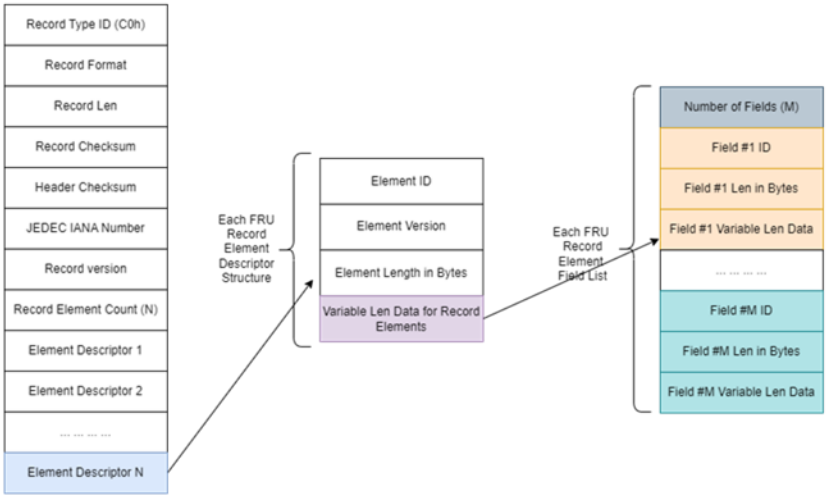

The first 9 bytes of the JEDEC Record for CXL Memory Devices shall be the FRU Record Header. Byte 10 onwards will be a linked list of individual FRU Records referred as Element Descriptor. Finally, each record or element descriptor shall be a TLV list of the record fields.

Table 15 — Structure Definition of JEDEC Record for CXL Memory Devices

Byte Offset	Factory Default Value	Description	Comments
00h	C0h	Record Type ID	Fixed for JEDEC compliant CXL Memory devices
01h	02h or 82h		Bit 7 - Set to 1 if last record in the multi-record list Bit [6:0] - Record Format Version set to 2h
02h	Implementation Specific	Record Length (RLEN)	Indicates the length of the Multi-Record Area in bytes
03h	Implementation Specific	Record Checksum	Used to give the record data a zero checksum
04h	Implementation Specific	Header Checksum	Used to give the record header a zero checksum
07h:05h	ED8Fh	JEDEC IANA Number	Internet Assigned Numbers Authority (IANA) manufacturer ID for LDAP, SNMP, etc.
08h	01h	Version Number	Indicates the version number of this multi-record

Table 15 — Structure Definition of JEDEC Record for CXL Memory Devices (cont'd)

Byte Offset	Factory Default Value	Description	Comments
09h	Implementation Specific	Element Count (N)	Indicates the number of Element Descriptors in this Topology Multi-Record
Implementation Specific	Implementation Specific	Element Descriptor 0	First Record Element Descriptor in this Topology Multi-Record
Implementation Specific	Implementation Specific	Element Descriptor 1	Second Record Element Descriptor in this Topology Multi-Record
.....	.....	Element Descriptor *N*	Applicable if count at offset 9 is > *N*

8.2.1 FRU Record Element Descriptor

The FRU Record Elements are based on the FRU Record Element Descriptor Format and are shown in Table 16.

Table 16 — CXL Memory Device FRU Record Element Descriptor Format

Byte Offset	Factory Default Value	Description	Notes
01h:00h	Implementation Specific	Element ID: Bits 15:12 = Reserved Bits 11:8 = Group ID • 00h = CXL 2.0 Base Specification for memory devices Bits 7:0 = Sub-Type: The exact element type defined for each of the records	CXL version of base specification – does not refer to specific features, etc.
02h	Implementation Specific	Version Number: This field indicates the version number of the Element Descriptor	
03h	Implementation Specific	Length: Number of bytes in the Element Descriptor	
04h	Implementation Specific	Element Data: This variable length area contains the Type-specific information associated with the Element Descriptor.	

8.2.2 FRU Record Element Data Field Format

FRU Record Element Data fields and offsets are defined in Table 17:

Table 17 — FRU Record Element Data Fields Definition

Byte Offset	Description
00h	Number of element fields in this record element (*M*)
01h	Field 1 - Field Type Number
Implementation Specific	Length: Number of bytes in the field
Implementation Specific	Field Data: This variable length area contains the element field data
Implementation Specific	Field 2 - Field Type Number
Implementation Specific	Length: Number of bytes in the field
Implementation Specific	Field Data: This variable length area contains the element field data
...	...
Implementation Specific	Field *M* - Field Type Number
Implementation Specific	Length: Number of bytes in the field
Implementation Specific	Field Data: This variable length area contains the element field data

8.2.3 Notes Regarding FRU Records

1. All Strings in the 0xC0 JEDEC Record for CXL Memory Devices are formatted as 8-Bit ASCII + Latin 1, as defined in the IPMI FRU Information Storage specification
2. All checksums should be calculated using the ‘Checksum8 Twos Complement’ method
3. For multi-port devices, FRU is common between all ports

8.2.4 CXL Memory Device FRU Content

The content of the CXL Memory Device VPD is divided into Element Record Type definitions and Element field definitions for each defined Record Type.

8.2.4.1 CXL Memory Device FRU Record Element Types

The CXL Memory Device FRU Record Element Types are defined in **Table 18**:

Table 18 — FRU Record Element Type Definitions

Element Type ID	Description	Optional / Mandatory	Description
0	Reserved		
1	General FRU Record	Mandatory	A CXL device shall contain one General FRU record where the fields match the JEDEC device label
2	Sub-FRU Record	Optional	A CXL device may optionally contain zero or more sub-FRU records to provide information about the sub-components of the CXL memory device. The description field shall be used to identify the type of sub-component.
3-253	Reserved		
254	OEM FRU Record	Optional	This record area to have optional write access for end-user
255	Reserved		

8.2.4.2 CXL Memory Device General FRU Record Element – Fields Definition

Table 19 defines the General FRU Record Element Fields:

Table 19 — CXL Memory Device General FRU Record Element Descriptor Definition

Field ID	Field Name	Data Type	Length in Bytes	Field Values
0	Reserved	N/A		
1	Media Type reference	Uint8	3	0: Reserved Refer to section 8.2.5 Media Type Field All others: Reserved
2	CXL Revision (base)	Uint8	1	0: Reserved 1: 2.0 2: 3.0 3: 3.1 All others: Reserved

Table 19 — CXL Memory Device General FRU Record Element Descriptor Definition (cont'd)

Field ID	Field Name	Data Type	Length in Bytes	Field Values
3	PCIe Revision	Uint8	1	0: Reserved 1: 5.0 2: 6.0 All others: Reserved
4	CXL Interface Type	Uint8	1	0: Reserved 1: Type 1 2: Type 2 3: Type 3 All others: Reserved
5	Module Capacity in GB	Uint32	4	Value Note: 00h designates devices which rely on sub-FRUs for media capacity
6	Module Form Factor and Thickness	Uint8	1	0: Reserved 1: E3.S-1T 2: E3.S-2T 3: E1.S-15mm 4: Other All others: Reserved
7	Connector Count	Uint8	1	0: Reserved 1: 1C/x4 2: 2C/x8 3: 4C/x16 4: 2X1C All others: Reserved
8	Port Configuration and IO Width	String	6	Example: 1Px8, 2Px4, 1Px16 Note: NULL characters (0x00) should be appended to field values when required to fully occupy 6-byte length
9	Serial Number <vid><mfgloc><mfgdate><serial>	String	24	Example: 800404A801250912345678 Note: Optional when PIA/BIA supported
10	Manufacturing Part Number, MPN	String	28	Example: MTA12ASF2G72PA-4H4A0 Note: Optional when PIA/BIA supported
11	Security Identifier, PSID	String	32	Example: 1A2B3C4D0918273600112233FFEEDDCC
12 - 255	Reserved	N/A		

8.2.4.3 CXL Memory Device Sub-FRU Record Element – Fields Definition

Table 20 defines the Sub-FRU Record Element Fields:

Table 20 — CXL Memory Device Sub-FRU Record Element Descriptor Definition

Field ID	Field Name	Data Type	Length in Bytes	Field Values	Examples
0	Reserved				
1	Type of Sub-FRUs	Uint8	1 Byte	0: RDIMM 1: MRDIMM All others: Reserved	RDIMM, LRDIMM, UDIMM, SODIMM, CAMM, etc.
2	Sub-FRU Bus	Uint8	1 Byte	0: DDR4 1: DDR5 All others: Reserved	DDR4, DDR5, LPDDR5, etc.
3	Number of Sub-FRU (DIMM) slots	Uint8	1 Byte		
4 - 255	Reserved				

8.2.5 CXL Memory Device Media Type Field of General Record

Table 21 defines the Media Type Field of the General Record Element:

Table 21 — CXL Memory Device Media Type Record Element Descriptor Definition Byte 0

Byte	Description	[7:4] Memory Type	[3:0] Memory Class
0	Memory Media	0h = DDR4 1h = DDR5 All others: Reserved	0h = volatile memory
		0h = NAND 1h = PCM All others: Reserved	1h = non-volatile memory
		0h = DDR4 + NAND 1h = DDR4 + PCM 2h = DDR5 + NAND 3h = DDR5 + PCM	2h = hybrid memory
		All others: Reserved	

8.2.5 CXL Memory Device Media Type Field of General Record (cont'd)

Table 22 — CXL Memory Device Media Type Record Element Descriptor Definition Byte 1

Byte	Description	Slot Type
1	Media Format	00h = Memory Down 01h = RDIMM All others: Reserved

Table 23 — CXL Memory Device Media Type Record Element Descriptor Definition Byte 2

Byte	Description	Slot Type
2	Reserved	Reserved

8.3 JEDEC Record Type 0xE1 – Thermal Record

This record provides information pertinent to the device's thermal/cooling requirements. The term "device" is used throughout this section to denote an assembly or part such as a PCI CEM adapter card or EDSFF module. Form factors are defined in Table 24.

Table 24 — Thermal Record Format

Byte Offset		Field Descriptor	Size in Bytes	Fixed Content	Comment
DEC	HEX				
0	0000	Record Type	1	0xE1	JEDEC Thermal Record for CXL Memory Devices
1	0001	End of List/ Record Format	1		
2	0002	Record Data Length	1	0x19	The data portion of this record has a length of 25 bytes.
3	0003	Record Checksum	1		
4	0004	Header Checksum	1		
5	0005	Manufacturer ID	3	0x8F	JEDEC IANA ID 0x00ED8F (LSB First)
6	0006			0xED	
7	0007			0x00	
8	0008	Version Number	1	0x01	Version 01
9	0009	Device Form Factor	2		Physical/Mechanical Form factor description (LSB,MSB)
10	000A				
11	000B	Airflow Impedance (AFI) Level	1		Classification group of device based on its impedance to airflow
12	000C	MaxTherm Level for Hot Aisle/Normal Airflow	1		Operational cooling requirement for hot-aisle installations
13	000D	MaxTherm Level for Cold Aisle/Reserve Airflow	1		Operational cooling requirement for cold-aisle installations
14	000E	Extended Operating Temperature (EOT) Thermal Level	1		Cooling Requirement for Extended Operating Temperature Range of 55-65 °C. May be the same or different from MaxTherm Level.
15	000F	Designed Thermal (DTherm) Level	1		Cooling level for failure of on-card air mover.
16	0010	Configuration	1		Device cooling capabilities and auxiliary power cable requirements
17	0011	Self-Protection Capabilities	1		Self-Protection Capabilities (i.e., throttling and shutdown)

Table 24 — Thermal Record Format (cont'd)

Byte Offset		Field Descriptor	Size in Bytes	Fixed Content	Comment
DEC	HEX				
18	0012	Minimum Local Inlet Temperature	2		Minimum allowable local inlet limit - deg C (LSB,MSB)
19	0013				
20	0014	Maximum Local Inlet Temperature Limit	2		Maximum allowable local inlet limit - deg C (LSB,MSB)
21	0015				
22	0016	Reserved	1		Reserved
23	0017	Minimum LFM	1		Minimum LFM to be provided by system fan cooling. Recommend used as needed in conjunction with MaxTherm Level.
24	0018	Minimum LFM Active Card Fault	1		Minimum LFM to be provided by system fan cooling for active cards in the event of the failure of the card air mover.
25	0019	Reserved	1		Reserved
26	001A	Reserved	1		Reserved
27	001B	Reserved	1		Reserved
28	001C	Reserved	1		Reserved
29	001D	Reserved	1		Reserved

8.3.1 Record Type

JEDEC Thermal Information Record for CXL Memory Devices has been designated as 0xE1.

8.3.2 Record Info

This field enables the software to determine Record Format Version. Please refer to the IPMI FRU specification for additional information. The bits are defined below.

- Bit - 7 End of List
- Bits - 6:4 Reserved
- Bits - 3:0 Record Format Version Number, 0x01 for IPMI FRU v1.0.

8.3.3 Record Data Length

This field states the length of the data portion of the record in 1-byte units (header fields are not included). Please refer to the IPMI FRU specification for details.

The data portion of this record has a length of 25 bytes for a hexadecimal value of 0x19.

8.3.4 Record Checksum

The checksum covers byte positions after the header (range: 0x0005 to 0x0028) and is calculated using Checksum8 + Two's Complement. This value must be calculated before Record Header Checksum. Please refer to the IPMI FRU specification for details.

8.3.5 Record Header Checksum

The checksum covers byte positions 0x0000 to 0x0004 and is calculated using Checksum8 + Two's Complement. This value must be calculated after Record Checksum. Please refer to the IPMI FRU specification for details.

8.3.6 Record Manufacturer ID

3-byte Manufacturer ID field, as specified in the GetDevice ID command in the IPMI v1.0 specification. This occupies the first 3 bytes of the data portion of the record. LSB first. JEDEC's IANA ID is 0x00ED8F. Please refer to the IPMI FRU specification for details.

8.3.7 JEDEC Record Version Number

JEDEC record version. This record has a version number of 0x01.

8.3.8 Form Factor Information

Two-byte field provides the form factor of the card or device. Location with the host system and thus cooling requirements will vary for different device form factors. The table below provides a list of potential CXL memory FRU form factors and the industry specifications that defines the thermal levels associated with each device form factor. The form factor information in combination with the thermal levels populated per the definitions in Sections X, Y, Z may be utilized by the host system to ensure proper cooling to the device. Since cooling requirements also vary for different form factors, it is necessary to populate fields for both the form factor and the required thermal level associated with that device form factor.

Format is LSB, MSB.

8.3.8 Form Factor Information (cont'd)

- Bits[15-8]
    - Reserved
- Bits[7:0]
    - Form Factors

Table 25 — Form Factors

Form Factor Code	Form Factors	Thermal Reporting Specification
0x00	PCIe Adapter Card: Low profile, half-length, single slot	PCI-SIG PCI Express Card Electromechanical Specification Revision 5.1
0x01	PCIe Adapter Card: Low profile, three quarter-length, single slot	Same as above
0x02	PCIe Adapter Card: Low profile, full-length, single slot	Same as above
0x03	PCIe Adapter Card: Low profile, half-length, dual slot	Same as above
0x04	PCIe Adapter Card: Low profile, three quarter-length, dual slot	Same as above
0x05	PCIe Adapter Card: Low profile, full-length, dual slot	Same as above
0x06	PCIe Adapter Card: Low profile, half-length, triple slot	Same as above
0x07	PCIe Adapter Card: Low profile, three quarter-length, triple slot	Same as above
0x08	PCIe Adapter Card: Low profile, full-length, triple slot	Same as above
0x09	PCIe Adapter Card: Full height, half-length, single slot	Same as above
0x0A	PCIe Adapter Card: Full height, three quarter-length, single slot	Same as above
0x0B	PCIe Adapter Card: Full height, full-length, single slot	Same as above
0x0C	PCIe Adapter Card: Full height, half-length, dual slot	Same as above
0x0D	PCIe Adapter Card: Full height, three quarter-length, dual slot	Same as above
0x0E	PCIe Adapter Card: Full height, full-length, dual slot	Same as above
0x0F	PCIe Adapter Card: Full height, half-length, triple slot	Same as above
0x10	PCIe Adapter Card: Full height, three quarter-length, triple slot	Same as above
0x11	PCIe Adapter Card: Full height, full-length, triple slot	Same as above
0x12-0x22	Reserved	N/A

Table 25 — Form Factors (cont'd)

Form Factor Code	Form Factors	Thermal Reporting Specification
0x23	SNIA EDSFF E3.S-1T	SNIA SFF-TA-1023
0x24	SNIA EDSFF E3.S-2T	Same as above
0x25	SNIA EDSFF E3.L-1T	Same as above
0x26	SNIA EDSFF E3.L-2T	Same as above
0x27	SNIA EDSFF E1.L 9.5mm	Same as above
0x28	SNIA EDSFF E1.L 18mm	Same as above
0x29	SNIA EDSFF E1.S 5.9mm	Same as above
0x2A	SNIA EDSFF E1.S 8mm	Same as above
0x2B	SNIA EDSFF E1.S 9.5mm (Symmetric Enclosure)	Same as above
0x2C	SNIA EDSFF E1.S 15mm (Asymmetric Enclosure)	Same as above
0x2D	SNIA EDSFF E1.S 25mm (Asymmetric Enclosure)	Same as above

8.3.9 Airflow Impedance (AFI) Level

This single byte provides the device's airflow impedance level as per the device's form factor thermal specification listed in the previous form factor section.

- Bits[7]
    - Reserved
- Bits[6:0]
    - Populated based on the AFI levels as defined in the device's form factor thermal specification.

8.3.10 MaxTherm Level for Hot Aisle/Normal Airflow

This single byte provides the device's operational cooling requirement in terms of a MaxTherm level for hot-aisle chassis configurations. In a hot-aisle configuration, devices are installed at the rear or exhaust (i.e., hot) end of the system such that IO connections are accessible to the user at the rear. Hot-aisle configurations are most prevalent for PCIe adaptors in server systems with the normal airflow direction having I/O connections on the exhaust (downstream) side of the card.

Reference the form factor specific thermal specification listed in the previous form factor section for MaxTherm level definition. For a particular device, the level that is loaded into this byte must match the defined cooling requirement that the device has been validated to meet. In general, the thermal levels relate the required inlet (approach) velocity of the device to the local device inlet air temperature.

In terms of power states, Operational in this context would correspond to a system level power state of S0. In S0 the system is completely operational, and all fans are turned on.

8.3.10 MaxTherm Level for Hot Aisle/Normal Airflow (cont'd)

- Bits[7]
    - Reserved
- Bits[6:0]
    - Populated based on the MaxTherm levels as defined in the device's form factor thermal specification.

8.3.11 MaxTherm Level for Cold Aisle/Reverse Airflow

This single byte provides the device's operational cooling requirement in terms of a MaxTherm level for cold-aisle chassis configurations. In a cold-aisle configuration, devices are installed at the front or inlet (i.e., cold) of the system such that IO connections are accessible to the user at the front of the system. Note that in cold-aisle configurations the airflow direction is reversed compared to hot-aisle configurations (described in previous section) with I/O connections on the upstream (inlet) side of the card.

Reference the form factor specific thermal specification listed in the previous form factor section for MaxTherm level definition. For a particular device, the level that is loaded into this byte must match the defined cooling requirement that the device has been validated to meet. In general, the thermal levels relate the required inlet (approach) velocity of the device to the local device inlet air temperature.

In terms of power states, Operational in this context would correspond to a system level power state of S0. In S0 the system is completely operational, and all fans are turned on.

The thermal specification associated with some form factors may not define multiple MaxTherm levels for cold-aisle configurations. In this case, the MaxTherm Level should be populated as Level 1.

For PCIe adapter card form factors, the MaxTherm Level for cold-aisle cooling may be determined by testing the card in reverse airflow using the level definitions in PCI Express Card Electromechanical Specification Revision 5.1.

In addition, some form factors thermal specifications may not define cooling level for cold-aisle configurations at all. In scenarios where the cold-aisle thermal level is not defined or required, populate a value of 0xFF for this field.

- Bits[7]
    - Reserved
- Bits[6:0]
    - Populated based on the MaxTherm levels as defined in the device's form factor thermal specification.

8.3.12 Extended Operating Temperature (EOT) Level

This field defines device's cooling requirement in terms of a thermal level if it supports an Extended Operating Temperature (EOT) range of 55-65 °C as defined in the respective form factor thermal reporting specification.

For cards that do not support operation in the EOT range (above 55 °C), populate a value of 0xFF for this field. For cards that do support EOT operation, the required thermal level must be populated. The expectation for the EOT Level is that it should either be equal to or greater than the worst case MaxTherm Level (normal vs reverse airflow).

- Bits[7]
    - Reserved
- Bits[6:0]
    - Populate based on the MaxTherm Level required for EOT operation (55-65 °C). Reference form factor section for external specification on thermal reporting.

8.3.13 Designed Thermal (DTherm) Level for Active Cooling Faults

Single byte that defines the designed thermal (DTherm) level required to cool a device with an on-board air mover (i.e., active) when a failure of the on-board fan occurs. The system fan control algorithms may set the system fan speeds to provide airflow equivalent to this cooling level when a fault for the device's onboard air mover is detected/flagged. This should NOT be populated with the cooling level for normal operation.

The value must be populated for active devices (active is designated when bit [5:4] of the 'Configuration' field is set to 01).

This value is not applicable for passive or liquid cooled devices which must use a value of 0x7F. (Passive is designated when bit [5:4] of the 'Configuration' field is set to 00. Liquid cooled is designated when bit [5:4] of the 'Configuration' field is set to 10).

- Bits[7]
    - Reserved
- Bits[6:0]
    - Populated based on the DTherm levels as defined in the device's form factor thermal specification.

8.3.14 Configuration

Single byte provides the configuration information.

- Bits[7:6]
    - Reserved
- Bit[5:4]
    - Denotes device cooling solution.
        - 00 = Passive
        - 01 = Active
        - 10 = Liquid Cooled
        - 11 = Reserved

8.3.14 Configuration (cont'd)

- Bit[3]
    - Auxiliary Power Cable Requirement - Denotes a cards AUX power cable requirement for maximum performance
        - No cable requirement defined = 0
        - Auxiliary cable required for maximum performance = 1
- Bits[2:0] – Reserved

8.3.15 Self-Protection Capability

Single byte provides information about the card's capability to throttle or shutdown/power-off to prevent damage due to overheating.

- Bits[7:3]
    - Reserved
- Bits[2:0]
    - Self-Protection Options
        - 000 = none
            - The device has no capability to throttle or shutdown to prevent damage due to overheating
        - 001 = throttle capable
            - The device can throttle (power, performance, etc.) to reduce the thermal load to help prevent overheating.
        - 010 = self-shutdown capable
            - The device has the ability to power off in order to prevent damage due to overheating
        - 011 = throttle and self-shutdown
            - The device has the capability to both throttle (power, performance, etc.) and power off in order to prevent damage due to overheating
        - 100 = Reserved
        - 101 = Reserved
        - 110 = Reserved
        - 111 = Reserved

8.3.16 Minimum Local Inlet Temperature Limit

Two bytes that provides the minimum allowable device inlet air temperature. This is defined as 'MinAmbient' in some form factor thermal specifications. Values are in units of degrees Celsius (°C). The minimum local inlet air temperature limit should only be used for devices that have a lower bound for the required local inlet air temperature. Most devices will not have this requirement and should be populated with 0xFFFF to denote that this field is not applicable for the particular device. Convert decimal to hexadecimal and enter in LSB, MSB order.

0xFFFF = Not applicable

Range: -32k to +32k

8.3.17 Maximum Local Inlet Temperature Limit

Two bytes that provide the maximum allowable device inlet air temperature. This is defined as 'MaxAmbient' in some form factor thermal specifications. Values are in units of degrees Celsius (°C). The maximum local inlet air temperature limit should only be used for devices that have an upper bound for the required local inlet air temperature. For example, a DIMM-expansion AIC might have a maximum inlet temperature of 45 °C. Most devices will not have this requirement and should be populated with 0xFFFF to denote that this field is not applicable for the particular device. Convert decimal to hexadecimal and enter in LSB, MSB order.

0xFFFF = Not applicable

Range: -32k to +32k

8.3.18 Minimum LFM

Single byte that provides the minimum allowable device local air velocity in LFM. The minimum allowable device local LFM should be used for devices that require an LFM that is greater than that provided by the assigned MaxTherm level. Minimum LFM is applicable to devices for which the device cooling requirement does not align with the LFM range from the MaxTherm level. In addition, if a device supports the upper end of the MaxTherm level but requires a higher LFM at the lower end of the operating temperature range, the Minimum LFM can be implemented to define a floor for the LFM that will be provided to the card. Cards that do not require a minimum LFM should be populated with 0xFF to denote that this field is not applicable for the particular device. Convert decimal to hexadecimal and enter in LSB, MSB order.

Min LFM required = (<Field Value> - 100) * 25

Hence the allowable LFM value supported is between 25 – 3850 and should be entered in the range of 101-254. An example table is shown in Table 26.

A decimal value from 101 – 254 will represent an encoded Minimum LFM value for all devices that require minimum LFM definition for their device cooling needs. Decimal value of 255 is not valid and equates to "Not Applicable". The formula to decode the field value (101 – 254) to Minimum LFM required is as follows:

Example: If a device complies with the upper end of the MaxTherm Level, for example Level 7, but does not comply to the lower end of the curve then the vendor would populate Level 17 for the MaxTherm level and the Min LFM field would be populated with the minimum LFM floor – e.g., 250 LFM (converts to decimal 108 and enter 0x6C in the field) for the Level 10 example.

Conversion Example: To define device LFM requirement of 500 LFM, divide the value by 25 (each single value above 100 is at a 25 LFM increment) and add 100 to it. Hence, for 500 LFM, it would be (500/25) + 100 = 120 (hex 0x78)

8.3.18 Minimum LFM (cont'd)

Table 26 — Example Table on Required LFM to Value to be Entered in the Field

Minimum LFM Desired	Decimal Value for the Field	Hex Value to be Entered in the Field
100	104	0x68
150	106	0x6A
200	108	0x6C
250	110	0x6E
300	112	0x70
350	114	0x72
400	116	0x74
450	118	0x76
500	120	0x78
550	122	0x7A
600	124	0x7C
650	126	0x7E
700	128	0x80
750	130	0x82
800	132	0x84
850	134	0x86
900	136	0x88
950	138	0x8A
1000	140	0x8C
.	.	.
.	.	.
3850	254	0xFE

8.3.19 Minimum Card LFM for Active Card Fault (Device's On-Board Air Mover Fails)

Single byte that provides the minimum allowable local air velocity in LFM required to cool a device with an on-board air mover (i.e., active) when a failure of the on-board fan occurs. The system fan control algorithms will set the system fan speeds to provide no lower than this LFM to the device when a fault for the device's on-board air mover is detected/flagged.

This field should be used for devices that require an LFM greater than that provided by the assigned DTherm level. Minimum LFM for Active Card Fault is applicable to devices for which the device is sufficiently cooled at the upper end of the DTherm level but requires a higher LFM at the lower end of the operating temperature range, the Minimum LFM for Active Card Fault can be implemented to define a floor for the LFM that will be provided to the card when a fault for the device's on-board air mover is detected/flagged.

Cards that do not require a Minimum Card LFM for Active Card Fault should be populated with 0xFF to denote that this field is not applicable for the particular device.

The Minimum Card LFM for Active Card Fault value is to be entered in the same format as the Minimum Card LFM definition in the previous section.

8.3.20 Reserved Fields

Reserved fields shall be populated with 0xFF.

8.4 JEDEC Record Type 0xE2 – Power Record

This record is used to describe the maximum power load required by a device at various states of operation and platform power states (S0 and S5). The term device is used throughout this section to denote an assembly or part such as a PCI CEM adapter card or EDSFF module. Devices that support sub-FRUs (such as DIMMs) shall make their measurements and report values assuming the highest power consuming components supported is populated.

The S0 power state as referred to in this record is defined by ACPI power state standards and is intended for when the server platform is in a fully-on and operational state.

The S5 power state is also defined by ACPI power state standards and is intended for when the server platform is in a soft-off state where typically only auxiliary power rails are available.

Table 27 — Power Record Format

Byte Offset		Field Descriptor	Size in Bytes	Fixed Content	Comment
DEC	HEX				
0	0000	Record Type	1	0xE2	JEDEC Power Record for CXL Memory Devices
1	0001	End of List/ Record Format	1		
2	0002	Record Data Length	1	0x0F	The data portion of this record has a length of 15 bytes
3	0003	Record Checksum	1		
4	0004	Header Checksum	1		
5	0005	Manufacturer ID	3	0x8F	JEDEC IANA ID 0x00ED8F (LSB First)
6	0006			0xED	
7	0007			0x00	
8	0008	Version Number	1	0x01	Version 01
9	0009	Max Sustained Power in S0	2		LSB
10	000A				MSB
11	000B	Max Peak Power in S0	2		LSB
12	000C				MSB
13	000D	Max Throttled Power in S0	2		LSB
14	000E				MSB
15	000F	Reserved	2		LSB
16	0010				MSB
17	0011	Reserved	2		LSB
18	0012				MSB
19	0013	Power Brake	1		

8.4.1 Record Type

JEDEC Power Information Record for CXL Memory Devices has been designated as 0xE2.

8.4.2 Record Info

This field enables the software to determine Record Format Version. Please refer to the IPMI FRU specification for additional information. The bits are defined below.

Bits	Meaning
7	End of List
6:4	Reserved
3:0	Record Format Version Number, 0x02 for IPMI FRU v1.0.

8.4.3 Record Data Length

This field states the length of the data portion of the record in 1-byte units (header fields are not included). Please refer to the IPMI FRU specification for details.

The data portion of this record has a length of 15 bytes and a hexadecimal value of 0x0F.

8.4.4 Record Checksum

The checksum covers byte positions after the header (range: 0x0005 to 0x0019) and is calculated using Checksum8 + Two's Complement. This value must be calculated before Record Header Checksum. Please refer to the IPMI FRU specification for details.

8.4.5 Record Header Checksum

The checksum covers byte positions 0x0000 to 0x0004 and is calculated using Checksum8 + Two's Complement. This value must be calculated after Record Checksum. Please refer to the IPMI FRU specification for details.

8.4.6 Manufacturer ID

3-byte Manufacturer ID field, as specified in the GetDevice ID command in the IPMI v1.0 specification. This occupies the first 3 bytes of the data portion of the record. LSB first. JEDEC's IANA ID is 0x00ED8F. Please refer to the IPMI FRU specification for details.

8.4.7 JEDEC Record Version Number

JEDEC record version. This record has a version number of 0x01.

8.4.8 Maximum Sustained Power (S0)

Maximum sustained power (in watts) consumed by the entire device while fully operational and platform is in an S0-power state. This power consumption value must be sustainable by the device indefinitely and would be typically achieved at 100% workload of the device. May be referred to as Thermal Design Point (TDP) and should be measured over a one second window with a sampling rate of 2 ms or better.

8.4.8 Maximum Sustained Power (S0) (cont'd)

- The values are stored as LSB in the first byte and MSB second byte.
- The values for this field are specific to the device implementation.
- Example: 0x000C = 12 watts

8.4.9 Maximum Peak Power (S0)

Maximum peak power (in watts) consumed by the entire device while fully operational and platform is in an S0-power state. This power consumption value may be sustained up to 10 µs, otherwise the maximum sustained power value should be used instead. May be referred to as Electrical Design Point (EDP) and should be measured over 100 µs window with a sampling rate of 4 µs or smaller.

- The values are stored as LSB in the first byte and MSB second byte.
- The values for this field are specific to the device implementation.
- Example: 0x000C = 12 watts
- This value should be the same as Maximum Sustained Power (S0) if there are no "peak" levels of power consumption by the entity.

8.4.10 Maximum Throttled Power (S0)

Maximum power (in watts) consumed by the entire device while fully operational but maximally throttled, and platform is in an S0-power state. This power consumption value would be typically achieved at 100% workload of the device but is externally forced to the lowest possible power consumption state (or internally forced by device's system management). Examples of externally triggered throttling may be signals such as GPUHOT#, PWRBRAKE#, etc. Maximum Throttled Power (S0) should be measured over a one second window with a sampling rate of 2 ms or better.

- The values are stored as LSB in the first byte and MSB second byte.
- The values for this field are specific to the device implementation.
- Example: 0x000C = 12 watts
- If this device does not support throttling, then this value should be set to 0xFFFF.

8.4.11 Reserved Fields

The reserved fields in this record use a value of 0x00.

8.4.12 Power Brake

Default value of 0xFF indicates no power brake support. Use a value of 0x00 to indicate support for power brake.

This page intentionally left blank

Standard Improvement Form **JEDEC Standard No. JESD325**

The purpose of this form is to provide the Technical Committees of JEDEC with input from the industry regarding usage of the subject standard. Individuals or companies are invited to submit comments to JEDEC. All comments will be collected and dispersed to the appropriate committee(s).

If you can provide input, please complete this form and return to:

JEDEC
Attn: Publications Department
3103 10th Street North
Suite 240S
Arlington, VA 22201

E-mail : angies@jedec.org

1. I recommend changes to the following:

☐ Requirement, clause number _____

☐ Test method number _____ Clause number _____

The referenced clause number has proven to be:

☐ Unclear ☐ Too Rigid ☐ In Error

☐ Other _____

2. Recommendations for correction:

_____

_____

_____

_____

3. Other suggestions for document improvement:

_____

_____

_____

_____

Submitted by

Name: _____ Phone: _____

Company: _____ E-mail: _____

Address: _____

City/State/Zip: _____ Date: _____


JEDEC®