A Comprehensive Verification Guide for Extended Metadata

Presented by: Po Chun, Wang

©2025 Conference Concepts, Inc. All Rights Reserved

Agenda

- Introduction
  - Background of CXL, metadata
  - Trailer and Extended Metadata transfer defined in CXL3.x
- The Challenge of Extended Metadata Implementation
  - CXL Link Layer
  - CXL Transaction Layer and System-Level Considerations
- Proposed Solution
  - Siemens Avery VIP Solution
- Conclusion

©2025 Conference Concepts, Inc. All Rights Reserved

Introduction

Background of CXL, metadata

Trailer and Extended Metadata transfer defined in CXL3.x

©2025 Conference Concepts, Inc. All Rights Reserved

Compute Express Link (CXL)

- CXL is an industry-standard, high-speed interconnect designed to support next-generation memory devices and accelerators with a **coherent, low-latency, and high-bandwidth interface**.
- **Key Capabilities:**
  - **Memory Management**: Efficient support for pooled, tiered, and disaggregated memory
  - **Coherence and Caching**: Enables seamless sharing and synchronization across CPUs, accelerators, and memory devices
- **Scalable Architecture:**
  CXL enables **rack-scale and row-scale system designs**, making it a critical enabler for workloads in:
  - **AI/ML**
  - **High-Performance Computing (HPC)**
  - **Cloud & Data Center Infrastructure**

©2025 Conference Concepts, Inc. All Rights Reserved

Metadata: Data That Describes Data

- **Metadata** is information about data - it provides context, structure, and meaning that enhances how data is managed and used
- It is **separate from the actual content**, but critical for its **interpretation, storage, access, and optimization**

- **Metadata Categories**[1]:
  - Structural - Data model, hierarchy, schema
  - Technical - Format, resolution, access speeds
  - Preservation - Backup info, version history
  - Administrative - Ownership, access rights

CXL3.x **Extended MetaData (EMD)**

    - **In CXL 1.1**, a field named **MetaField** was introduced to store **2 bits** of metadata per cache line — known simply as **MetaData**

Ref 1: IBM site – What is metadata
https://www.ibm.com/think/topics/metadata#:~:text=Metadata%20is%20information%E2%80%94such%20as,for%2C%20organize%20and%20use%20data.

©2025 Conference Concepts, Inc. All Rights Reserved

Extended Metadata – Concept Overview

- Key Concepts:
  - Each **64B CXL.mem transaction** can carry up to **32 bits of metadata**

- EMD Can Be Carried In:
  - **M2S RwD** (Memory Write):
    - From Master (Host) → Subordinate (Device)
  - **S2M DRS** (Memory Read Data):
    - From Subordinate (Device) → Master (Host)

- Packing Efficiency:
  - A specialized packing mechanism allows **up to 3 EMDs** to be bundled within a single DRS

©2025 Conference Concepts, Inc. All Rights Reserved

Extended Metadata – Integration Across Layers

- How EMD Is Transferred in CXL 3.x
  - The **CXL Transaction Layer** includes indicators that inform the **Link Layer** whether the corresponding Slot should append a Trailer, and whether that Trailer carries Extended Metadata (EMD)

- Transaction Layer
  - EMD Indicators

Field	Width (bit)	Encoding
TRP	1	1: Trailer Present
MetaField	2	01: Extended Meta State

- CXL Link Layer
  - Flit Construction

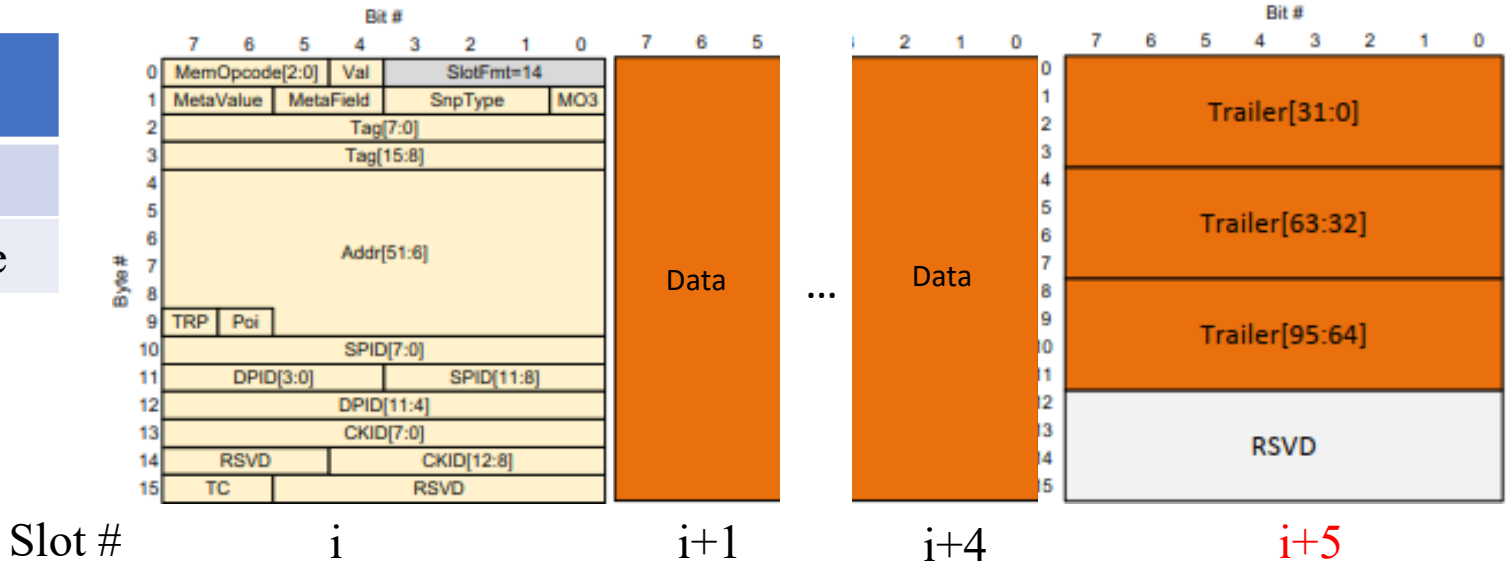

Extended Metadata – Multi-Header Packing

- Special Packing Rule for DRS
  - CXL 3.x introduces a new packing mechanism that **optimizes trailer placement** in **S2M DRS** Flits
  - The Trailer field supports the **packing of up to 3 Extended Metadata (EMD) units**—one for each DRS header

- Key Behavior
  - The **Trailer** is placed **immediately after the first 64B data segment** (associated with **Header 0**)
  - This Trailer can contain **EMD for Header 0, Header 1, and Header 2**, **even if Header 0 itself has no EMD**
  - This design improves bandwidth efficiency and minimizes fragmentation, contributing to **~5–7% higher data transfer efficiency**

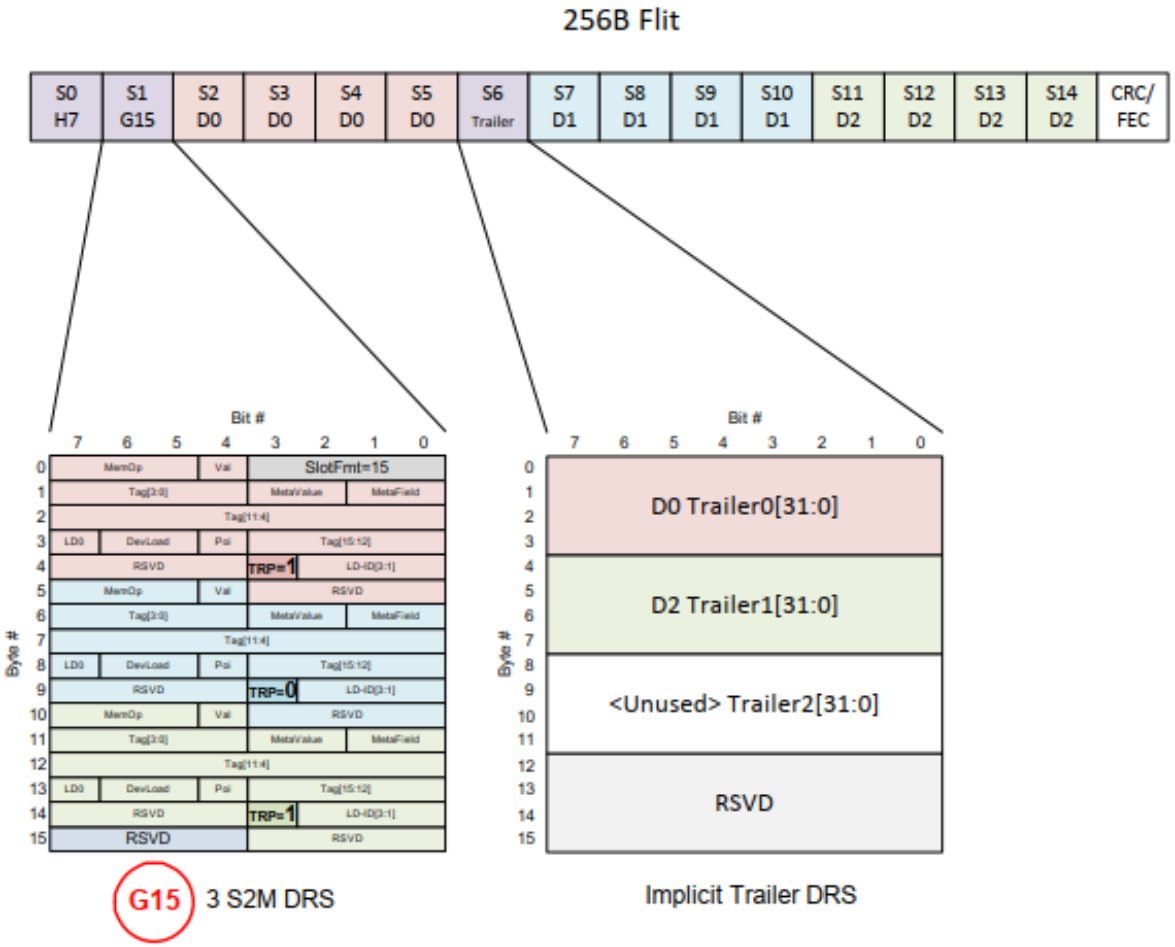

©2025 Conference Concepts, Inc. All Rights Reserved

Implementation Challenges

CXL Link Layer Challenges

CXL Transaction Layer and System-Level Considerations

©2025 Conference Concepts, Inc. All Rights Reserved

Implementation Challenges

- Adopting EMD in CXL 3.x has revealed practical challenges during implementation.
  - These challenges are based on both our development experience and feedback from early customer integrations

- **CXL Link Layer**
  - Risk of Data Corruption
    - Improper unpacking due to implicit slot-type encoding
  - Late Poison Data Mismatch
    - Failure to detect Trailer correctly can result in incorrect message poisoning
- **CXL Transaction Layer and above**
  - EMD Read-Back Mismatches
    - Caused by misconfigured EMD size in capabilities
  - Case Study: Error Reporting Ambiguity
    - CXL Spec defines 4 EMD error cases—but only one is reliably detectable in practice

©2025 Conference Concepts, Inc. All Rights Reserved

Challenge – Risk of Data Corruption

- There is no **explicit slot-type encoding** to distinguish between **data slots** and **trailer slots**
  - The receiver must rely on rollover tracking and previously received message headers to determine whether a received slot is a data slot or a trailer slot
- If this unpacking logic fails (e.g., header mismatch, misaligned rollover), it could cause:
  - Cascading errors across multiple transactions
  - Falsely decoding a data slot as an unexpected slot format may lead to flit corruption and protocol violations

- Figure 1: Transmitted a 256B-Flit (TX Side)

- Figure 2: Received a 256B-Flit (RX Side – Missed Trailer Detection)

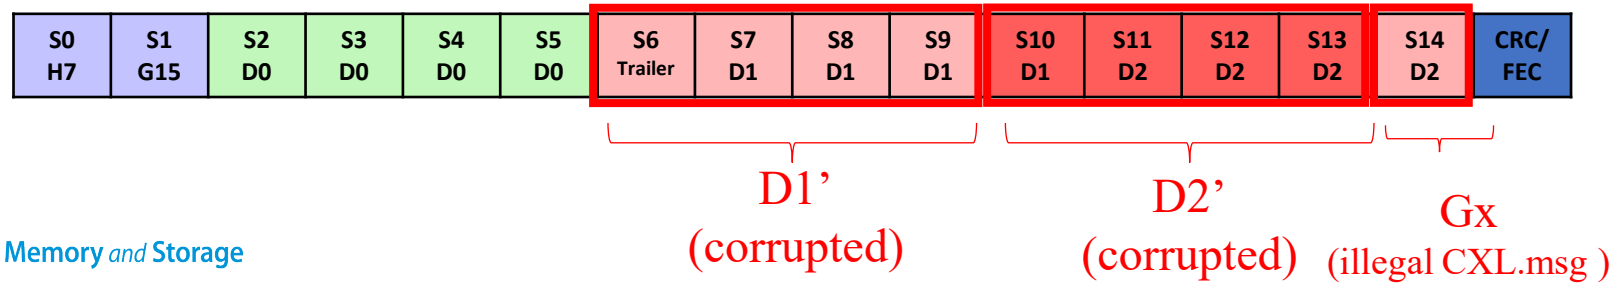

Challenge – Late Poison Mismatch

- Late poison is applied **after headers are sent**, via a **Control Flit** with a **message offset**
  - msg_offset = 0 → Poison the first CXL.msg in the next Flit
  - msg_offset = 1 → Poison the second CXL.msg in the next Flit
- If the receiver **fails to detect a Trailer**, it may associate the poison with the **wrong message**

- Figure 3: Transmitted 3 256B-Flits (TX Side)

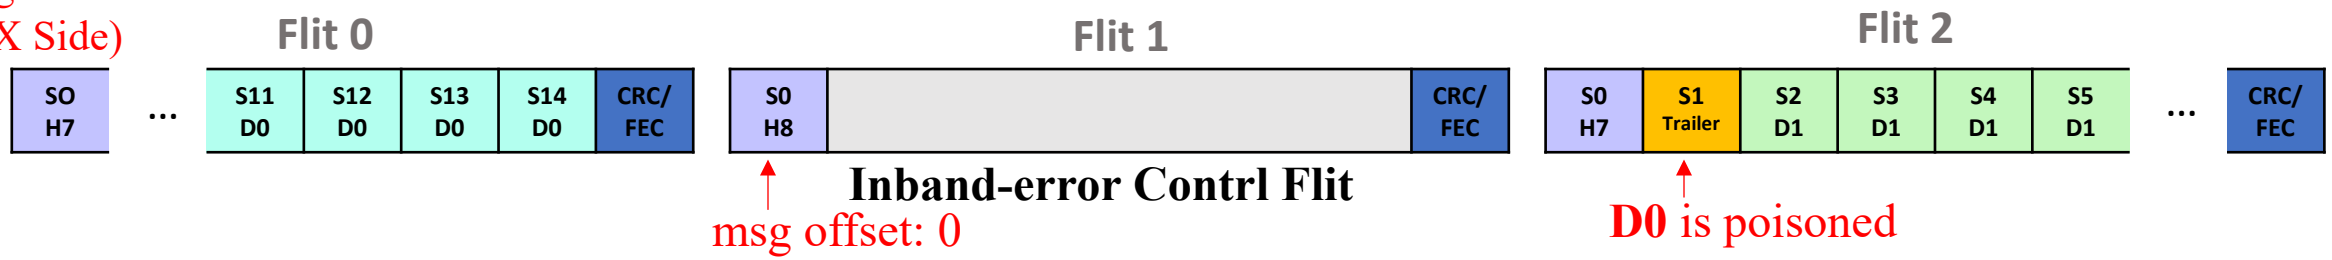

- Figure 4: Received 3 256B-Flits (RX Side – Missed Trailer Detection)

©2025 Conference Concepts, Inc. All Rights Reserved

Challenge – Incorrect EMD Configuration

- While each transaction can carry **up to 32-bit EMD**, the **actual bit-width** is determined by:
  - The "**Max Size of Extended Metadata**" field in **CXL Extended Metadata Capability Register**
  - The **"Size of Extended Metadata"** field in the **CXL Extended Metadata Control Register**
- Common issue: Customers set this incorrectly (e.g., 0 or smaller than expected)
  - EMD content dropped or truncated
  - Mismatch between write and read-back EMD

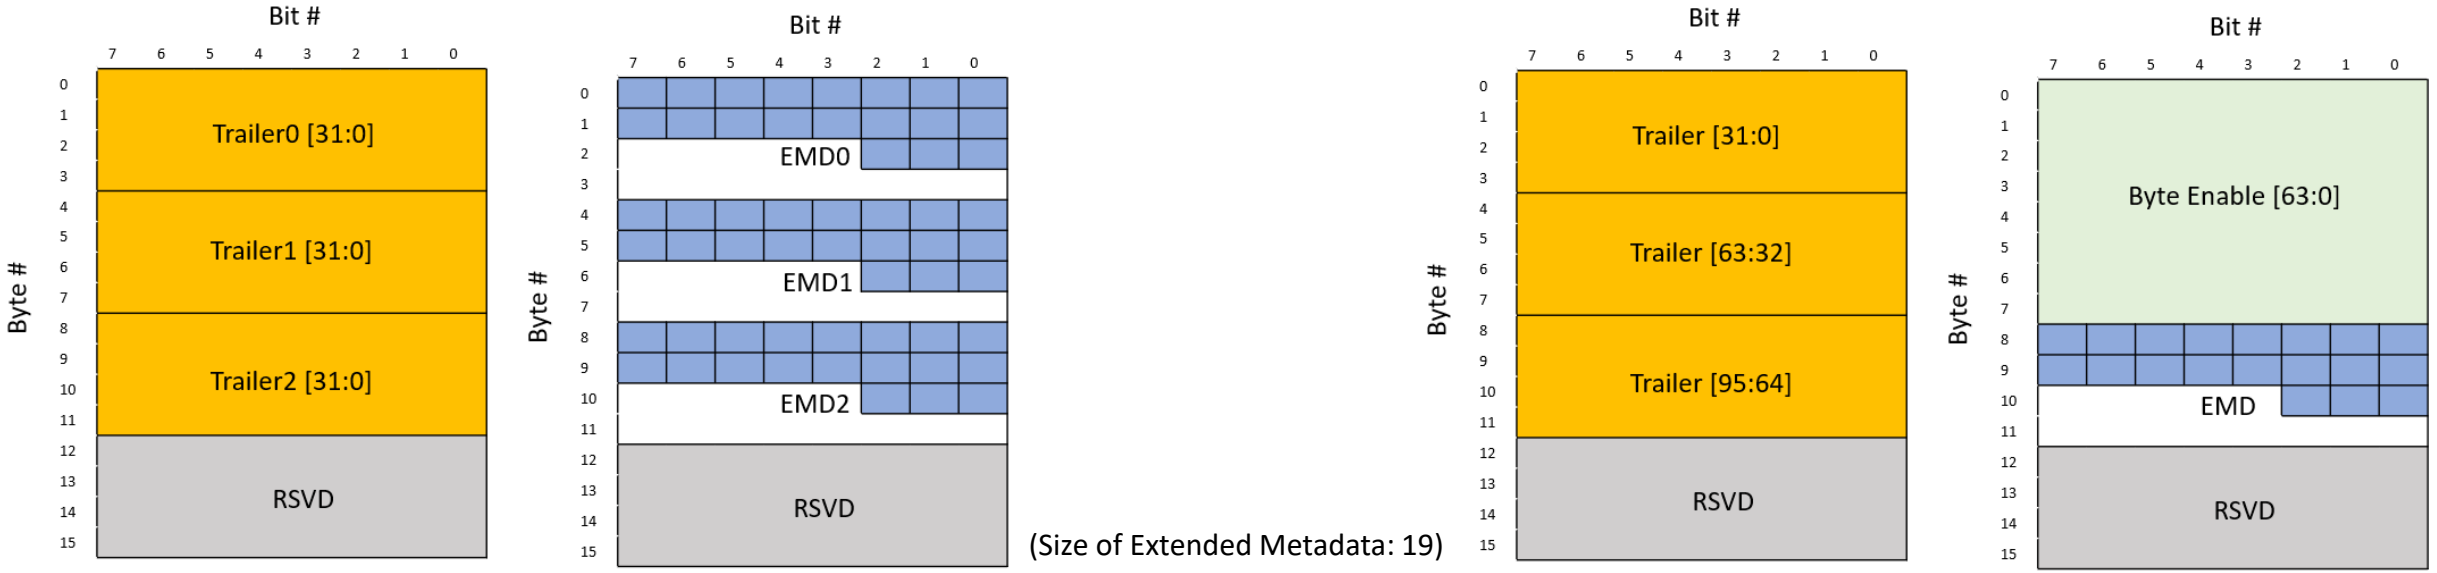

Challenge – EMD Error Reporting Case Study

- CXL spec defines **4 types** Extended Metadata Error
- In practice:
  - Only type3 leads to a detectable EMD-specific error by the receiver
  - Blindly relying on all 4 spec-defined cases may result in false confidence or misdiagnosis

Uncorrectable Error Status Register (Offset 00h) (Sheet 4 of 4)

Bit Location	Attributes	Description
17	RW1CS	**Extended_Metadata Error**: An error associated with Extended Metadata field.[2] DWORD 0 of the Header Log register captures the type of error: • 0 = A Root Port in an Extended Metadata-aware host received unexpected Extended Metadata on S2M DRS. • 1 = An Extended Metadata-aware device received unexpected Extended Metadata on M2S RwD. • 2 = A Root Port in an Extended Metadata-aware host expected but did not receive Extended Metadata on S2M DRS. • 3 = An Extended Metadata-aware device expected but did not receive Extended Metadata on M2S RwD. DWORD 1 of the Header Log register contains the following: • Bits[15:0]: Tag field associated with the value of the transaction with the EMDErr. • Bits[17:16]: MetaField value of the transaction with the EMDErr. • Bits[19:18]: MetaValue value of the transaction with the EMDErr. • Bit[20]: Indicates that an EMD value was captured with the EMDErr and is stored in DWORD 1. Must be 0 if the Enable Extended Metadata Error Logging bit is 0 in the CXL Extended Metadata Control register (see Table 8-170). DWORD 2 of the Header Log register captures the Extended Metadata field value if bit[20] of DWORD 2[1] is 1. This bit must be 0 if the Enable Extended Metadata Error Logging bit is 0.

Case	Enabled EMD	WrPTL	MetaField	TRP	Error Type
1	0	0	EMS	0	None
2	0	0	EMS	1	(Correctable Error)
3	0	0	non_EMS	0	None
4	0	0	non_EMS	1	None
5	0	1	EMS	0	None
6	0	1	EMS	1	(Correctable Error)
7	0	1	non_EMS	0	None
8	0	1	non_EMS	1	None
9	1	0	EMS	0	None
10	1	0	EMS	1	None
11	1	0	non_EMS	0	3
12	1	0	non_EMS	1	None
13	1	1	EMS	0	None
14	1	1	EMS	1	None
15	1	1	non_EMS	0	None
16	1	1	non_EMS	1	3

Solution

Siemens Avery VIP

©2025 Conference Concepts, Inc. All Rights Reserved

Siemens Avery VIP – Verification Solution for EMD

- **Verification Coverage**

Over **2000** checklist items in our CXL BFM, including deep validation of EMD-related features

Feature area	Checklist Items
Load/Store Extended Metadata	2
Flit Packing / Unpacking	31+
Late Poison Handling	45
RAS – Extended Metadata	2
EMD Configuration Validation	22

- **Compliance Test suite**
  - EMD Consistency Check
    - Configure CXL Extended Metadata Capability
    - Initiate MemWr/MemWrPtrl with EMD
    - Record EMD for every transaction
    - Initiate MemRd to retrieve data + EMD
    - Perform golden check for both data and metadata
  - EMD Error Reporting Test (Enabled upon CXL Spec RAS update)
    - Configure CXL Extended Metadata Capability
    - Inject EMD-related protocol errors
    - Validate RAS capability registers
    - Verify that DUT logs EMD errors correctly

©2025 Conference Concepts, Inc. All Rights Reserved

Flexible Topology Support – Full Stack or Layered Integration

- Siemens Avery VIP is built to support a wide range of DUT configurations through its **layered and configurable architecture**

DUT Configuration	Supported DUT-VIP topology
Full-stack DUT	Use full-stack BFM
DUT without TL	Use full-stack BFM + BFM (PHY+ARBMUX+DL)
DUT with TL only	Use full-stack BFM + APP BFM (TL)

- **Full Stack BFM**

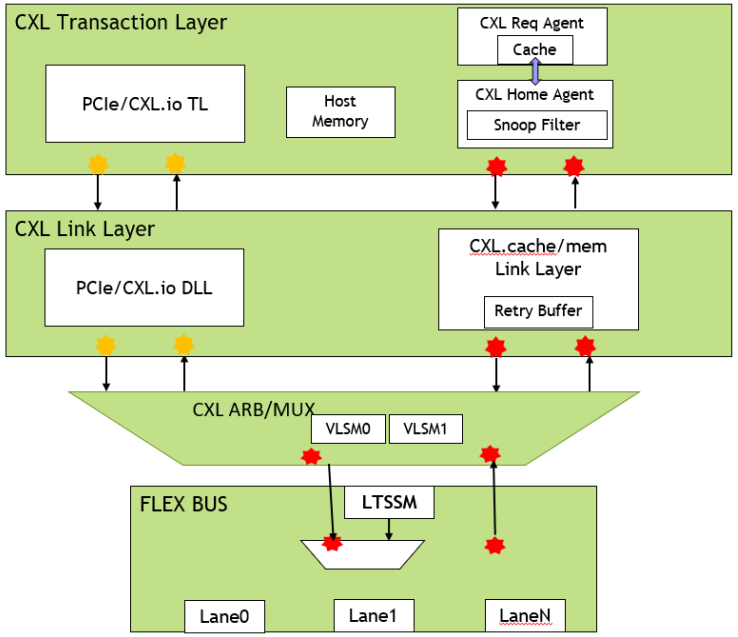

- **BFM (PHY+ ARBMUX + DL)**

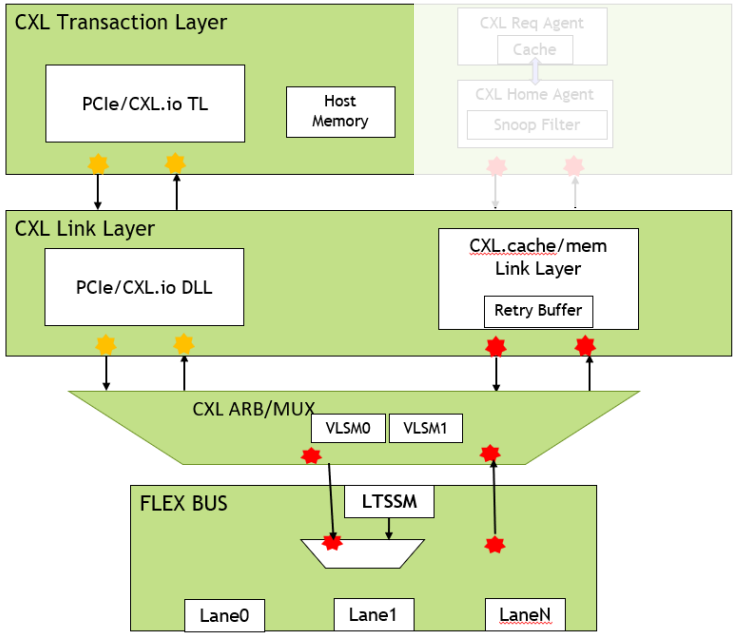

- **APP BFM (TL)**

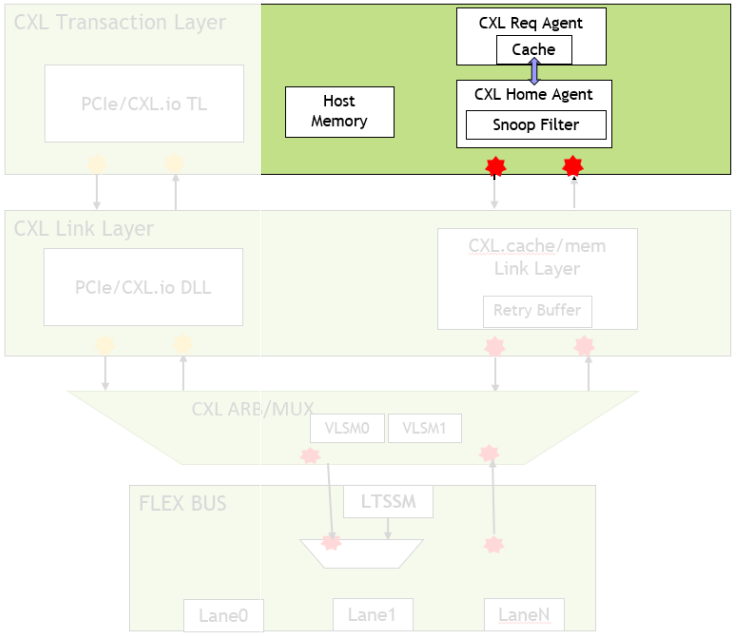

Summary

©2025 Conference Concepts, Inc. All Rights Reserved

Summary

- The **Extended Metadata (EMD)** feature in CXL 3.x offers a **flexible and scalable mechanism** to transfer metadata alongside memory transactions
- While conceptually straightforward, **EMD implementation involves non-trivial challenges** across both the Transaction and Link Layers—such as packing rules, trailer handling, late poison alignment, and configuration consistency
- **Siemens Avery VIP** provides a **comprehensive verification solution** to address these challenges—covering:
  - Protocol correctness
  - Edge-case detection
  - Compliance testing
  - Layer-specific integration

©2025 Conference Concepts, Inc. All Rights Reserved

QA

©2025 Conference Concepts, Inc. All Rights Reserved

Appendix

©2025 Conference Concepts, Inc. All Rights Reserved

Appendix - A

Case	Enabled EMD	WrPTL	MetaField	TRP	Error Type	Protocol Violation
1	0	0	EMS	0	None	v
2	0	0	EMS	1	None	INMO – this is a correctable error, but the unpacking logic should correctly unpack this
3	0	0	non_EMS	0	None	
4	0	0	non_EMS	1	None	v
5	0	1	EMS	0	None	v
6	0	1	EMS	1	None	INMO – this is a correctable error, but the unpacking logic should correctly unpack this
7	0	1	non_EMS	0	None	v
8	0	1	non_EMS	1	None	
9	1	0	EMS	0	None	v
10	1	0	EMS	1	None	
11	1	0	non_EMS	0	3	Uncorrectable Error
12	1	0	non_EMS	1	None	v
13	1	1	EMS	0	None	v
14	1	1	EMS	1	None	
15	1	1	non_EMS	0	None	v
16	1	1	non_EMS	1	3	Uncorrectable Error

©2025 Conference Concepts, Inc. All Rights Reserved

Appendix - B

Table 8-169. CXL Extended Metadata Capability Register (Offset 00h)

Bit Location	Attributes	Description
6:0	RO	**Max Size of Extended Metadata**: Defines the maximum size of the Extended Metadata Field within the EMD trailer. Valid values are from 1 to 32. • 1 = 1 bit of EMD • ... • 32 = 32 bits of EMD

Table 8-170. CXL Extended Metadata Control Register (Offset 04h)

Bit Location	Attributes	Description
6:0	RWL	**Size of Extended Metadata**: Defines the Extended Metadata Field size of a transfer. The device behavior is undefined if this register is set to a value that exceeds the Max Size of Extended Metadata reported via the CXL Extended Metadata Capability register[1]. • 1 = 1-bit EMD field. Corresponds to the LSB of the EMD Trailer. • ... • 31 = 31-bit EMD field, Corresponds to the 31 least significant bits in the EMD Trailer. • 32 = 32-bit EMD field Locked by the CONFIG_LOCK bit[2].
7	RO	**Reserved**
8	RWL/RO	**Enable Extended Metadata Error Logging**: This bit must be RWL if the Support for Extended Metadata Error Logging bit in the CXL Extended Metadata Capability register[1] is set; otherwise, this bit is permitted to be hardwired to 0. Software must not set this bit unless the Support for Extended Metadata Error Logging bit is set. If set, the device logs Extended Metadata content associated with the error, if possible, in the Header Log (see Table 8-95 for details). Locked by the CONFIG_LOCK bit[2].
30:9	RO	**Reserved**
31	RWL	**Enable Extended Metadata Field Transfers**: If set, the CXL device expects to receive and send Extended Metadata on data transfers via the trailer. Locked by the CONFIG_LOCK bit[2].

Reserved

Appendix - C

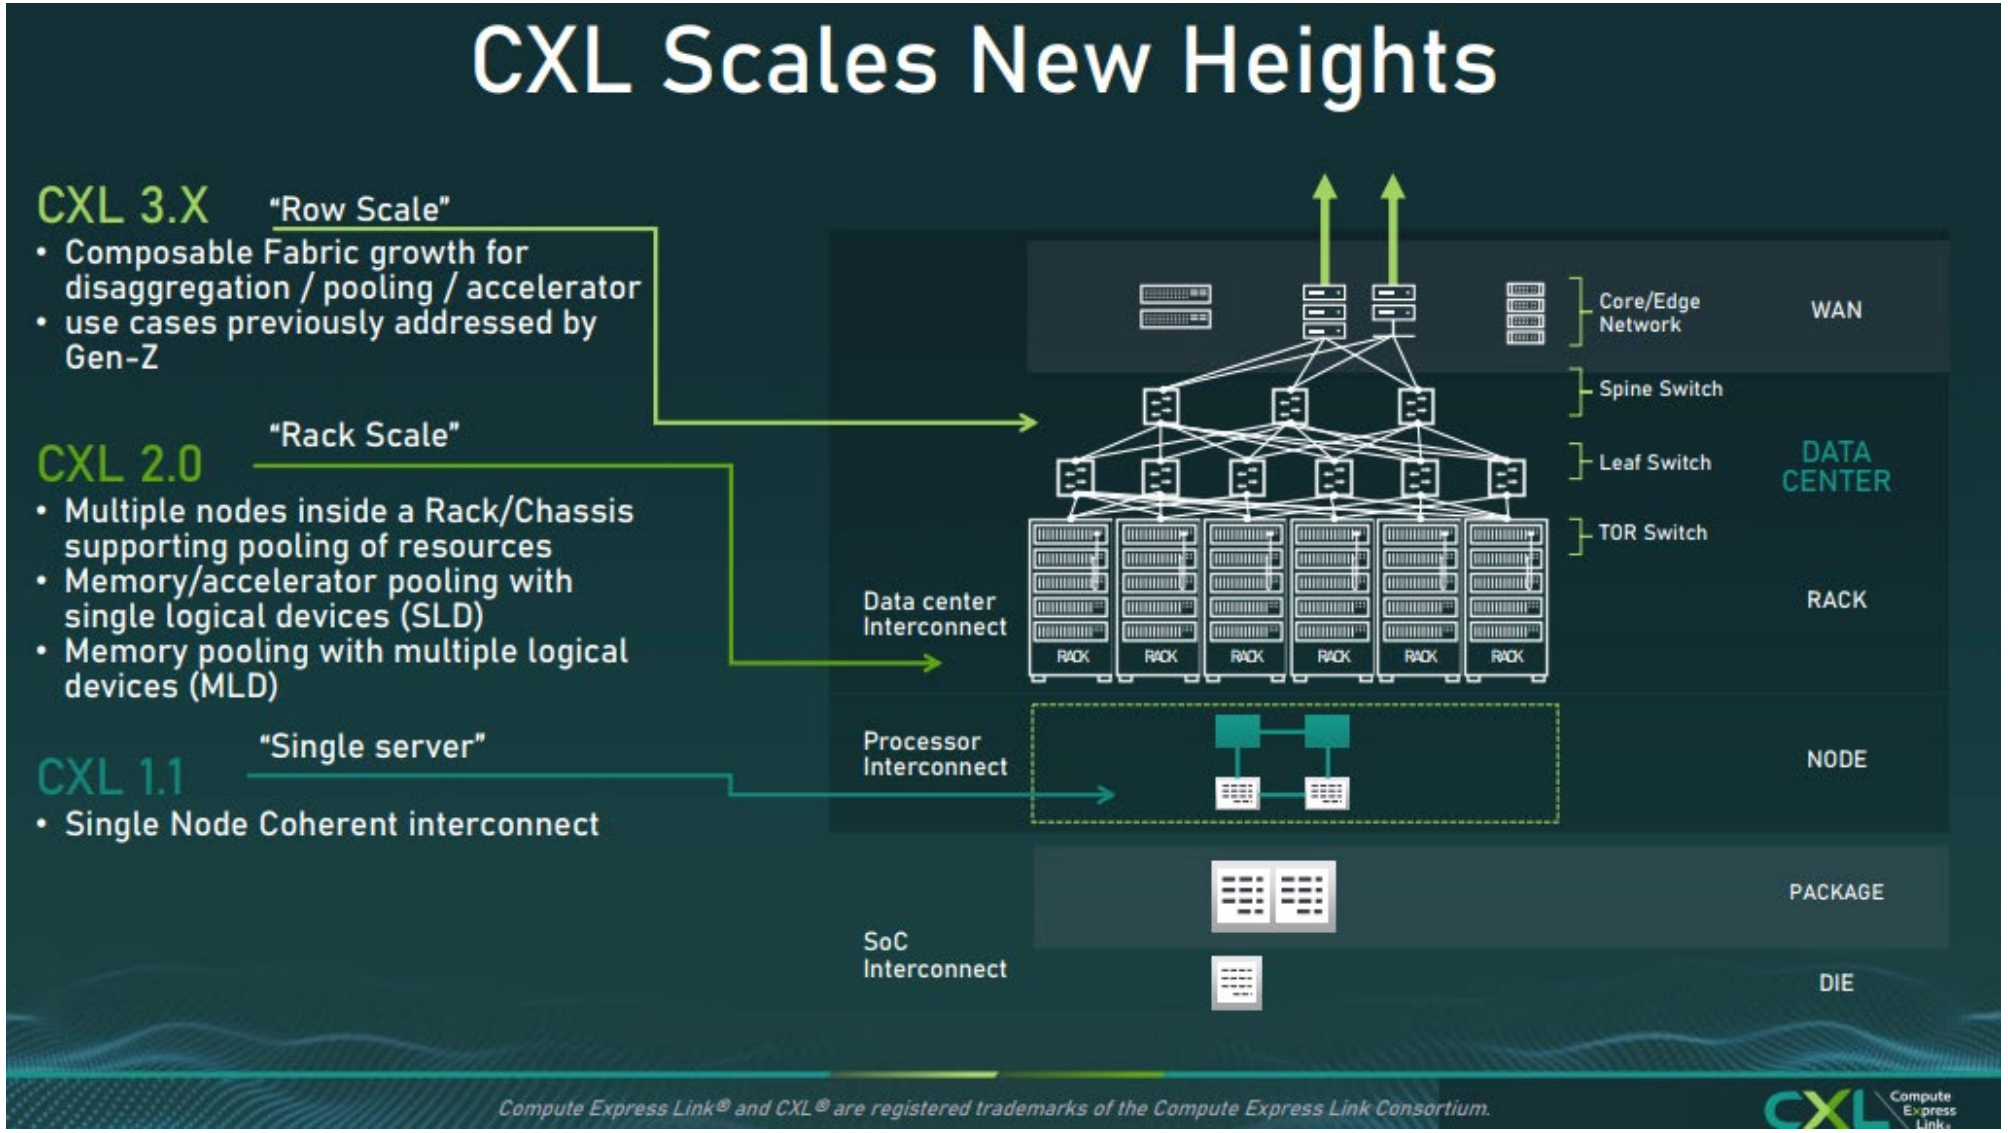
CXL Scales New Heights
CXL 3.X
"Row Scale"
• Composable Fabric growth for disaggregation / pooling / accelerator
• use cases previously addressed by Gen-Z
CXL 2.0
"Rack Scale"
• Multiple nodes inside a Rack/Chassis supporting pooling of resources
• Memory/accelerator pooling with single logical devices (SLD)
• Memory pooling with multiple logical devices (MLD)
CXL 1.1
"Single server"
• Single Node Coherent interconnect
Core/Edge Network
WAN
Spine Switch
Leaf Switch
DATA CENTER
TOR Switch
Data center Interconnect
RACK
RACK
RACK
RACK
RACK
RACK
RACK
Processor Interconnect
NODE
PACKAGE
SoC Interconnect
DIE
Compute Express Link® and CXL® are registered trademarks of the Compute Express Link Consortium.
CXL Compute Express Link