




Redfish CXL Device Management Models

November 23, 2021



Goals & Scope

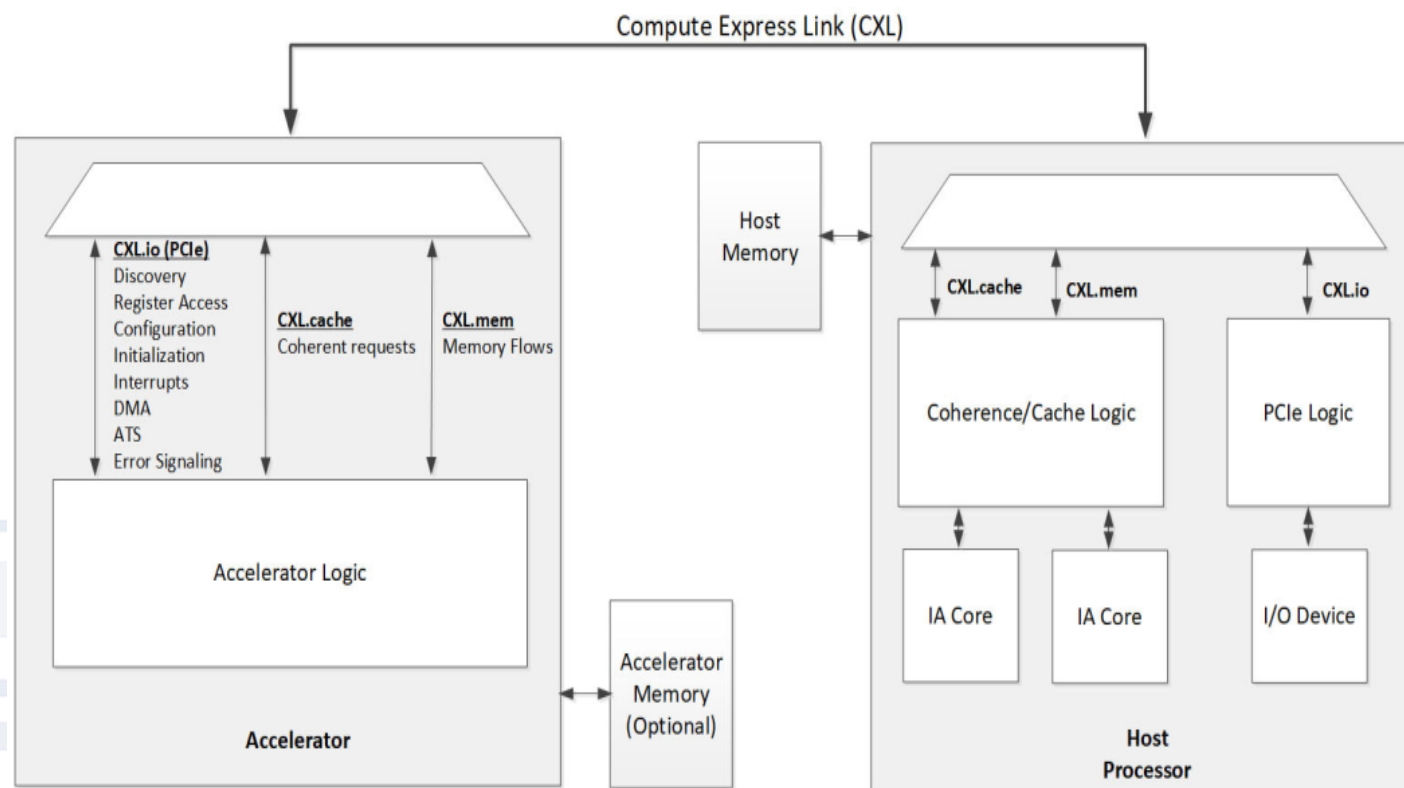
- Define Redfish management models for CXL devices
- Support all CXL device use cases
- Support local and fabric attached CXL devices
- Extend Redfish fabric management model with CXL devices support



Redfish



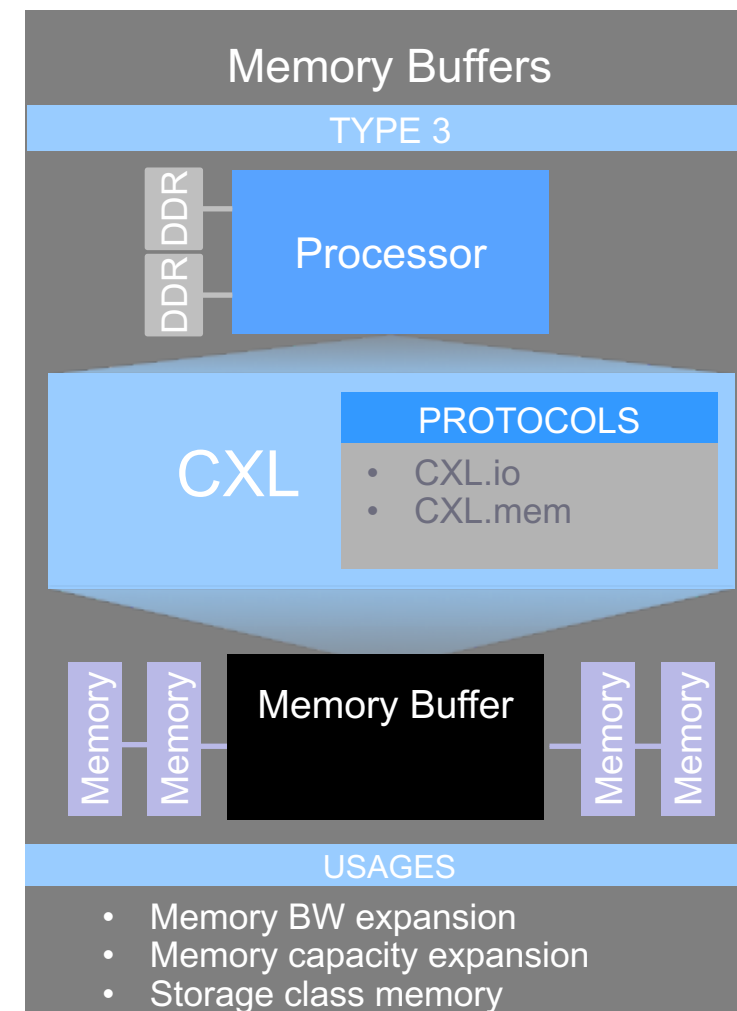

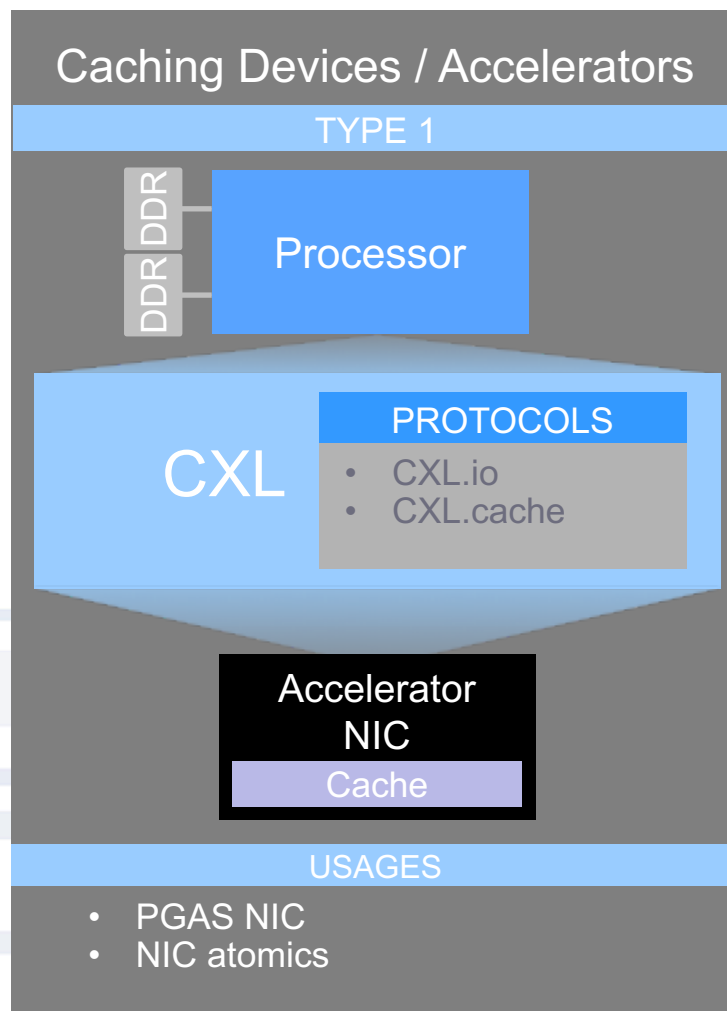
Compute Express Link



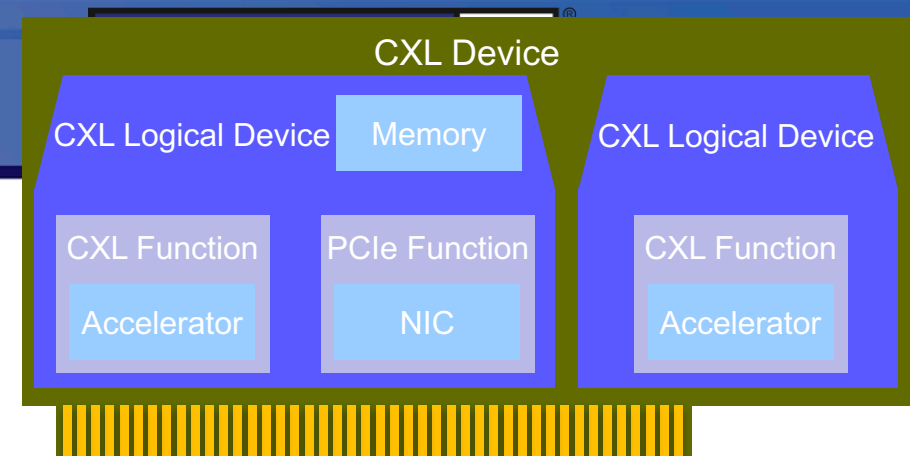
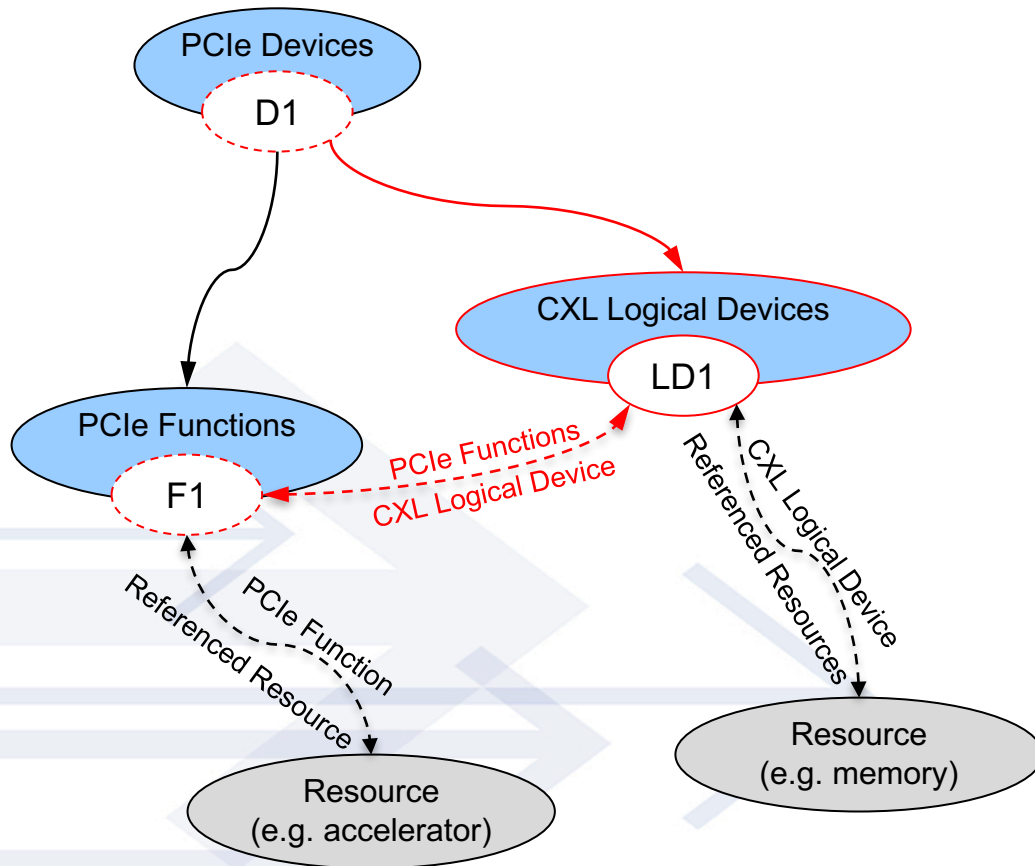
- Open industry standard for high bandwidth, low-latency interconnect
- Connectivity between host processor and accelerators/ memory device/ smart NIC
- Based on PCIe® 5.0 PHY infrastructure
 - Leverages channel, retimers, PHY, Logical, Protocols
 - Pin-to-pin compatible with PCIe connectors
 - CXL.io – 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



CXL Device Base Redfish Model




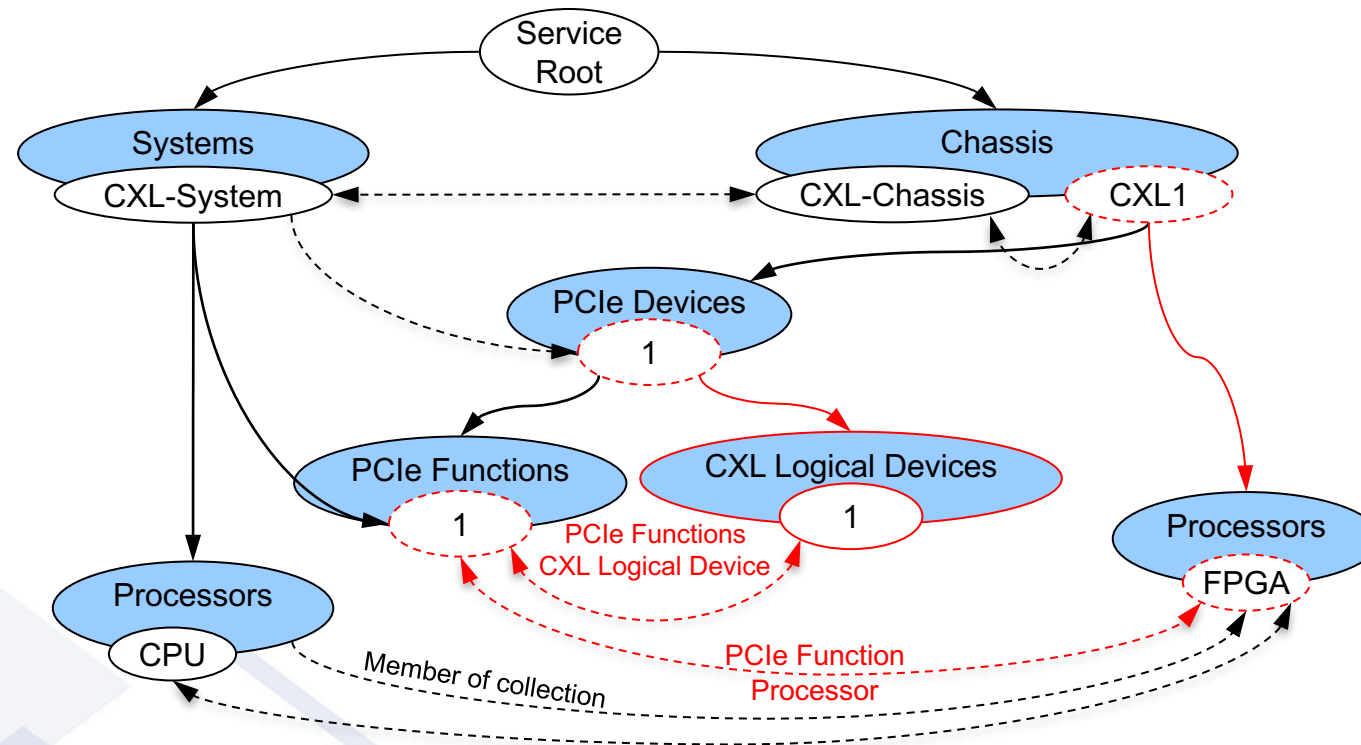
- CXL management model built on top of Redfish PCIe model
- CXL Device can provide standard PCIe Functions as well as extended functions supporting new CXL cache and memory semantics
- New CXL Logical Device allows partitioning of device resources and binding to different hosts through CXL switch
- PCIe 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



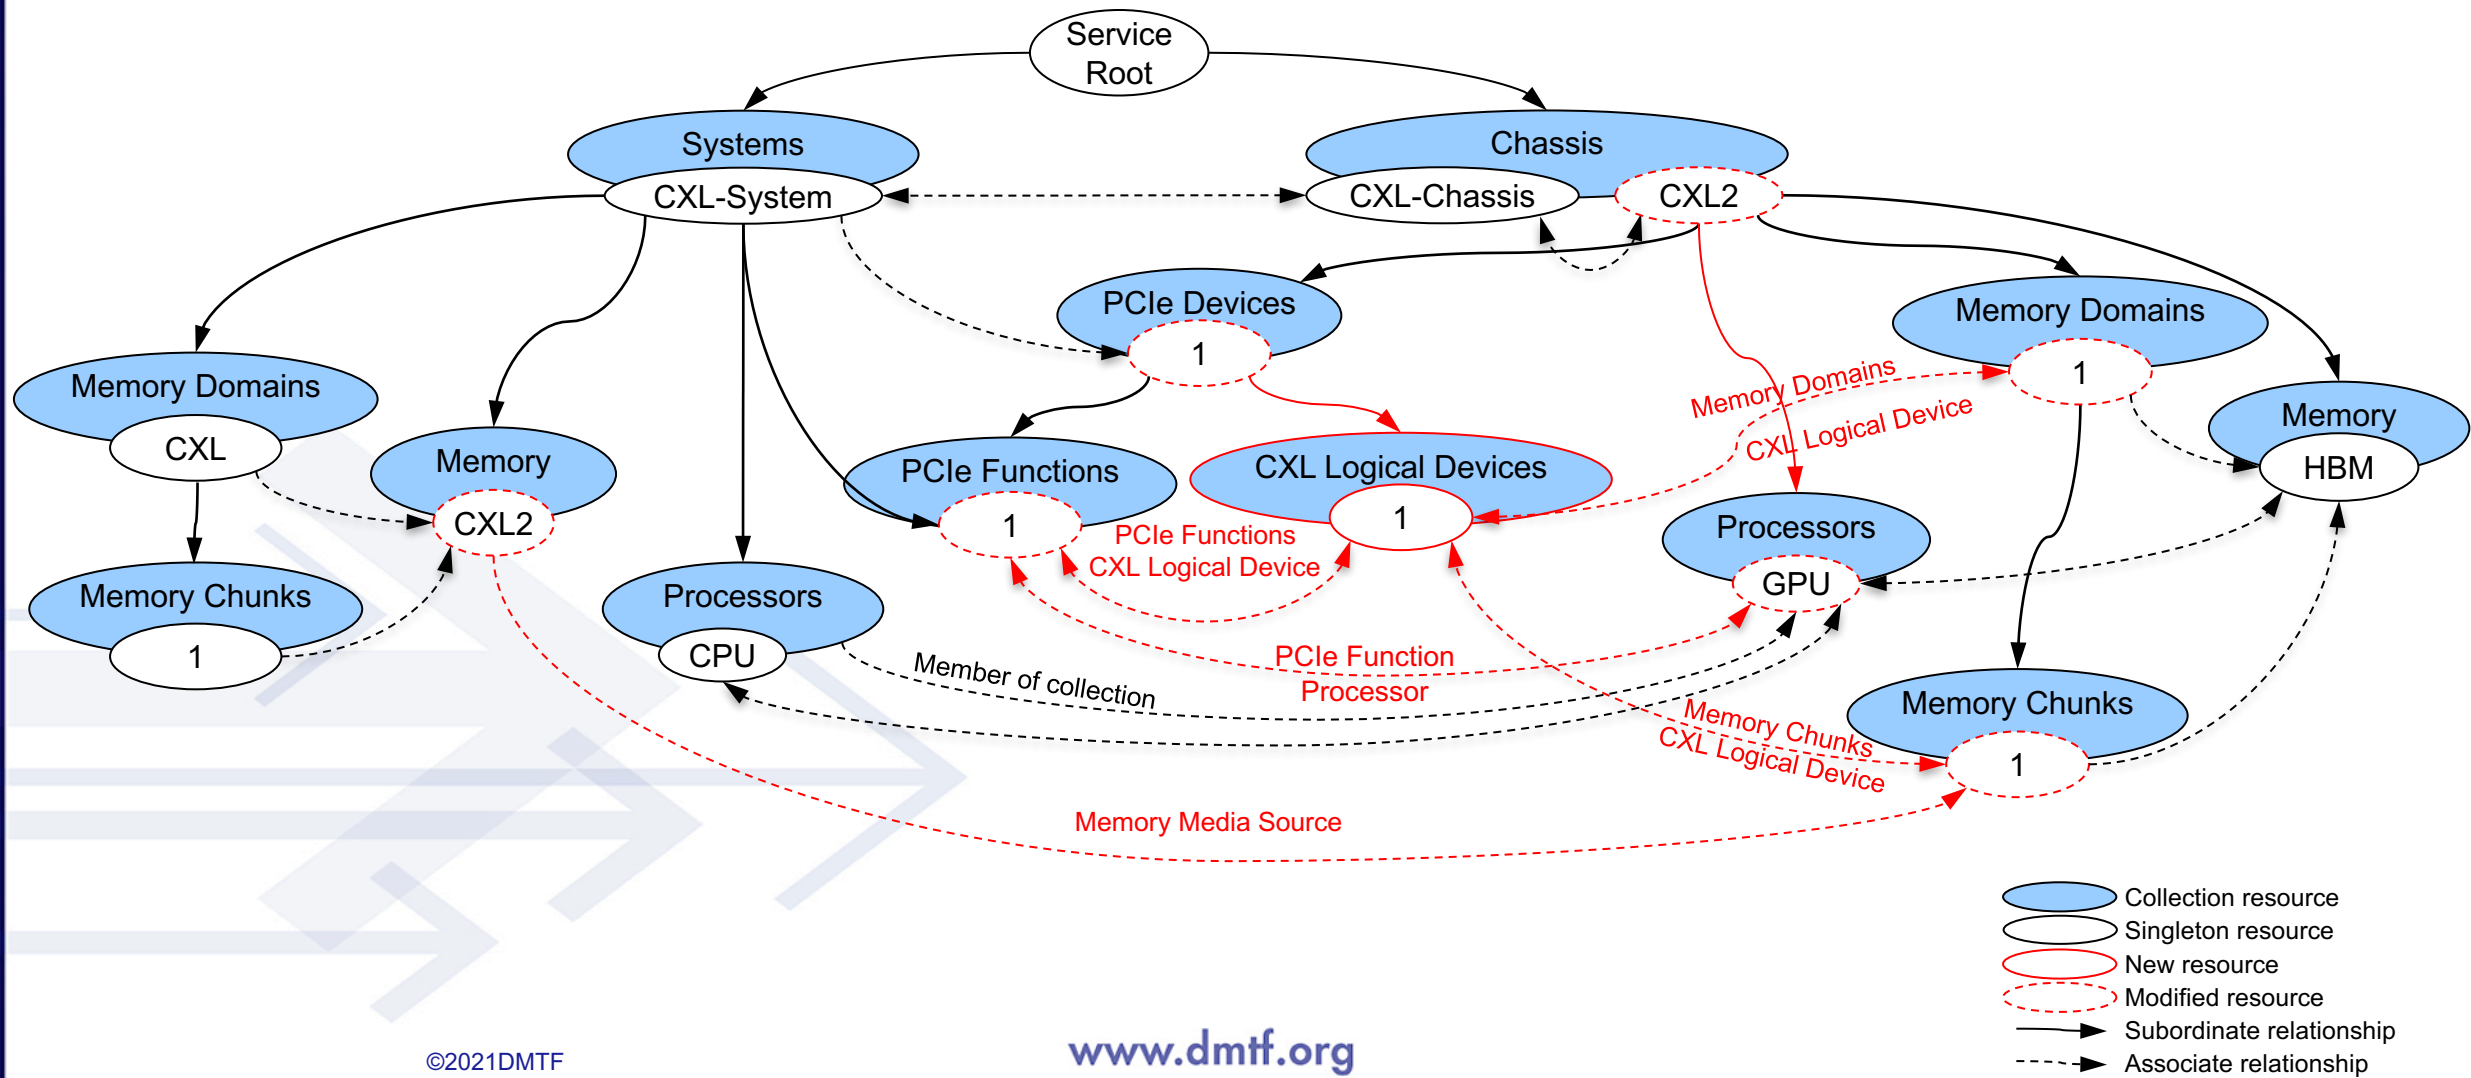
Attached directly to Computer System

LOCAL CXL DEVICE MODELS

Local CXL Type 1 Device Model

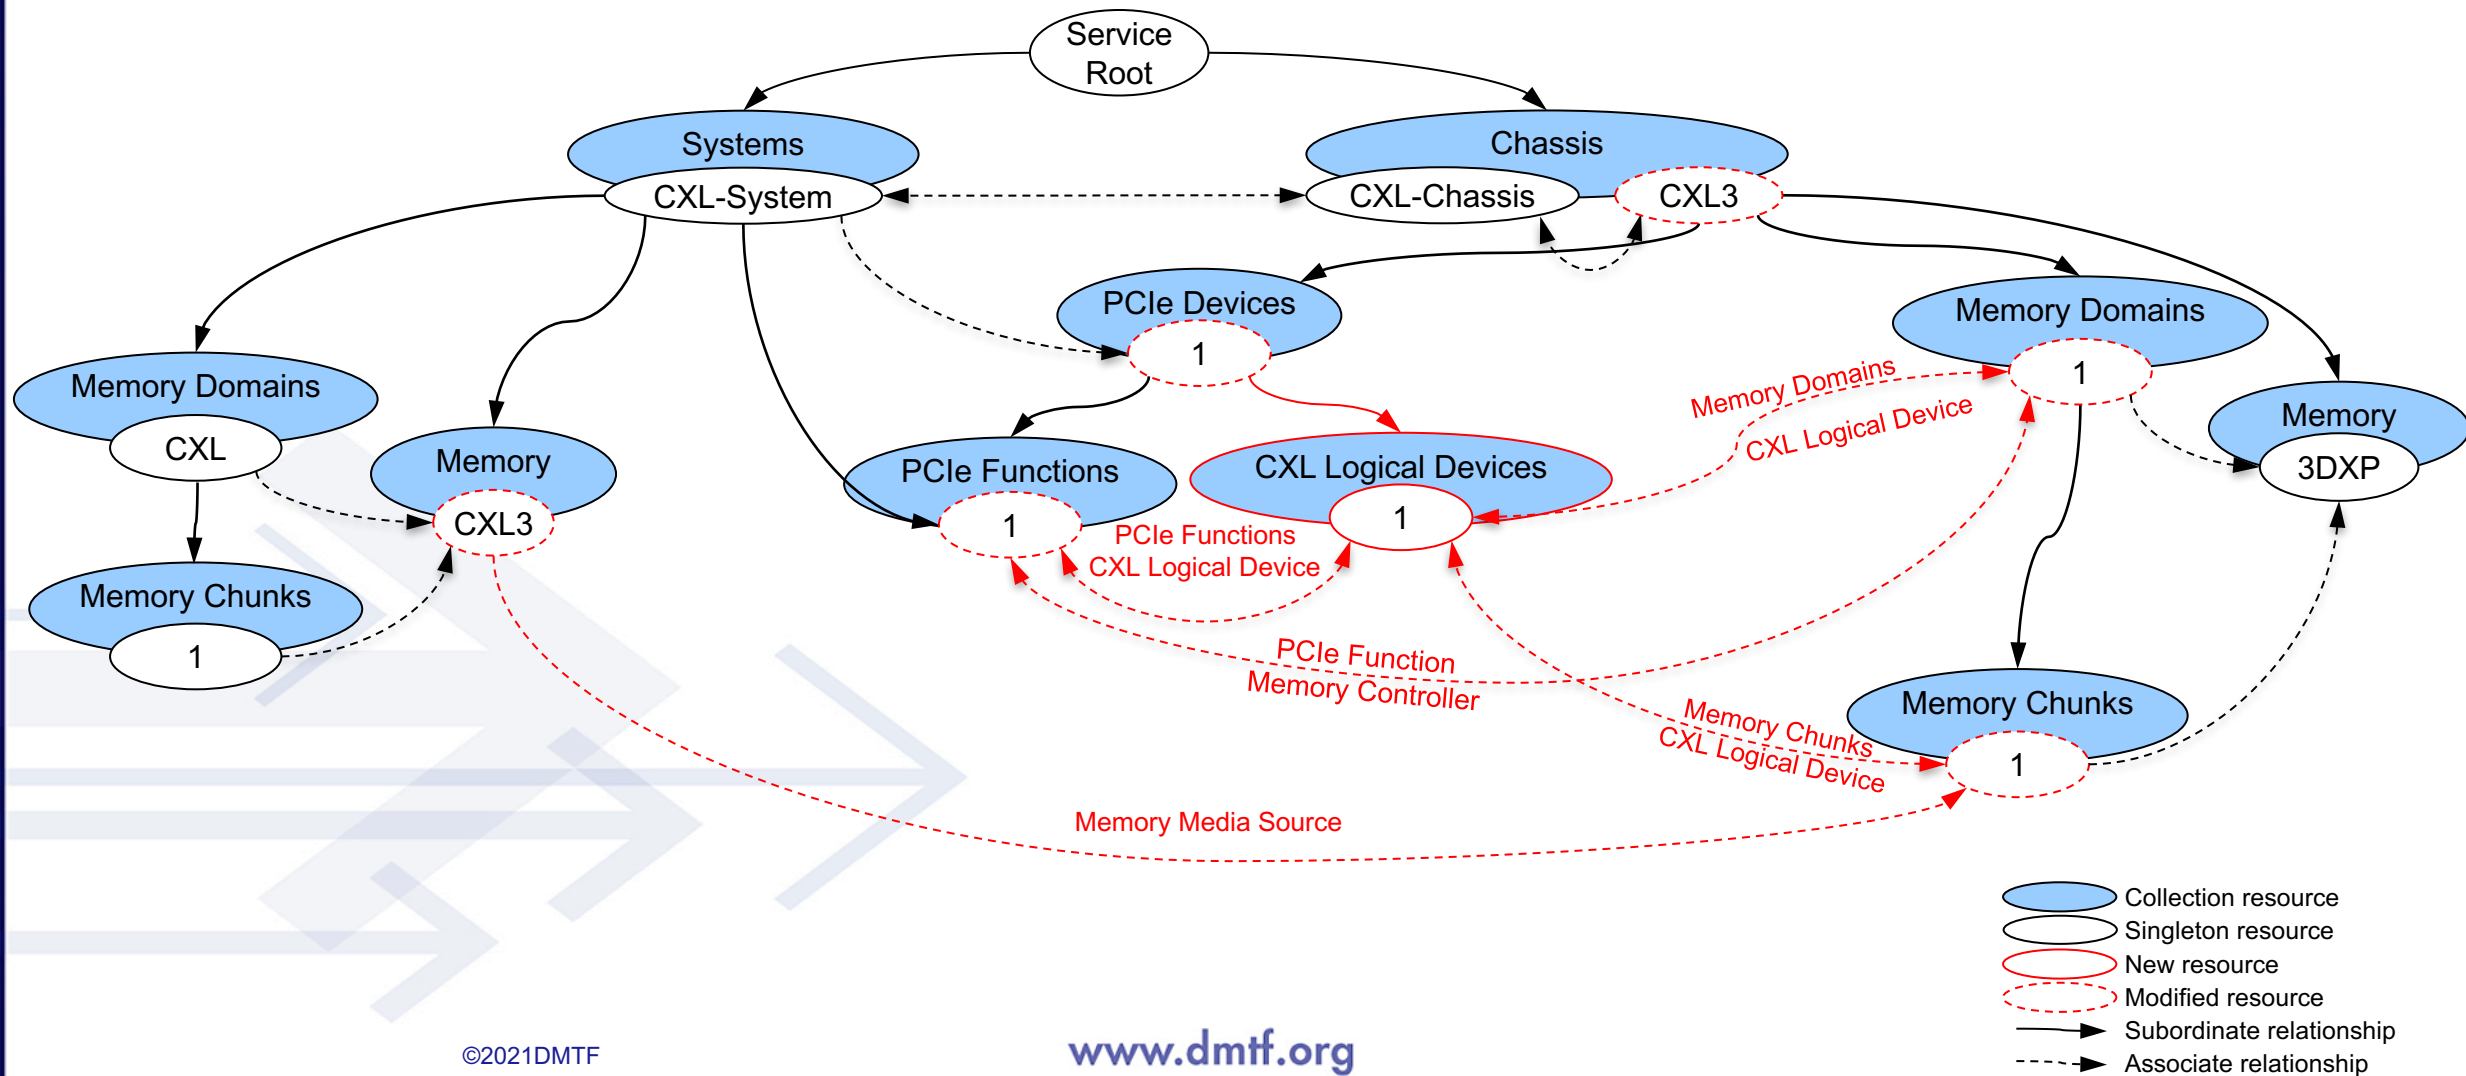


Local CXL Type 2 Device Model





Local CXL Type 3 Device Model



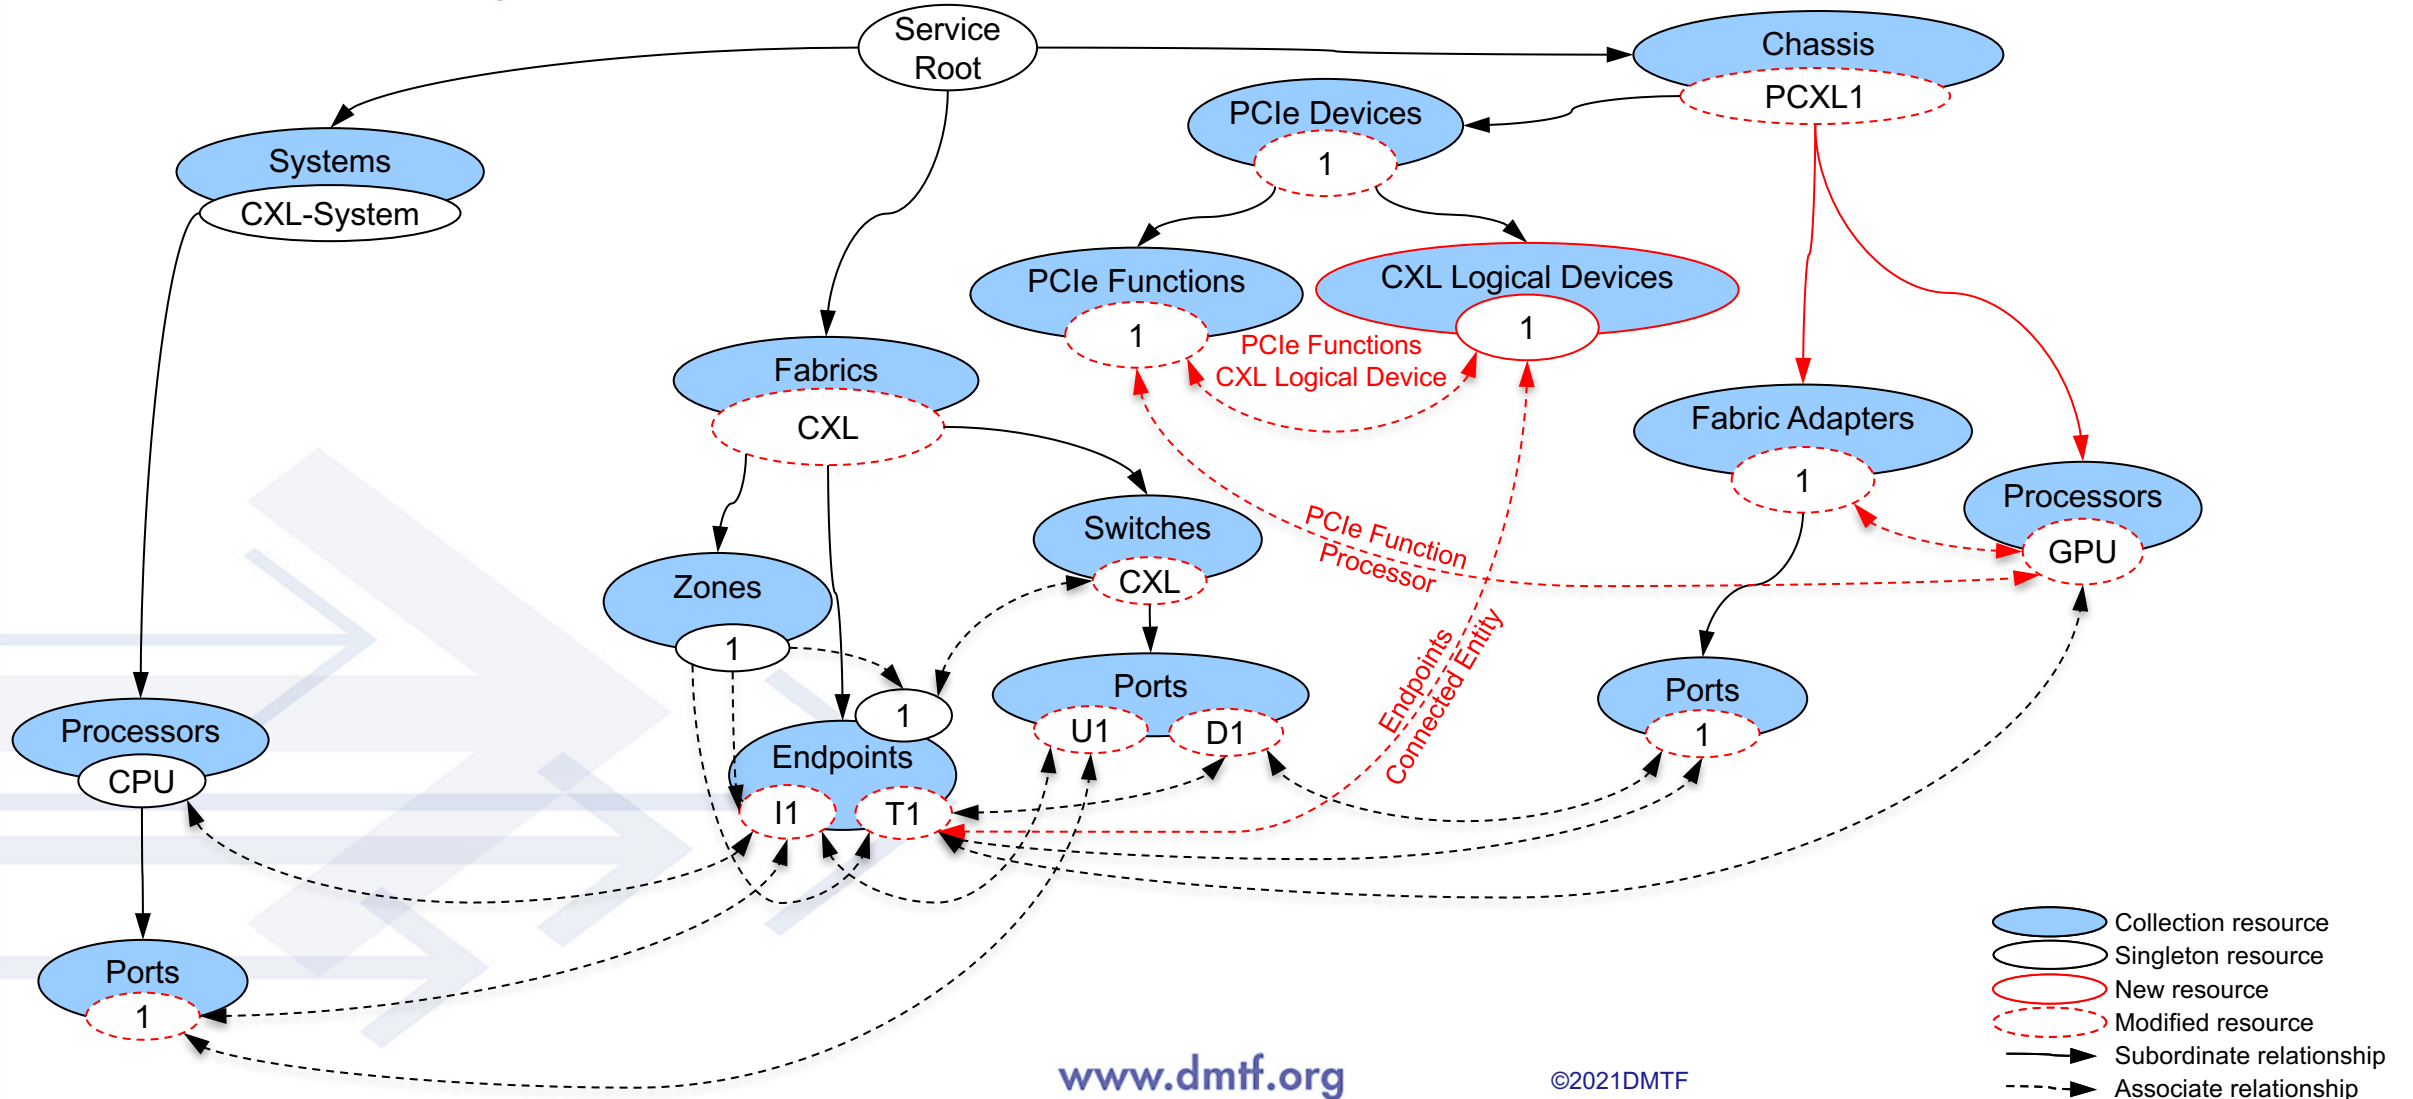


Attached to Computer System through Fabric

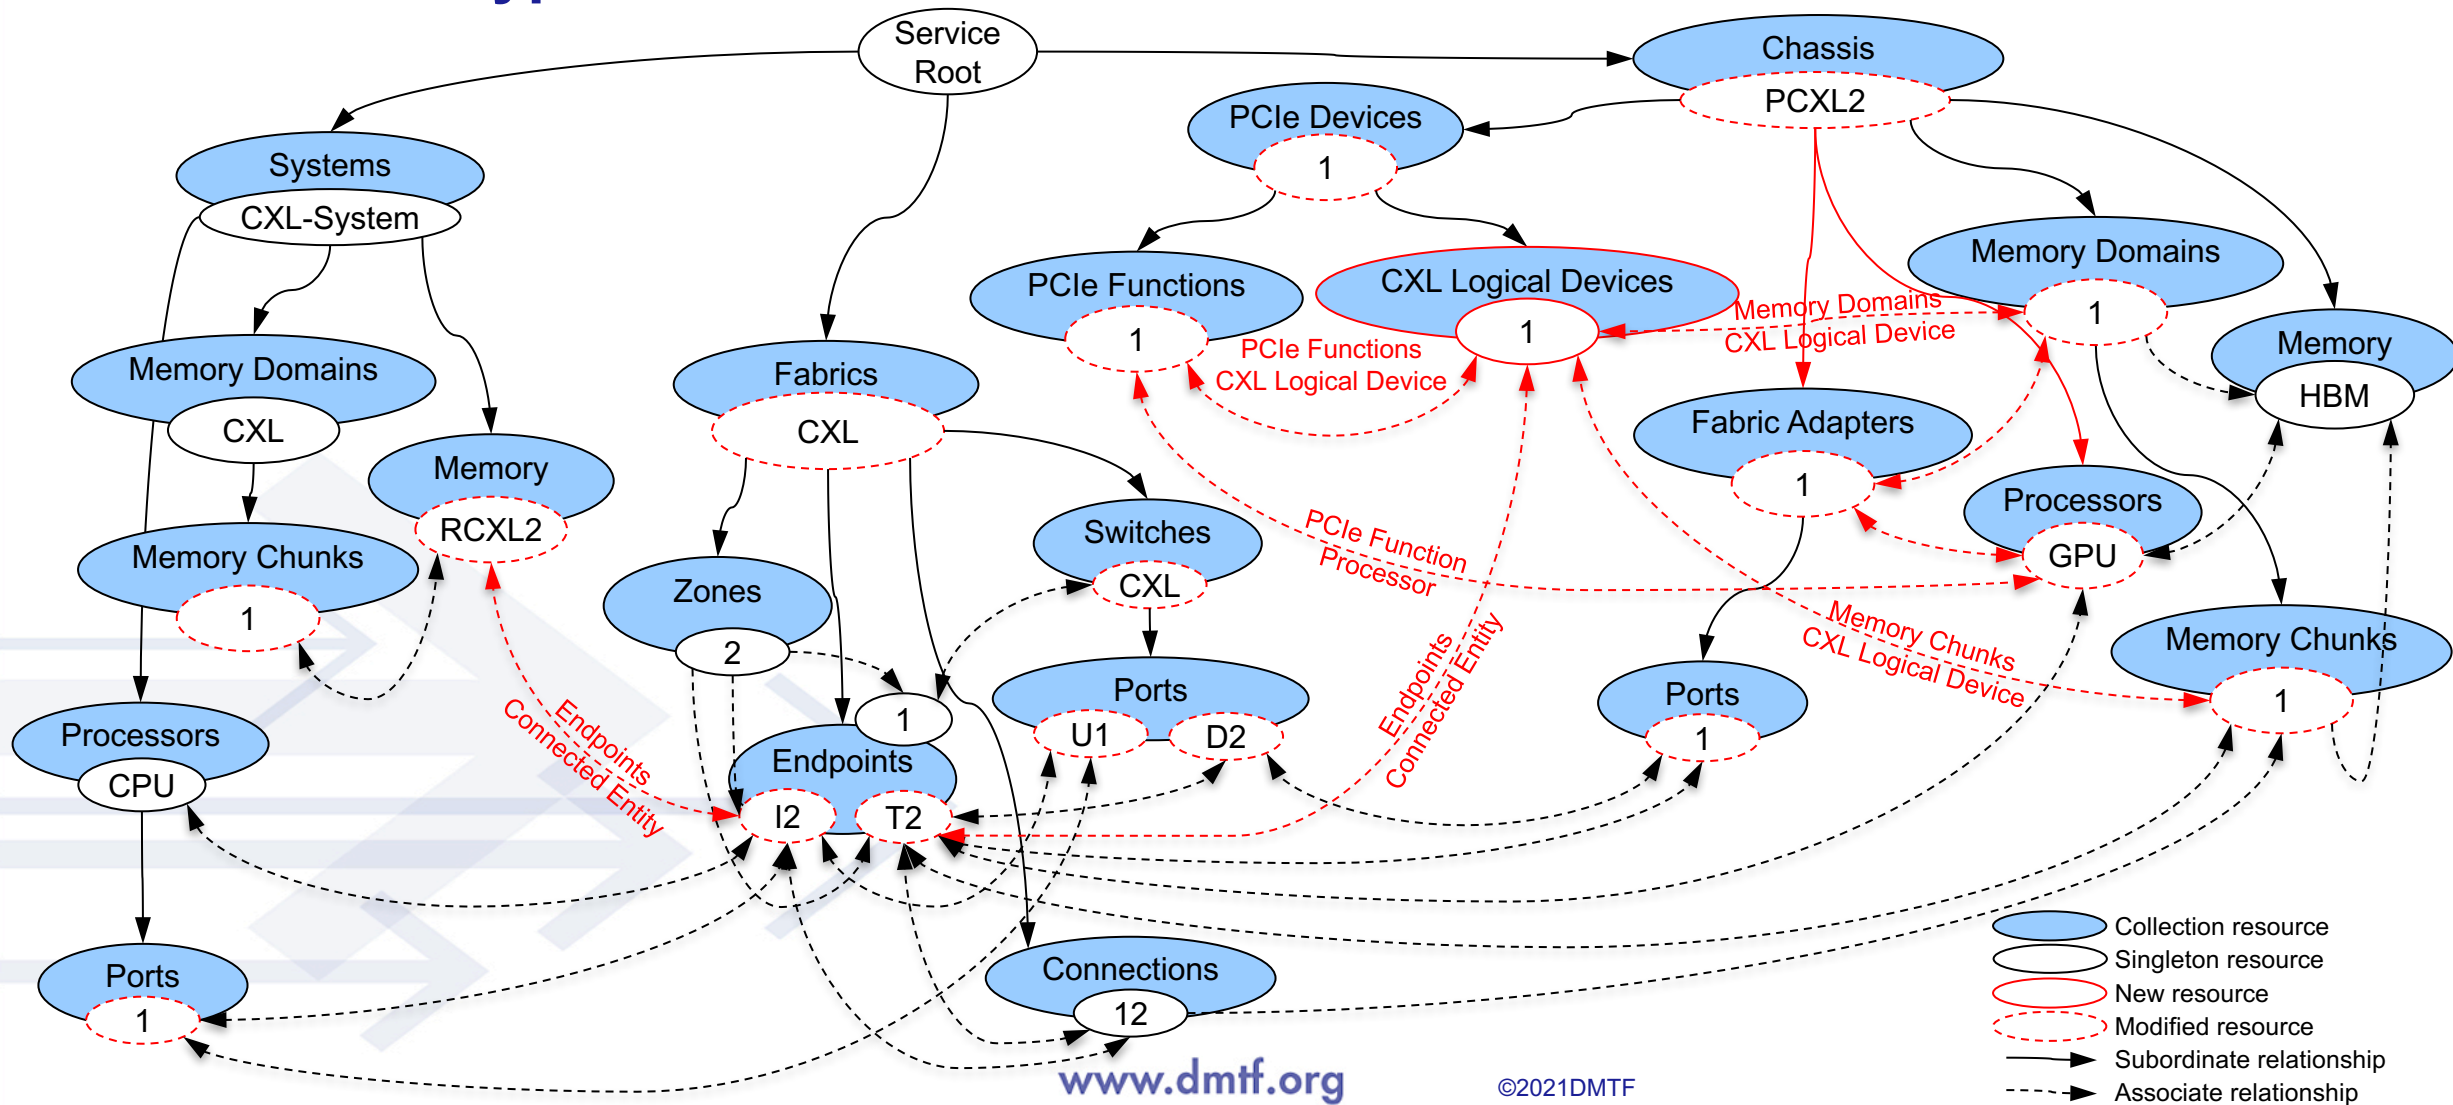
REMOTE CXL DEVICE MODELS



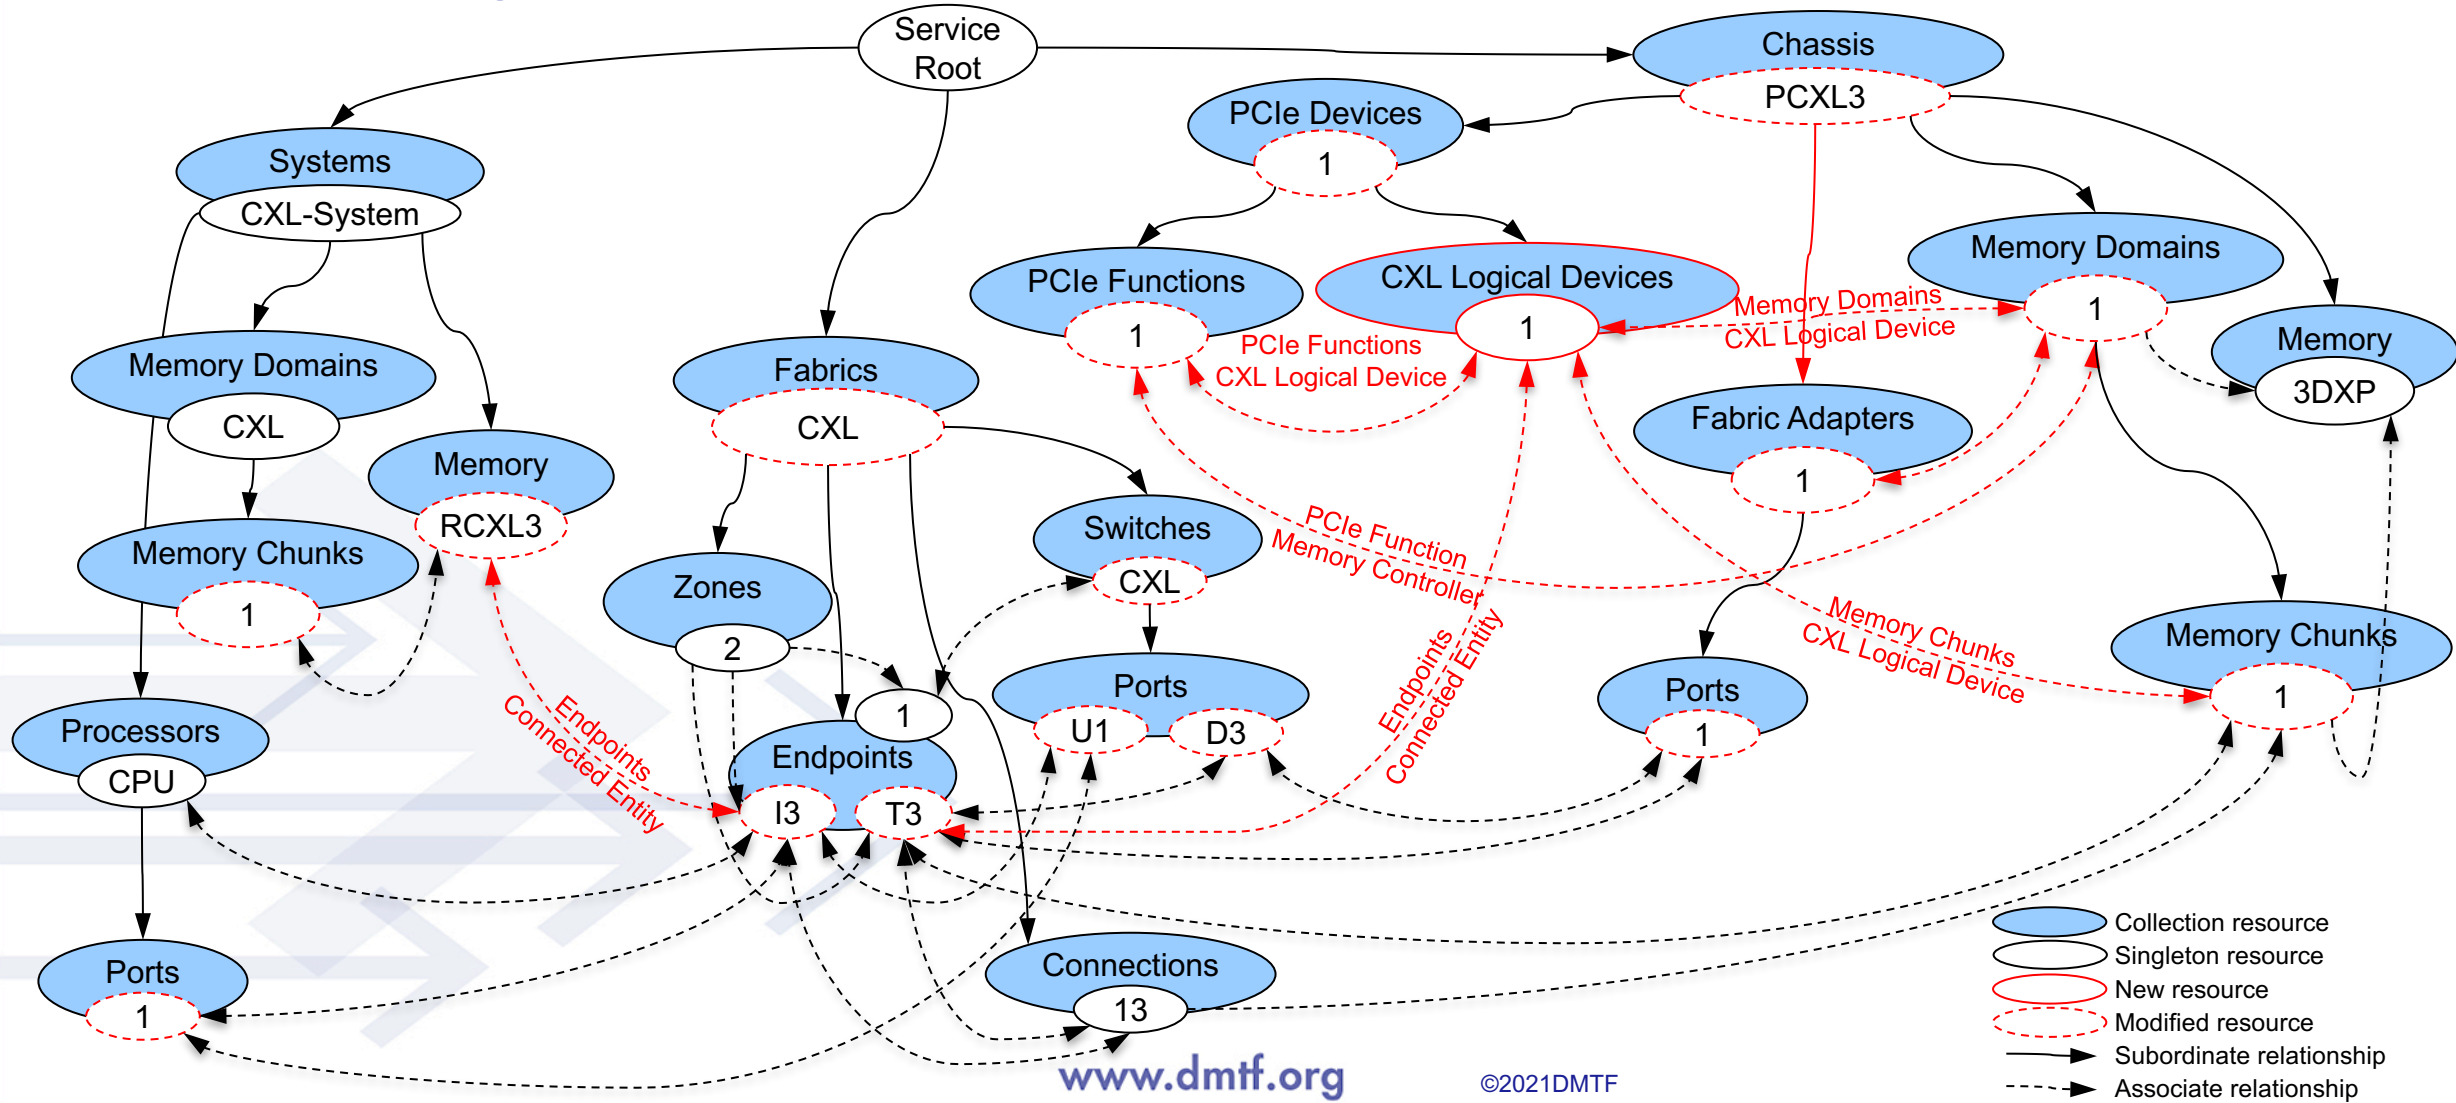
Remote CXL Type 1 Device Model



Remote CXL Type 2 Device Model



Remote CXL Type 3 Device Model





REDFISH SCHEMAS



Redfish Schemas Changes

New schemas:

- CXL Logical Device
- CXL Logical Device Collection

Modified schemas:

- PCIe Device
- PCIe Function
- Chassis
- Processor
- Port
- Fabric Adapter
- Memory
- Memory Domain
- Memory Chunks
- Endpoint
- Protocol
- Resource



PCIe Device Schema

- Existing schema updated with new properties
 - CXL protocol negotiated with root complex (computer system or switch)
 - CXL protocol supported by device
 - CXL device type (Type1, Type2 or Type3)
 - Flag indicating whether device is single or multi logical device
 - Link to collection of CXL Logical Devices contained in PCIe Device
- CXL model will use existing properties (e.g. PCIe interface properties and link to collection of PCIe Functions)

```
{
  "@odata.type": "#PCIeDevice.v1_10_0.PCIeDevice",
  "@odata.id": "/redfish/v1/Chassis/CXL/PCIeDevices/1",
  "Description": "CXL Device",
  "Id": "1",
  "CXLDevice": {
    "CurrentProtocolVersion": "v1_1 | v2_0",
    "CapableProtocolVersion": "v1_1 | v2_0",
    "DeviceType": "Type1 | Type2 | Type3",
    "MultiLogicalDevice": false
  },
  "CXLLogicalDevices": {
    "@odata.id": "/redfish/v1/Chassis/CXL/PCIeDevices/1/CXLLogicalDevices"
  }
}
```



PCIe Function Schema

- Existing schema updated with new properties
 - Function Protocol (PCIe or CXL)
 - Links to Processor resources that PCIe Function represents
 - Links to Memory Controller resources that PCIe Function represents
 - Link to CXL Logical Device that PCIe Function is assigned to
- CXL model will use existing properties (e.g. Device Class, Device ID and Vendor Id)

```
{
  "@odata.id": "/redfish/v1/Chassis/CXL/PCIeDevices/CXL/PCIeFunctions/1",
  "@odata.type": "#PCIeFunction.v1_4_0.PCIeFunction",
  "Description": "CXL Function",
  "Id": "1",
  "FunctionProtocol": "CXL | PCIe",
  "Links": {
    "Processors": [
      {
        "@odata.id": "/redfish/v1/Chassis/CXL/Processors/FPGA"
      }
    ],
    "MemoryControllers": [
      {
        "@odata.id": "/redfish/v1/Chassis/CXL/MemoryDomains/1"
      }
    ],
    "CXLLogicalDevice": {
      "@odata.id": "/redfish/v1/Chassis/CXL/PCIeDevices/1/CXLLogicalDevices/1"
    }
  }
}
```

CXL Logical Device Schema

- New schema defining CXL Logical Device with following properties
 - CXL Logical Device globally unique identifier
 - List of supported CXL semantics
 - CXLio – CXL.io I/O semantics similar to PCIe
 - CXLcache – CXL.cache caching protocol semantics
 - CXLmem – CXL.mem memory access semantics
 - Link to PCIe Device containing CXL Logical Device
 - Array of links to PCIe Functions assigned to CXL Logical Device
 - Array of links to Memory Domains owned by CXL Logical Device
 - Array of links to Memory Chunks referenced by CXL Logical Device
 - Array of links to Endpoints representing CXL Logical Device on fabrics

```
{
  "@odata.id": "/redfish/v1/Chassis/CXL/PCIDevices/1/CXLLogicalDevices/1",
  "@odata.type": "#CXLLogicalDevice.v1_0_0.CXLLogicalDevice",
  "Description": "CXL Logical Device",
  "Id": "1",
  "Name": "CXL Logical Device",
  "Status": {
    "State": "Enabled",
    "Health": "OK",
    "HealthRollup": "OK"
  },
  "Identifiers": [
    {
      "DurableName": "4C-1D-96-FF-FE-DD-D8-35:1",
      "DurableNameFormat": "GCXLID"
    }
  ],
  "CXLSemanticsSupported": [
    "CXLio",
    "CXLcache",
    "CXLmem"
  ],
  "Links": {
    "PCIeDevice": {
      "@odata.id": "/redfish/v1/Chassis/CXL/PCIDevices/1"
    },
    "PCIeFunctions": [
      {
        "@odata.id": "/redfish/v1/Chassis/CXL/PCIDevices/1/PCIeFunctions/1"
      }
    ],
    "MemoryDomains": [
      {
        "@odata.id": "/redfish/v1/Chassis/CXL/MemoryDomains/1"
      }
    ],
    "MemoryChunks": [
      {
        "@odata.id": "/redfish/v1/Chassis/CXL/MemoryDomains/1/MemoryChunks/1"
      }
    ],
    "Endpoints": [
      {
        "@odata.id": "/redfish/v1/Fabrics/CXL/Endpoints/1"
      }
    ]
  }
}
```



Chassis Schema

- Existing schema updated with new properties
 - Link to collection of Processors present in chassis
 - Link to collection of Fabric Adapters present in chassis
- CXL model will use existing properties (e.g. Memory, Memory Domains and Media Controllers)

```
{
  "@odata.type": "#Chassis.v1_20_0.Chassis",
  "@odata.id": "/redfish/v1/Chassis/CXL",
  "Description": "CXL Chassis",
  "Id": "CXL",
  "Processors": {
    "@odata.id": "/redfish/v1/Chassis/CXL/Processors"
  },
  "FabricAdapters": {
    "@odata.id": "/redfish/v1/Chassis/CXL/FabricAdapters"
  }
}
```

Processor Schema

- Existing schema update with new properties
 - Links to Fabric Adapter resources that processor uses
 - Port Collection long description annotation modifications
- CXL model will use existing properties (e.g. information about port collection, embedded memory, internal cache for example)

```
{
  "@odata.type": "#Processor.v1_15_0.Processor",
  "@odata.id": "/redfish/v1/Chassis/CXL/Processors/1",
  "Description": "Processor with CXL ports",
  "Id": "1",
  "Links": {
    "FabricAdapters": [
      {
        "@odata.id": "/redfish/v1/Chassis/CXL/FabricAdapters/1"
      }
    ]
  }
}
```




Port Schema

- Existing schema updated with new properties
 - Remote Port Id allowing physical connection topology building
- New Port Protocol and Link Network Technology value (defined in Protocol schema)
- CXL model will use existing properties (e.g. Port Id, Port Type, Speed, and Width)

```
{  
  "@odata.id": "/redfish/v1/Systems/CXL/FabricAdapters/CXL/Ports/1",  
  "@odata.type": "#Port.v1_6_0.Port",  
  "Description": "CXL Port",  
  "Id": "1",  
  "RemotePortId": "4C-1D-96-FF-FE-DD-D8-35",  
  "PortProtocol": "CXL",  
  "LinkNetworkTechnology": "PCIe"  
}
```



Memory Schema

- Existing schema updated with new properties
 - New Memory Media type values: “CXL” and “Fabric”
 - Array of links to Memory Media Sources
 - Array of links to Endpoints representing Memory on fabrics
- CXL model will use existing properties (e.g. capacity and rank count)

```
{
  "@odata.type": "#Memory.v1_15_0.Memory",
  "@odata.id": "/redfish/v1/Systems/CXL/Memory/1",
  "Description": "CXL device memory",
  "Id": "1",
  "MemoryMedia": [
    "CXL", "Fabric"
  ],
  "MemoryMediaSources": [
    {
      "@odata.id": "/redfish/v1/Chassis/CXL/MemoryDomains/CXL/MemoryChunks/1"
    }
  ],
  "Links": {
    "Endpoints": [
      {
        "@odata.id": "/redfish/v1/Fabrics/CXL/Endpoints/I1"
      }
    ]
  }
}
```



Memory Domain Schema

- Existing schema updated with new properties
 - Min Memory Chunk Size in domain
 - Memory Chunk Allocation Increment unit size in domain
 - Array of links to CXL Logical Device associated with Memory Domain
 - Array of links to PCIe Function representing Memory Domain
 - Array of links to Fabric Adapters
- CXL model will use existing properties (e.g. flag indicating whether creation of memory chunks is allowed and link to collection of memory chunks created in domain)

```
{
  "@odata.id": "/redfish/v1/Chassis/CXL/MemoryDomains/1",
  "@odata.type": "#MemoryDomain.v1_4_0.MemoryDomain",
  "Description": "Pooled Memory Domain accessible through CXL fabric",
  "Id": "1",
  "MinMemoryChunkSizeMiB": 4096,
  "MemoryChunkIncrementMiB": 4096,
  "Links": {
    "CXLLogicalDevices": [
      {
        "@odata.id": "/redfish/v1/Chassis/CXL/PCIeDevices/1/CXLLogicalDevices/1"
      }
    ],
    "PCIeFunctions": [
      {
        "@odata.id": "/redfish/v1/Chassis/CXL/PCIeDevices/1/PCIeFunctions/1"
      }
    ],
    "FabricAdapters": [
      {
        "@odata.id": "/redfish/v1/Chassis/CXL/FabricAdapters/1"
      }
    ]
  }
}
```



Memory Chunks Schema

- Existing schema updated with new properties
 - Memory Media Location
 - Operational State allowing graceful memory chunk disconnection
 - Array of links to CXL Logical Device associated with Memory Chunks
- CXL model will use existing properties (e.g. Memory Chunk Size and Address Range Offset)

```
{
  "@odata.type": "#MemoryChunks.v1_5_0.MemoryChunks",
  "@odata.id": "/redfish/v1/Chassis/CXL/MemoryDomains/1/MemoryChunks/1",
  "Description": "CXL memory chunk",
  "Id": "1",
  "Status": {
    "State": "Enabled | StandbyOffline",
    "Health": "OK",
    "HealthRollup": "OK"
  },
  "MediaLocation": "Local | Remote | Mixed",
  "OperationalState": "Online | Offline",
  "Links": {
    "CXLLogicalDevices": [
      {
        "@odata.id": "/redfish/v1/Chassis/CXL/PCIeDevices/1/CXLLogicalDevices/1"
      }
    ]
  }
}
```



Endpoint Schema

- Existing schema updated with new properties
 - New Entity Type value - Memory
 - New Entity Type value – CXL Device
- New Durable Name Format for CXL Logical Device (defined in Resource schema)

```
{
  "@odata.type": "#Endpoint.v1_7_0.Endpoint",
  "@odata.id": "/redfish/v1/Fabrics/CXL/Endpoints/T1",
  "Description": "CXL Logical Device endpoint",
  "Id": "T1",
  "EndpointProtocol": "CXL",
  "ConnectedEntities": [
    {
      "EntityType": "CXLDevice",
      "EntityRole": "Target",
      "Identifiers": [
        {
          "DurableName": "4C-1D-96-FF-FE-DD-D8-35:1",
          "DurableNameFormat": "GCXLID"
        }
      ],
      "EntityLink": {
        "@odata.id": "/redfish/v1/Chassis/PCIEDevices/1/CXLLogicalDevices/1"
      }
    },
    {
      "EntityType": "Memory",
      "EntityRole": "Initiator",
      "EntityLink": {
        "@odata.id": "/redfish/v1/Systems/CXL/Memory/1"
      }
    }
  ]
}
```



Fabric Adapter Schema

- Existing schema updated with new properties to replace Media Controller
 - Array of links to Memory Domain associated with Fabric Adapter
 - Array of links to Processor associated with Fabric Adapter
- CXL model will use existing properties (e.g. Port Collection and Endpoint links)

```
{
  "@odata.id": "/redfish/v1/Chassis/CXL/FabricAdapters/1",
  "@odata.type": "#FabricAdapter.v1_3_0.FabricAdapter",
  "Description": "Fabric Adapter allowing access through CXL fabric",
  "Id": "1",
  "Links": {
    "Processors": [
      {
        "@odata.id": "/redfish/v1/Chassis/CXL/Processors/1"
      }
    ],
    "MemoryDomains": [
      {
        "@odata.id": "/redfish/v1/Chassis/CXL/MemoryDomains/1"
      }
    ]
  }
}
```




Protocol Schema

- Existing schema updated with new protocol – CXL
- Used by Fabric, Switch and Port schemas

```
{
  "$id": "http://redfish.dmtf.org/schemas/v1/Protocol.json",
  "$schema": "http://redfish.dmtf.org/schemas/v1/redfish-schema-v1.json",
  "definitions": {
    "Protocol": {
      "enum": [
        "CXL"
      ],
      "enumDescriptions": {
        "CXL": "Compute Express Link."
      },
      "enumLongDescriptions": {
        "CXL": "This value shall indicate conformance to the Compute Express Link Consortium Specification."
      }
    }
  }
}
```



Resource Schema

- Existing schema updated with new durable name format definition – GCXLID
- Used by CXL Logical Device and Endpoint schemas

```
{
  "$id": "http://redfish.dmtf.org/schemas/v1/Resource.v1_15_0.json",
  "$schema": "http://redfish.dmtf.org/schemas/v1/redfish-schema-v1.json",
  "definitions": {
    "DurableNameFormat": {
      "enum": [
        "GCXLID"
      ],
      "enumDescriptions": {
        "GCXLID": "The Globally Unique CXL Logical Device Identifier (GCXLID).",
      },
      "enumLongDescriptions": {
        "GCXLID": "This durable name shall be in the Globally Unique CXL Logical Device Identifier (GCXLID) defined as {PCIe Serial Number}:{CXL Logical Device Identifier}. The DurableName property shall follow the pattern '^([0-9A-Fa-f]{2}[:-]){8}([0-9]{2})$', where the most significant octet is first."
      },
      "enumVersionAdded": {
        "GCXLID": "v1_15_0"
      }
    }
  }
}
```