

Redfish

## Redfish support for Compute Express Link (CXL)

Copyright © 2023 DMTF

### Agenda

Overview of CXL The Chassis and System Model for CXL

- PCIe Devices/Functions
- CXL Logical Devices
- Memory Domains and Memory Chunks for CXL Memory

#### Modeling Local CXL Devices

- Type 1 Devices (SmartNICs)
- Type 2 Devices (Accelerators)
- Type 3 Devices (Memory Buffers)

- Modeling Remote CXL Devices
  - Type 1 Devices

DMTF

- Type 2 Devices
- Type 3 Devices
- The fabric model for all CXL Devices
  - Switches/Ports
  - Connections

DMTF Redfish



## **Compute Express Link Overview**



Figure above is © 2019-2022 COMPUTE EXPRESS LINK CONSORTIUM, INC.

Open industry standard for high bandwidth, low-latency interconnect

Connectivity between host processor and accelerators/ memory device/ smart NIC

#### Based on PCIe® 5.0 PHY infrastructure

- Leverages channel, retimers, PHY, Logical, Protocols
- Pin-to-pin compatible with PCIe connectors
- CXL.io I/O semantics, similar to PCIe mandatory
- CXL.cache Caching semantics optional
- CXL.mem Memory semantics optional
- CXL Link supports both standard PCIe devices as well as multi-protocol CXL devices
- CXL Logical Devices allow resources partitioning and binding to different hosts



#### **Representative CXL Use Cases**







Figure above is © 2019-2022 COMPUTE EXPRESS LINK CONSORTIUM, INC. www.dmtf.org

## **CXL Device Base Redfish Model**



# CXL DeviceCXL Logical DeviceMemoryCXL Logical DeviceCXL FunctionPCle FunctionCXL FunctionAcceleratorNICAccelerator

- CXL management model built on top of Redfish
  PCIe model
- CXL Device can provide standard PCIe Functions as well as extended functions supporting new CXL cache and memory semantics
- New CXL Logical Device allows partitioning of device resources and binding to different hosts through CXL switch
- PCIe Functions may be flexibly associated with CXL Logical Devices thus bound to different compute host (future functionality not defined in current specification)
- All PCIe Functions supporting CXL extensions associated with CXL Logical Device can use resources (e.g., memory) referenced by these devices



## **CHASSIS AND SYSTEM MODEL FOR CXL**



## **Chassis Model for CXL Device**

- CXL devices are not always bound to a specific host physically
- CXL device objects are located in the Chassis Model
  - Chassis describes resources within the chassis
  - Remote accelerators on a CXL fabric can also reside in Chassis
- PCIeDevice in Chassis describes CXL devices in the chassis
  - Fabric Manager can create Memory Chunks to be available for assignment
  - Assignment is done through the CXL Logical
    Device
- Devices in Chassis describes a remote resource
  - Shown here is a processor however other devices such as memory or I/O could also be represented here



## 

## **System Model for CXL**

- System describes a host that is attached to a CXL fabric
  - Ports of processor describes how the system is connected to the fabric
  - 'Memory' can contain remote memory that describes memory bound to the host over the fabric
- MemoryDomain in system can describe either locally attached CXL memory or remote CXL memory
- MemoryChunk in system describes a section of address space backed by local or remote memory
  - Can be interleaved or a contiguous range across local or remote devices
- Memory describes a logical memory device that represents the memory that is assigned to this host over the CXL Fabric





Attached directly to Computer System

## LOCAL CXL DEVICE MODELS















Attached to Computer System through Fabric

## **REMOTE CXL DEVICE MODELS**



#### **Remote CXL Type 1 Device Model**





Remote CXL Type 2 Device Model







13





## **EXAMPLE RESOURCES**



## **CXL Logical Device**

Remote CXL Logical Device





## **Remote Memory at Host**

Showing memory in the host from a Remote source





## **Memory Chunk**

Memory range accessible by Host

{  "@odata.id": "/redfish/v1/Chassis/PCXL3/MemorvDomains/1/MemorvChunks/1",	Memory chunk is in remote chassis PCXL3
"Id": "1",	
"Name": "Memory Chunk 1",	
"Description": "Memory chunk accessible through CXL fabric",	
"MemoryChunkSizeMiB": 4096,	Total Size of this Memory Chunk
"AddressRangeType": "PMEM",	Momony Type and offect within this CXL Device
"AddressRangeOffsetMiB": 1024,	Memory Type and onset within this CAL Device
"MediaLocation": "Local",	
"RequestedOperationalState": "Online",	
"Links": {	
"CXLLogicalDevices": [	Logical Device that is exporting this
{ "@odata.id": "/redfish/v1/Chassis/PCXL3/PCIeDevices/1/CXLLogicalDevic	Memory Chunk
{ "@odata.id": "/redfish/v1/Fabrics/CXL/Endpoints/T3" }	Endpoint of this device on the CXL Fabric
www.dmtf.org	



### **Additional Resources**

- More details of the CXL model can be found at the CXL Public Mockup on the DMTF Website <u>https://redfish.dmtf.org</u>
- The Fabrics Whitepaper provides details of the redfish fabrics model and example fabric types
  - <u>https://www.dmtf.org/sites/default/files/standards/documents/DSP2066</u>
    <u>1.0.0.pdf</u>



## Thank you!

#### For more information, visit us online at dmtf.org

Visit the Redfish Developer Hub at redfish.dmtf.org

