

# **Universal Serial Bus 4 (USB4®) Connection Manager Guide**

**Version 2.0**

**November 2022**

## Release History

Version	Comments	Issue Date
1.0	First release	April 2021
2.0	<p>Aligned with USB4 v2 features, changes include</p> <ul style="list-style-type: none"><li>• 80G lane initialization</li><li>• Asymmetric link configuration</li><li>• Bandwidth management</li><li>• TMU adaptive mode</li><li>• DP2.1 tunnel support</li><li>• Host Router reset</li><li>• PCIe extended encapsulation</li><li>• Completion timeouts</li><li>• Arbitration configuration</li><li>• CLx support</li><li>• In band PCIe wake</li><li>• Error Handling</li><li>• New Notifications</li><li>• DP Rx and DP Tx discovery</li><li>• Vendor Specific Adapter Path Setup and Teardown</li><li>• GenT support</li><li>• USB3 Gen X bandwidth negotiation</li><li>• TMU configuration for Enhanced Uni-Directional Mode</li><li>• Bandwidth Allocation during Path Setup</li></ul>	November 2022

**NOTE:** Adopters may only use this USB specification to implement USB or third party functionality as expressly described in this Specification; all other uses are prohibited.

**LIMITED COPYRIGHT LICENSE:** The Promoters grant a conditional copyright license under the copyrights embodied in this USB Specification to use and reproduce the Specification for the sole purpose of, and solely to the extent necessary for, evaluating whether to implement the Specification in products that would comply with the specification. Without limiting the foregoing, use of the Specification for the purpose of filing or modifying any patent application to target the Specification or USB compliant products is not authorized. Except for this express copyright license, no other rights or licenses are granted, including without limitation any patent licenses. In order to obtain any additional intellectual property licenses or licensing commitments associated with the Specification a party must execute the USB Adopters Agreement. **NOTE:** By using the Specification, you accept these license terms on your own behalf and, in the case where you are doing this as an employee, on behalf of your employer.

#### **INTELLECTUAL PROPERTY DISCLAIMER**

THIS SPECIFICATION IS PROVIDED TO YOU “AS IS” WITH NO WARRANTIES WHATSOEVER INCLUDING ANY WARRANTY OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE. THE AUTHORS OF THIS SPECIFICATION DISCLAIM ALL LIABILITY FOR INFRINGEMENT OF ANY PROPRIETARY RIGHTS, RELATING TO THE USE OR IMPLEMENTATION OF INFORMATION IN THIS SPECIFICATION. THE PROVISION OF THIS SPECIFICATION TO YOU DOES NOT PROVIDE YOU WITH ANY LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS.

Please send comments to [techsup@usb.org](mailto:techsup@usb.org).

For industry information, refer to the USB Implementers Forum web page at <http://www.usb.org>.

USB Type-C®, USB-C®, USB 2.0 Type-C™, and USB4® are trademarks of the Universal Serial Bus Implementers Forum (USB-IF). DisplayPort™ is a trademark of VESA. All product names are trademarks, registered trademarks, or service marks of their respective owners.

Thunderbolt™ is a trademark of Intel Corporation. You may only use the Thunderbolt™ trademark or logo in conjunction with products designed to this specification that complete proper certification and executing a Thunderbolt™ trademark license – see [usb.org/compliance](http://usb.org/compliance) for further information.

Copyright © 2022, USB Promoter Group (Apple Inc., HP Inc., Intel Corporation, Microsoft Corporation, Renesas Corporation, STMicroelectronics, and Texas Instruments).

## Table of Contents

<b><u>1. INTRODUCTION</u></b>	<b>8</b>
1.1. SCOPE	8
1.2. REFERENCE DOCUMENTS	8
1.3. OVERVIEW	8
<b><u>2. HOST INTERFACE INITIALIZATION</u></b>	<b>10</b>
2.1. RING 0 INITIALIZATION	10
2.2. CONNECTION MANAGER TRANSMIT FLOW	11
2.3. CONNECTION MANAGER RECEIVE FLOW	11
2.4. HOST INTERFACE RESET	12
2.5. DROM ACCESS	12
2.6. PCIe MEMORY BAR ACCESS	12
<b><u>3. ROUTER CONFIGURATION</u></b>	<b>13</b>
3.1. CONNECTION EVENT DETECTION	13
3.1.1. USB4 CONNECTION EVENTS	13
3.1.2. DISPLAY CONNECTION EVENTS	13
3.2. ROUTER INITIALIZATION	14
3.2.1. ROUTER ENUMERATION AND CONFIGURATION	14
3.2.2. ADAPTER ENUMERATION	17
3.2.3. ENABLING PROTOCOL TUNNELING	18
3.3. IDENTIFYING LOOPS AND INTER-DOMAIN LINKS	19
3.4. DETECTING A REMOVAL EVENT	22
3.4.1. LANE ADAPTER	22
3.4.2. DP ADAPTER	23
3.5. HOST ROUTER RESET	23
<b><u>4. CONTROL PACKETS</u></b>	<b>25</b>
4.1. ROUTING	25
4.2. PROTOCOL RULES	25
4.3. HOT PLUG	25
<b><u>5. PATH CONFIGURATION</u></b>	<b>27</b>
5.1. FLOW CONTROL	28
5.1.1. FLOW CONTROL DISABLED	29
5.1.2. DEDICATED FLOW CONTROL	29
5.1.3. SHARED FLOW CONTROL	29
5.1.4. BUFFER ALLOCATION	29
5.2. PATH SETUP AND TEARDOWN	31
5.2.1. PATH SETUP	31
5.2.2. PATH TEARDOWN	33

<b>5.3. PCIe PATH SETUP AND TEARDOWN</b>	<b>33</b>
5.3.1. PCIe PATH SETUP	33
5.3.2. PCIe PATH TEARDOWN	35
<b>5.4. DISPLAYPORT PATH SETUP AND TEARDOWN</b>	<b>35</b>
5.4.1. CONNECTION MANAGER DISCOVERY	35
5.4.2. PAIRING DP ADAPTERS	42
5.4.3. DISPLAY PORT PATH SETUP AND TEARDOWN	43
<b>5.5. USB3 GEN X PATH SETUP AND TEARDOWN</b>	<b>48</b>
5.5.1. USB3 GEN X PATH SETUP	48
5.5.2. USB3 GEN X PATH TEARDOWN	50
5.5.3. USB3 GEN X ADAPTER DISCONNECT	50
<b>5.6. USB3 GEN T PATH SETUP AND TEARDOWN</b>	<b>50</b>
5.6.1. USB3 GEN T PATH SETUP	50
5.6.2. USB3 GEN T PATH TEARDOWN	53
<b>5.7. INTER-DOMAIN PATH SETUP AND TEARDOWN</b>	<b>54</b>
5.7.1. TRANSMIT RING SETUP IN A PCIe HOST INTERFACE ADAPTER LAYER	55
5.7.2. RECEIVE RING SETUP IN A PCIe HOST INTERFACE ADAPTER LAYER	56
5.7.3. RING DISABLE IN A PCIe HOST INTERFACE ADAPTER LAYER	57
<b>5.8. VENDOR SPECIFIC ADAPTER PATH SETUP AND TEARDOWN</b>	<b>57</b>
<b>5.9. PATH COUNTERS</b>	<b>57</b>
<b>6. BANDWIDTH MANAGEMENT</b>	<b>59</b>
<b>6.1. BANDWIDTH ALLOCATION</b>	<b>59</b>
6.1.1. SCHEDULING PRIORITIES	59
6.1.2. BANDWIDTH ALLOCATION	60
<b>6.2. BANDWIDTH CLIENTS</b>	<b>63</b>
6.2.1. DISPLAYPORT TUNNELING	63
6.2.2. USB3 GEN X TUNNELING	67
6.2.3. USB3 GEN T TUNNELING	69
6.2.4. PCIe TUNNELING	70
6.2.5. HOST-TO-HOST TUNNELING	71
<b>7. USB4 PORT MANAGEMENT</b>	<b>72</b>
<b>7.1. LANE BONDING (GEN 2 OR GEN 3 ONLY)</b>	<b>72</b>
<b>7.2. ASYMMETRIC LINK (VERSION 2.0 ROUTER)</b>	<b>73</b>
7.2.1. TRANSITION FROM SYMMETRIC LINK TO ASYMMETRIC LINK	74
7.2.2. TRANSITION FROM ASYMMETRIC LINK TO SYMMETRIC LINK	75
<b>7.3. DOWNSTREAM PORT RESET AND CHANGE OF LINK PARAMETERS</b>	<b>76</b>
<b>7.4. TIME SYNCHRONIZATION</b>	<b>76</b>
7.4.1. TMU MODES	77
7.4.2. ROUTER ACCURACY LEVEL REQUIREMENTS	77
7.4.3. TMU CONFIGURATION	77
<b>7.5. LANE DISABLE AND ENABLE</b>	<b>82</b>
7.5.1. INTRA-DOMAIN LINKS	83

7.5.2. INTER-DOMAIN LINKS	84
<b>7.6. LINK ERRORS</b>	<b>84</b>
<b>7.7. NOTIFICATIONS</b>	<b>84</b>
7.7.1. NOTIFICATION ENABLE	84
7.7.2. NOTIFICATION ACKNOWLEDGEMENT	85
<b>8. POWER MANAGEMENT</b>	<b>86</b>
<b>8.1. LOW-POWER STATES</b>	<b>86</b>
<b>8.2. SLEEP AND WAKE</b>	<b>87</b>
8.2.1. ENTRY TO SLEEP STATE	87
8.2.2. EXIT FROM SLEEP STATE	89
8.2.3. BEHAVIOR WITH A NEIGHBOR DOMAIN IN SLEEP STATE	90
<b>9. OPERATIONS</b>	<b>93</b>
<b>9.1. ROUTER OPERATIONS</b>	<b>93</b>
<b>9.2. PORT OPERATIONS</b>	<b>93</b>

## Figures

Figure 1-1: High-Level PCIe SW Architecture with a Connection Manager	9
Figure 3-1: Flow for Identifying a Loop in a Domain	21
Figure 3-2: Flow for identifying an Inter-Domain Link	22
Figure 4-1: TopologyID Format	25
Figure 5-1: Path Segments	28
Figure 5-2: PCIe Tunneling Scheme	34
Figure 5-3: Example - Full and Partial Implementation of DPRX Discovery Feature	37
Figure 5-4: Example - Virtual AUX Tunneled Path for Full DPRX Implementation	39
Figure 5-5: Example - Virtual AUX Tunneled Path for Partial DPRX Implementation	40
Figure 5-6: Example topology with DPTX discovery feature	41
Figure 5-7: DPTX Discovery flow	42
Figure 5-8: DP Tunneling Setup	45
Figure 5-9: USB3 Gen X Tunneling Setup	49
Figure 5-10: USB3 Gen T Tunneling Setup	51
Figure 5-11: Inter-Domain Tunneling Scheme	54
Figure 6-1: Example - WRR, Bundle Weight and BW Allocation	63
Figure 6-2: DP BW Allocation Interaction with Connection Manager	67
Figure 6-3. Bandwidth Negotiation by the Connection Manager	69
Figure 7-1: Symmetric Link	74
Figure 7-2: Asymmetric Link	74
Figure 8-1: Monitoring Sleep State Status in a Neighbor Domain	91

## Tables

Table 5-1: PCIe Path Attributes	35
Table 5-2: DP Path Attributes	47
Table 5-3: USB3 Gen X Path Attributes	50
Table 5-4: USB3 Gen T Path Attributes	53
Table 6-1: Tunneled Protocol Path Priority (Recommended)	60
Table 6-2: USB4 Raw Bandwidth	60

Table 6-3: DisplayPort Required Bandwidth (Gbps).....	64
Table 7-1: TMU Modes .....	77
Table 7-2: TMU Mode Configuration .....	78
Table 7-3: TMU Mode Parameters .....	78
Table 8-1: Wake Enable Bits .....	87
Table 8-2: Wake Status Bits.....	89
Table 8-3: Monitoring Sleep state status in a neighbor Domain.....	91

## **1. Introduction**

### **1.1. Scope**

The USB4 Connection Manager Guide is an informative document. It is intended to give the reader an understanding of how a Connection Manager (CM) functions as well as examples of Connection Manager flows and usages. This document targets Version 2.0 Connection Managers operating with devices (Host Routers, Device Routers, Re-timers, Cables) that comply with Version 2.0 or Version 1.0 of the USB4 specification. Unless specified otherwise, any reference to a Version 2.0 device also applies to future devices that comply with future versions of the USB4 specification (forward compatibility).

This document describes Connection Managers that use a PCIe host interface. Connection Managers that use other interfaces are outside the scope of this document. However, many of the flows and information in this document are applicable to all Connection Managers.

This document is targeted at USB4 host developers, but also contains useful information for USB4 device developers hoping to better understand how a Connection Manager configures and manages USB4 devices.

This document should be read in conjunction with the Connection Manager Notes in the USB4 Specification. The Connection Manager Notes give requirements and recommendations for Connection Manager implementation.

This document assumes that the reader has a basic understanding of the USB 3.2, DisplayPort, and PCIe protocols.

### **1.2. Reference Documents**

- Universal Serial Bus (USB4®) Specification, Revision 2.0 October 2022 (USB4 Specification v2.0)
- Universal Serial Bus (USB4®) Inter-Domain Service Protocol, Version 2.0, (USB4 Inter-Domain Service Specification)

### **1.3. Overview**

A Connection Manager is part of a USB4 host system. It is the entity that discovers, manages, and configures connected USB4 devices. When a Connection Manager enumerates a USB4 device, that device is considered part of the Connection Manager's Domain.

The Connection Manager executes the following configuration tasks within its Domain:

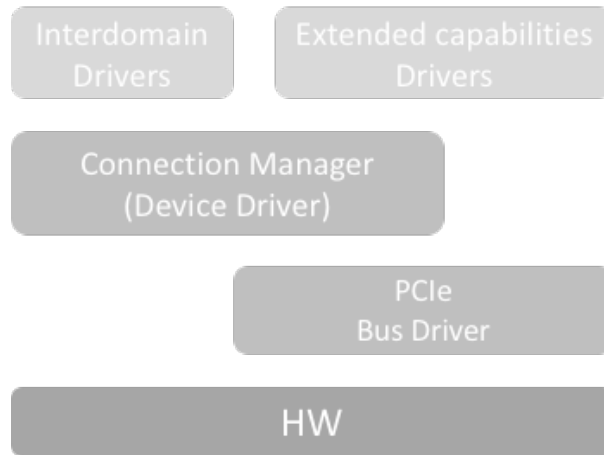
- Enumeration of Routers
- Setup and tear-down of Paths
- Initialization of a hot plugged Router and removal of a hot unplugged Router
- Configuration of QoS behavior including flow control and bandwidth allocation

A Connection Manager can optionally perform inter-Domain configuration and management functions. Inter-Domain configuration and management is outside the scope of this document and described in the USB4 Inter Domain Specification.

A Connection Manager uses Ring 0 to transmit and receive Control Packets. Any additional Rings are used for Inter-Domain Packets.

A Connection Manager can be implemented in software or firmware. Figure 1-1 shows a Connection Manager as part of a USB4 host interface that is implemented as a PCIe function. The interface between the Connection Manager and the Operating System is system specific and not described in this document.

**Figure 1-1: High-Level PCIe SW Architecture with a Connection Manager**



## 2. Host Interface Initialization

Prior to initialization, the Connection Manager may reset the host interface either via a Host Router Reset (Version 2.0 Host Router or higher, see Section 3.5), or via a Host Interface Reset (see Section 2.4).

### 2.1. Ring 0 Initialization

To properly operate and enable Connection Manager functionality, the host interface driver initializes a Transmit Descriptor Ring for HopID 0 (referred to as “Transmit Ring 0”) and a Receive Descriptor Ring for HopID 0 (referred to as “Receive Ring 0”).

A Connection Manager injects Control Packets into a Domain through Transmit Ring 0. A Connection Manager receives Control Packets from a Domain through Receive Ring 0.

The host interface driver performs the following steps to initialize Transmit Ring 0 and Receive Ring 0:

1. Configure interrupts and interrupt controls.
2. Allocate a buffer for Transmit Ring 0 in host memory.
3. Initialize Transmit Ring 0 as follows:
  - Set the Ring base address in the Base Address Low Register and the Base Address High Register of Transmit Ring 0.
  - Set the Ring size in the Ring Size Register of Transmit Ring 0.
  - Set the following Ring control attributes in the Ring Control Register of Transmit Ring 0:
    - *E2E Flow Control Enable* bit = 0b.
    - *No-Snoop Flag* bit = 0b.
    - *Raw Mode* bit = 1b.
    - *Ring Valid* bit = 1b.
  - As necessary, allocate new transmit buffers in host memory for outgoing Control Packets and queue descriptors with valid pointers to these buffers.
4. Allocate a buffer for Receive Ring 0 in host memory.
5. Initialize Receive Ring 0 as follows:
  - Set the Ring base address in the Base Address Low Register and the Base Address High Register of Receive Ring 0.
  - Set the Ring size and data buffer size for the Ring in the Ring Size Register of Receive Ring 0.
    - The Ring size needs to be at least 256 bytes.
  - Set the following Ring control attributes in the Ring Control Register of Receive Ring 0:
    - *E2E Flow Control Enable* bit = 0b.
    - *No-Snoop Flag* bit = 0b.
    - *Raw Mode* bit = 1b.
    - *Ring Valid* = 1b.

- As necessary, allocate new receive buffers in host memory for incoming Control Packets and queue pointers to these buffers. Update the *Consumer Index* field accordingly.
6. After initializing Transmit Ring 0 and Receive Ring 0, the Connection Manager is ready to both send and receive USB4 packets to/from the Host Router.

## 2.2. Connection Manager Transmit Flow

A Connection Manager configures Transmit Ring 0, generating Transmit Descriptors (the data structures intended for transmission), filling and queuing Transmit Descriptors, and processing transmit done messages.

A Connection Manager performs the following steps to issue a Control Packet:

1. Create a Transmit Descriptor with the following fields:
  - The *Address* field points to the Data Buffer in Host Memory that holds the Control Packet data to be transmitted.
  - The *Offset* field contains the first location of data in the Data Buffer.
  - The *Data Length* field contains the number of bytes to be transmitted.
  - The *EOF PDF* field contains the PDF value to be inserted into the Control Packet carrying the data.
  - The *Request Status* field is set to 1b.
  - The *Interrupt Enable* field is set to the desired value.
2. Compare the local copies of the *Producer Index* and *Consumer Index* fields for Transmit Ring 0 to determine if the transmit Ring is full.
  - If  $Consumer\ Index \neq (Producer\ Index + 1) \% Ring\ Size$ , Transmit Ring 0 is not full. Write the Transmit Descriptor into the Transmit Ring 0 data area in host memory.
  - If  $Consumer\ Index == (Producer\ Index + 1) \% Ring\ Size$ , Transmit Ring 0 is full. Wait until space is available, then write the Transmit Descriptor into the Transmit Ring 0 data area in host memory.
3. Increment the Producer Index and write the resulting value (modulo *Ring Size*) to the *Producer Index* field in the *Producer and Consumer Indexes* Register for Transmit Ring 0.
4. If the *Interrupt Enable* bit was set to 1b in the Transmit Descriptor, when a Transmit Done interrupt is received, scan the Transmit Descriptors in host memory starting from the *Consumer Index* location. For each Transmit Descriptor with the *Descriptor Done* bit set to 1b, either free or recycle the buffer as appropriate.
5. If the *Interrupt Enable* bit was not set to 1b in the Transmit Descriptor, poll the *Consumer Index* field. When the *Consumer Index* field changes, scan the Transmit Descriptors in host memory starting from the *Consumer Index* location. For each Transmit Descriptor with the *Descriptor Done* bit set to 1b, either free or recycle the buffer as appropriate.

## 2.3. Connection Manager Receive Flow

A Connection Manager configures Receive Ring 0, allocating memory data buffers for incoming data, creating descriptors pointing to the buffers and queuing these descriptors to the Receive Ring. The Connection Manager also processes the data posted into Host Memory buffers by the Host Router.

A Connection Manager performs the following steps when it receives an interrupt indicating that a Host Router has posted Control Packet data into a Data Buffer for Receive Ring 0:

1. Read the *Producer Index* and *Consumer Index* fields for Receive Ring 0 to determine the number of Receive Descriptors ready to be processed.

*Note: A Receive Descriptor is ready to be processed when the Descriptor Done bit in the Receive Descriptor posted by the Router is 1b.*

2. For each Receive Descriptor ready to be processed:
  - Read the *Address Low* and *Address High* fields to get the location of the associated Data Buffer.
  - Consume the data in the Data Buffer. Data buffers should be consumed in the order they are received from the Host Router.
3. After processing, either free or recycle the data buffers as appropriate, increment the Consumer Index accordingly and write the result (modulo *Ring Size*) to the *Consumer Index* field in the *Producer and Consumer Indexes* Register for Receive Ring 0.

#### **2.4. Host Interface Reset**

A Connection Manager resets the Transmit Rings and Receive Rings within the Host interface by setting the *RST* bit in the Host Interface PCI Configuration Space to 1b. Setting the *RST* bit to 1b brings the registers in the memory BAR to their default state and clears the End-to-End Flow Control state.

After setting the *RST* bit to 1b, the Connection Manager waits for a minimum of 10 milliseconds for the reset to complete. Note that transition of the *RST* bit to 0b does not indicate completion of a reset.

#### **2.5. DROM Access**

A USB4 Device Router contains Device ROM (DROM). A USB4 Host Router may also contain DROM. The DROM provides data about the product and the Router within it, which the Connection Manager uses to properly enumerate and allocate resources. The DROM is customized for each model of the product. See the USB4 DROM Specification for more information on how the DROM in a USB4 device is formatted.

A Connection Manager uses the DROM READ Router Operation, defined in Chapter 8 of the USB4 Specification to access the DROM of a USB4 Router.

#### **2.6. PCIe Memory BAR Access**

A Connection Manager limits the size of a read or write access from the PCIe Memory BAR to aligned 32 bits only.

### 3. Router Configuration

#### 3.1. Connection Event Detection

A Router detects and reports connection events on its USB4 Ports and its DP Adapters. USB3 Adapters, PCIe Adapters, and Host Interface Adapters are not pluggable.

##### 3.1.1. USB4 Connection Events

A Connection Manager can add a Router to its Domain when any of the following USB4 connection events occur:

- The Domain powers up (including the Host Router).
- A Router is hot plugged into the DFP of another Router in the Domain.
- A Router completes a Downstream Port Reset.
- The UFP of a Router is disabled, then re-enabled.
- Routers exit sleep state.

The Connection Manager can use any of the following mechanisms to detect a newly connected Router:

- A Hot Plug Event Packet.
  - The Connection Manager enables (or disabled) the generation of Hot Plug Event Packets in a Downstream Facing Port by setting the *Disable Hot Plug Events* bit to the desired value. The bit shall be set before the *Lock* bit for the Port is set to 0b.
  - When a USB4 connection event occurs, the Router that detects the connection sends the Connection Manager a Hot Plug Event Packet (if enabled) with the *UPG* bit set to 0b. The Router periodically resends the Hot Plug Event Packet until the Connection Manager responds with a Hot Plug Acknowledgment Notification Packet. If the Connection Manager receives multiple Hot Plug Event Packets for a hot plug event, it sends at least one Notification Packet for the event
- Periodically scanning the Downstream Facing Ports of enumerated Routers. For example, the Connection Manager can read the *LANE\_ADAP\_CS\_1.Adapter State* field of a Downstream Facing Lane Adapter. If the Adapter State is CLd, the Connection Manager can conclude that nothing is connected downstream. If the Adapter State is CL0, the Connection Manager can conclude that another Router is connected downstream.

After detecting that a new Router is connected to the Domain, a Connection Manager can either initialize the Router right away, delay Router initialization (for any period of time) or chose not to initialize the Router. Section 3.2 describes Router initialization.

##### 3.1.2. Display Connection Events

A DP Source or DP Sink can be added when any of the following display connection events occur:

- The Domain powers up.
- A DP Source or DP Sink is connected to an enumerated Router.
- A Router with a connected DP Source or DP Sink exits sleep state and completes enumeration.

- A Router with a connected DP Source or DP Sink exits reset and completes enumeration.

When a display connection event occurs, the Router that detects the connection sends the Connection Manager a Hot Plug Event Packet with the *UPG* bit set to 0b. The Router periodically resends the Hot Plug Event Packet until the Connection Manager responds with a Notification Packet with a Hot Plug Acknowledgment.

The Connection Manager can use any of the following mechanisms to detect a display connection event:

- A Hot Plug Event Packet.
- Read the *HPD Status* bit in ADP\_DP\_CS\_2 – relevant only for DP Sink (monitor).

Section 5.4 describes how a Connection Manager configures a DP Path after a display connection event.

### 3.2. Router Initialization

Router Initialization consists of the following steps:

1. The Connection Manager reads the Router Configuration Space of the Router to determine the Router's capabilities (Section 3.2.1).
2. The Connection Manager enumerates the Router by assigning it a Topology ID (Section 3.2.1).
3. The Connection Manager scans the Router Ports to determine how many Ports and Adapters the Router has. It reads the Adapter type and Adapter number of each Adapter (Section 3.2.2).
4. The Connection Manager enables the appropriate protocol tunnels and waits for the *Configuration Ready* bit to be set to 1b (Section 3.2.3)
  - a. Based on the accuracy level required by the Routers and the implementation specific Connection Manager policy, the Connection Manager configures TMU capabilities as described in Section 7.4.3.
5. The Connection Manager creates and configures one or more Paths through the Router (Chapter 5).
6. The Connection Manager configures Protocol Adapters and enables the associated Paths (Chapter 5).

Each initialization step is described in more detail in the sub-sections below.

Notes:

1. If the Connection Manager receives a Notification Packet with an *Event Code* = ERR\_CONN at any time during operation, then the Adapter listed in the Notification Packet is disconnected.
2. If the Connection Manager receives a Notification Packet with an *Event Code* = ERR\_LOCK at any time during operation, the Adapter listed in the Notification Packet was reconnected (i.e. a disconnect happened and the ADP\_CS\_4.Lock bit set to 1b, followed by a connect event).

#### 3.2.1. Router Enumeration and Configuration

Prior to enumerating a Domain, the Connection Manager may reset a Version 2.0 (or higher) Host Router and its Domain by issuing a Host Router Reset. See Section 3.5.

A Connection Manager reads from and writes to the Configuration Space of a Router in order to enumerate it. Before accessing the Configuration Space of a newly connected Device Router, the Connection Manager “unlocks” the DFP that the Router is connected to. The Connection Manager does this by writing 0b to the ADP\_CS\_4.Lock bit in the DFP.

After unlocking the DFP, a Connection Manager performs the following steps to enumerate the Router:

1. Send a Read Request that targets DW0 through DW4 (and optionally additional DWs) in the Router Configuration Space of the Router.
  - If the Connection Manager receives a Notification Packet with an *Event Code* = ERR\_ENUM, then the target Router is already enumerated. In this case, the Connection Manager does not enumerate the Router. Instead, it performs the flow described in Section 3.3 to determine whether the Router belongs to a different Domain or whether a loop exists in this Domain’s topology.
  - If the Connection Manager receives a Notification Packet with an *Event Code* = ERR\_NUA, then the target Router is uninitialized, but is connected to the Domain via its DFP. If the UFP is later connected, the Router will send a Hot Plug Event Packet, which can be used to determine whether the Router belongs to a different Domain or whether a loop exists in the Domain’s topology. See case 1 in Figure 3-1 for more information.
  - If Router sends a Read Response, then the Connection Manager can continue to Step 2).
2. Determine the Depth of the Router by reading the Depth of the upstream Router and adding 1.
  - If the Router does not exceed the maximum depth (maximum depth = 5), then continue to Step 3.
  - Else, the Router is not enumerated, and the enumeration flow ends here.
3. Parse the Read Response from the Router and look at the ROUTER\_CS\_4.USB4 Version field:
  - If the *USB4 Version* field is set to Version 2.0 or higher, then the Router supports USB4 Version 2.
  - If the *USB4 Version* field is set to Version 1.0, then the Router supports USB4 Version 1.
  - If the *USB4 Version* field is set to 00h, then the Router is a TBT3 Router. If the Connection Manager is TBT3-compatible, it shall continue with TBT3 Router configuration.
4. The Connection Manager sends a single Write Request to the Router that writes the following information to the Router Configuration Space:
  - *Depth* field = Router depth
  - *TopologyID* field = Router Topology ID
  - *Upstream Adapter* field = the value read earlier (i.e., value does not change)
  - *Connection Manager USB4 Version* field
    - = Version 1.0 if the Router is a USB4 Version 1.0 Router

- = Version 2.0 if the Router is a USB4 Version 2.0 Router
- *TopologyID Valid* bit = 1b

*Note: See the USB4 Specification for more information on what values to write to the Depth field and TopologyID field.*

*Note: The Connection Manager shall not issue a Write Request to an Uninitialized Router other than the Write Request that enumerates the Router.*

*Note: Unless specified otherwise, a Connection Manager shall not change the value in any of the fields in DW1 through DW3 after the TopologyID Valid bit is set to 1b.*

*Note: A Version 2.0 Connection Manager shall not change the value in the Connection Manager USB4 Version field after setting the TopologyID Valid bit to 1b*

After the Connection Manager enumerates a Router in a USB4 device, it can read the DROM from the device. If, after reading the DROM, the Connection Manager decides to exclude the Router from its Domain, it can un-enumerate the Router by issuing a Downstream Port Reset to the DFP that the Router is connected on.

A Connection Manager may keep a list of Router UUIDs (read from the *UUID* field in the Router Configuration Space) so that on subsequent connection events, the Connection Manager can compare the UUID value of a connected Router to the list to determine whether to enumerate the Router.

The Connection Manager shall execute a Get Capabilities Operation to Router to identify capabilities supported by the Router:

- Hot Plug Failure Indication
- Sequence bit in Notification Packet

The Connection Manager enables Hot Plug Failure Indication, if desired, by issuing a Set Capabilities Operation to the Router.

The Connection Manager sets the mechanism for Notifications as described in Section 7.7.

After enumerating a Router, the Connection Manager configures the Router by writing the desired parameters and settings to Router Configuration Space. See the USB4 Specification for Router Configuration Space fields and values.

The Connection Manager can optionally set the following fields in Router Configuration Space to a value different than the Router default:

- *Notification Timeout* field (to a value > 0)

If the *TBT3 Not Supported* bit in the Router Configuration Space is set to 0b, then

- After configuring a Version 1.0 Router, the Connection Manager polls the *ROUTER\_CS\_6.Router Ready* bit until the bit is set to 1b.
- After configuring a Version 2.0 Router, the Connection Manager waits for reception of a Notification Packet with *Event Code* = *ROP\_CMPLT* and the *Event Info* field set to 01h. The Connection Manager reads the *ROUTER\_CS\_6.Router Ready* bit to verify it is set to 1b. Alternatively, the Connection Manager may poll the *ROUTER\_CS\_6.Router Ready* bit until the bit is set to 1b.

It then scans the Router's Adapters as described in Section 3.2.2.

*Note:*

*A TBT3-compatible Router sets the Router Ready bit to 1b within 50ms after the TopologyID Valid bit is set to 1b. It is recommended that the Connection Manager waits for 500ms before it times out.*

### **3.2.2. Adapter Enumeration**

The Connection Manager enumerates the Adapters of a Router by reading the Adapter Configuration Space of each Adapter. The Connection Manager sends a Read Request to each Adapter Configuration Space starting from Adapter number 1 up to Adapter number {Max Adapter}. The Read Requests target ADP\_CS\_0 to ADP\_CS\_5 of each Adapter Configuration Space.

If the response to a Read Request is a Notification Packet with Event Code = ERR\_ADDR, the Connection Manager makes a note that the Adapter is unused. Else, the Connection Manager parses the Read Data in the Read Response and looks at the *Adapter Type Protocol*, *Adapter Type Version*, and *Adapter Type Sub-type* fields to determine the Adapter Type. See Table 8-10 in the USB4 Specification for how to identify the different Adapter Types.

Lane Adapters come in pairs. Each pair of Lane Adapters corresponds to a USB4 Port. The Lane Adapter with the lowest Adapter Number is the Lane 0 Adapter. The Lane Adapter with the highest number is the Lane 1 Adapter. Lane 1 needs to either be aggregated with Lane 0 or be disabled before the Connection Manager can setup any Paths through the USB4 Port. A Link may operate as a Gen 4 Link (Version 2.0 Routers only, either Symmetric or Asymmetric, see Section 7.2), as an Aggregated Gen 2 / Gen 3 Link (see in Section 7.1), or as a Single-Lane Link with Lane 1 disabled (see Section 7.5).

If the Router has an Upstream USB3 Gen X Adapter, it supports USB3 Gen X Tunneling. The Connection Manager makes a note which Adapter is the Upstream USB3 Gen X Adapter. The Connection Manager can setup a Path for USB3 Gen X Tunneling as described in Section 5.5. If the Connection Manager does not setup a Path to a USB3 Gen X Adapter, it disconnects the Adapter as defined in Section 5.5.3. Note that each USB3 Gen X Adapter only supports one USB3 Gen X Path in each direction.

If the Router has one or more Downstream USB3 Gen X Adapters, the Connection Manager maps each Downstream USB3 Gen X Adapter to a USB4 Port. USB4 Ports and USB3 Gen X Adapters are ordered in pairs in increasing Adapter numbers (see example in the USB4 Specification).

If the Device Router has an Upstream USB3 Gen T Adapter, it supports USB3 Gen T Tunneling. The Connection Manager makes a note which Adapter is the Upstream USB3 Gen T Adapter. Note that the USB3 Gen T Adapter may support more than one USB3 Gen T Path. The Connection Manager can setup Paths for USB3 Gen T Tunneling as described in Section 5.6.

If the Host Router has one or more Downstream USB3 Gen T Adapters, the Connection Manager maps each Downstream USB3 Gen T Adapter to a USB4 Port. USB4 Ports and USB3 Gen T Adapters are ordered in pairs in increasing Adapter numbers (see example in the USB4 Specification).

If the Router has an Upstream PCIe Adapter, then PCIe tunneling is supported by the Router. The Connection Manager makes a note which Adapter is the Upstream PCIe Adapter. The Connection Manager can setup a Path for PCIe tunneling as described in Section 5.3. Note that each PCIe Adapter only supports one PCIe Path in each direction.

If the Router has one or more Downstream PCIe Adapters, the Connection Manager maps each Downstream PCIe Adapter to a USB4 Port. USB4 Ports and PCIe Adapters are ordered in pairs in increasing Adapter numbers (see example in the USB4 Specification).

If the Router has a DP IN Adapter, the Connection Manager makes a note that the Router supports DP Tunneling for a connected DP Source.

If the Router has a DP OUT Adapter, the Connection Manager makes a note that the Router Supports DP Tunneling for a connected DP Sink.

If Router has any Adapters with Adapter Type = “Unsupported Adapter”, the Connection Manager makes a note that those Adapters are not used.

After the Connection Manager reads the DROM and gets the list of Unused Adapters it makes a note that those Adapters are not used.

The Connection Manager ignores any Adapters with an unknown Adapter Type.

For each Adapter, a Connection Manager reads its Adapter Configuration Space by following the Linked List as described in section 8.2.2 in the USB4 Specification. The Connection Manager skips undefined Capabilities. It should be robust and handle Capabilities that are not ordered as defined in the USB4 Specification.

### **3.2.3. Enabling Protocol Tunneling**

The last step in Device Router initialization is to enable Tunneled Protocols. By default, USB3 Gen T Tunneling, USB3 Gen X Tunneling and PCIe Tunneling are disabled in a Device Router.

The Connection Manager enables USB3 Gen T Tunneling in a Device Router if all conditions are true:

- USB3 Gen T Tunneling is supported in the Host Router.
- The Device Router has an Upstream USB3 Gen T Adapter.
- Link is a USB4 Link.

The Connection Manager sets the *ROUTER\_CS\_5.Host Supports USB3 Gen T* bit to 1b to enable USB3 Gen T Tunneling. The Connection Manager sets the *ROUTER\_CS\_5.Host Supports USB3 Gen T* bit to 0b to disable USB3 Gen T Tunneling.

The Connection Manager enables USB3 Gen X Tunneling in a Device Router if all conditions are true:

- USB3 Gen X Tunneling is enabled in the Router that is connected to the Upstream Facing Port of the Device Router.
- The Device Router has an Upstream USB3 Gen X Adapter.
- Link is a USB4 Link.

The Connection Manager sets the *ROUTER\_CS\_5.USB3 Tunneling On* bit to 1b to enable USB3 Gen X Tunneling. The Connection Manager sets the *ROUTER\_CS\_5.USB3 Tunneling On* bit to 0b to disable USB3 Gen X Tunneling.

The Connection Manager must set USB3 Tunneling if the Device Router supports USB3 tunneling.

The Connection Manager enables PCIe Tunneling for the Device Router if both conditions are true:

- PCIe Tunneling is enabled in the Router that is connected to the Upstream Facing Port of the Device Router.
- The Device Router has an Upstream PCIe Adapter.

The Connection Manager sets the ROUTER\_CS\_5.*PCIe Tunneling On* bit to 1b to enable PCIe Tunneling. The Connection Manager sets the ROUTER\_CS\_5.*PCIe Tunneling On* bit to 0b to disable PCIe Tunneling.

If USB3 Gen X Tunneling is disabled, if the *Internal Host Controller Implemented* bit is set to 1b, and if PCIe Tunneling is enabled in the Router, then the Connection Manager may enable the Internal Host Controller in the Device Router by setting the ROUTER\_CS\_5.*Internal Host Controller On* bit to 1b.

The Connection Manager may configure its wake policy by setting the *Wake Enable* bits to the desired values.

The Connection Manager sets the ROUTER\_CS\_5.*Enumerated State PCIe Wake* bit in a Version 2.0 Router to 1b if the system supports wake by a connected PCIe device in D3cold state while the USB4 Domain is in Enumerated state.

*Note: See Section 8.2.1.1 for more information on how to configure the Wake Enable bits and the USB4 Port is Configured bit.*

After enabling USB3 Gen T Tunneling, USB3 Gen X Tunneling and PCIe tunneling in a Device Router, the Connection Manager executes the following steps:

1. Set the ROUTER\_CS\_5.*Configuration Valid* bit to 1b.
2. Version 1.0 Router: Poll the ROUTER\_CS\_6.*Configuration Ready* bit until it is set to 1b, which means that the Device Router is ready to operate the enabled protocol tunneling.
3. Version 2.0 Router: Wait for reception of a Notification Packet with *Event Code* = ROP\_CMPLT and the *Event Info* field set to 02h. Read the ROUTER\_CS\_6.*Configuration Ready* bit to verify that it is set to 1b, which means that the Device Router is ready to operate the enabled protocol tunneling. Alternatively, the Connection Manager may poll the ROUTER\_CS\_6.*Configuration Ready* bit until the bit is set to 1b.
  - a. Based on the accuracy level required by the Routers and the implementation specific Connection Manager policy, the Connection Manager configures TMU capabilities as described in Section 7.4.3.

*Note: The Router sets the Configuration Ready bit to 1b within 50ms from setting of the Configuration Valid bit to 1b. It is recommended that the Connection Manager waits for 500ms before it times out.*

4. Setup Paths for the enabled Tunneled Protocols as described in Section 5.

### **3.3. Identifying Loops and Inter-Domain Links**

A loop occurs when the DFP of a Router is connected to the DFP of another Router in the same Domain. An Inter-Domain Link occurs when the DFP of a Router is connected to the DFP of another Router in a different Domain.

When attempting to enumerate a Router, a Connection Manager might receive a Notification Packet with an *Event Code* = ERR\_ENUM or an *Event Code* = ERR\_NUA. When the Connection Manager receives a Notification Packet with an *Event Code* = ERR\_ENUM or an *Event Code* = ERR\_NUA, it needs to identify whether the event is the result of a loop in its Domain topology or that the Router is part of another Domain.

The following steps describe how a Connection Manager identifies a DFP-to-DFP connection and determines whether it is an Inter-Domain Link or a loop within its own Domain:

1. A Router in the Connection Manager's Domain sends a Hot Plug Event Packet, which indicates a new connection on its DFP.
2. After unlocking the DFP, the Connection Manager sends a Read Request to the Router. The Read Request targets the Adapter referenced in the *Adapter Num* field of the Hot Plug Event Packet.
3. The hot-plugged Router receiving the Read Request responds with a Notification Packet with an *Event Code* = *ERR\_ENUM* (if the Router is enumerated) or an *Event Code* = *ERR\_NUA* (if the Router is not enumerated). The Notification Packet contains The Route String High (24 bits of Topology ID) and Route String Low (32 bits of the Topology ID) indicating the Topology ID that was used by the Connection Manager to access the Router.
4. The Connection Manager sends an Inter-Domain UUID Request Packet as defined in the USB4 Inter-Domain Service Protocol Specification with the *Route String* field pointing to the Router that was the target of the Read Request.
5. The Router that receives the Inter-Domain UUID Request Packet on its DFP forwards the Packet on its UFP to its Connection Manager.
6. The Connection Manager that receives the Inter-Domain UUID Request Packet responds with an Inter-Domain UUID Response Packet as defined in the USB4 Inter-Domain Service Protocol Specification. The Inter-Domain UUID Response Packet targets the Router that sent the Inter-Domain UUID Request Packet. The Response Packet contains a 128-bit UUID value that identifies the Connection Manager.
7. The Router that receives the Inter-Domain UUID Response Packet on its DFP forwards the Packet on its UFP to its Connection Manager.
8. The Connection Manager compares the UUID value in the Inter-Domain UUID Response Packet to its own UUID. If the UUIDs are the same, there is a loop in the Domain. If the UUIDs are different, there is an Inter-Domain Link.

When the Connection Manager detects an Inter-Domain Link, it does the following:

- Set the *PORT\_CS\_19.USB4 Port is Inter-Domain* bit to 1b in the Port that connects to the Inter-Domain Link. The Connection Manager sets the *PORT\_CS\_19.USB4 Port is Inter-Domain* bit to 0b for all other USB4 Ports.

See the USB4 Inter-Domain Service Protocol Specification for details on how Connection Manager manages the Inter-Domain Link.

If the Connection Manager detects a topology loop it does the following in the Ports on both sides of the Link that created the loop in the topology:

- Disable Time Sync Handshakes by setting the *TMU\_ADP\_CS\_6.Disable Time Sync* bit to 1b.
- Do not set the *PORT\_CS\_19.USB4 Port is Configured* bit to 1b.
- Do not set any of the *PORT\_CS\_19.Wake on Connect* bits to 1b.

Figure 3-1 illustrates the steps above when there is a loop in the Domain. In Case 1, the Connection Manager handles the Hot Plug Event Packet sent by Router B. In Case 2, the Connection Manager handles the Hot Plug Event Packet sent by Router A.

**Figure 3-1: Flow for Identifying a Loop in a Domain**

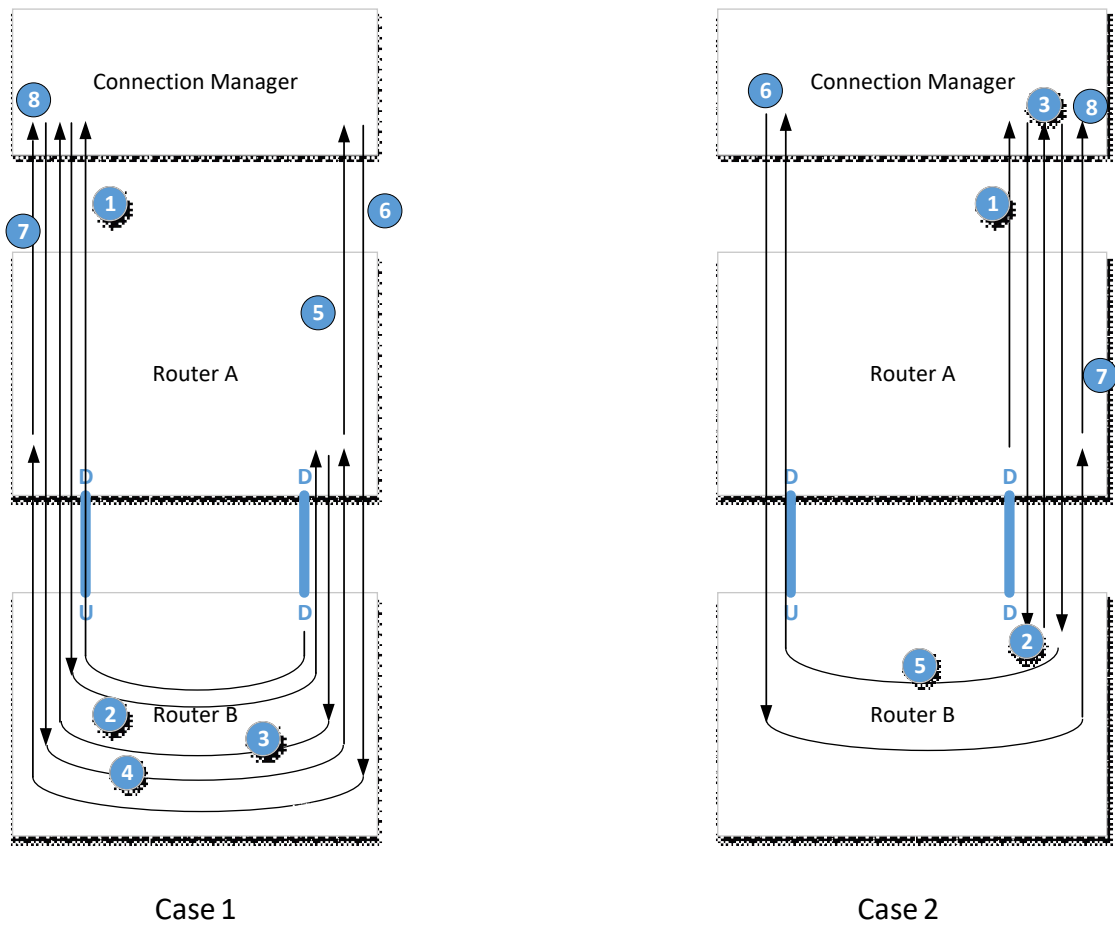
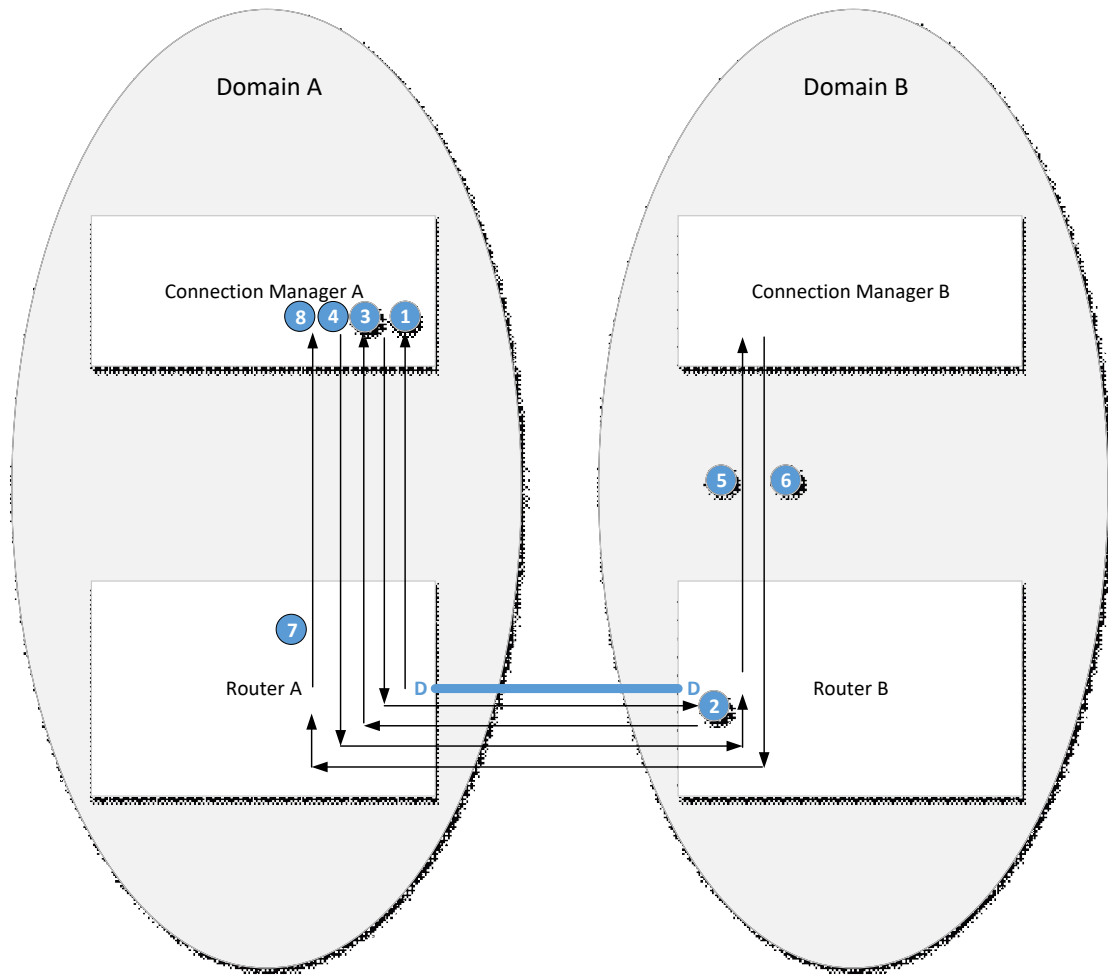


Figure 3-2 illustrates the steps above when there is in Inter-Domain Link.

**Figure 3-2: Flow for identifying an Inter-Domain Link**



### 3.4. Detecting a Removal Event

A Router can detect and report removal events on its USB4 Ports and its DP Adapters. USB3 Adapters, PCIe Adapters, and Host Interface Adapters are not removable.

#### 3.4.1. Lane Adapter

A Router is removed from a Domain upon one of the following removal events:

- The Router's UFP is physically unplugged
- The Router's UFP goes through a Port disconnect as defined in the USB4 specification (Section 4.4.5)

When a removal event occurs, the Router that detects the removal event sends the Connection Manager a Hot Plug Event Packet with the *UPG* bit set to 1b. The Router periodically resends the Hot Plug Event packet until the Connection Manager responds with a Notification Packet with a Hot Plug Acknowledgment.

The Connection Manager can use any of the following mechanisms to detect a removal event on a USB4 Port:

- A Hot Plug Event Packet.
- Periodically scanning the Downstream Facing Ports of enumerated Routers. For example, the Connection Manager can read the *LANE\_ADP\_CS\_1.Adapter State*

field in a Downstream Facing Lane Adapter. If the Adapter State is CLd, the Connection Manager can conclude that nothing is connected downstream. If the Adapter State is CL0, the Connection Manager can conclude that another Router is connected downstream.

After a removal event, the Connection Manager tears down all Paths (other than Path 0) that either terminate at the removed Router or traverse the removed Router. Section 5.2.2 defines the steps for Path teardown. The Connection Manager also restores the following fields to their default values in both Lane Adapters of the USB4 Port that detected the removal event:

- *TMU Uni-Directional Mode*
- *USB4 Port is Configured*
- *Link Credits*

### **3.4.2. DP Adapter**

A DP Source or DP Sink is removed from a Domain when a DP Source or DP Sink is unplugged from a Router in the Domain

When a removal event occurs, the Router that detects the removal event sends the Connection Manager a Hot Plug Event Packet with the *UPG* bit set to 1b. The Router continues to resend the Hot Plug Event packet until the Connection Manager responds with a Notification Packet with a Hot Plug Acknowledgment. The Connection Manager responds to incoming Hot Plug Events from a Router in the order received (i.e. first-in, first-out (FIFO)).

The Connection Manager can use any of the following mechanisms to detect a removal event on a DP Adapter:

- A Hot Plug Event Packet.
- An HPD indication in the *ADP\_DP\_CS\_2.HDP Status* field of a DP OUT Adapter.

After a removal event, the Connection Manager tears down any DP Tunneling Paths as described in Section 5.4.3.4.

### **3.5. Host Router Reset**

If the Host Router is a Version 2.0 Router, the Connection Manager may reset the Domain by asserting a Host Router reset.

The Connection Manager checks support for Host Router reset by reading the *Host Interface Version* field in the PCIe Memory BAR within the host interface of the Host Router. If the value is 2.0 or higher, then Host Router reset is supported.

The Connection Manager sets the *Host Router Reset* bit to 1b to initiate a Host Router reset. The Connection Manager periodically reads the *Host Router Reset* bit till its value is 0b, indicating completion of reset. Polling starts 50ms after the *Host Router Reset* bit is set to 1b.

*Notes:*

1. *The Connection Manager shall disable all Transmit Descriptor Rings and wait for at least 1 millisecond prior to setting the Host Router Reset bit to 1b. After the Connection Manager sets the Host Router Reset bit to 1b, it shall not access the Receive Descriptor Rings until the Host Router Reset bit is set to 0b.*
2. *The Connection Manager shall not access the Host Router Reset bit for 50ms after setting it to 1b.*

3. *The Host Router is required to complete its reset within 500ms after the Host Router Reset bit is set to 1b.*
4. *After a Connection Manager initiates a Host Router reset, it is not guaranteed that responses will arrive for any outstanding transactions.*
5. *A Connection Manager shall not write to any registers in the Host Interface Memory BAR registers while a Host Router reset is in progress. Router behavior is not defined if the Memory BAR registers are written to during a Host Router reset.*
6. *While the Host Router Reset bit is set to 1b, the read value of registers in the Host Router Memory BAR space (with the exception of the Host Router Reset register) is undefined.*

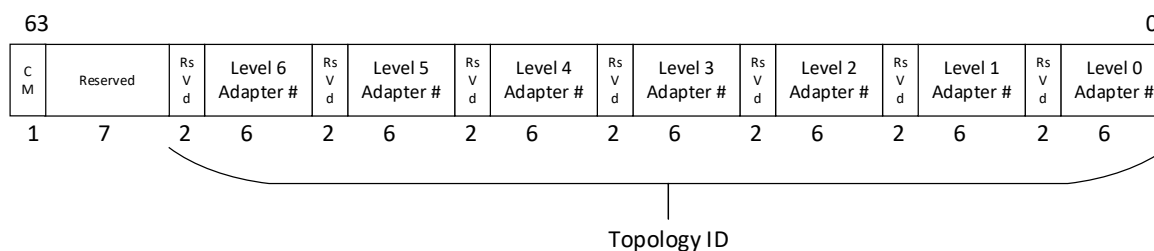
## 4. Control Packets

### 4.1. Routing

A Connection Manager uses Control Packets to manage the Routers within its Domain. Control Packets use HopID 0.

A Control Packet contains a Route String in its payload. The format of the Route String is shown in Figure 4-1.

**Figure 4-1: TopologyID Format**



For Control Packets flowing downstream (i.e. away from the Connection Manager), the Route String contains the TopologyID of the Router that is the final recipient of the Packet. The *CM* bit in the Route String is set to 0b. Control Packets flowing downstream are routed based on Route String.

For Control Packets flowing upstream (i.e. towards the Connection Manager), the Route String contains the TopologyID of the Router from which the Control Packet originates (for Enumerated Routers), or the Route String from the request packet (Uninitialized Router). The *CM* bit in the Route String is set to 1b. Control Packets flowing upstream are forwarded over the upstream Port of each Router and do not rely on the Route String for routing.

### 4.2. Protocol Rules

A Connection Manager shall only have one outstanding Read Request or Write Request at a time. The Connection Manager shall wait for a response to the previous Read or Write Request before sending the next Read or Write Request.

A Connection Manager can optionally implement a timeout mechanism and may retry a Control Packet if no response is received within the timeout interval. If a Connection Manager implements a timeout, it is recommended that the timeout is at least 10 milliseconds for Control Packets within a Domain and is at least 1 second for Inter-Domain Control Packets.

A Connection Manager should use the *Sequence Number* in a Response Packet to match it to the corresponding Request Packet.

### 4.3. Hot Plug

A Connection Manager shall send a Hot Plug Acknowledgment Packet in response to a Hot Plug Event Packet. It does not have to send the Hot Plug Acknowledgment Packet within a specific time after receiving the Hot Plug Event Packet. However, a Router will continue to resend a Hot Plug Event Packet and will not send any new Hot Plug Event Packets until it receives a Hot Plug Acknowledgment Packet.

A Connection Manager can send one Hot Plug Acknowledgement Packet per Hot Plug or Hot Unplug Event (regardless of the number of Hot Plug or Hot Unplug Event Packets received for that event) or it can send a Hot Plug Acknowledgement Packet for each Hot

Plug/Unplug Packet. A Connection Manager shall respond to Hot Plug Event Packets from a Router in the order received.

## 5. Path Configuration

This section describes how a Connection Manager configures a Path (Path Setup) and removes a Path (Path Teardown).

- Section 5.1 only applies to Path Setup in a USB4 Port.
- Section 5.2 applies to all Paths, regardless of the protocol that is tunneled.
- Section 5.3 describes the additional steps required for a Path that tunnels PCIe traffic.
- Section 5.4 describes the additional steps required for a Path that tunnels DisplayPort traffic.
- Section 5.5 describes the additional steps required for a Path that tunnels USB3 Gen X traffic.
- Section 5.6 describes the additional steps required for a Path that tunnels USB3 Gen T traffic.
- Section 5.7 describes the additional steps required for a Path that tunnels Inter-Domain traffic.

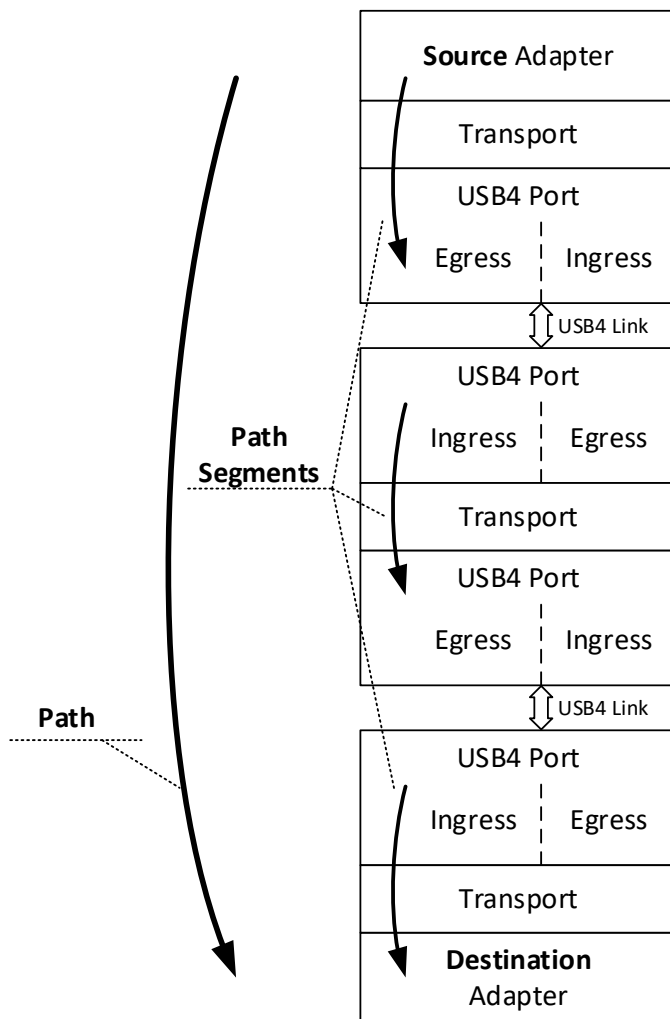
*Note: This section does not apply to the Control Path (Path 0). The Control Path is set up by a Router upon enumeration and torn down by the Router when it is disconnected.*

A Connection Manager does not configure a Path until all Device Routers that the Path traverses are enumerated and have the ROUTER\_CS\_6.Configuration Ready bit set to 1b.

A Connection Manager does not enable any Path in a USB4 Port with two Single-Lane Links that have not yet been bonded.

A Connection Manager configures and removes a Path in segments. The portion of the Path within a Router is called a Path segment. Figure 5-1 illustrates the Path segments for a Path that traverses three Routers.

**Figure 5-1: Path Segments**



There are two types of shared resources a Connection Manager considers when configuring a Path:

- Bandwidth – Bandwidth management maps into two Path settings: Path Priority and Path weight (see Section 6).
- Flow Control buffers – See Section 5.1.

### 5.1. Flow Control

There are four types of Flow Control schemes defined in the USB4 Specification. The type of Flow Control scheme to be used depends on which protocol is being tunneled and is detailed within each specific protocol Path Setup section.

A Connection Manager ensures that a Path has the same flow control scheme at each side of a Link (i.e., Path has same flow control scheme in both the Egress Adapter and Ingress Adapter). A Connection Manager ensures that the same flow control scheme is used throughout all the Links on a Path.

The subsections below describe how a Connection Manager configures each of the Flow Control schemes.

#### **5.1.1. Flow Control Disabled**

A Connection Manager sets the number of buffers that an Ingress Adapter may use for Paths that have flow control disabled. The Connection Manager does this by writing to the *ADP\_CS\_4.Non-Flow Controlled Buffers* field. A Connection Manager can optionally increase the number of buffers in the *Non-Flow Controlled Buffers* field when Paths with flow control disabled are configured and decrease the number of buffers in the *Non-Flow Controlled Buffers* field when Paths with flow control disabled are removed.

A Connection Manager does not change the *Non-Flow Controlled Buffers* field in a Protocol Adapter.

A Connection Manager shall configure DisplayPort Main-Link Path with the Flow Control Disabled scheme.

#### **5.1.2. Dedicated Flow Control**

When a Path with Dedicated Flow Control is setup, the Connection Manager allocates buffers for the Path. The Connection Manager does this by writing to the *PATH\_CS\_0.Path Credits Allocated* field.

After a Connection Manager tears down a Path, the buffers for the Dedicated Flow Control buffers are automatically deallocated from the Path and can be reallocated to a different Path.

A Connection Manager shall configure a USB3 Gen X Path with the Dedicated Flow Control Buffer Allocation scheme.

A Connection Manager may configure a USB3 Gen T Path with the Dedicated Flow Control Buffer Allocation scheme.

A Connection Manager shall configure PCIe Paths with the Dedicated Flow Control Buffer Allocation scheme.

A Connection Manager shall configure DP AUX Path with the Dedicated Flow Control Buffer Allocation scheme.

#### **5.1.3. Shared Flow Control**

A Connection Manager sets the number of buffers that an Ingress Adapter may use for Paths that use the Shared Flow Control scheme. The Connection Manager does this by writing to the *Link Credits Allocated* field in the Adapter Configuration Space.

A Connection Manager shall not enable Shared Flow Control or Restricted Shared Flow Control in any Link for more than one Tunneled Protocol.

A Connection Manager may configure a USB3 Gen T Path with the Shared Flow Control or the Restricted Shared Flow Control Buffer Allocation scheme.

A Connection Manager does not change the *Link Credits Allocated* field if there is an enabled Path through the Ingress Adapter that has *PATH\_CS\_1.ISE Flag* field set to 1b.

#### **5.1.4. Buffer Allocation**

The following steps can be used by a Connection Manager to allocate the buffers at an Ingress Lane Adapter:

1. Read the total number of buffers supported by the USB4 Port from the *ADP\_CS\_4.Total Buffers* field in the Lane 0 Adapter.

2. Read the number of buffers that are reserved for the Control Path from the `PATH_CS_0.Path Credits Allocated` field for Path 0.
3. Get the preferred buffer configuration for the Router by initiating Buffer Allocation Request Router Operations.
  - If a Router supports *Buffer Allocation Per USB4 Port* feature, and the Connection Manager enabled this feature through the Set Capabilities Router Operation, then the Connection Manager shall issue a Buffer Allocation Request Router Operation per USB4 Port setting the *Lane Adapter Number* field to the Lane 0 Adapter number.
  - Else, the Connection Manager issues a single Buffer Allocation Request Router Operation (where *Lane Adapter Number* field is inapplicable), for all USB4 Ports in the Router.
4. Calculate the number of buffers per each tunnel type as follows:
  - DP Paths:  $\#DP\ Streams * (baMinDPaux + baMinDPmain)$   
 where #DP Streams is the total number of current or future DP streams through that Lane Adapter based on the implementation specific Connection Manager Policy, and `baMinDPaux`, `baMinDPmain` are provided in the Buffer Allocation Request Operation
  - USB3 Gen X Path: `baMaxUSB3GenX`  
 where `baMaxUSB3GenX` is provided in the Buffer Allocation Request Operation
  - USB3 Gen T Paths: `baMaxUSB3GenT`  
 where `baMaxUSB3GenT` is provided in the Buffer Allocation Request Operation
  - PCIe Path: `baMaxPCIe`  
 where `baMaxPCIe` is provided in the Buffer Allocation Request Operation
  - Host-to-Host Path: `baMaxHI`  
 where `baMaxHI` is provided in the Buffer Allocation Request Operation
5. Allocate buffers to Paths based on the calculations in the previous step, while not exceeding the total number of available buffers
  - A Connection Manager shall setup a DP Path if it can allocate at least `baMinDPmain` buffers for the DP Main-Link Path and at least `baMinDPaux` buffers for the DP AUX Path.
  - It is recommended that a Connection Manager allocates the maximum available number of buffers for the Gen X Path, not exceeding `baMaxUSB3GenX`. However, the Connection Manager may setup a USB3 Gen X Path with less buffers as long as the number of buffers supports a minimal bandwidth of 1.5Gbps.

The number of buffers required for a given bandwidth is calculated by:

$$\text{Number of buffers} = (\text{BW} * \text{RT\_latency}) / \text{USB4\_packet\_size}$$

where:

BW is the desired bandwidth (in bits/sec)

RT\_latency is the round-trip latency (in sec) for flow-control credits. It is recommended to use a value of 2 microseconds for worst-case latency

USB4\_packet\_size is the average payload size (in bits) of USB4 Packets for the Path

Example: If desired bandwidth is 1.5Gbps and the average USB Packet size for a Gen X Path is 256B, then

$$\text{Number of buffers} = 1.5 * 10^9 * 2 * 10^{-6} / (256 * 8) < 2$$

- It is recommended that a Connection Manager allocates the maximum available number of buffers for the Gen T Paths, not exceeding baMaxUSB3GenT. However, the Connection Manager may setup the USB3 Gen T Paths with less buffers as long as the number of buffers supports a minimal bandwidth of 1.5Gbps for each supported Gen T path.

To calculate the number of buffers, see above formula.

- The Connection Manager may setup a PCIe Path with less than the maximum requested buffers as long as the number of buffers support a minimal bandwidth of 1.5Gbps.

To calculate the number of buffers, see above formula.

- The Connection Manager may setup a Host-to-Host Path with less than the maximum requested buffers.

## 5.2. Path Setup and Teardown

This section applies to all Paths that Tunnel protocol traffic.

### 5.2.1. Path Setup

When configuring a Path to a Protocol Adapter, a Connection Manager first reads the target Path Entry in Path Configuration Space. This is to ensure that the Connection Manager knows the existing values in Path Configuration Space.

The Connection Manager then writes to the target Path Entry in Path Configuration Space. The Connection Manager must use a single Read Request to read the Path Entry and a single Write Request to write to the Path Entry. The Connection Manager shall set a valid Lane Adapter or a Protocol Adapter field when issuing a Write Request to a Path Entry.

A Connection Manager can access registers for multiple contiguous Paths by increasing the read/write length in a Read or Write Request. For example, a Read Request targeting the entries for the Paths with Ingress HopIDs 8 and 9 would contain values of 16 in the *Address* field and 4 in the *Length* field.

A Connection Manager shall not send a Read or a Write Request that targets a Path Configuration Space entry that is either reserved or does not exist.

The only fields in the Path Configuration Space entry for Path 0 that a Connection Manager can change are the *Counter ID* field and the *Counter Enable* bit. To avoid changing the value of the other fields, it is recommended that a Connection Manager first read the Path Configuration Space entry for Path 0, make the necessary changes to the

*Counter ID* field and the *Counter Enable* bit, then write the results back to Path Configuration Space.

A Connection Manager does not change the value of any of the following fields in the Path Configuration Space of the Protocol Adapter:

- *Path Credits Allocated*
- *Ingress Flow Control*
- *Ingress Shared Buffering Enable*

A Connection Manager shall follow the rules below when allocating HopID values in a Router:

- An Input HopID can only be used once per Ingress Adapter.
- A HopID can only be used once per Egress Adapter.
- An Output HopID can be used more than once in an Ingress Adapter as long as the Output HopID targets a different Egress Adapter.
- The HopIDs for a Path can be different in each Router that the Path traverses.
- When the Egress Adapter is not a Host Interface Adapter:
  - Output HopID cannot be less than 8.
  - Output HopID cannot exceed the Max Output HopID field in the Adapter Configuration Space of the Egress Adapter.
  - If the Egress Adapter is a Lane Adapter, Output HopID cannot exceed the Max Input HopID field in the Adapter Configuration Space of the next Ingress Adapter (i.e. the Adapter that will be receiving the packet from the Egress Adapter).
- When the Egress Adapter is a Host Interface Adapter:
  - Output HopID cannot be less than 1.
  - Output HopID cannot exceed the Max Input HopID field in the Adapter Configuration Space of the Host Interface Adapter.

A Connection Manager performs the following steps to establish the Paths for a Tunneled Protocol:

1. For each Path, the Connection Manager writes to the Path Configuration Space of each Adapter that the Path traverses (including the Source and Destination Adapters). The Paths may be set up in any order.
  - The Connection Manager writes to the Path Configuration Space of each of the Ingress Adapters along the Path. The Connection Manager configures the Output HopID, Output Adapter, and QoS parameters of the Path and sets the `PATH_CS_0.Valid` bit to 1b.

*Note: When setting up a Path, it is recommended that the Connection Manager configure the Adapters along the Path in sequence from the Source Adapter to the Destination Adapter. This sequence ensures that the initial Credit Grant Packet, which is sent following buffer allocation, is received and processed by the Egress Adapter at the other Router.*

2. The Connection Manager sets the *Path Enable* bit to 1b in the Configuration Spaces of the Source Adapter and the Destination Adapter, which allows the Adapters to send and receive Transport Layer Packets on the Path.

### 5.2.2. Path Teardown

A Connection Manager performs Path Teardown after a Router is disconnected. A Connection Manager can also optionally perform Path Teardown any time after a Path is configured.

Each Tunneled Protocol has at least one Path in each direction and all Paths are torn down when the Tunneled Protocol is removed.

The Connection Manager performs the following steps to tear down a Path for a Tunneled Protocol:

*Note: The following steps compose the full list of steps. If Path teardown is due to a Router disconnect, the relevant steps that target the removed Router are skipped.*

1. Disable sending Transport Layer Packets on the Path by setting the *Path Enable* bit to 0b in both Source and Destination Adapters.
2. For each hop along the Path:
  - a. Set the *PATH\_CS\_0.Valid* bit to 0b in the Ingress Adapter. The Connection Manager does not change any other fields in the Path Configuration Space at this time.
  - b. Periodically poll the *PATH\_CS\_1.Pending Packets* bit in the Ingress Adapter until it is 0b, which indicates that all Transport Layer Packets belonging to the Path have been dequeued.
  - c. Wait at least 2 microseconds after a Router sets the *Pending Packets* bit to 0b before setting the *Valid* bit for the Path to 1b again.

*Note: It is recommended that a Connection Manager tear down a Path from the Source Adapter to the Destination Adapter. This sequence ensures that Tunneled Packets are stopped first at the origin.*

## 5.3. PCIe Path Setup and Teardown

This section describes how a Connection Manager sets up and tears down PCIe Paths between two Routers.

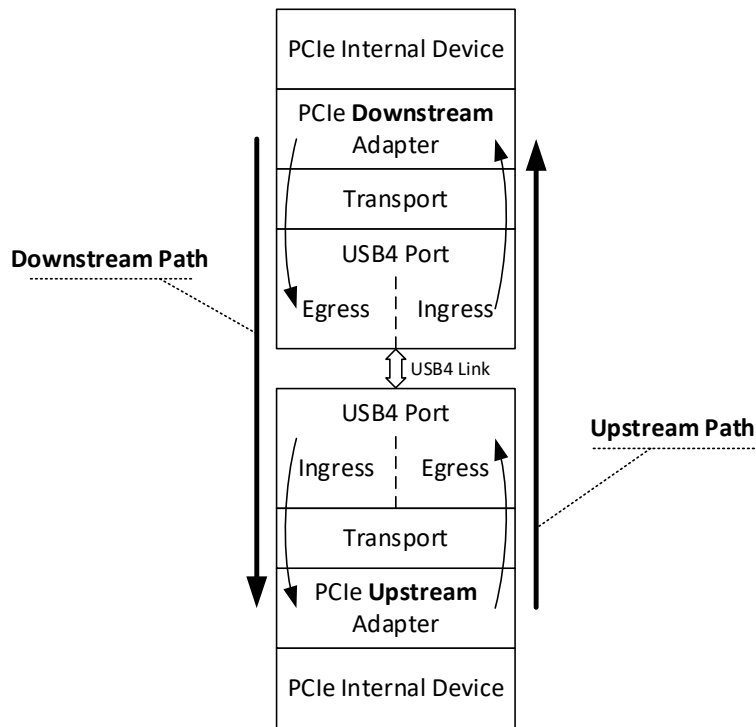
### 5.3.1. PCIe Path Setup

A Connection Manager can set up a PCIe Path between the PCIe Upstream Adapter and a PCIe Downstream Adapter of two physically connected Routers within the same Domain. A Connection Manager does not set up a PCIe Path between two Routers unless they are directly connected to each other.

A Connection Manager can set up a Path for PCIe Tunneling when it first enumerates a Router that supports PCIe Tunneling. It can also wait until some or all Routers in the Domain are enumerated.

A Connection Manager needs to establish two Paths, one upstream and one downstream, for PCIe tunneling. Figure 5-2 shows the two PCIe Tunneling Paths.

**Figure 5-2: PCIe Tunneling Scheme**



A Connection Manager sets up the PCIe Paths as follows:

1. Poll the `ADP_PCIE_CS_0.LTSSM` field in both PCIe Adapters until both are 0h (Indicating Detect state).
2. Set up the downstream Path.
  - Set Path Attributes in the Path Configuration Space of the Downstream PCIe Adapter. The Path Attributes are defined in Table 5-1 (Path segment is from the Downstream PCIe Adapter to the Lane Adapter).
  - Set Path Attributes in the Path Configuration Space of the Lane 0 Adapter of the USB4 Port. The Path Attributes are defined in Table 5-1 (Path segment is from the Lane Adapter to the Upstream PCIe Adapter).
3. Set up the upstream Path.
  - Set Path Attributes in the Path Configuration Space of the Upstream PCIe Adapter. The Path Attributes are defined in Table 5-1 (Path segment is from the Upstream PCIe Adapter to the Lane Adapter).
  - Set Path Attributes in the Path Configuration Space of the Lane 0 Adapter of the USB4 Port. The Path Attributes are defined in Table 5-1 (Path segment is from the Lane Adapter to the Downstream PCIe Adapter).
4. If the *USB4 Version* field in the Router Configuration Space of the Router and the Link Partner are both set to 2.0 or higher, then set the `ADP_PCIE_CS_1.Extended Encapsulation` bit to 1b in both the Upstream PCIe Adapter and the Downstream PCIe Adapter.

5. Enable Transport Layer Packets through the Upstream PCIe Adapter by setting the ADP\_PCIE\_CS\_0.Path Enable bit to 1b in the Upstream PCIe Adapter.
6. Enable Transport Layer Packets through the Downstream PCIe Adapter by setting the ADP\_PCIE\_CS\_0.Path Enable bit to 1b in the Downstream PCIe Adapter.

**Table 5-1: PCIe Path Attributes**

Path Segment	Input HopID	Output HopID	Buffers Allocation: Dedicated	Priority	Weight	IFC	EFC	ISE	ESE
PCIe Adapter to USB4 Port	8		NA <sup>1</sup>	See Section 6.1.1	See Section 6.1.2.3	NA <sup>1</sup>	1	NA <sup>1</sup>	0
USB4 Port to PCIe Adapter		8	baMaxPCle from the Buffer Allocation Request Operation <sup>2</sup>	See Section 6.1.1	See Section 6.1.2.3	1	0	0	0
<ol style="list-style-type: none"> <li>1. A Connection Manager performs Read/Modify/Write to Path Configuration Space and does not change the PATH_CS_0.Path Credits Allocated, PATH_CS_1.IFC Flag and PATH_CS_1.ISE Flag fields at the PCIe Adapter</li> <li>2. For more information regarding Buffer allocation, see Section 5.1.4</li> </ol>									

### 5.3.2. PCIe Path Teardown

A Connection Manager performs the following steps to tear down the PCIe Paths between two Routers:

1. Disable Transport Layer Packets in the Downstream PCIe Adapter by setting the ADP\_PCIE\_CS\_0.Path Enable bit to 0b in the Adapter.
2. Disable Transport Layer Packets in the Upstream PCIe Adapter by setting the ADP\_PCIE\_CS\_0.Path Enable bit to 0b in the Adapter.
3. If the Router is a Version 2.0 Router and Extended Encapsulation is enabled, set the ADP\_PCIE\_CS\_1.Extended Encapsulation bit to 0b in both the Upstream PCIe Adapter and the Downstream PCIe Adapter to disable Extended Encapsulation.
4. Teardown the downstream Path as defined in Section 5.2.2.
5. Teardown the upstream Path as defined in Section 5.2.2.

## 5.4. DisplayPort Path Setup and Teardown

### 5.4.1. Connection Manager Discovery

Before setting up a Path between a DP IN Adapter and a DP OUT Adapter, a Connection Manager needs to make sure that the DP Adapters are available. Section 5.4.1.1 describes how a Connection Manager checks DP IN Adapter Availability. Section 5.4.1.2 describes how a Connection Manager checks DP OUT Adapter availability.

#### 5.4.1.1. DP IN Adapters

A DP IN Adapter is available for DP tunneling after all of the following occur:

1. The Connection Manager has either:
  - Received a Hot Plug Event Packet with *UPG* = 0b for a DP IN Adapter;

- The Connection Manager determines whether a Hot Plug Event Packet is sent by the Adapter by setting the *Disable Hot Plug Events* bit in the Adapter Configuration Space to the desired value. The bit shall be set before the *AUX Enable* bit and the *Video Enable* bit in the DP IN Adapter Configuration Capability are set to 1b

or

- Confirmed that the DP IN Adapter has an available DP resource. The Connection Manager uses a QUERY\_DP\_RESOURCE command with the DisplayPort Number parameter equal to the DP IN Adapter number to query DP resource availability.

2. The Connection Manager has allocated a DP stream resource to the DP IN Adapter:

- The Connection Manager uses an ALLOCATE\_DP\_RESOURCE command with a DisplayPort Number parameter equal to the DP IN Adapter number to allocate a DP resource.

Once a DP IN Adapter is available, all of the fields in its Adapter Configuration Capability Field are valid.

#### 5.4.1.2. DP OUT Adapters

A DP OUT Adapter is available for DP tunneling when either:

- The Connection Manager has received a Hot Plug Event Packet with UPG = 0b for the DP OUT Adapter;
  - The Connection Manager determines whether a Hot Plug Event Packet is sent by the Adapter by setting the *Disable Hot Plug Events* bit in the Adapter Configuration Space to the desired value. The bit shall be set before the *AUX Enable* bit and the *Video Enable* bit in the DP OUT Adapter Configuration Capability are set to 1b

or

- The ADP\_DP\_CS\_2.HPD Status field is set to 1b in the DP OUT Adapter Configuration Capability.

Once a DP OUT Adapter is available, all of the fields in its Adapter Configuration Capability are valid.

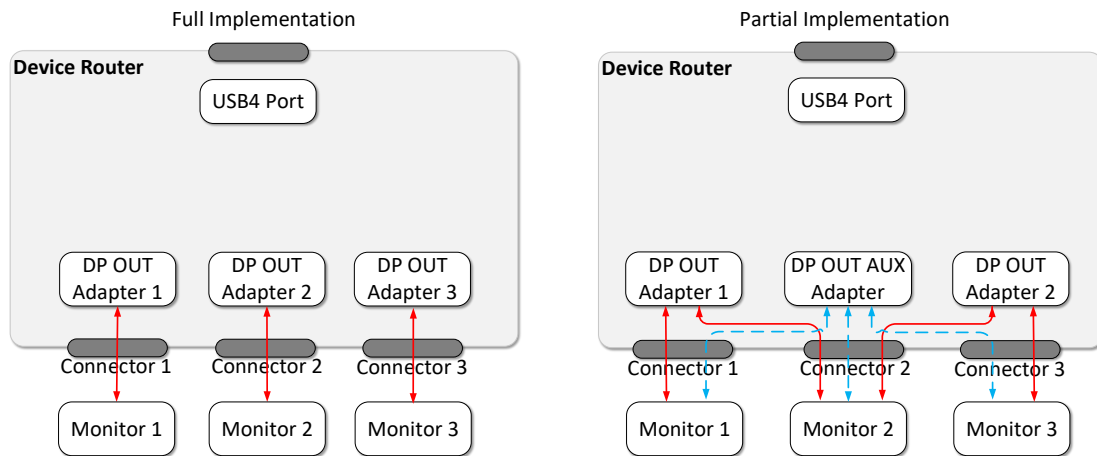
#### 5.4.1.3. DPRX Discovery

DPRX Discovery feature allows the System Software to discover all connected monitors irrespective of the availability of DP IN and DP OUT Adapters. This feature enables the System Software to determine which DP IN and DP OUT Adapters should be paired and which monitors should be functional.

DPRX Discovery feature can either be:

- A Full Implementation where there is 1 to 1 mapping between the Connectors and the DP OUT Adapters (refer Figure 5-3).
- A Partial Implementation without 1 to 1 mapping between the Connectors and the DP OUT Adapters with lesser number of DP OUT Adapters and implementation of DP OUT AUX Adapter with connectivity to all Connectors (refer Figure 5-3).

**Figure 5-3: Example - Full and Partial Implementation of DPRX Discovery Feature**



A Connection Manager that supports USB4 Version 2.0, may use DPRX Discovery feature to:

- Read the basic capabilities of DPRX before pairing DP IN and DP OUT Adapters.
- Read the DPRX capabilities when a monitor is connected, and no more DP IN Adapters are available.
- Read the DPRX capabilities when a monitor is connected, and no more DP OUT Adapters are available in case the Router supports partial implementation of DPRX Discovery (ROUTER\_CS\_6.Partial DP Connectivity Implementation bit is set to 1b).

The Connection Manager performs the steps listed below to discover the capabilities of a connected monitor by using the DPRX Discovery feature:

- The Connection Manager reads ROUTER\_CS\_6.Partial DP Connectivity Implementation bit of Device Router to determine if it supports Full or Partial implementation of DPRX Discovery.
- If the Device Router has Full Implementation of DPRX feature, that is if ROUTER\_CS\_6.Partial DP Connectivity Implementation bit is set to 0b, the Connection Manager sets ROUTER\_CS\_5.SW Mapping bit in Device Router to 0b for Router/HW controlled mapping of connector to DP OUT Adapter.
- If the Device Router has Partial Implementation of DPRX feature that is if ROUTER\_CS\_6.Partial DP Connectivity Implementation bit is set to 1b,
  - The Connection Manager retrieves the information about the Downstream facing USB Type-C ports and native DP connectors (referred to collectively as “Connectors”) that are capable of being connected to a DP OUT or DP OUT AUX Adapter, by using a Get Connectors Information Operation.
  - The Connection Manager may set ROUTER\_CS\_5.SW Mapping bit in Device Router to 0b for Router/HW controlled mapping of connector to DP OUT Adapter or may set ROUTER\_CS\_5.SW Mapping bit in Device Router to 1b for SW controlled mapping of connector to DP OUT Adapter.
- If the Connection Manager has set the ROUTER\_CS\_5.SW Mapping bit in Device Router to 0b (HW Mapping), upon a monitor connection or IRQ\_HPD, the Connection Manager may discover the capabilities of a monitor connected to a connector:

- For a Full Implementation the Router sends a Plug Packet (containing the TopologyID and Adapter Number) for the relevant DP OUT Adapter.
  - For a Partial Implementation if the DP OUT Adapter is available the Router connects the DP OUT Adapter to the connector and sends a Plug Packet for the relevant DP OUT Adapter. If DP OUT Adapter is not available, the Router does not send any Plug Packet for the flow to proceed further.
- If the Connection Manager has set the ROUTER\_CS\_5.SW Mapping bit in Device Router to 1b (SW Mapping, applicable for Partial Implementation only), upon a monitor connection to a connector or IRQ, the Device Router sends the Connection Manager a Notification Packet with *Event Code* = DP\_CON\_CHANGE.
  - Connection Manager notifies the SW that a monitor is ready for discovery.
  - Upon trigger from SW, the Connection Manager prepares for a monitor discovery.
  - For a Full Implementation, the Connection Manager sets up two Virtual AUX Tunneled Paths between the Host Interface DMA Ring and the DP OUT Adapter (refer Figure 5-4).
  - For a Partial Implementation, the Connection Manager
    - Connects a DP OUT AUX Adapter to the connector using the Connect DP OUT Adapter Operation.
    - After successfully mapping the Connectors to the DP OUT AUX Adapter in the Device Router, the Connection Manager sets two Virtual AUX Tunneled Paths (refer Figure 5-5).
  - The two Virtual AUX Tunneled Paths include
    - An AUX outbound Path from a Host Interface Transmit Ring to the DP OUT Adapter or the DP OUT AUX Adapter.
    - An AUX inbound Path from the DP OUT Adapter or DP OUT AUX Adapter to a Host Interface Receive Ring.
  - After receiving an HPD Packet (from DP OUT or DP OUT AUX Adapter), the SW knows that the Virtual Aux path is set. The SW, mimics a DP IN Adapter behavior, and sends the SET\_CONFIG Packets as defined in the USB4 Base specification. The SW then issues AUX Requests and receives AUX Responses and IRQ HDP (from DP OUT or DP OUT AUX Adapter) over the Virtual AUX Tunneled path.
  - After DPRX discovery is complete the SW notifies the Connection Manager the completion of the monitor discovery process.
  - The Connection Manager tears down the two Virtual AUX Tunneled Paths. In case of Partial Implementation after tearing down the Virtual AUX paths, the Connection Manager disconnects the DP OUT AUX Adapter from the connector by using Connect DP OUT Adapter Router Operation.
  - When directed by the SW, the Connection Manager sets up the full DP Tunneling path to a specific monitor. In case of Partial Implementation, the Connection Manager, and the Router use Connect DP OUT Adapter Router Operation to Connect DP OUT Adapter to the connector before setting up the DP Tunneling path. Refer Section 5.4.2 for pairing of DP IN and DP OUT Adapters and refer Section 5.4.3.3 for DP Path Setup.

**Figure 5-4: Example - Virtual AUX Tunneled Path for Full DPRX Implementation**

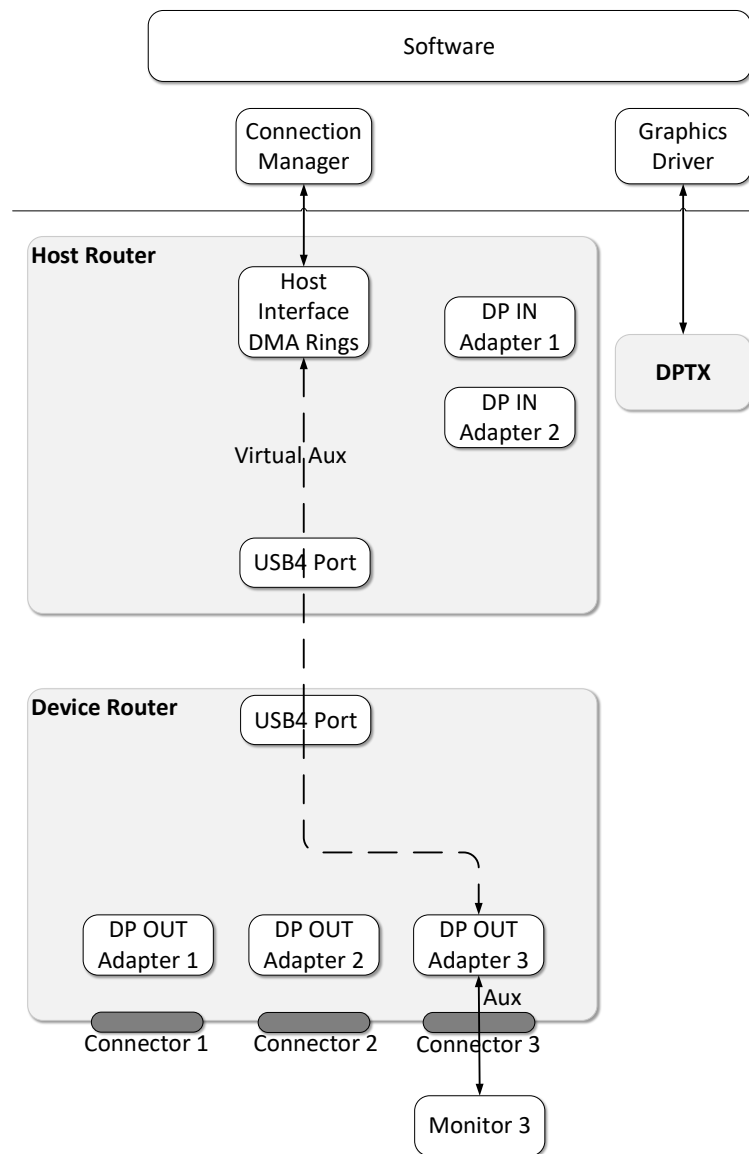
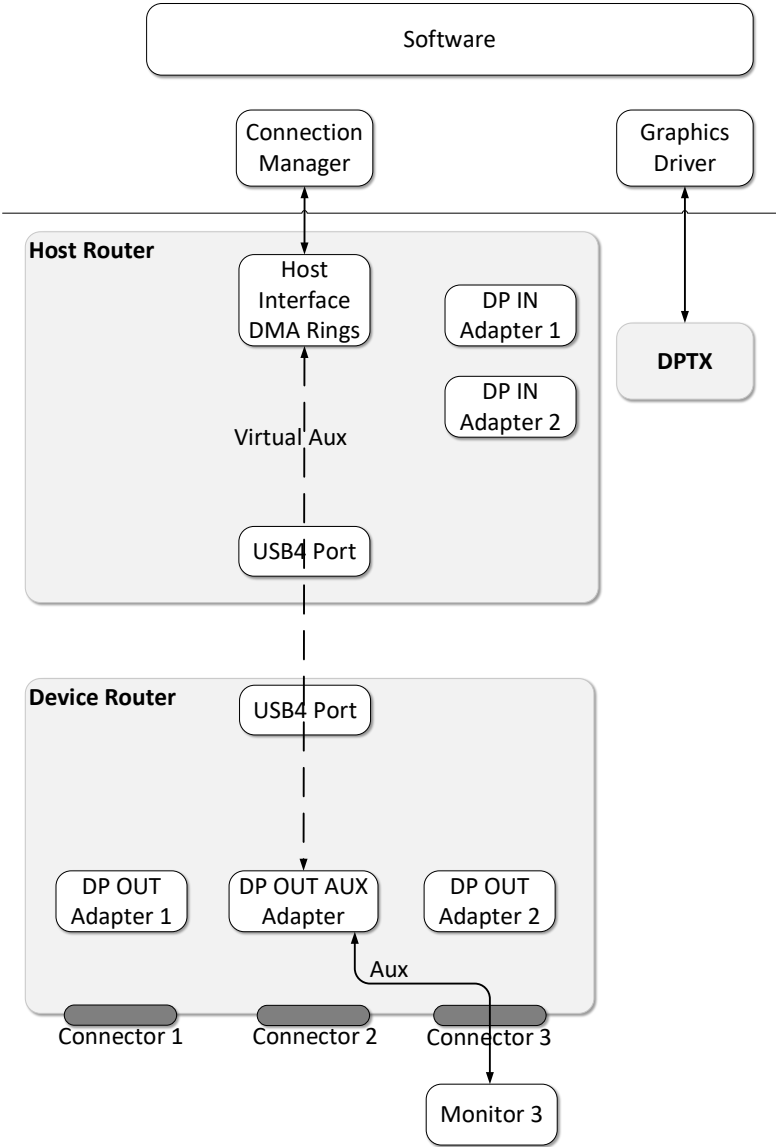


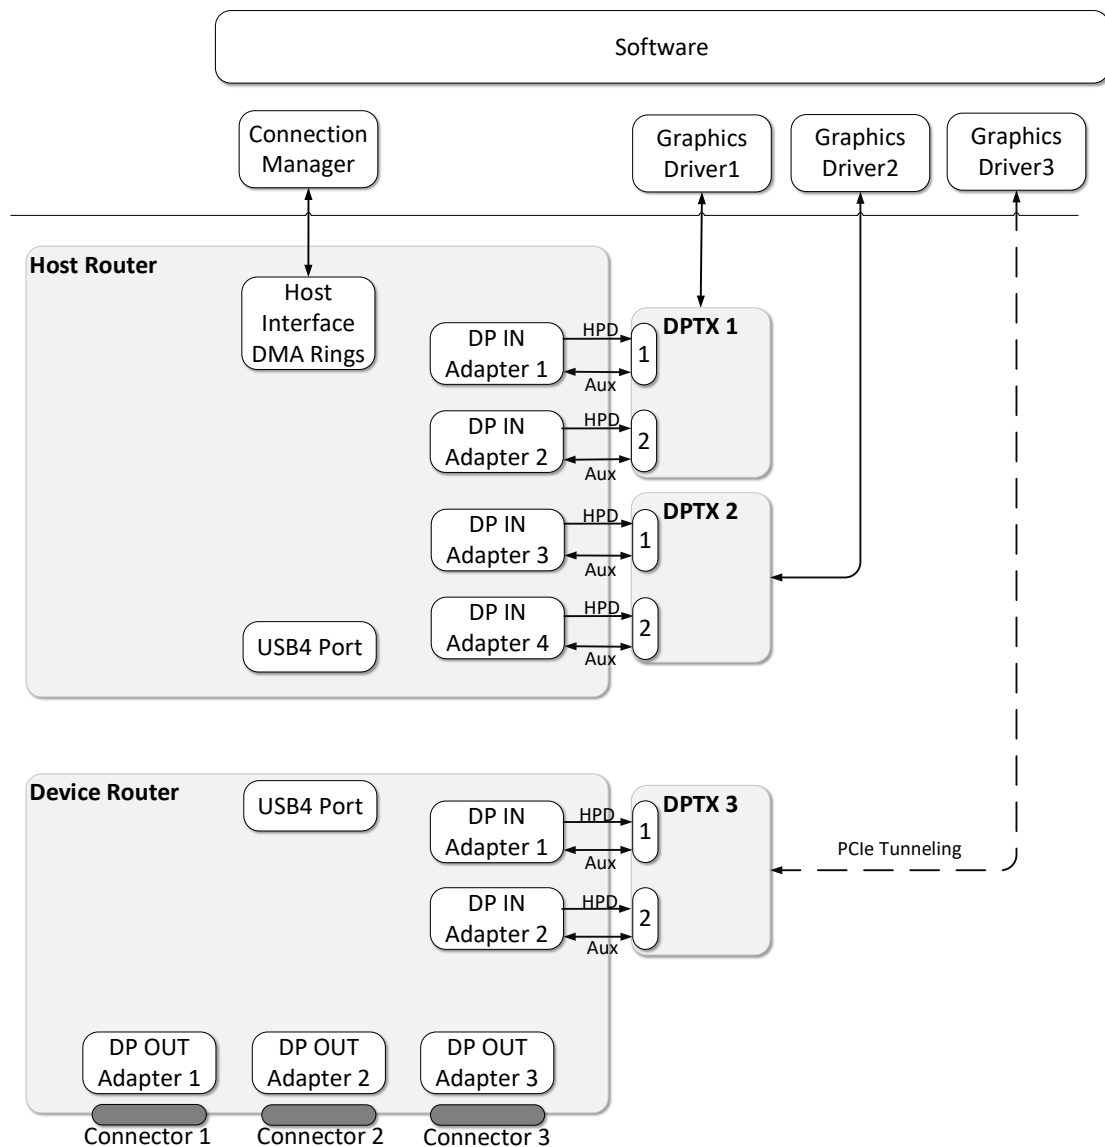
Figure 5-5: Example - Virtual AUX Tunneled Path for Partial DPRX Implementation



5.4.1.4. DPTX Discovery

The DPTX Discovery feature enables system software to discover the connectivity between a DPTX and a DP IN Adapter. System Software can use this information to choose the pairing of a DPTX and a DPRX in case there are multiple Graphics Processing Units (GPUs). Figure 5-6 illustrates an example topology with multiple GPUs and DPTX discovery feature.

**Figure 5-6: Example topology with DPTX discovery feature**

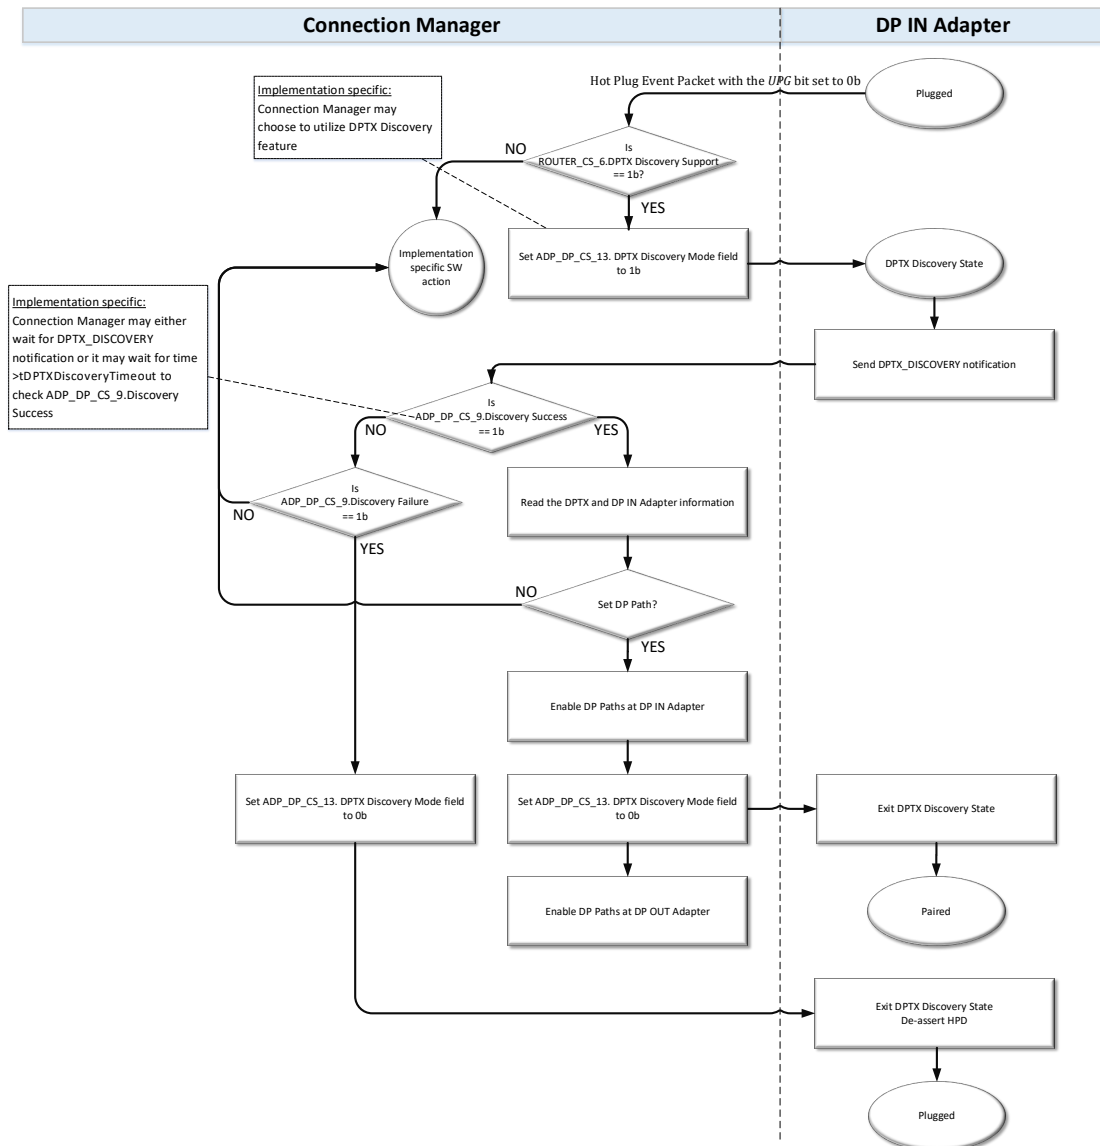


Refer Figure 5-7 for the DPTX discovery flow.

- When the Connection Manager receives a Hot Plug Event Packet with the *UPG* bit set to 0b for a DP IN Adapter, if *ROUTER\_CS\_6.DPTX Discovery Support* is set to 1b, the Connection Manager may set *ADP\_DP\_CS\_13.DPTX Discovery Mode* field to 1b in the DP IN Adapter to enable DPTX discovery mode.
- The Connection Manager may either wait for the Notification Packet with Event Code = *DPTX\_DISCOVERY* or it may wait for a time > *tDPTXDiscoveryTimeout* to check if *ADP\_DP\_CS\_9.Discovery Success*, bit is set to 1b or if the *ADP\_DP\_CS\_9.Discovery Failure*, bit is set to 1b.
  - If *ADP\_DP\_CS\_9.Discovery Success*, bit is set to 1b, the Connection Manager reads and stores the DPTX information.
    - If the Connection Manager chooses to set the DP Paths while in DPTX Discovery mode, it:
      - Sets up all 3 DP Paths.

- Enables the DP Paths at the DP IN Adapter (sets AE and VE to 1b).
  - Sets ADP\_DP\_CS\_13.DPTX Discovery Mode bit is set to 0b.
  - Enables the DP Paths at the DP OUT Adapter (sets AE and VE to 1b).
- If ADP\_DP\_CS\_9.Discovery Failure, bit is set to 1b, the Connection Manager sets ADP\_DP\_CS\_13.DPTX Discovery Mode bit to 0b to exit DPTX Discovery Mode. The DP IN Adapter de-asserts HPD before a new DP Path is set up.

**Figure 5-7: DPTX Discovery flow**



#### 5.4.2. Pairing DP Adapters

A Connection Manager decides which DP IN Adapter associates with which DP OUT Adapter. This process is called “pairing”. The method for pairing a DP IN Adapter with a DP OUT Adapter is implementation specific.

The following is an example of how a Connection Manager can pair DP Adapters:

- DP Adapters are discovered during Adapter Enumeration (see Section 3.2.2). Each DP Adapter found during Adapter Enumeration is stored as a potential DP resource to be used.
- If there is a change in the DP resources (plug or unplug), the Connection Manager goes over the stored display resources and searches for a free pair of DP IN Adapter and a DP OUT Adapter
- Once the Connection Manager finds a DP IN Adapter and a DP OUT Adapter with the desired characteristics, and decides to pair the two Adapters, it will setup Paths as described in Section 5.4.3.
- A Connection Manager may optionally use DPRX Discovery (refer Section 5.4.1.3) feature to discover DPRX and its capabilities before pairing DP IN and DP OUT Adapters.

#### **5.4.3. Display Port Path Setup and Teardown**

Before configuring a Path between a set of paired DP Adapters, the Connection Manager must do the following:

1. Determine the available USB4 bandwidth as described in Section 6.2.1.
2. Determine the relevant DP IN and DP OUT Adapter Capabilities as described in Section 5.4.3.1.
3. Configure the Adapters along the Path as described in Section 5.4.3.2.

Section 5.4.3.3 describes how a Path is configured.

##### **5.4.3.1. Capabilities Exchange**

A Connection Manager determines the capabilities of the DP Adapters as follows:

1. Read the DP\_LOCAL\_CAP.*Protocol Adapter Version*:
  - If the value is less than 4, DP Tunneling over USB4 is not supported and a DP Path cannot be set up.
  - Else, continue with steps below.
2. Set the DP\_STATUS\_CTRL.*UF* field to 1b in the DP OUT Adapter.
3. Set the DP\_STATUS\_CTRL.*CMHS* (Connection Manager Handshake) field to 1b in the DP OUT Adapter.
4. Poll the DP\_STATUS\_CTRL.*CMHS* field in the DP OUT Adapter until it is reset to zero.

*Note: The Router sets DP\_STATUS\_CTRL.CMHS to 0b within 5ms after DP\_STATUS\_CTRL.CMHS is set to 1b. It is recommended that the Connection Manager waits for 50ms before it times out.*

5. Read the DP\_LOCAL\_CAP register of the DP Adapters at each end of the Path.
6. Copy the value read from the DP\_LOCAL\_CAP register of the DP IN Adapter to the DP\_REMOTE\_CAP register of the DP OUT Adapter.
7. Copy the value read from the DP\_LOCAL\_CAP register of the DP OUT Adapter to the DP\_REMOTE\_CAP register of the DP IN Adapter.

- Limit DP BW – If the Connection Manager determines that the DP BW needs to be limited, (see Section 6.2.1), it limits the DisplayPort bandwidth by writing to the *DP\_REMOTE\_CAP.Remote Maximal Link Rate* and *DP\_REMOTE\_CAP.Remote Maximal Lane Count* fields of the DP IN Adapter. Besides the *Remote Maximal Link Rate* and *Remote Maximal Lane Count* fields, the Connection Manager does not change any other fields in the *DP\_REMOTE\_CAP*.
8. Set the *ADP\_DP\_CS\_2.CM\_ID* field in the DP IN Adapter to the Connection Manager index, which is the unique Connection Manager ID where the USB4 Host has multiple Host Routers.
  9. If the Connection Manager supports DP BW Allocation Mode, it does the following:
    - Checks if the DP IN Adapter supports DP BW Allocation Mode by reading *DP\_LOCAL\_CAP.DP\_IN\_BW\_Allocation Mode Support*.
    - If the DP IN Adapter supports DP BW Allocation Mode, the Connection Manager does the following in the DP IN Adapter Configuration Space:
      - Sets *ADP\_DP\_CS\_2.CM BW Allocation Mode Support* to 1b.
      - If the Path being setup between the two Adapters travels through the same USB4 Links as another DP Path, then it sets *ADP\_DP\_CS\_2.Group\_ID* to a non-zero value for all the Paths which travel through the same USB4 Links, otherwise it sets *ADP\_DP\_CS\_2.Group\_ID* to 0h.
      - Sets *ADP\_DP\_CS\_2.NRD\_Maximal\_Link\_Rate* to the *DP\_LOCAL\_CAP.Maximal\_Link\_Rate* of the DP IN Adapter or DP OUT Adapter (whichever is lower).
      - Sets *ADP\_DP\_CS\_2.NRD\_Maximal\_Lane\_Count* to the *DP\_LOCAL\_CAP.Maximal\_Lane\_Count* of the DP IN Adapter or DP OUT Adapter (whichever is lower).
      - Sets *ADP\_DP\_CS\_2.Estimated BW* according to Section 6.2.1.3.1.
      - Sets *ADP\_DP\_CS\_2.Granularity* to any of the valid values.
      - Sets *DP\_STATUS.Allocated BW* to 0h.
      - Sets *ADP\_DP\_CS\_2.CM Ack* to 0b.
  10. If ALPM is supported, set the *ADP\_DP\_CS\_13.CL1 Exit Time* field in the DP IN Adapter Configuration Space to the time it takes the USB4 Links, that the DP Path traverses through, to exit CL1 state.

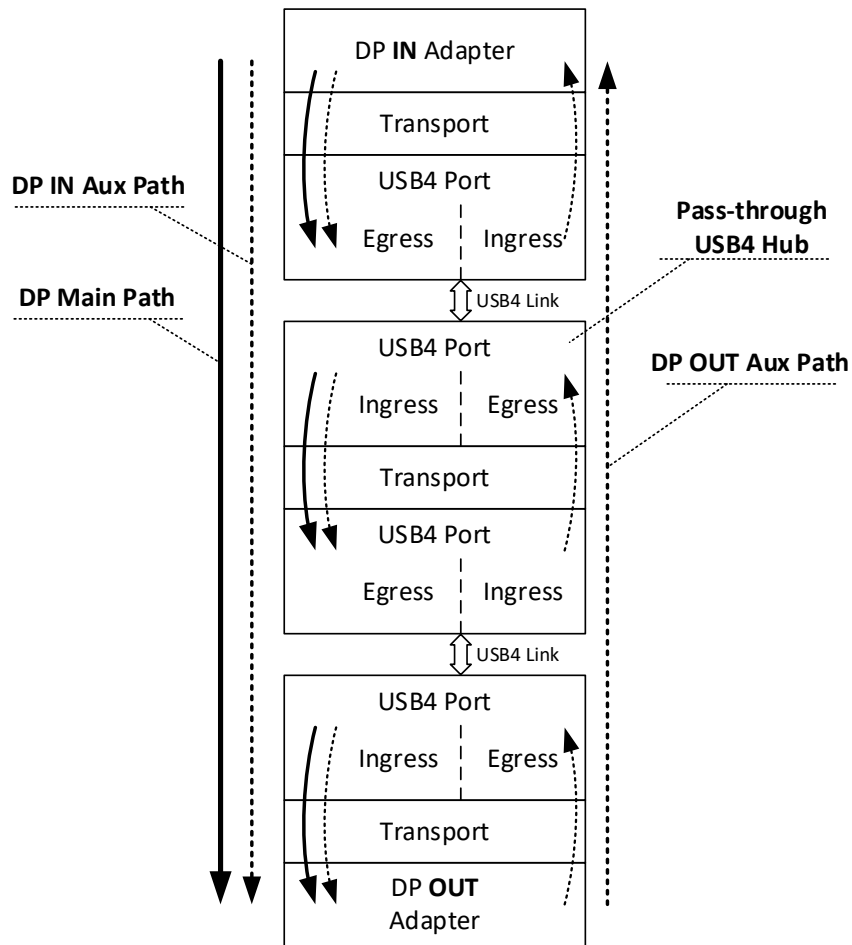
#### **5.4.3.2. Adapter Configuration**

A Connection Manager needs to increase the number of Non-Flow Control Buffers at each Ingress Lane Adapter that the Main-Link Path goes through. The number of buffers to be added by the Connection Manager is listed in Table 5-2 and is programmed to the *Non-Flow Control* field in the Adapter Configuration Space of the Lane Adapter.

#### **5.4.3.3. DP Path Setup**

A Connection Manager establishes three Paths for DP Tunneling: a DP Main Path, a DP IN Aux Path, and a DP OUT Aux Path. Figure 5-8 shows an example of a DP Tunneling scheme with its basic constructs.

**Figure 5-8: DP Tunneling Setup**



A Connection Manager sets up the three DP Paths as follows. The Paths may be configured in any order:

1. Setup DP Main Path:
  - In the Router that interfaces with the DP Source, configure the DP Main Path segment that goes from the DP IN Adapter to the Lane Adapter. This is done by writing the Path Attributes in Table 5-2 to the Path Configuration Space of the DP IN Adapter.
  - For each USB4 Hub that the Path traverses, configure the DP Main Path segment between the Ingress and Egress Ports on the hub's Router. This is done by writing the Path Attributes in Table 5-2 to the Path Configuration Space of the Ingress Lane Adapter.
  - In the Router that interfaces with the DP Sink, configure the DP Main Path segment that goes from the Ingress Lane Adapter to the DP Sink. This is done by writing the Path Attributes in Table 5-2 to the Path Configuration Space of the Ingress Lane Adapter.
  - If the *USB4 Version* field in the Router Configuration Space of the Router that implements the DP IN Adapter is set to 2.0 or higher, and the Connection Manager supports USB4 Version 2.0, the Connection Manager sets the *PATH\_CS\_0.PM Packet Support* to 1b in the Path Configuration

Space of the Ingress Lane Adapter in USB4 Version 2.0 hubs along the DP Main Path.

*Note: The PATH\_CS\_0.PM Packet Support bit cannot be changed after the path valid (PATH\_CS\_0.Valid) bit is set to 1b.*

## 2. Setup DP IN Aux Path:

- In the Router that interfaces with the DP Source, configure the DP IN Aux Path segment that goes from the DP IN Adapter to Lane Adapter. This is done by writing the Path Attributes in Table 5-2 to the Path Configuration Space of the DP IN Adapter.
- For each USB4 Hub that the Path traverses, configure the DP IN Aux Path segment that goes from the Ingress Lane Adapter to the Egress Lane Adapter. This is done by writing the Path Attributes in Table 5-2 to the Path Configuration Space of the Ingress Lane Adapter.
- In the Router that interfaces with the DP Sink, configure the DP IN Aux Path segment that goes from the Ingress Lane Adapter to the DP OUT Adapter. This is done by writing the Path Attributes in Table 5-2 to the Path Configuration Space of the Ingress Lane Adapter.
- If the *USB4 Version* field in the Router Configuration Space of the Router that implements the DP IN Adapter is set to 2.0 or higher, and the Connection Manager supports USB4 Version 2.0, the Connection Manager sets the *PATH\_CS\_0.PM Packet Support* to 1b in the Path Configuration Space of the Ingress Lane Adapter in USB4 Version 2.0 hubs along the DP IN Aux Path.

*Note: The PATH\_CS\_0.PM Packet Support bit cannot be changed after the path valid (PATH\_CS\_0.Valid) bit is set to 1b.*

## 3. Setup DP OUT Aux Path:

- In the Router that interfaces with the DP Sink, configure the DP OUT Aux Path segment that goes from the DP OUT Adapter to Lane Adapter. This is done by writing the Lane Attributes in Table 5-2 to the Path Configuration Space of the DP OUT Adapter.
- For each USB4 Hub that the Path traverses, configure the DP OUT Aux Path segment that goes from Ingress Lane Adapter to the Egress Lane Adapter. This is done by writing the Lane Attributes in Table 5-2 to the Path Configuration Space of the Lane Adapter.
- In the Router that interfaces with the DP Source, configure the DP OUT Aux Path segment that goes from the Ingress Lane Adapter to DP IN Adapter. This is done by writing the Lane Attributes in Table 5-2 to the Path Configuration Space of the Lane Adapter.
- If the *USB4 Version* field in the Router Configuration Space of the Router that implements the DP IN Adapter is set to 2.0 or higher, and the Connection Manager supports USB4 Version 2.0, the Connection Manager sets the *PATH\_CS\_0.PM Packet Support* to 1b in the Path Configuration Space of the Ingress Lane Adapter in USB4 Version 2.0 hubs along the DP OUT Aux Path.

*Note: The PATH\_CS\_0.PM Packet Support bit cannot be changed after the path valid (PATH\_CS\_0.Valid) bit is set to 1b.*

A Connection Manager enables the Adapters according to the exact order as follows. (It allows the HPD Tunneled Packet which is sent by the DP OUT Adapter to be received at the DP IN Adapter):

1. Set ADP\_DP\_CS\_0.AE to 1b and set ADP\_DP\_CS\_0.VE to 1b in the Adapter Configuration Space of the DP IN Adapter to enable the Paths.
  - A Connection Manager shall set the values in the DP IN Adapter Configuration Capability before it sets the *Video Enable* bit and the *Aux Enable* bit in the Adapter to 1b. However, the Connection Manager may set the Bandwidth Management fields (i.e., *CM Ack*, *Granularity*, *Estimated BW*, and *Allocated BW*) at any time after the DP resource is allocated via an Allocate DP Resource Router Operation.
2. If the Connection Manager enables the DP Path while in DPTX Discovery Mode it shall set the DPTX Discovery Mode bit to 0b before setting ADP\_DP\_CS\_0.AE to 1b and set ADP\_DP\_CS\_0.VE to 1b in the Adapter Configuration Space of the DP OUT Adapter to enable the Paths.
  - A Connection Manager shall set the values in the DP OUT Adapter Configuration Capability after it receives a Hot Plug Event Packet with the *UPG* bit set to 0b from the Adapter, and before it sets the *Video Enable* bit and the *Aux Enable* bit in the Adapter to 1b.

**Table 5-2: DP Path Attributes**

Path Segment	Input HopID	Output HopID	Buffers Allocation	Priority	Weight	IFC	EFC	ISE	ESE
DP AUX: DP Adapter to USB4	8		NA <sup>1</sup>	See Section 6.1.1	See Section 6.1.2.3	NA <sup>1</sup>	1	NA <sup>1</sup>	0
DP AUX: USB4 to USB4			Dedicated Flow Control: baMinDPaux from the Buffer Allocation Request Operation <sup>2</sup>	See Section 6.1.1	See Section 6.1.2.3	1	1	0	0
DP AUX: USB4 to DP Adapter		8		See Section 6.1.1	See Section 6.1.2.3	1	0	0	0
DP Main: DP IN Adapter to USB4	9		NA <sup>1</sup>	See Section 6.1.1	See Section 6.1.2.3	NA <sup>1</sup>	0	NA <sup>1</sup>	0
DP Main: USB4 to USB4			Flow Control Disabled: baMinDPmain from the Buffer Allocation Request Operation <sup>2</sup>	See Section 6.1.1	See Section 6.1.2.3	0	0	0	0
DP Main: USB4 to DP OUT Adapter		9		See Section 6.1.1	See Section 6.1.2.3	0	0	0	0
1. A Connection Manager performs Read/Modify/Write to Path Configuration Space and does not change the PATH_CS_0.Path Credits Allocated, PATH_CS_1.IFC Flag and PATH_CS_1.ISE Flag fields at the DP Adapter.									
2. For more information regarding Buffer allocation, see Section 5.1.4.									

#### 5.4.3.4. DP Path Teardown

A Connection Manager performs the following steps to teardown the three DP Paths:

1. Disable the Paths in the DP OUT Adapter by setting ADP\_DP\_CS\_0.AE bit to 0b and the ADP\_DP\_CS\_0.VE bit to 0b

2. Disable the Paths in the DP IN Adapter by setting ADP\_DP\_CS\_0.AE bit to 0b and the ADP\_DP\_CS\_0.VE bit to 0b
3. Teardown the DP Main Path according to Section 5.2.2
4. Teardown the DP OUT Aux Path according to Section 5.2.2
5. Teardown the DP IN Aux Path according to Section 5.2.2

After tearing down a DP Path, a Connection Manager initiates a DEALLOCATE\_DP\_RESOURCE Router Operation to release the DP stream resource. The DEALLOCATE\_DP\_RESOURCE Operation has the DisplayPort Number parameter equal to the DP IN Adapter number being released.

## 5.5. USB3 Gen X Path Setup and Teardown

This section describes how a Connection Manager sets up and tears down the USB3 Gen X Paths between two Routers.

### 5.5.1. USB3 Gen X Path Setup

A USB3 Gen X Path is created between the Upstream USB3 Gen X Adapter and a Downstream USB3 Gen X Adapter of two physically connected Routers within the same Domain. A Connection Manager does not set up a USB3 Gen X Path between two Routers unless they are directly connected to each other.

A Connection Manager can set up a Path for USB3 Gen X Tunneling when it first enumerates a Router that supports USB3 Gen X Tunneling. It can also wait until some or all Routers in the Domain are enumerated.

A Connection Manager does not set up a Gen X Path to the Upstream USB3 Gen X Adapter in a Device Router if the following conditions are true:

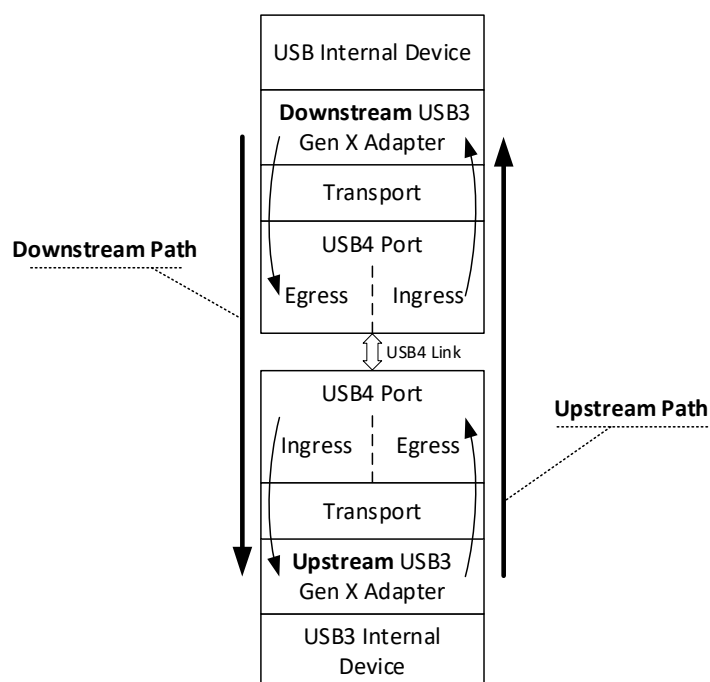
1. The Connection Manager has set up a USB3 Gen T Path to the Upstream USB3 Gen T Adapter for each of the Gen T Ports, and
2. The ADP\_USB3\_GT\_CS\_0.Gen X Adapter Coupled bit is set to 1b in that Upstream USB3 Gen T Adapter,

If ADP\_USB3\_GT\_CS\_0.Gen X Adapter Coupled bit is set to 0b in the Upstream USB3 Gen T Adapter, the Connection Manager shall set a USB3 Gen X Path to the Upstream USB3 Gen X Adapter, regardless if a USB3 Gen T Path is set or not.

If the Connection Manager does not setup a Path to a USB3 Gen X Adapter, it disconnects the Adapter as defined in Section 5.5.3.

A Connection Manager needs to establish two Paths, one upstream and one downstream, for USB3 Gen X Tunneling. Figure 5-9 shows a USB3 Gen X Tunneling scheme with its basic constructs.

**Figure 5-9: USB3 Gen X Tunneling Setup**



A Connection Manager performs the following steps to establish a USB3 Gen X Path:

1. Wait at least 500ms from previous Path tear down.
2. Setup Downstream Path:
  - Set Path Attributes in the Path Configuration Space of the Downstream USB3 Gen X Adapter. The Path Attributes are defined in Table 5-3 (Path segment is from the Downstream USB3 Gen X Adapter to the USB4 Port).
  - Set Path Attributes in the Path Configuration Space of the Lane Adapter. The Path Attributes are defined in Table 5-3 (Path segment is from the USB4 Port to the Upstream USB3 Gen X Adapter).
3. Setup Upstream Path:
  - Set Path Attributes in the Path Configuration Space of the USB3 Upstream Adapter. The Path Attributes are defined in Table 5-3 (Path segment is from the USB3 Upstream Adapter to the USB4 Port).
  - Set Path Attributes in the Path Configuration Space of the Lane Adapter. The Path Attributes are defined in Table 5-3 (Path segment is from the USB4 Port to the USB3 Downstream Adapter).
4. Enable Transport Layer Packets through the Downstream USB3 Gen X Adapter by setting the `ADP_USB3_GX_CS_0.Path Enable` bit to 1b and the `ADP_USB3_GX_CS_0.Valid` bit to 1b in the Adapter. The *Valid* bit should not be set before the *Path Enable* bit is set.
5. Enable Transport Layer Packets through the Upstream USB3 Gen X Adapter by setting the `ADP_USB3_GX_CS_0.Path Enable` bit to 1b and the `ADP_USB3_GX_CS_0.Valid` bit to 1b in the Adapter. The *Valid* bit should not be set before the *Path Enable* bit is set.

**Table 5-3: USB3 Gen X Path Attributes**

Path Segment	Input HopID	Output HopID	Buffers Allocation: Dedicated Flow Control	Priority	Weight	IFC	EFC	ISE	ESE
USB3 Gen X Adapter to USB4 Port	8		NA	See Section 6.1.1	See Section 6.1.2.3	NA <sup>1</sup>	1	NA <sup>1</sup>	0
USB4 Port to USB3 Gen X Adapter		8	baMaxUSB3GenX from the Buffer Allocation Request Operation <sup>2</sup>	See Section 6.1.1	See Section 6.1.2.3	1	0	0	0
<ol style="list-style-type: none"> <li>1. A Connection Manager performs Read/Modify/Write to Path Configuration Space and does not change the PATH_CS_0.Path Credits Allocated, PATH_CS_1.IFC Flag and PATH_CS_1.ISE Flag fields at the USB3 Gen X Adapter.</li> <li>2. For more information regarding Buffer allocation, see Section 5.1.4.</li> </ol>									

### 5.5.2. USB3 Gen X Path Teardown

A Connection Manager performs the following steps to tear down the USB3 Gen X Paths between two Routers:

1. Disable the Upstream USB3 Gen X Adapter by setting the ADP\_USB3\_GX\_CS\_0.Path Enable bit to 0b and the ADP\_USB3\_GX\_CS\_0.Valid bit to 1b in the Adapter.
2. Disable the Downstream USB3 Gen X Adapter by setting the ADP\_USB3\_GX\_CS\_0.Path Enable bit to 0b and the ADP\_USB3\_GX\_CS\_0.Valid bit to 1b in the Adapter.
3. Teardown the downstream Path according to Section 5.2.2.
4. Teardown the upstream Path according to Section 5.2.2.

### 5.5.3. USB3 Gen X Adapter Disconnect

When the Connection Manager sets the ADP\_USB3\_GX\_CS\_0.Path Enable bit to 0b and the ADP\_USB3\_GX\_CS\_0.Valid bit to 1b, it initiates a USB3 Gen X Adapter Disconnect. A USB3 Gen X Adapter disconnect causes the USB3 Gen X Adapter to remove far-end receiver termination to the internal USB3 device.

## 5.6. USB3 Gen T Path Setup and Teardown

This section describes how a Connection Manager sets up and tears down the USB3 Gen T Paths between two Routers.

### 5.6.1. USB3 Gen T Path Setup

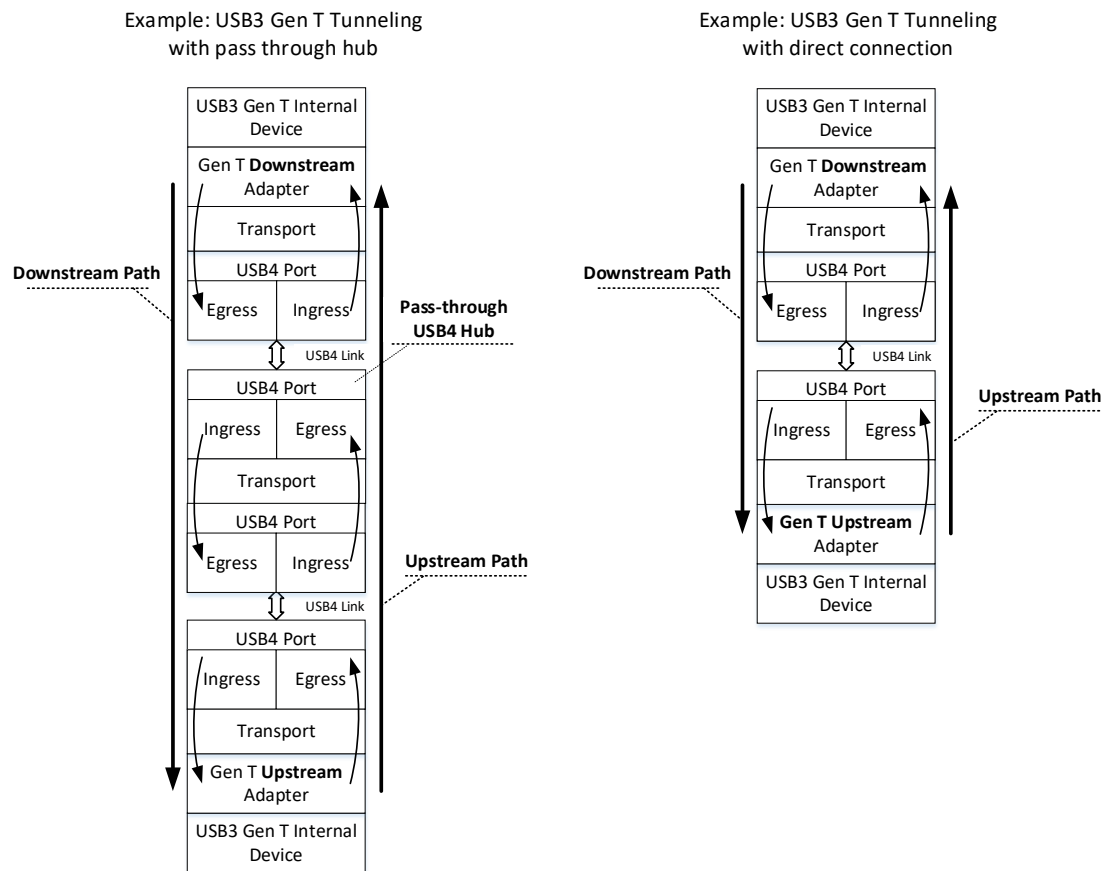
A Connection Manager can set up a point-to-point Gen T Path between a Gen T Upstream Adapter and a Gen T Downstream Adapter. There may be one or more pass through hubs between the two Routers (with Gen T Adapters) within the same Domain or these Routers may be directly connected to each other (refer Figure 5-10).

A Gen T Adapter may interface to one or more USB3 Gen T Ports. Each USB3 Gen T Port supports one USB3 Gen T Path in each direction. Each Gen T Port is assigned a unique index number, the Gen T Port Index. Though a USB3 Gen T Port may be mapped to multiple USB3 Gen T Adapters, a USB3 Gen T Port can work with one USB3 Gen T Adapter at a time (see example in the USB4 Specification).

A Connection Manager can set up a Path for Gen T Tunneling when it first enumerates a Device Router that supports Gen T Tunneling provided the Host Router supports Gen T Tunneling. It can also wait until some, or all Routers in the Domain are enumerated.

A Connection Manager needs to establish two Paths, one upstream and one downstream, for Gen T tunneling. Figure 5-10 shows the two Gen T Tunneling Paths.

**Figure 5-10: USB3 Gen T Tunneling Setup**



Before setting up a USB3 Gen T Path to a USB3 Gen T Port, a Connection Manager requests to allocate the Gen T Port by interfacing with the Host Controller software. A Connection Manager uniquely identifies a Gen T Port according to the following fields:

- `ADP_USB3_GT_PORT_CS_0.USB3 Host Controller Index`
- `ADP_USB3_GT_PORT_CS_0.Gen T Port Number`

If the Gen T Port cannot be allocated, the Connection Manager continues to check the availability of the other Gen T Ports that are mapped to the USB3 Gen T Downstream. If a Gen T Port is allocated, the Connection Manager proceeds with USB3 Gen T Path setup. If none of the Gen T Ports are available, the Connection Manager sets the `ADP_USB3_GT_PORT_CS_1.Gen T Port Not Available` bit to 1b in the Upstream USB3 Gen T Adapter Configuration Capability. If the `ADP_USB3_GT_CS_0.Gen X Adapter Coupled` bit is set to 1b in the Upstream USB3 Gen T Adapter, it sets up a Gen X Path (refer Section 5.5.1).

The Connection Manager performs the following steps to establish a USB3 Gen T Path:

1. Waits at least 500ms from previous Path tear down to the same Gen T Port
2. Setup two paths:
  - i. Setup Downstream Path:

- In the Router with Downstream Gen T Adapter, configure the Downstream Path segment that goes from the Downstream Gen T Adapter to Lane Adapter. This is done by writing the Path Attributes in Table 5-4 to the Path Configuration Space for the Downstream Gen T Adapter.
- For each USB4 Hub that the Path traverses, configure the Downstream Path segment that goes from the Ingress Lane Adapter to the Egress Lane Adapter. This is done by writing the Path Attributes in Table 5-4 to the Path Configuration Space for the Ingress Lane Adapter.
  - The Connection Manager sets the *PATH\_CS\_0.PM Packet Support* to 1b in the Path Configuration Space of the Ingress Lane Adapter in USB4 Version 2.0 hubs along the Gen T Path.

*Note: The PATH\_CS\_0.PM Packet Support bit cannot be changed after the path valid (PATH\_CS\_0.Valid) bit is set to 1b*

- In the Router with Upstream Gen T Adapter, configure the Downstream Path segment that goes from the Ingress Lane Adapter to the Upstream Gen T Adapter. This is done by writing the Path Attributes in Table 5-4 to the Path Configuration Space of the Ingress Lane Adapter.

## ii. Setup Upstream Path

- In the Router that Upstream Gen T Adapter, configure the Upstream Path segment that goes from the Upstream Gen T Adapter to Lane Adapter. This is done by writing the Lane Attributes in Table 5-4 to the Path Configuration Space for the Upstream Gen T Adapter.
- For each USB4 Hub that the Path traverses, configure the Upstream Path segment that goes from Ingress Lane Adapter to the Egress Lane Adapter. This is done by writing the Lane Attributes in Table 5-4 to the Path Configuration Space for the Lane Adapter.
  - The Connection Manager sets the *PATH\_CS\_0.PM Packet Support* to 1b in the Path Configuration Space of the Ingress Lane Adapter in USB4 Version 2.0 hubs along the Gen T Path.

*Note: The PATH\_CS\_0.PM Packet Support bit cannot be changed after the path valid (PATH\_CS\_0.Valid) bit is set to 1b*

- In the Router with Downstream Gen T Adapter, configure the Upstream Path segment that goes from the Ingress Lane Adapter to Downstream Gen T Adapter. This is done by writing the Lane Attributes in Table 5-4 to the Path Configuration Space for the Lane Adapter.

## iii. The Connection Manager sets the ADP\_USB3\_GT\_PORT\_CS\_1.Active Tx Link Rate field in both the Downstream and Upstream USB3 Gen T Adapters to lowest link rate among:

- The local USB3 Gen T Port Maximum Tx Supported Link Rate (ADP\_USB3\_GT\_PORT\_CS\_0. *Maximum Tx Supported Link Rate*).
- The remote USB 3 Gen T Port Maximum Rx Supported Link Rate (ADP\_USB3\_GT\_PORT\_CS\_0. *Maximum Rx Supported Link Rate*).
- The USB4 Link speed that the USB3 Gen T Path traverses.

- iv. The Connection Manager sets the ADP\_USB3\_GT\_PORT\_CS\_1.Active Rx Link Rate field in both the Downstream and Upstream USB3 Gen T Adapters to the lowest link rate among:
  - The local USB3 Gen T Port Maximum Rx Supported Link Rate (ADP\_USB3\_GT\_PORT\_CS\_0. Maximum Rx Supported Link Rate).
  - The remote USB 3 Gen T Port Maximum Tx Supported Link Rate (ADP\_USB3\_GT\_PORT\_CS\_0. Maximum Tx Supported Link Rate).
  - The USB4 Link speed that the USB3 Gen T Path traverses
- v. Set ADP\_USB3\_GT\_PORT\_CS\_1.Link Commands Aggregation Enable bit to 1b in both the Downstream and Upstream USB3 Gen T Adapters which belong to the same Path if there are less than 3 hubs between the Host Router and the Device Router.
- vi. Set ADP\_USB3\_GT\_PORT\_CS\_1.Path Enable bit in the USB3 Gen T Adapter Configuration Capability to 1b in both the Downstream and Upstream USB3 Gen T Adapters.
- vii. The Connection Manager may set the ADP\_USB3\_GT\_PORT\_CS\_1.U2CL2 Enable bit to 1b, if either of conditions listed below are true:
  - The endpoint reports a BELT value greater than the U2 exit latency plus tConvergeTime.
  - The endpoint reports a BELT value greater than the U2 exit latency and it sets the Accurate ITP Required bit to 0b in the Device BOS Descriptor (refer Section 9.4.3.2 of USB4 Version 2.0 Specification).

*Note: The Connection Manager may obtain the BELT value and the Accurate ITP Required by interacting with the USB Host Controller Software.*

**Table 5-4: USB3 Gen T Path Attributes**

Path Segment	Input HopID	Output HopID	Buffers Allocation: Shared Flow Control	Priority	Weight	IFC	EFC	ISE	ESE
USB3 Gen T Adapter to USB4 Port	8+n <sup>3</sup>		NA	See Section 6.1.1	See Section 6.1.2.3	0	1	0	1
USB4 Port to USB4 Port			baMaxUSB3GenT from the Buffer Allocation Request Operation <sup>2</sup>	See Section 6.1.1	See Section 6.1.2.3	1	1	1	1
USB4 Port to USB3 Gen T Adapter		8+n <sup>3</sup>		See Section 6.1.1	See Section 6.1.2.3	1	0	1	0
Notes: 1. The table details the configuration for the recommended Shared Buffer scheme. In case of Dedicated Buffer scheme is used, the ISE and ESE Flags shall be set to 0b. 2. For more information regarding Buffer allocation, see Section 5.1.4. 3. Where n = the Gen T Index of the USB3 Gen T Port									

### 5.6.2. USB3 Gen T Path Teardown

A Connection Manager performs the following steps to tear down the USB3 Gen T Paths between two Routers:

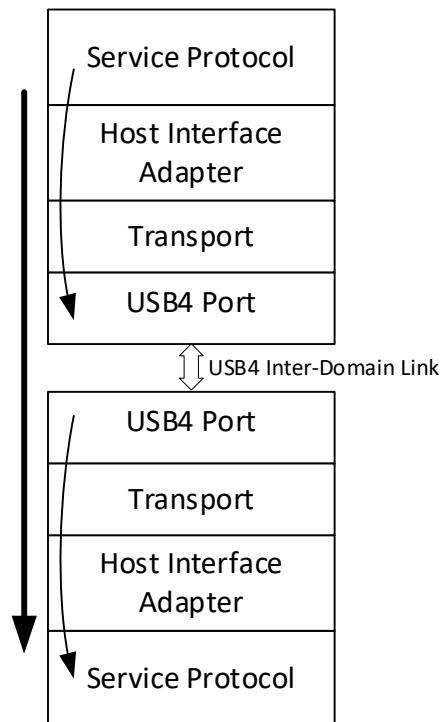
1. Disable the Upstream USB3 Gen T Adapter by setting the ADP\_USB3\_GT\_PORT\_CS\_1.Path Enable bit to 0b in the Adapter.
2. Disable the Downstream USB3 Gen T Adapter by setting the ADP\_USB3\_GT\_PORT\_CS\_1.Path Enable bit to 0b in the Adapter.
3. Teardown the downstream Path according to Section 5.2.2.
4. Teardown the upstream Path according to Section 5.2.2.

After tearing down the USB3 Gen T Path, the Connection Manager shall request to de-allocate the Gen T Port by interfacing with the Host Controller software.

### 5.7. Inter-Domain Path Setup and Teardown

An Inter-Domain Tunnel is a set of Inter-Domain Paths that transport traffic between the hosts in two adjacent Domains. After the Inter-Domain Tunnel is established, a Service Protocol uses the Tunnel to exchange traffic between the two hosts. Tunneling of Ethernet traffic is an example of a Service Protocol using the Inter-Domain Tunnel. Figure 5-11 shows an Inter-Domain Tunneling scheme between two hosts connected by an Inter-Domain Link.

**Figure 5-11: Inter-Domain Tunneling Scheme**



This section defines how a Connection Manager sets up and tears down the Inter-Domain Paths for an Inter-Domain Tunnel. This section should be read in combination with the USB4 Inter-Domain Service Protocol Specification, which describes Connection Manager discovery, capability exchange, Inter-Domain Link management, and Inter-Domain Service Protocol configuration.

A Connection Manager sets up an Inter-Domain Path as follows:

1. Determine the existence of an Inter-Domain Link per Section 3.3.
2. Exchange identification information with the peer Connection Manager.
3. Discover the capabilities of the peer Connection Manager.

4. If the Inter-Domain Link operates at Gen 2 or Gen 3 speeds and supports Lane Bonding, then together with the peer Connection Manager set the Inter-Domain Link to operate in aggregated mode.
5. If the Inter-Domain Link operates at Gen 4 speed and comes up with an asymmetric setting, optionally transition the Link to symmetric operation.
6. To enable a Service Protocol, exchange protocol-specific Login Packets with the peer Connection Manager.
7. Configure and enable an inbound, an outbound Path, and the associated Descriptor Rings for the Service Protocol:
  - Enable Transmit Descriptor Rings and the Receive Descriptor Rings for the service. The setting in a Host Router with a PCIe Host Interface Adapter Layer is defined in Section 5.7.1 and in Section 5.7.2. The order of setting the Rings is not important.
  - For an outbound Path, set the Path from a Transmit Descriptor Ring in the Host Router's Host Interface to the Inter-Domain Link.
  - For an inbound Path, set the Path from the Inter-Domain Link to a Receive Descriptor Ring in the Host Router's Host Interface.
  - Set the attributes for each Path:
    - Table 6-1 defines how to set the *Priority* fields.
    - Link Flow Control attributes are set according to the selected Service Protocol.
      - The EFC and ESE bits shall be set to 0b in the Path segment from the Lane Adapter to the Host Interface Adapter.

A Connection Manager performs the following steps to stop a Protocol Service with a peer Connection Manager:

1. Exchange protocol-specific Logout Packets with the peer Connection Manager.
2. Tear down the inbound and outbound Paths and disable the Descriptor Rings associated with the defunct Service Protocol. Section 5.7.3 defines how to disable a Ring in a PCIe Host Interface Adapter Layer.

#### **5.7.1. Transmit Ring Setup in a PCIe Host Interface Adapter Layer**

A Connection Manager selects an available Transmit Descriptor Ring and performs the following steps to set it:

1. Configure interrupts and interrupt controls.
2. Allocate a buffer in host memory for the Transmit Descriptor Ring.
3. Initialize the Transmit Descriptor Ring as follows:
  - a. Set the Ring base address in the *Base Address Low* Register and the *Base Address High* Register of the Transmit Descriptor Ring.
  - b. Set the Ring size in the *Ring Size* Register of the Transmit Descriptor Ring.

- c. Set the following Ring control attributes in the *Ring Control* Register of the Transmit Descriptor Ring:
  - i. *E2E Flow Control Enable* bit = set to 1b if E2E flow control is desired.
  - ii. *No-Snoop flag* bit = set to a value that is compatible with the software model.
  - iii. *Raw Mode* bit = set to 0b to operate the Descriptor Ring in Frame mode. Set to 1b to operate the Descriptor Ring in Raw mode.
  - iv. *Ring Valid* bit = 1b.

*Note: A Host Interface Adapter that generates Transport Layer Packets from Transmit Ring N, sets the HopID field in the Transport Layer Packet to N.*

### **5.7.2. Receive Ring Setup in a PCIe Host Interface Adapter Layer**

The Connection Manager selects an available Receive Descriptor Ring and performs the following steps to set it:

1. Configure interrupts and interrupt controls.
2. Allocate a buffer in host memory for the Receive Descriptor Ring.
3. Initialize the Receive Descriptor Ring as follows:
  - a. Set the Ring base address in the *Base Address Low* Register and the *Base Address High* Register of the Receive Descriptor Ring.
  - b. Set the Ring size and data buffer size for the Ring in the *Ring Size* Register of the Receive Descriptor Ring.
  - c. Set the *SOF PDF Bitmask* field and the *EOF PDF Bitmask* field in the *PDF Bit Masks* Register.
  - d. Set the following Ring control attributes in the *Ring Control* Register of the Receive Descriptor Ring:
    - i. *Transmit E2E HopID* field = set to the HopID value of the Transmit Descriptor Ring associated with the service provided by the Receive Descriptor Ring. In most cases, the Receive Descriptor Ring and the Transmit Descriptor Ring have the same HopID.
    - ii. *E2E Flow Control Enable* bit = set to 1b if E2E flow control is desired.
    - iii. *No-Snoop Flag* bit = set to a value that is compatible with the software model.
    - iv. *Raw Mode* bit = set to 0b to operate the Descriptor Ring in Frame mode. Set to 1b to operate the Descriptor Ring in Raw mode.
    - v. *Ring Valid* bit = 1b.

*Note: A Host Interface Adapter that receives a Transport Layer Packet with HopID N, directs this packet to Receive Ring N.*

### 5.7.3. Ring Disable in a PCIe Host Interface Adapter Layer

The Connection Manager disables a Descriptor Ring by setting all fields in the Ring Control register, including the *Ring Valid* bit, to their default values. The Connection Manager then sets the following fields to their default values:

- *Base Address Low*
- *Base Address High*
- *Ring Size*

### 5.8. Vendor Specific Adapter Path Setup and Teardown

A USB4 Host and Device Router may contain Vendor Specific Adapters which can be uniquely identified by the Connection Manager. The Connection Manager reads ADP\_CS\_2. *Adapter Type Sub-type*, ADP\_CS\_2. *Adapter Type Version*, ADP\_VS\_CS\_1. *VSA Type* and ADP\_VS\_CS\_1. *VSA Vendor ID* fields to identify a Vendor Specific Adapter. The Connection Manager can read registers ADP\_VS\_CS\_2 to ADP\_VS\_CS\_n to determine the rest of the Vendor Specific Adapter Capabilities.

If the Connection Manager recognizes a Vendor Specific Adapter, it should be able to establish Paths between the Vendor Specific Adapter of the same type in a USB4 Host and the Vendor Specific Adapter in a USB4 Device and configure any applicable QoS parameters.

The Connection Manager cannot configure any Paths between the Vendor Specific Adapters in two USB4 Devices.

### 5.9. Path Counters

A Connection Manager configures the performance counters in an Adapter while the performance counters are disabled for all Paths in the Adapter. The performance counters are disabled for a Path when the *Counter Enable* bit in the Path Configuration Space is 0b.

1. Read the *Counters Configuration Space (CCS)* Flag in the Lane 0 Adapter Configuration Space. If the value is 0, then a Counters Configuration Space is not supported by the Adapter. Else, continue.
2. Read the *Max Counter Sets* field in the Lane 0 Adapter Configuration Space. This is the number of counter sets in the Counters Configuration Space of the Adapter.
3. In a Version 2.0 Router, read the *Bytes Counter Supported* bit in the Lane 0 Adapter Configuration Space. If its value is 1b, the counters support byte counting. If supported:
  - a. Set the *Received Bytes Counter Enable* bit in the Lane 0 Adapter Configuration Space to the desired counting policy for all counter sets of the Adapter:
    - i. 0b - CNT\_CS\_0 and CNT\_CS\_1 count number of Transport Layer Packets routed by the Path
    - ii. 1b - CNT\_CS\_0 and CNT\_CS\_1 count total number of payload bytes in Transport Layer Packets routed by the Path
  - b. If counting bytes, read the *Lock Bytes Counter with TimeOffsetFromHR Low Supported* bit in the Lane 0 Adapter Configuration Space. If set to 1b, the Adapter supports locking the byte counters by reading the

*TimeOffsetFromHR Low* counter in the TMU Router Configuration Capability. If supported:

- i. Set the *Lock Bytes Counter with TimeOffsetFromHR Low Enable* bit in the Lane 0 Adapter Configuration Space to the desired sampling policy of the byte counters:

0b – *Received Counter Low* field and *Received Counter High* field in the Counters Configuration Space reflect current counter values.

1b - *Received Counter Low* field and *Received Counter High* field in the Counters Configuration Space are updated when the *TimeOffsetFromHR Low* field is read.

A Connection Manager performs the following steps to set performance counters to a Path in an Adapter:

1. Select a Counter ID for the Path. The Counter ID shall be less than the *Max Counter Sets* field of the Adapter. A Counter may be shared among multiple Paths. Write the Counter ID value into the *Counter ID* field in the Path Configuration Space.
2. Set the *Counter Enable* bit in the Path Configuration Space to 1b. The counter set is now counting events for the Path.

Note: The only fields in the Path Configuration Space entry for Path 0 that a Connection Manager can change are the Counter ID field and the Counter Enable bit. To avoid changing the value of the other fields, it is recommended that a Connection Manager first read the Path Configuration Space entry for Path 0, make the necessary changes to the Counter ID field and the Counter Enable bit, then write the results back to Path Configuration Space.

3. Each Counter Set includes a 32-bit *Dropped Packets* register. The Connection Manager may read the counter values at any time.
4. Each Counter Set includes a 64-bit *Received Counter* register.
  - a. If counting packets, the Connection Manager may read the counter values at any time
  - b. If counting bytes (Version 2.0 Router only),
    - i. If the *Lock Bytes Counter with TimeOffsetFromHR Low Enable* bit in the Lane 0 Adapter Configuration Space is set to 0b, the Connection Manager may read the counter values at any time.
      1. It is recommended that a Connection Manager read the *Received Packets Low* field and the *Received Packets High* field in a single Read Request to guarantee that the received packets counter values are from the same counter snapshot.
    - ii. If the *Lock Bytes Counter with TimeOffsetFromHR Low Enable* bit in the Lane 0 Adapter Configuration Space is set to 1b, the Connection Manager reads the *TimeOffsetFromHR Low* counter in the TMU Router Configuration Capability to upload the Received Counter value into the *Received Counter* register. The value loaded corresponds to time values captured by reading the *TimeOffsetFromHR Low* counter. The Connection Manager may

then use the Host Router's Local Clock for accurate bandwidth calculations.

5. Each counter stops counting at FFFF ... FFFFh. The Connection Manager may restart a counter from 0 by writing into the counter.

A Connection Manager performs the following steps to detach performance counters from a Path in an Adapter:

1. Set the Counter Enable bit in the Path Configuration Space to 0b. The counter set stops counting events for the Path.

## 6. Bandwidth Management

A Connection Manager calculates and manages the available USB4 bandwidth and allocates it to the various Tunneled Protocols within its Domain.

Tunneled Protocols utilize two types of USB4 bandwidth: isochronous and non-isochronous. A Connection Manager must reserve 10% of the maximum USB4 bandwidth as "guard band bandwidth" for Link overhead and for non-isochronous Tunneled Protocols. This helps prevent starvation to the non-isochronous protocols, which can lead to system visible failures.

The Connection Manager can use up to the remaining 90% of the maximum USB4 bandwidth for isochronous Tunneled Protocols. Any bandwidth not used for isochronous Tunneled Protocols is used for Non-Isochronous Tunneled Protocols.

Section 6.1 describes how a Connection Manager allocates USB4 bandwidth between the Tunneled Protocols. Section 6.2 describes specifics per each tunneled protocol.

### 6.1. Bandwidth Allocation

#### 6.1.1. Scheduling Priorities

A Router uses a 3-layer scheduling scheme to manage traffic at each egress Port. The Connection Manager uses this scheme to allocate USB4 Link bandwidth to each egress Path. A Connection Manager sets the *PATH\_CS\_1.Priority* and *PATH\_CS\_1.Weight* fields to set the USB4 bandwidth for a Path. The *Priority* field determines which priority group the Path belongs to. The *Weight* field determines the weight of the Path within its Priority Group with respect to the Weighted Round Robin (WRR) scheduling protocol. See the USB4 Specification for more details.

The following considerations are recommended when assigning priorities to tunneled protocols:

- Tunneled Protocols that carry isochronous traffic only (e.g., DP Main-Link tunnel), are assigned highest priorities, so that isochronous contracts can be guaranteed
- Tunneled Protocols that carry a mixture of isochronous and bulk traffic (e.g., USB3 Gen T Paths, USB3 Gen X Path) are assigned the next highest priority, and all such protocols share the same priority. This ensures that isochronous contracts can be guaranteed
- Tunneled Protocols that carry bulk traffic only (e.g., PCIe) are assigned to either the same priority as Tunneled Protocols that carry a mixture of isochronous and bulk traffic, or to lower priorities.
  - Tunneled Protocols that carry very little traffic (e.g., DP Aux) may be assigned higher priority as long as its bandwidth is minimal and is bounded.

Table 6-1 shows recommended Priority values for each type of Tunneled Protocol.

**Table 6-1: Tunneled Protocol Path Priority (Recommended)**

Path	Priority
DP Main-Link	1
DP AUX	2
USB3 Gen X	3
USB3 Gen T	3
PCIe	3
Host-to-Host	4

#### 6.1.2. Bandwidth Allocation

This section describes how a Connection Manager calculates the bandwidth requested for a Path and how it determines the actual bandwidth allocated to the Path.

The Connection Manager may negotiate with the USB controller the allocation of bandwidth for USB3 Paths, and with DPTX the allocation of bandwidth to DP Paths (optionally supported). Bandwidth requirements of tunneled protocols are considered when distributing the available bandwidth to Paths. At the same time, the available bandwidth of Links along a Path affects the setting of tunneled protocols (e.g., choice of DP mode).

##### 6.1.2.1. Raw USB4 Bandwidth

The maximum USB4 bandwidth of a Link is the raw bandwidth produced by the operating link rate and the number of Lanes used. Table 6-2 shows the raw bandwidth for each type of USB4 Link.

**Table 6-2: USB4 Raw Bandwidth**

USB4 Link	Raw Bandwidth (USB4)
Gen 2x1	10 Gbps in each direction
Gen 2x2	20 Gbps in each direction
Gen 3x1	20 Gbps in each direction
Gen 3x2	40 Gbps in each direction
Gen 4 Symmetric	80 Gbps in each direction
Gen 4 Asymmetric	120 Gbps in one direction 40 Gbps in the other direction

*Note: The prefix Gbps represents  $10^9$  bits-per-second and not  $2^{30}$  bits-per-second*

When a Path traverses a single USB4 Link, the raw bandwidth of the Link is the maximum USB4 bandwidth for the Path. When a Path traverses multiple USB4 Links, the raw bandwidth of all Links along the Path needs to be considered.

##### 6.1.2.2. Bandwidth Management Events

When a protocol tunnel is established, the Connection Manager sets aside bandwidth for each Path of the tunnel.

- Initial bandwidth allocation is done per the general guidelines in Section 6.1.2.3
- Initial bandwidth allocation for a DP tunnel is described in Section 6.2.1.1

- Initial bandwidth allocation for a USB3 Gen X Path is described in Section 6.2.2.1
- Initial bandwidth allocation for a USB3 Gen T Path is described in Section 6.2.3.1

Bandwidth allocated to a DP tunnel may change as described in Section 6.2.1.2 and Section 6.2.1.3. The Connection Manager may allocate more bandwidth to a DP tunnel as bandwidth becomes available. A DP tunnel may increase or decrease its consumed bandwidth (within its allocation) as its needs change.

Bandwidth allocated to a USB3 Gen X Path may change as described in Section 6.2.2.2. The Connection Manager may allocate more bandwidth to a USB3 Gen X Path as bandwidth becomes available. A USB host controller may increase or decrease its USB3 Gen X consumed bandwidth (within its allocation) as its needs change.

Bandwidth allocated to USB3 Gen T Paths may change as described in Section 6.2.3.2. The Connection Manager may allocate more bandwidth to a USB3 Gen T Path as bandwidth becomes available. A USB host controller may request allocation of bandwidth for isochronous traffic, and it may increase or decrease its consumed bandwidth (within its allocation) as its needs change.

#### **6.1.2.3. Bandwidth Allocation**

Once scheduling priorities are set (see Section 6.1.1), the Connection Manager determines the value of the *weight* field for each Path within each priority. The assigned values are calculated from the allocated bandwidth to each Path.

The following considerations are recommended when assigning bandwidth to tunneled protocols:

- Allocate bandwidth to the Paths in each priority group, starting with the highest priority, and until all bandwidth has been allocated.
- Ensure that the sum of bandwidths committed to isochronous Paths does not exceed 90% of the raw bandwidth in each Link along the Path, and at least minimum bandwidth is allocated to each USB and PCIe Path as described below.
- Within each priority group, distribute to bulk traffic any remaining bandwidth after all isochronous requirements have been met. A possible policy is to allocate equal bulk bandwidth to each Path that sends bulk traffic.
- Allocate at least 1.5Gbps bandwidth to each USB3 Gen X, USB3 Gen T, and PCIe Paths. Allocating a minimum bandwidth to a Path allows a Router implementation to ensure that Path latency does not exceed a certain value.
- When a Path is torn down, its allocated bandwidth is reclaimed and can be allocated to existing Paths
- In each priority group, calculate the value of the *weight* field for each Path relative to the sum of allocated isochronous and bulk bandwidth for the Path.

#### **6.1.2.4. An example of bandwidth allocation**

This section describes an example of an arbitration scheme and how the Connection Manager may allocate bandwidth and program the weights for the different Paths. Refer to Figure 6-1 for Path assignments and Weighted Round Robin (WRR) weights.

In this example, the following conditions are assumed:

- Link raw bandwidth = 80Gbps.
- Available bandwidth = 90% of raw bandwidth = 72Gbps
- Assigned Priorities

Path	Priority
DP	1
USB3 Gen X	3
USB3 Gen T	3
PCIe	3

- Isochronous bandwidth requested by each Path
  - DP Path: 36Gbps (50% of available bandwidth)
  - USB3 Gen X Port isochronous bandwidth: 10Gbps
  - USB3 Gen T Port 1 isochronous bandwidth: 5Gbps
  - USB3 Gen T Port 2 isochronous bandwidth: 5Gbps

*Note: Isochronous bandwidth requested on Gen X and Gen T ports is total bandwidth requested by multiple isochronous endpoints.*

Based on the above conditions:

- Total isochronous bandwidth = 36Gbps + 10Gbps + 5Gbps + 5Gbps = 56Gbps
- Total bulk bandwidth = Available bandwidth - Total isochronous bandwidth  
= 72Gbps – 56Gbps = 16 Gbps
- Bulk bandwidth assigned per Path = 16Gbps / 4 = 4 Gbps

*Note: In this example (refer Figure 6-1), there are 4 such Paths and we recommend same bulk bandwidth for each Path (B1, B2, B3 and B4 in Figure 6-1).*

**Bandwidth assigned to each Path:**

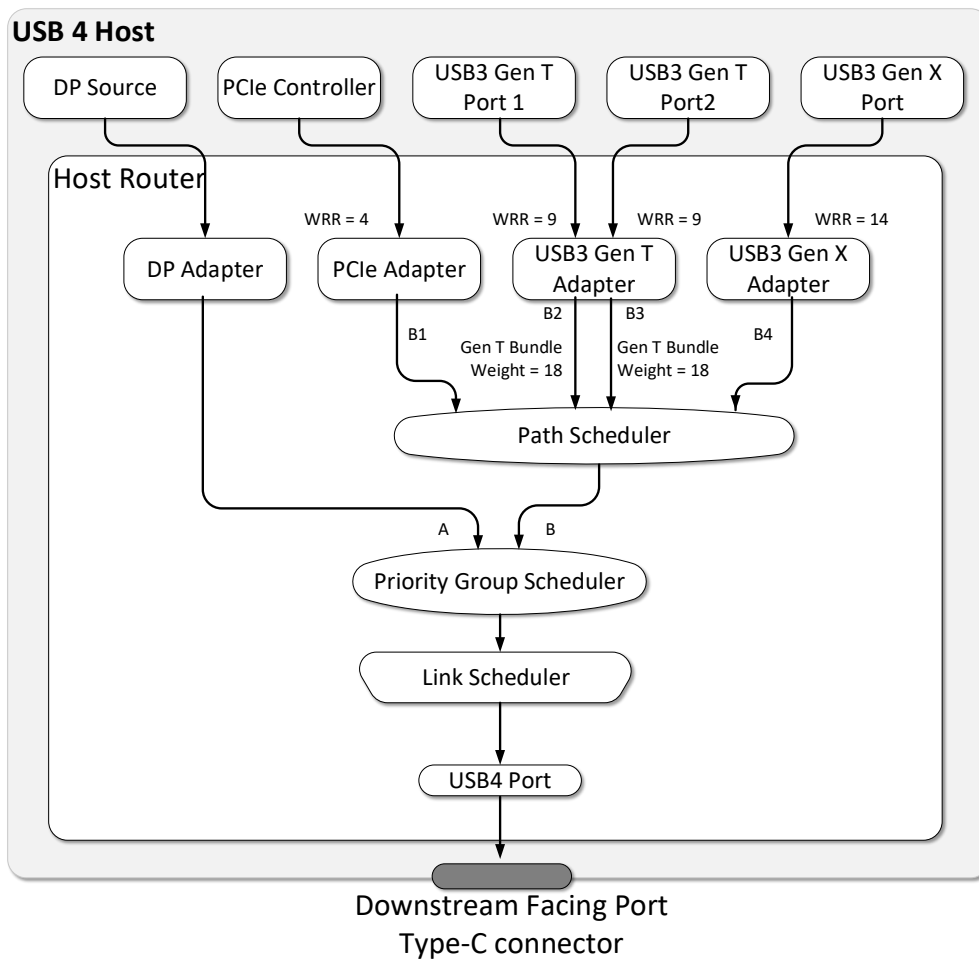
- DP Path: 36Gbps
- USB3 Gen X Path
  - Isochronous traffic: 10Gbps
  - Bulk traffic: 4Gbps
  - Total bandwidth for Path: 14Gbps
- USB3 Gen T Port 1
  - Isochronous traffic: 5Gbps
  - Bulk traffic: 4Gbps
  - Total bandwidth for Path: 9Gbps
- USB3 Gen T Port 2
  - Isochronous traffic: 5Gbps
  - Bulk traffic: 4Gbps
  - Total bandwidth for Path: 9Gbps
- PCIe Path

- Bulk traffic: 4 Gbps
- Total bandwidth for Path:  
4Gbps

### **Weight assignments for Priority Group with priority = 3**

- USB3 Gen X Path: 14
- USB3 Gen T Port 1: 9
- USB3 Gen T Port 2: 9
- PCIe Path: 4
- If Gen T bundle mode is enabled:
  - Gen T Bundle Weight = 18

**Figure 6-1: Example - WRR, Bundle Weight and BW Allocation**



## **6.2. Bandwidth Clients**

### **6.2.1. DisplayPort Tunneling**

DisplayPort Tunneling uses Isochronous bandwidth, and it has two methods for Bandwidth Allocation:

1. *Max Link BW* – This method is supported by default by a DP IN Adapter and a DPTX. According to the maximum Link Rate and maximum Lane Count, the Connection Manager allocate 80% of that Bandwidth. This method does not support 128b/132 DP Link.
2. *DP BW Allocation Mode* – This method is supported by a DP IN Adapter in a Version 2.0 Router, but not necessarily supported by the DPTX. The Connection Manager and the DPTX negotiate for the allocated bandwidth through registers located at the DP IN Adapter, both at the DPCD space and the Adapter Configuration Space. This method supports 8b/10b and 128b/132b DP Links.

During Path setup, the Connection Manager initialized the two allocation methods as described in Section 6.2.1.1. By default, the *Max Link BW* method is used according to Sections 6.2.1.1 and 6.2.1.2. If a DPTX enables the *DP BW Allocation Mode*, then the bandwidth allocation method is according to Section 6.2.1.3.

#### 6.2.1.1. Initial Configuration

When a Connection Manager sets up the Path it initialized both bandwidth allocation methods, while the bandwidth which is allocated through the *Max Link BW* method is guaranteed.

##### **Max Link BW**

The Isochronous bandwidth required for a DisplayPort traffic varies with Link clock frequency and the number of DP lanes it utilizes.

The initial DisplayPort Link setup uses 8b/10b encoding, which is removed by the DP IN Adapter, lowering the required DisplayPort bandwidth over a USB4 Link to 80% of the original DisplayPort Link.

The Connection Manager calculates the USB4 Link bandwidth required for DisplayPort traffic according to the formula below:

$$\text{DP Required BW} = \text{DP Link Rate} * \text{DP Lane Count} * 0.8$$

When a Connection Manager sets up a DP Tunneling Path, it allocates the maximum available USB4 bandwidth to the Path. If needed, it adjusts the maximum bandwidth capability of the DisplayPort Link to be smaller than the maximum available USB4 bandwidth for the Path. The video format, resolution and blanking scheme used by the DPTX does not affect the bandwidth calculation.

Table 6-3 shows the Asynchronous Bandwidth needed for each DP Link according to lane count and link rate.

**Table 6-3: DisplayPort Required Bandwidth (Gbps)**

	4 Lanes	2 Lanes	1 Lane
HBR3 (8.1Gbps)	25.92	12.96	6.48
HBR2 (5.4Gbps)	17.28	8.64	4.32
HBR (2.7Gbps)	8.64	4.32	2.16
RBR (1.62Gbps)	5.18	2.59	1.30

Example:

- USB4 raw bandwidth = 40.0 Gbps
- Guard band bandwidth = 4.0 Gbps
- The first DisplayPort Link uses the maximum configuration:  
$$\text{HBR3x4Lanes} = 8.1\text{G bps} * 4 * 0.8 = 25.92\text{ Gbps}$$
- Remaining Bandwidth =  $40.0 - 4.0 - 25.92 = 10.08\text{ Gbps}$
- The second DisplayPort Link is constrained by the Connection Manager, so it does not consume more than the remaining bandwidth:  
$$\text{HBR2x2Lanes} = 5.4\text{Gbps} * 2 * 0.8 = 8.64\text{ Gbps}$$

#### **DP BW Allocation Mode**

During Path setup, if DP IN Adapter supports DP BW Allocation Mode (DP\_LOCAL\_CAP.DP\_IN\_BW\_Allocation Mode Support is set to 1b), the Connection Manager calculates the estimated bandwidth for the DP IN Adapter. The estimated bandwidth for a DP IN Adapter includes any throughput that can be assigned to the DP IN Adapter:

- The bandwidth that has already been allocated to the DP IN Adapter.
- The available BW along the Path.
- Bandwidth that is allocated to the internal Host Controller, but not consumed

After the estimated bandwidth is calculated, a Connection Manager updates the ADP\_DP\_CS\_2.Estimated BW field.

##### **6.2.1.2. Additional Configuration for Max Link BW**

When the DP Paths are first enabled, the Maximal Lane Count and Link Rate are based on the minimum values advertised by the DP IN and DP OUT Adapters. After the DP Paths are enabled, the DP IN Adapter may reduce the Maximal Link Rate and Lane Count based on the information read from the DPRX during DPTX discovery. The Connection Manager can determine when the updated Maximal Lane Count and Link Rate values are valid as follows:

1. After the DP Paths are enabled, the Connection Manager polls the DPRX DP\_COMMON\_CAP.Capabilities Read Done field of the DP IN Adapter. The Connection Manager polls the field until it is set to 1b.
2. The Connection Manager reads the DP\_COMMON\_CAP register of the DP IN Adapter to get the updated *Maximal Link Rate*, and *Maximal Lane Count* fields. If the *Maximal Lane Count* or the *Maximal Link Rate* were reduced, it means that BW was freed by this DP Path.

### **6.2.1.3. DP BW Allocation Mode**

#### **6.2.1.3.1. BW Estimation**

When there is a change in bandwidth across its Domain, the Connection Manager recalculates the estimated bandwidth for each DP IN Adapter in the Domain. The estimated bandwidth for a DP IN Adapter includes any throughput that can be assigned to the DP IN Adapter:

- The bandwidth that has already been allocated to the DP IN Adapter.
- The available BW along the Path.
- Bandwidth that is allocated to the internal Host Controller, but not consumed

After the estimated bandwidth is recalculated, a Connection Manager updates the ADP\_DP\_CS\_2.*Estimated BW* field.

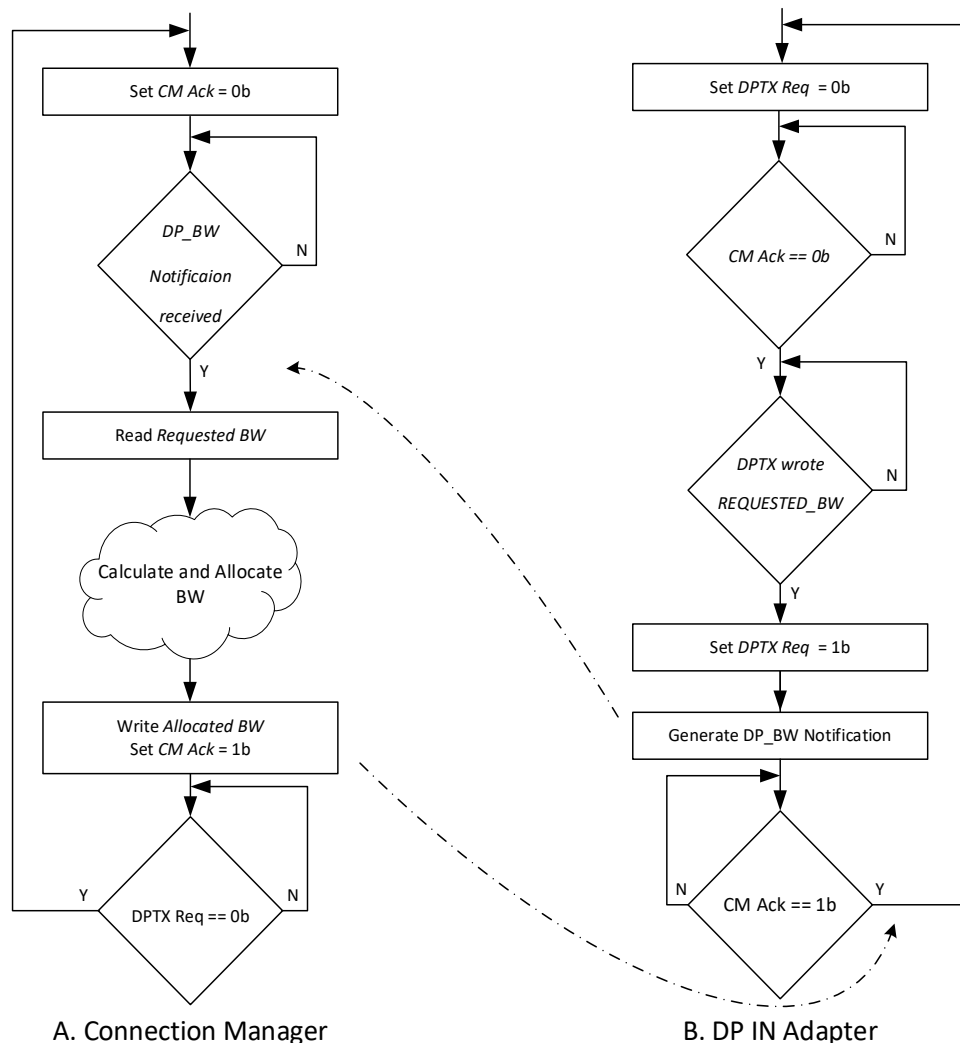
#### **6.2.1.3.2. BW Negotiation**

When the Connection Manager receives a DP\_BW Notification Packet, it completes the following steps. The Connection Manager completes step 4 within 100ms from reception of the Notification Packet:

1. Reads ADP\_DP\_CS\_8.*Requested BW*.
2. Allocates bandwidth:
  - a. If the Request bandwidth is less than or equal to the currently allocated bandwidth, then the new allocation is equal to the new requested bandwidth.
    - If the DP IN Adapter has a Group\_ID different than zero, then for at least 10 seconds, the Connection Manager reserves the released bandwidth to be allocated to DP IN Adapters with the same Group\_ID.
  - b. Else, the Connection Manager tries to allocate bandwidth from the available bandwidth and by freeing bandwidth from other clients.
3. Writes the new allocated bandwidth to the DP\_STATUS.*Allocated BW* field.
4. Sets the ADP\_DP\_CS\_2.*CM Ack* bit to 1b.
5. Waits for the DP IN Adapter to set the ADP\_DP\_CS\_8.*DPTX Req* bit to 0b.
6. Sets the ADP\_DP\_CS\_2.*CM Ack* bit to 0b.

A DP IN Adapter interacts with a Connection Manager according to Figure 6-2.

**Figure 6-2: DP BW Allocation Interaction with Connection Manager**



### 6.2.2. USB3 Gen X Tunneling

USB3 Gen X Tunneling uses both Isochronous and Non-Isochronous bandwidth, depending on the connected USB 3.2 endpoints. The Connection Manager and the internal host controller share a bandwidth negotiation mechanism to manage the allocation of the Isochronous bandwidth. The Connection Manager can allocate separate bandwidth for the downstream and upstream directions. See Section 6.2.2.2 for more details on the operation of bandwidth negotiation.

USB 3.2 Isochronous and USB 3.2 Non-Isochronous traffic share a single Gen X Path in each direction and therefore are treated the same way over the USB3 Gen X Path. It is up to the internal host controller to use the allocated USB4 Link bandwidth for the Isochronous traffic first.

The Connection Manager allocates at least 1.5Gbps bandwidth to a USB3 Gen X Path. Allocating a minimum bandwidth to a Path allows a Router implementation to ensure that Path latency does not exceed a certain value.

#### 6.2.2.1. Bandwidth Allocation during Path Setup

When setting a new USB3 Gen X Path, the Connection Manager assigns an initial bandwidth allocation to the Path based on available bandwidth and the needs of other Paths. Bandwidth allocation for the USB3 Gen X Path does not exceed the value of *Maximum Supported Link Rate* field at both upstream and downstream Adapters. Bandwidth allocation for isochronous traffic, reported to the USB host controller, does not exceed 90% of the initial allocation for the Path at both upstream and downstream Adapters. The Connection Manager writes the initial isochronous allocation to the *Allocated Upstream Bandwidth* and *Allocated Downstream Bandwidth* fields at each Adapter of the Host Router. The Connection Manager sets the *Scale* field in the USB3 Gen X Configuration Capability to the desired granularity for bandwidth negotiation.

*Note: After a USB3 Gen X Path is established, the Allocated Upstream Bandwidth and Allocated Downstream Bandwidth can be decreased according to the Actual Link Rate field.*

*Note: After a USB3 Gen X Path is established, the USB host controller updates the Connection Manager about the actual isochronous traffic consumed by writing to the Consumed Downstream Bandwidth field and the Consumed Upstream Bandwidth field in the Host Router's Adapter.*

#### 6.2.2.2. Dynamic Bandwidth Allocation

USB3 Gen X Tunneling bandwidth is dynamic, with both Connection Manager and USB host controller continuously updating each other.

The Connection Manager updates the allocated bandwidth to the USB3 Gen X Path as conditions change. The Connection Manager may increase the allocated bandwidth at any time. The Connection Manager may also decrease the allocated bandwidth, but it never decreases allocation for isochronous traffic to a value lower than the consumed bandwidth at that time. As mentioned above, the Connection Manager also ensures that it allocates at least 1.5Gbps to the total of isochronous and bulk bandwidth.

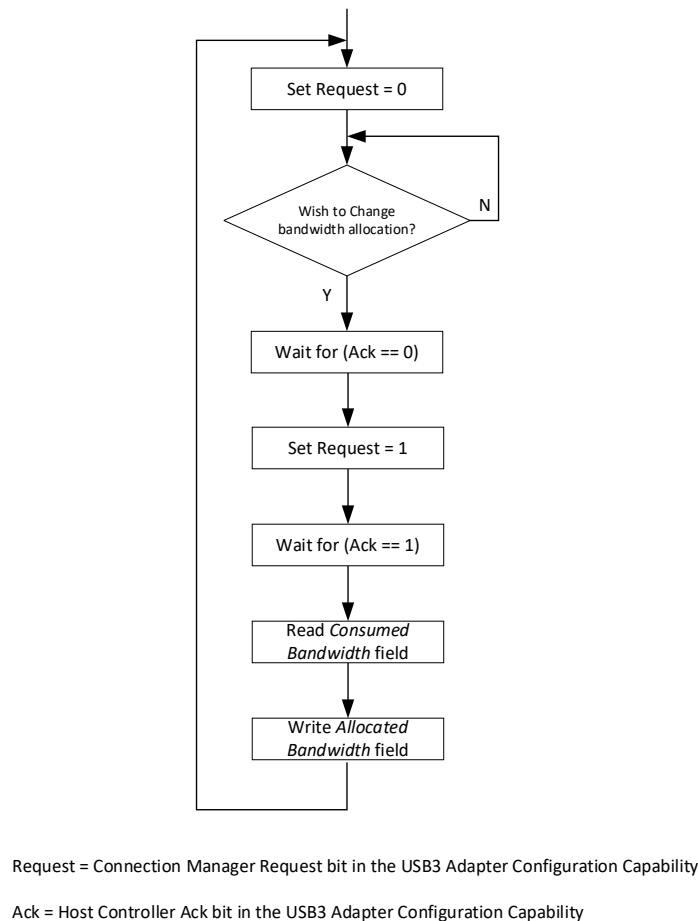
The USB host controller updates the Connection Manager with its consumed isochronous bandwidth as its needs change. The USB host controller may decrease the consumed bandwidth at any time. The USB host controller may also increase the consumed bandwidth, but never to a value higher than the allocated bandwidth at that time.

When a Connection Manager changes the bandwidth allocation of the USB3 Gen X Path, it updates the *weight* field in all Paths that share an Egress Port with the USB3 Gen X Path.

A Connection Manager uses the *Connection Manager Request* and *Host Controller Ack* fields as follows (see Figure 6-3):

- When the *Host Controller Ack* field = 0 (The default value):
  - The Connection Manager shall not update the *Allocated Upstream Bandwidth* or *Allocated Downstream Bandwidth* fields.
  - The Connection Manager shall not read the *Consumed Upstream Bandwidth* or *Consumed Downstream Bandwidth* fields.
- When the *Host Controller Ack* field = 1:
  - The Connection Manager can read the *Consumed Upstream Bandwidth* and/or *Consumed Downstream Bandwidth* fields.
  - The Connection Manager can change the *Allocated Upstream Bandwidth* and/or *Allocated Downstream Bandwidth* fields.

**Figure 6-3. Bandwidth Negotiation by the Connection Manager**



### 6.2.3. USB3 Gen T Tunneling

USB3 Gen T Tunneling uses both isochronous and non-isochronous bandwidth, depending on the connected USB 3.2 endpoints. The Connection Manager and the internal host controller share a bandwidth negotiation mechanism to manage the allocation of the Isochronous bandwidth of each USB3 Gen T Path. The Connection Manager can allocate separate bandwidth for the downstream and upstream directions.

The details of bandwidth negotiation for USB3 Gen T Paths is beyond the scope of this document.

USB 3.2 Isochronous and USB 3.2 Non-Isochronous traffic share a Gen T Path in each direction and are treated the same way over the USB3 Gen T Path. It is up to the internal host controller to use the allocated USB4 Link bandwidth for the Isochronous traffic first.

The Connection Manager allocates at least 1.5Gbps bandwidth to each USB3 Gen T Path. Allocating a minimum bandwidth to a Path allows a Router implementation to ensure that Path latency does not exceed a certain value.

#### **6.2.3.1. Bandwidth Allocation during Path Setup**

When setting a new USB3 Gen T Path, the Connection Manager assigns an initial bandwidth to the Path based on available bandwidth and the needs of other Paths. Bandwidth allocation to a Path shall not exceed the value of the *Active Tx Link Rate* field that is associated with the Path in the USB3 Gen T Adapter Configuration Capability (see Section 5.6.1). Initial bandwidth allocation is for bulk traffic only.

Once a USB3 Gen T Path is established, the USB3 host controller may request allocation of bandwidth for isochronous traffic using dynamic bandwidth allocation mentioned in the next section.

#### **6.2.3.2. Dynamic Bandwidth Allocation**

USB3 Gen T Tunneling bandwidth is dynamic, with both Connection Manager and USB host controller continuously updating each other.

The Connection Manager updates the allocated bandwidth to the USB3 Gen T Path as conditions change. The Connection Manager may increase the allocated bandwidth at any time. The Connection Manager may also decrease the allocated bandwidth, but it never decreases allocation to a value lower than the consumed isochronous bandwidth at that time.

The USB host controller updates the Connection Manager with its consumed isochronous bandwidth as its needs change. The USB host controller may decrease the required isochronous bandwidth at any time. The USB host controller may also increase the allocated isochronous bandwidth, but never to a value higher than the allocated isochronous bandwidth at that time.

The USB host controller may request an increase in allocation of isochronous bandwidth. The Connection Manager may increase the allocation of isochronous bandwidth if bandwidth is available.

For each USB3 Gen T Path, bandwidth allocation for isochronous traffic, reported to the USB host controller, does not exceed 70% of the total bandwidth allocated to the Path.

When a Connection Manager changes the bandwidth allocation of a USB3 Gen T Path, it reprograms the *weight* fields in all Paths that share the same Priority Group as the USB3 Gen T Path.

If ROUTER\_CS\_6.Gen T Bundle Weight Mode field in the Host Router is set to 1b, the Connection Manager sets the ADP\_USB3\_GT\_CS\_1. Bundle Weight field in the USB3 Gen T Adapter to the sum of weights of all Gen T Paths that are routed via the Adapter.

#### **6.2.4. PCIe Tunneling**

PCIe Tunneling is treated as non-isochronous and is not guaranteed bandwidth beyond a minimal allocation mentioned below.

A PCIe Path, USB3 Gen T Path and a USB3 Gen X Path are commonly set to the same Priority group. Therefore, if the Connection Manager allocates bandwidth to a USB3 Gen T Path or to the USB3 Gen X Path, it indirectly affects bandwidth allocation to the PCIe

Path. The WRR ratio between the PCIe Path, USB3 Gen T Path and the USB3 Gen X Path determine how much bandwidth is allocated to the PCIe Path.

The Connection Manager shall allocate at least 1.5Gbps bandwidth to a PCIe Path. Allocating a minimum bandwidth to a Path allows a Router implementation to ensure that Path latency does not exceed a certain value.

#### **6.2.5. Host-to-Host Tunneling**

Host-to-Host Tunneling is assumed to be using non-isochronous bandwidth only. The Connection Manager configures a Host-to-Host Tunneling Path to use the lowest Priority. The Host Interface Path uses any remaining bandwidth not used by high-priority Paths.

When a Connection Manager configures a Path for a tunneled protocol, it needs to make sure that it has enough available bandwidth for the tunneled protocol along each USB4 Link from the Source Adapter to the Destination Adapter. When the Path travels through more than one USB4 Link, the maximum available bandwidth for the Path is dictated by the USB4 Link which has the lowest available bandwidth.

## 7. USB4 Port Management

A USB4 Port comes up in one of the following configurations following USB4 Link Initialization:

- Symmetric Gen 4 Link (Version 2.0 Router) - Both Lane Adapters are in CL0 state. Lanes are aggregated into a single Link running Gen 4 speed and protocol. The Connection Manager may transition the Link to act as an Asymmetric Link (see Section 7.2.1).
- Two Single-Lane Gen 2 or Gen 3 Links - Both Lane Adapters are in CL0 state. The Connection Manager can either bond the Lanes into one Aggregated Link (see Section 7.1) or it can disable the Lane 1 Adapter in order to operate with one Single-Lane Link (see Section 7.5). The Connection Manager does not configure any Paths over the Port while the Port has two Single-Lane Links.
- One Single-Lane Gen 2 or Gen 3 Link - Only Lane 0 Adapter is in CL0 state. In this case, the Lane 1 Adapter is either unused, is not enabled in the AT Register Space, is in CLd state, or is in Training state. The Connection Manager can define Paths over the Single-Lane Link that corresponds with Lane 0 Adapter. The Connection Manager disables the Lane 1 Adapter (see Section 7.4).
- No Link - Neither Lane Adapter is in CL0 state. This is the case when the Lane Adapters are either unused, are not enabled in the AT Register Space, or are in CLd or Training state. The Connection Manager disables the Lane Adapters (see Section 7.4).

Unless specified otherwise, a Connection Manager always accesses the configuration space associated with a USB4 Port via the Port's Lane 0 Adapter Configuration Space.

### 7.1. Lane Bonding (Gen 2 or Gen 3 only)

A Connection Manager does not initiate Lane bonding when a Path (other than Path 0) is enabled in the Port. After Lane Bonding, the Connection Manager can set up additional Paths.

Before initiating Lane Bonding, the Connection Manager reads the *Negotiated Link Width* field. If the value read is 000010b or higher, the Lanes are already bonded, the Connection Manager does not initiate Lane Bonding and continues as if the Lane Bonding completed successfully.

The Connection Manager performs the following steps to initiate Lane Bonding of a Gen 2 or a Gen 3 Link:

1. On both sides of the Link, read the `LANE_ADP_CS_1.Adapter State` field of each Lane Adapter and verify that all Lane Adapters are in CL0 state.
2. Ensure that the Routers at both ends of the Link are configured as described in Section 3.
3. On both sides of the Link, set the `LANE_ADP_CS_1.Target Link Width` field to 11b.
  - (Version 1.0 Router) in both Lane Adapters
  - (Version 2.0 Router) In Lane 0 Adapter
4. On the DFP of the Link, set the `LANE_ADP_CS_1.Lane Bonding` bit to 1b in the Lane 0 Adapter.
5. Verify that Lane Bonding was successful by either:

- Waiting for a Hot Unplug Event Packet for Lane 1 with the *UPG* bit set to 1b.
- Polling the *LANE\_ADAP\_CS\_1.Negotiated Link Width* field in the Lane 0 Adapter until it indicates a USB4 Link width of Symmetric Link (x2).

When Lane Bonding is successful, an Aggregated Gen2/3 Link is established. The Connection Manager uses the Path Configuration Space of the Lane 0 Adapter to setup any Paths over the Link. Path setup is described in Section 5. If after tBonding time the *Negotiated Link Width* field in the Lane Adapter Configuration Capability of Lane 0 Adapter does not indicate a USB4 Link width of x2, then Lane bonding has failed. The Connection Manager can either retry Lane bonding or disable the Lane 1 Adapter and operate with one Single-Lane Link.

## **7.2. Asymmetric Link (Version 2.0 Router)**

This section describes how a Connection Manager enables and disables Asymmetric Link when the Link is capable of operating in Asymmetric Link configuration. The Connection Manager shall not initiate an Asymmetric transition on Inter-Domain Links.

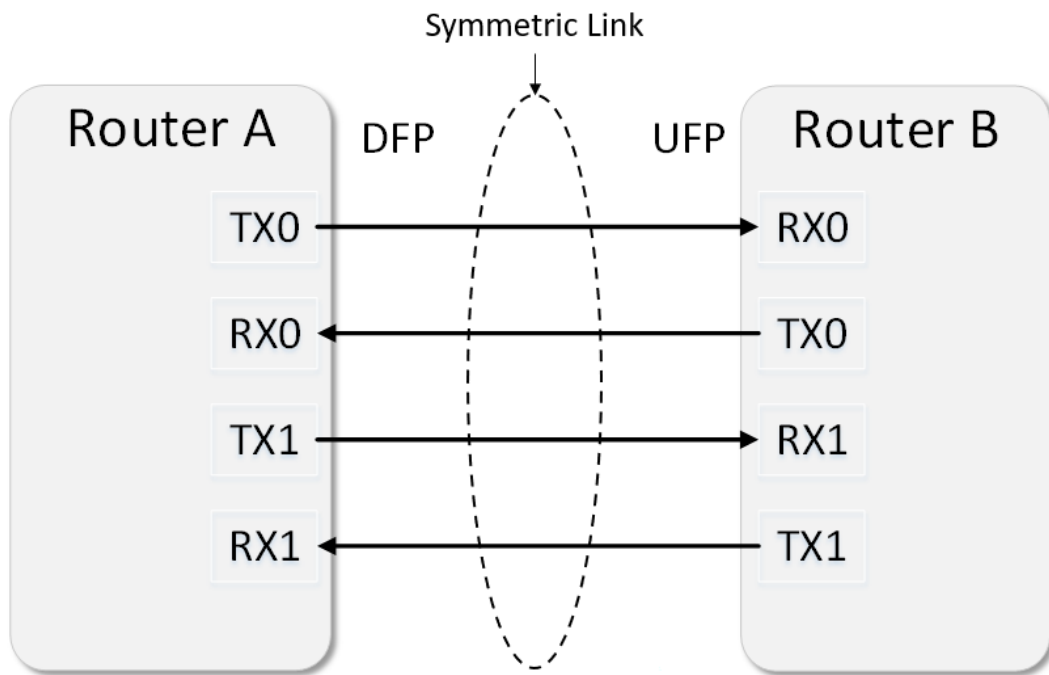
*Note: Stimulus and Conditions used by Connection Manager to determine when to enable or disable Asymmetric Link is beyond the scope of this document.*

Before initiating the configuration of Asymmetric Link, the Connection Manager reads the *Negotiated Link Width*. If the Lanes are already in the desired Asymmetric Link, the Connection Manager does not initiate transition to Asymmetric Link and continues as if the Asymmetric Link configuration completed successfully.

Similarly, before initiating the transition to Symmetric Link from an Asymmetric Link, the Connection Manager reads the *Negotiated Link Width*. If the Lanes are already in the Symmetric 2 Lanes link, the Connection Manager does not initiate transition to Dual-Lane Link and continues as if the Dual-Lane Link configuration completed successfully.

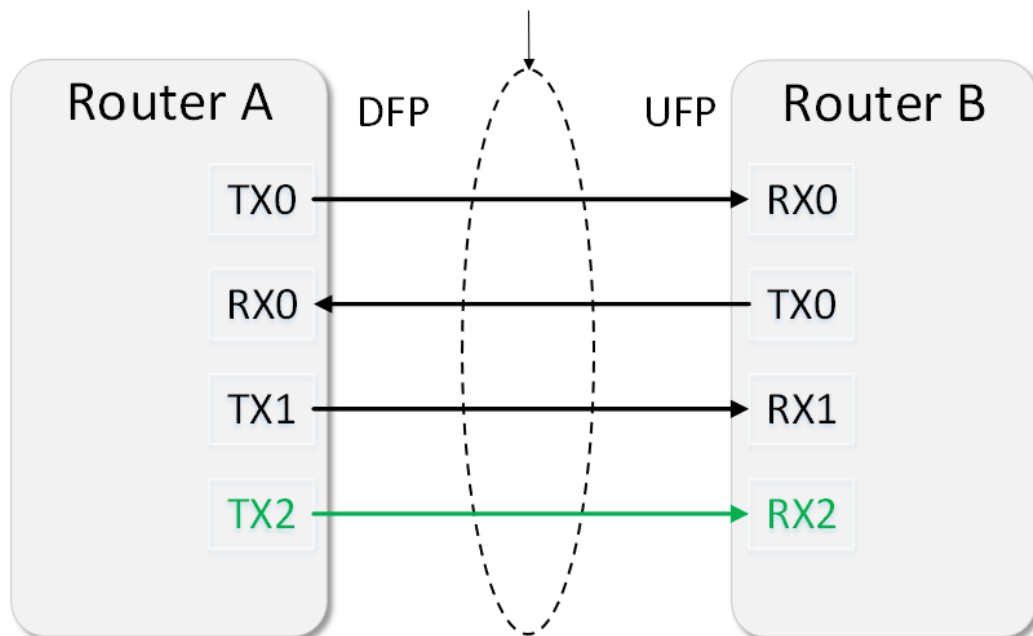
Section 7.2.1 describes the transition to Asymmetric Link from Symmetric Link between Router A and Router B. Section 7.2.2 describes transition to Symmetric Link from Asymmetric Link between Router A and Router B. The two routers referred to as Router A and Router B in these sections are two routers at each end of the link. Refer Figure 7-1 and Figure 7-2 for example Symmetric and Asymmetric Link between two routers Router A with DFP and Router B with UFP.

**Figure 7-1: Symmetric Link**



**Figure 7-2: Asymmetric Link**

Asymmetric Link with  
3 Tx in Router A and 3 Rx in Router B



#### **7.2.1. Transition from Symmetric Link to Asymmetric Link**

Connection Manager performs the following steps to enable Asymmetric Link between the two routers Router A and Router B:

1. Read *LANE\_ADAP\_CS\_0.Gen 4 Asymmetric Support* field (*Support 3 Tx* for 3 Transmitters, *Support 3 Rx* for 3 Receivers) of Lane 0 of USB4 Port in Router A and Router B and verify the USB4 Ports, all On-Board Re-timers connected between the USB4 Ports and the Cable support supports Gen 4 speed and Asymmetric Link with 3 transmitters or 3 receivers.
2. Read *PORT\_CS\_18.Cable Asymmetric Support* field in the Router transitioning to Asymmetric Link to ensure USB4 Port is connected to a Cable that supports an Asymmetric Link. This information does not indicate if the cable specifically supports Asymmetric Link with 3 Tx or 3 Rx.
3. Disable CLx on both ends of the Link by setting *LANE\_ADAP\_CS\_1.CL0s Enable*, *LANE\_ADAP\_CS\_1.CL1 Enable* & *LANE\_ADAP\_CS\_1.CL2 Enable* fields of Lane 0 Adapters to 0b.
4. On both sides of the Link, read the *LANE\_ADAP\_CS\_1.Adapter State* field of Lane 0 Adapter and verify that Lane Adapters are in CL0 state.
5. In the Router A and Router B set the *LANE\_ADAP\_CS\_1.Target Asymmetric Link* field of Lane 0 Adapter to the requested value. For example, set 01b for Asymmetric Link with 3 transmitters in Router A and correspondingly 10b for Asymmetric Link with 3 receivers in Router B or 10b for Asymmetric Link with 3 receivers in Router A and the correspondingly 01b for Asymmetric Link with 3 transmitters in Router B.
6. In the router which is changing one of its transmitter to receiver set *PORT\_CS\_19.StartAsymmetricFlow* field to 1b to initiate transition from Symmetric Link to Asymmetric Link. The Connection Manager does not write to USB4 Port Capability Register 19 while *StartAsymmetricFlow* is set to 1b.
7. Connection Manager waits at least 2 seconds to receive Notification Packet with Event Code = ASYM\_LINK from the DFP of the Link which confirms successful configuration of Asymmetric Link, else Connection Manager assumes the configuration of Asymmetric Link failed.
8. Connection Manager may read *PORT\_CS\_19.AsymmetricTransitionInProgress* field from the UFP of the Link a value of 1b indicates the transition to Asymmetric Link has not completed.
9. Read *LANE\_ADAP\_CS\_1.Negotiated Link Width* at both ends of the Link to ensure the link is active in the desired asymmetric configuration.  
Connection Manager may enable CLx on both ends of the Link by setting *LANE\_ADAP\_CS\_1.CL0s Enable*, *LANE\_ADAP\_CS\_1.CL1 Enable* & *LANE\_ADAP\_CS\_1.CL2 Enable* fields of Lane 0 Adapters to 1b.

### 7.2.2. Transition from Asymmetric Link to Symmetric Link

Connection Manager performs the following steps to enable Symmetric Link from an Asymmetric Link between the two routers Router A and Router B:

1. Disable CLx on both ends of the Link by setting *LANE\_ADAP\_CS\_1.CL0s Enable*, *LANE\_ADAP\_CS\_1.CL1 Enable* & *LANE\_ADAP\_CS\_1.CL2 Enable* fields of Lane 0 Adapters to 0b.
2. On both sides of the Link, read the *LANE\_ADAP\_CS\_1.Adapter State* field of Lane 0 Adapter and verify that all Lane Adapters are in CL0 state.
3. In the Router A and Router B set the *LANE\_ADAP\_CS\_1.Target Asymmetric Link* field of Lane 0 Adapter at both ends of the Link to 00b for Symmetric Link.
4. In the router which has 3 transmitters and is changing one of its transmitter to receiver set *PORT\_CS\_19.StartAsymmetricFlow* field to 1b to initiate transition from Asymmetric Link to Symmetric Link. The Connection Manager does not write to USB4 Port Capability Register 19 while *StartAsymmetricFlow* is set to 1b.

5. Connection Manager waits at least 2 seconds to receive Notification Packet with Event Code = ASYM\_LINK from the DFP of the Link which confirms successful configuration of Symmetric Link, else Connection Manager assumes the configuration of Symmetric Link failed.
6. Connection Manager may read PORT\_CS\_19. *AsymmetricTransitionInProgress* field from Router A a value of 1b indicates the transition to Symmetric Link has not completed.
7. Read LANE\_ADAP\_CS\_1. *Negotiated Link Width* at both ends of the Link to ensure the link is active in the Symmetric Link configuration.
8. Connection Manager may enable CLx on both ends of the Link by setting LANE\_ADAP\_CS\_1. *CL0s Enable*, LANE\_ADAP\_CS\_1. *CL1 Enable* & LANE\_ADAP\_CS\_1. *CL2 Enable* fields of all Lane Adapters to 1b.

### 7.3. Downstream Port Reset and Change of Link Parameters

The Connection Manager uses the PORT\_CS\_19. *Downstream Port Reset* bit in a Downstream Facing Port (DFP) to initiate a Downstream Port Reset. A Downstream Port Reset causes Lane Initialization to restart on the DFP. If the Port on the other side of the Link is an UFP, then the Router with the UFP enters the Uninitialized Unplugged state.

Because a Downstream Port Reset restarts Lane Initialization, it can be used to reconfigure the USB4 Link.

The Connection Manager performs the following steps to execute a Downstream Port Reset:

1. The Connection Manager can optionally change one or more of the following fields in the Adapter Configuration Space of the DFP:
  - *Target Link Speed* field (to initiate a change in Link speed).

*Note: (Version 1.0 Router) Both Lane Adapters in a Port must be programmed to the same Link speed.*

  - *Request RS-FEC Gen 2* bit (to initiate a change in RS-FEC at Gen 2 speeds).
  - *Request RS-FEC Gen 3* bit (to initiate a change in RS-FEC at Gen 3 speeds).

2. Set the PORT\_CS\_19. *Downstream Port Reset* bit of the DFP to 1b.

3. Wait for at least 10ms.

4. Set the PORT\_CS\_19. *Downstream Port Reset* bit of the DFP to 0b to allow Lane Initialization to take place.

*Note: After Lane Initialization of a Gen 2 or Gen 3 Link, the Lanes of the USB4 Port are not bonded. The Connection Manager needs to re-initiate Lane Bonding for the Link to operate as an Aggregated Gen 2 or Gen 3 Link.*

### 7.4. Time Synchronization

A Router contains a configurable Time Management Unit (TMU). The TMU implements a Time Synchronization Protocol, which provides a mechanism for synchronizing the real-time clocks and absolute time of connected Routers to a high degree of accuracy and precision.

#### 7.4.1. TMU Modes

The TMU in a Router can operate in multiple Modes. Each mode has a different accuracy level. Table 7-1 defines the possible TMU modes of operation and the accuracy level, where 0 is the lowest level of accuracy and 3 is the highest.

**Table 7-1: TMU Modes**

TMU Mode	Accuracy Level
Off	0
LowRes	1
HiFi-Uni	2
HiFi-Bi / HiFi-EnhancedUni / MedRes-EnhancedUni	3

#### 7.4.2. Router Accuracy Level Requirements

A Connection Manager determines the accuracy level requirements for a Router as follows:

- When a Router tunnels USB3 traffic, it requires an accuracy level of 1 or greater.
- When a Router tunnels DP traffic, it requires an accuracy level of 2 or greater.

A Connection Manager also reads the values in the TMU Minimum Requested Mode Entry in the DROM to determine the Router's minimum accuracy level Requirements.

*Note: A Router's accuracy level requirements can change dynamically when a DP Tunnel is setup or torn down.*

*Note: There are no accuracy level requirements for tunneling PCIe traffic.*

#### 7.4.3. TMU Configuration

A Connection Manager chooses the TMU Mode for each Link in a Domain as follows:

- If both Routers support Version 2.0 of the USB4 Specification, use MedRes Enhanced Uni-Directional Mode.
- Else
  - If CLx states are not enabled on the Link, use HiFi-Bi Mode.
  - On all Links on which CLx states are enabled, the Router with the maximum accuracy level requirements in the Domain determines the TMU Mode. The Connection Manager does the following:
    - Determine the accuracy level required by each Router in the Domain as described in 7.3.2.
    - Determine which Router has the maximum accuracy level requirements.
    - Select the TMU Mode that provides the maximum accuracy level.

For other TMU Modes the Connection Manager writes to the following fields:

- TMU\_RTR\_CS\_3.TSPacketInterval
- TMU\_ADP\_CS\_3.EnableUniDirectionalMode

- TMU\_ADP\_CS\_8.Enable Enhanced Uni-Directional Mode
- TMU\_ADP\_CS\_9.DirSwitchN

Table 7-2 defines the values that are used to configure each TMU Mode.

**Table 7-2: TMU Mode Configuration**

TMU Mode	TSPacket Interval	EnableUniDirectional Mode	AdapterTimeSync Interval	Enable Enhanced Uni-Directional Mode	DirSwitchN
Off	0	0	0	0	0
LowRes	1000	1	0	0	0
HiFi-Uni	16	1	0	0	0
HiFi-Bi	16	0	0	0	0
HiFi-EnhancedUni /MedRes-Enhanced Uni	NA	0	16	1	255

A Connection Manager also sets the parameters defined in Table 7-3 when it configures TMU Mode.

A Version 2.0 Connection Manager sets the values in the TMU Adapter Configuration Capability of an Upstream Facing Port before it sets the *Configuration Valid* bit of the Router to 1b.

**Table 7-3: TMU Mode Parameters**

TMU Mode	Downstream Router			DFP in Upstream Router		
	Freq Measurement Window	ErrorAvgConst, OffsetAvgConst, DelayAvgConst, FreqAvgConst	DeltaAvgCont	Replenish Timeout	Replenish Threshold	ReplenishN
Off	NA	NA	NA	NA	NA	NA
LowRes	30	4	NA	NA	NA	NA
HiFi-Uni	800	8	NA	NA	NA	NA
HiFi-Bi	800	8	NA	NA	NA	NA
MedRes-EnhancedUni	800	4	0	3125	25	128
HiFi-EnhancedUni	800	8	8	1875	0	222

#### 7.4.3.1. Router Connection

When a Router is first connected to a Domain, the TMU Mode in the Router is Off by default.

After detecting a new connection, the Connection Manager sets the TMU\_ADP\_CS\_6.Disable Time Sync bit to 1b in the Ports on both sides of the Link that connects the new Router to the Domain. This ensures that the Routers do not send any TMU Packets on the Link until the TMU in the new Router is configured.

A Connection Manager gathers the information needed to determine the accuracy level required for the Router during Router Enumeration and by reading DROM.

A Connection Manager does not enable TMU operation when Upstream Router and the Downstream Routers are connected via two Single-Lane Link.

The Connection Manager enables TMU operation on both sides of the Link after the Aggregated Link is established.

A Connection Manager performs the following steps to enable the TMU in a Router:

1. Set the Local Time of the Router using the Time Posting registers:
  - a. Read the Host Router Time of the Domain.
  - b. Write the Host Router Time to the *Post Local Time* registers in the TMU Router Configuration Capability of the Router.
  - c. Write a value of 0x1 to the TMU\_RTR\_CS\_24.*Post Time Low* field and a value of 0xFFFFFFFF to the TMU\_RTR\_CS\_25.*Post Time High* field of the Router.
  - d. Write a value of 0x0 to the TMU\_RTR\_CS\_25.*Post Time High* field, which causes the Router to immediately update its local time to the value in the *Post Local Time* registers.
  - e. Periodically read the *Post Time* field of the Router to determine when the Router is done updating its local time. The Router sets the *Post Time* field to 0x0 after it has updated its local time.
2. Optionally change the TMU Mode for the Routers in the Domain according to the TMU Mode Switch routine defined in Section 7.4.3.4.
3. Set the TMU Mode in the Router using the *TMU Enable* routine defined in Section 7.4.3.4.1.
4. Wait 50ms to allow the TMU to converge before Path Setup.

If the Link is an Inter-Domain Link, see USB4 Inter-Domain Service Protocol Specification for more details.

#### **7.4.3.2. Router Disconnect**

When a Router is disconnected from a Domain, the Connection Manager may decide to change the TMU Mode in the Domain. The TMU Mode change Routine in Section 7.3.3.4 describes how to change the TMU Mode in a Domain.

#### **7.4.3.3. DisplayPort Plug/Unplug**

When a DisplayPort Sink or Source is Plugged or Unplugged to a Domain, a Connection Manager may decide to change the TMU Mode in the Domain. The TMU Mode change Routine in Section 7.3.3.4 describes how to change the TMU Mode in a Domain.

#### **7.4.3.4. TMU Mode Change Routines**

A Connection Manager performs the following steps to change the TMU Mode in a Domain:

1. Teardown all DP Paths.
2. Indicate to the Routers in the Domain that a Time Disruption Event is occurring:
  - a. Set the TMU\_RTR\_CS\_0.*Time Disruption* bit to 1b in each Router in the Domain.
  - b. Start with the Router(s) at the highest Depth and end with the Router with the lowest Depth (i.e. the Host Router).

3. Disable Time Synchronization in the Domain:
  - a. Set the TMU\_ADAP\_CS\_6.Disable Time Sync bit to 1b in each Port in each Router in the Domain.
  - b. Start with the Router(s) at the highest Depth and end with the Router with the lowest Depth (i.e. the Host Router).

*Note: If there is an Inter-Domain Link, addition of any new Router to the Domain that is the Inter-Domain Time Source results in a Time Disruption Event, see USB4 Inter-Domain Service Protocol Specification for more details.*
4. Change the TMU Mode on each Link in the Domain. Start with the Link(s) between the Host Router and any connected Device Routers, then down the tree ending with the Links to the Device Routers at the maximum depth:
  - a. If changing the TMU Mode from Off to another Mode, follow the *TMU Enable* routine in Section 7.4.3.4.1.
  - b. If changing the TMU Mode to Off follow the TMU Disable routine in Section 7.4.3.4.2.
  - c. Else, follow the TMU Mode Switch routine in Section 7.4.3.4.3

*Note: If there is an Inter-Domain Link in the Domain, see USB4 Inter-Domain Service Protocol Specification for more details.*
5. Wait 50ms.
6. Connection Manager may enable CLx by setting LANE\_ADAP\_CS\_1.CL0s Enable, LANE\_ADAP\_CS\_1.CL1 Enable & LANE\_ADAP\_CS\_1.CL2 Enable fields of Lane 0 Adapters to 1b.
7. Indicate to the Routers in the Domain that the Time Disruption Event has ended:
  - a. Set the TMU\_RTR\_CS\_0 Time Disruption bit to 0b in each Router in the Domain.
  - b. Start with the Router(s) at the highest Depth and end with the Router with the lowest Depth (i.e., the Host Router).
8. Set up the DP Paths that were torn down in step 1.

#### **7.4.3.4.1. TMU Enable**

When the current TMU Mode is Off, a Connection Manager changes the TMU Mode on a Link as follows:

1. Set the TMU Parameters according to Table 7-3 in the Downstream Router
2. If the new TMU Mode is Enhanced Uni-Directional:
  - a. If CLx states are enabled, disable CLx by setting LANE\_ADAP\_CS\_1.CL0s Enable, LANE\_ADAP\_CS\_1.CL1 Enable & LANE\_ADAP\_CS\_1.CL2 Enable fields of Lane 0 Adapters in both sides of the Link to 0b.
  - b. Set the TMU\_ADAP\_CS\_9.AdapterTimeSyncInterval field to 16 and TMU\_ADAP\_CS\_8.Enable Enhanced Uni-Directional Mode bit to 1b in the TMU Adapter Configuration of the Downstream Router.

- c. Set the *TMU\_ADP\_CS\_9.AdapterTimeSyncInterval* field to 16 and *TMU\_ADP\_CS\_8.Enable Enhanced Uni-Directional Mode* bit to 1b in the TMU Adapter Configuration of the Upstream Router.
  - d. Set the *TMU\_ADP\_CS\_6.Disable Time Sync* bit to 0b in the TMU Adapter Configuration of the Downstream Router.
  - e. Set the *TMU\_ADP\_CS\_6.Disable Time Sync* bit to 0b in the TMU Adapter Configuration of the Upstream Router.
3. Else If the new TMU Mode is HiFi-Bi:
- a. Set the *TMU\_ADP\_CS\_3.EnableUniDirectionalMode* bit to 0b in both Ports.
  - b. Set the *TMU\_RTR\_CS\_3.TSPacketInterval* field according to Table 7-2 in the Downstream Router.
  - c. Set the *TMU\_ADP\_CS\_6.Disable Time Sync* bit in the TMU Adapter Configuration to 0b in the Upstream Router.
  - d. Set the *TMU\_ADP\_CS\_6.Disable Time Sync* bit in the TMU Adapter Configuration to 0b in the Downstream Router.
4. Else:
- a. Set the *TMU\_RTR\_CS\_3.TSPacketInterval* field according to Table 7-2 in the Upstream Router.
  - b. Set the *TMU\_ADP\_CS\_3.EnableUniDirectionalMode* bit to 1b and the *Disable Time-Sync* bit to 0b in the TMU Adapter Configuration of the UFP.
  - c. Set the *TMU\_ADP\_CS\_3.EnableUniDirectionalMode* bit to 1b and the *Disable Time-Sync* bit to 0b in the TMU Adapter Configuration of the DFP.

#### **7.4.3.4.2. TMU Disable**

A Connection Manager changes the TMU Mode to Off on a Link as follows:

- 1. If current TMU Mode is Enhanced Uni-Directional:
  - a. Set the *TMU\_ADP\_CS\_9.AdapterTimeSyncInterval* field to 0 and *TMU\_ADP\_CS\_8.Enable Enhanced Uni-Directional Mode* bit to 0b in the TMU Adapter Configuration of both the Downstream Router and the Upstream Router.
  - b. Set the *TMU\_ADP\_CS\_6.Disable Time Sync* bit to 1b in the TMU Adapter Configuration of both the Downstream Router and the Upstream Router.
- 2. Else If current TMU Mode is HiFi-Bi:
  - a. Set the *TMU\_RTR\_CS\_3.TSPacketInterval* field to 0b in the Downstream Router
  - b. Set the *TMU\_ADP\_CS\_6.Disable Time Sync* bit to 1b in in both Ports.
- 3. Else:
  - a. Set the *TMU\_RTR\_CS\_3.TSPacketInterval* field to 0b in the Upstream Router
  - b. Set the *TMU\_ADP\_CS\_3.EnableUniDirectionalMode* bit to 0b and the *TMU\_ADP\_CS\_6.Disable Time Sync* bit to 1b in both Ports.

#### 7.4.3.4.3. TMU Mode Switch

TMU Mode switch can happen only when the TMU is in Bi-Directional Mode or Uni-Directional Mode. When the TMU is in Enhanced Uni-Directional Mode, there is no need to switch to a different TMU Mode. When the current TMU Mode is not Off, a Connection Manager changes the TMU Mode on a Link to a new Mode (other than Off) as follows:

1. Set the *TMU\_ADP\_CS\_3.EnableUniDirectionalMode* bit (0b for Bi-Directional, 1b for Uni-Directional) in the TMU Adapter Configuration of the DFP to indicate the new TMU Mode.
2. Set the *TMU\_RTR\_CS\_3.TSPacketInterval* or *TMU\_ADP\_CS\_9.AdapterTimeSyncInterval* field to indicate the new TMU Mode (see Table 7-1 for values).
  - a. If the new TMU Mode is a Bi-Directional mode, set the *TMU\_RTR\_CS\_3.TSPacketInterval* field in the Downstream Router.
  - b. If the new TMU Mode is a Uni-Directional mode, set the *TMU\_RTR\_CS\_3.TSPacketInterval* field in the Upstream Router.
3. Set the TMU Parameters for the new TMU Mode in the Downstream Router according to Table 7-3.
4. Set the *TMU\_ADP\_CS\_3.EnableUniDirectionalMode* bit (0b for Bi-Directional, 1b for Uni-Directional) in the TMU Adapter Configuration of the UFP to indicate the new TMU Mode.
5. Set the *TMU\_ADP\_CS\_6.Disable Time Sync* bit in the DFP to 0b.
6. Set the *TMU\_ADP\_CS\_6.Disable Time Sync* bit in the UFP to 0b.

#### 7.5. Lane Disable and Enable

The Connection Manager writes to the *LANE\_ADP\_CS\_1.Lane Disable* bit in a Lane Adapter Configuration Capability to enable and disable that Lane. Setting the *Lane Disable* bit to 0b enables a Lane. Setting the *Lane Disable* bit to 1b disables a Lane.

The Connection Manager does not change the value of the *Lane Disable* bit when one or more of the following are true:

- One or more CLx states are enabled over the Link.
- The TMU Mode is not Off.
- Port is operating as an Aggregated Link.

When a Lane is disabled, the associated Lane Adapter transitions to the Disabled state. The Lane Adapter at the other end of the Link transitions to CLd state.

A Connection Manager must disable Lane 1 before disabling Lane 0 in a Port. When both Lanes are disabled, the Connection Manager must enable Lane 0 before enabling Lane 1.

A Connection Manager must disable the Lane 1 Adapter if, after Lane Initialization, the Lane 1 Adapter is in CL0 state, but the Lane 0 Adapter is not in CL0 state.

Before disabling or enabling a Lane, the Connection Manager disables all Paths on the Link and sets all fields in the Lane Adapter Configuration Capability to their default values except the following:

- Lane Disable

- Target Link Speed
- Adapter State

Section 7.5.1 describes how a Connection Manager disables and enables a Lane on a Link that is not an Inter-Domain Link. Section 7.5.2 describes how a Connection Manager disables and enables a Lane on an Inter-Domain Link.

#### 7.5.1. Intra-Domain Links

A Connection Manager does the following to disable Lane 1 in a Port:

1. If Port is not in two single-Link mode, initiate a Downstream Facing Port Reset to transition to Gen 2/3 Link speed and to two Single-Lane Links and wait for Lane Initialization to complete.
  - A Connection Manager shall not change the value of the *Lane Disable* bit for a Lane Adapter unless all of the fields in its Lane Adapter Configuration Capability except the following are set to their default values: *Lane Disable*, *Target Link Speed*, *Adapter State*.
2. When the Port has two Single-Lane Links, disable Lane 1 by setting the `LANE_ADP_CS_1.Lane Disable` bit to 1b in the Lane 1 Adapter in the Downstream Facing Port. Note that the Router then sends a Hot Plug Event Packet with the *UPG* bit set to 1b for the Lane 1 Adapter.

A Connection Manager does the following to disable Lane 0 in a Port:

1. If Lane 1 is not disabled, disable Lane 1 as described above.
2. After Lane 1 is disabled, disable Lane 0 by setting the `LANE_ADP_CS_1.Lane Disable` bit to 1b in the Lane 0 Adapter in the Downstream Facing Port.

A Connection Manager does the following to enable only Lane 0 in a Port:

1. Set the `LANE_ADP_CS_1.Lane Disable` bit to 0b for Lane 0 in the Downstream Facing Port.
2. Initiate a Downstream Facing Port Reset as described in Section 7.2.

A Connection Manager does the following to enable both Lane 0 and Lane 1 in a Port:

1. Set the `LANE_ADP_CS_1.Lane Disable` bit to 0b for both Lane 1 and Lane 0 in the Downstream Facing Port.
2. Initiate a Downstream Facing Port Reset as described in Section 7.2.

A Connection Manager does the following to enable only Lane 1 in a Port:

1. Confirm that Lane 0 is enabled.
2. Set the `LANE_ADP_CS_1.Lane Disable` bit to 0b for Lane 1 in the Downstream Facing Port.
3. Initiate a Downstream Facing Port Reset as described in Section 7.2.

*Note: The Downstream Facing Port Reset restarts Lane Initialization on both Lane 1 and Lane 0.*

### 7.5.2. Inter-Domain Links

A Connection Manager disables Lane 1 in an Inter-Domain Link if the Link operates in Gen 2/3 speed and either of the following are true:

- Aggregated (x2) operation over an Inter-Domain Link is not supported.
- Lane Bonding failed.

The Connection Manager sets the *LANE\_AD\_P\_CS\_1.Lane Disable* bit to 1b in the Downstream Facing Port that is in its Domain. This is a steady state that does not change until the Port is disconnected.

### 7.6. Link Errors

Following Link Initialization, a Lane Adapter may report logical layer errors in the *Logical Layer Errors* field of its Lane Adapter Configuration Capability. A Connection Manager that uses the *Logical Layer Errors* field needs to clear it between Link Initializations. Reading the *Logical Layer Errors* field clears it. Note that it is only possible to read the *Logical Layer Errors* field in a UFP after the Router with the UFP is enumerated.

(Version 2.0 Router) For a Gen 4 Link the Connection Manager may enable Link Recovery from uncorrectable error events by setting *PORT\_CS\_19.Enable Gen 4 Link Recovery* bit to 1b at both ends of the Link when the Link is in CL0.

(Version 2.0 Router) For a Gen4 Link the Connection Manager may trigger Link Recovery by setting *PORT\_CS\_19.Initiate Gen 4 Link Recovery* bit to 1b at either end (either the Downstream Facing Port or the Upstream Facing Port) of the Link. Reception of a Notification Packet with Event Code = LINK\_RECOVERY indicates that the Link is back to Active state, after which the Connection Manager shall verify that all tunneled protocols using the recovered Link are in an operational state. Connection Manager may check fields such as *ADP\_USB3\_GX\_CS\_4.Port Link State*, *ADP\_USB3\_GT\_PORT\_CS\_0.Port Link State*, *ADP\_PCIE\_CS\_0.Link*, DP IN and DP OUT Adapter Configuration Capability Fields to ensure the recovered Link is in operational state.

A Connection Manager can optionally configure a Router to send a Notification Packet for certain Link errors. Section 7.7.1 describes how a Connection Manager enables Notification Packets. Section 7.7.2 describes how a Connection Manager acknowledges a Notification Packet from a Router.

### 7.7. Notifications

#### 7.7.1. Notification Enable

A Connection Manager enables Notification Packets from a Router as follows:

1. A Version 2.0 Connection Manager shall enable the “Sequence bit in Notification Packet” capability if supported by the Router.
2. Set the desired bits to 1b in the Adapter Configuration Space:
  - a. Set the corresponding bits to 1b in the *LANE\_AD\_P\_CS\_2.Logical Layer Errors Enable* field to enable Notification Packets for specific Link Layer errors.
    - i. When a Port operates with two single-Lane Links, error reporting and notification is enabled independently for each Lane. Else, error reporting and notification is enabled for Lane 0 only.

- b. Set the *ADP\_CS\_5.HEC Error Enable* bit to 1b to enable Notification Packets for HEC errors.
  - c. Set the *ADP\_CS\_5.Flow Control Error Enable* bit to 1b to enable Notification Packets for flow control errors.
- 3. Read back the fields in Adapter Configuration Space that were modified:
  - a. If any of the bits are 0b, then the Router does not support Notification Packets for that error type.
  - b. Else, the Router supports that error type and will send Notification Packets when an error occurs.

#### **7.7.2. Notification Acknowledgement**

A Connection Manager shall process Notification Packets from a specific Router in the order received and needs to send a Notification Acknowledgment in response to the following Notification Packets:

- A Notification Packet with Event Code = ERR\_LINK
- A Notification Packet with Event Code = ERR\_HEC
- A Notification Packet with Event Code = ERR\_FC
- A Notification Packet with Event Code = ERR\_PLUG
- A Notification Packet with Event Code = DP\_BW
- A Notification Packet with Event Code = PCIE\_WAKE
- A Notification Packet with Event Code = DP\_CON\_CHANGE
- A Notification Packet with Event Code = DPTX\_DISCOVERY

The Connection Manager shall not send Notification Acknowledgement for any other packets.

If the “Sequence bit in Notification Packet” capability is disabled, a Connection Manager shall respond to a Notification Packet with a Notification Acknowledgment Packet as defined in Section 6.4.2.8 of the USB4 Base Specification. If the “Sequence bit in Notification Packet” capability is enabled, a Connection Manager shall respond to a Notification Packet with an Enhanced Notification Acknowledgment Packet as defined in Section 6.4.2.9 of the USB4 Base Specification.

## 8. Power Management

### 8.1. Low-Power States

The Connection Manager follows the rules below when it enables CL2, CL1, and CL0s states:

- Do not enable any CLx state in a Port with two Single-Lane Links. The Connection Manager either bonds the Lanes or disables Lane 1 Adapter before enabling any CLx state.
- Do not enable any CLx state when the PORT\_CS\_18. *CLx Protocol Support* bit is set to 0b, indicating that CLx states are not supported on the Link.
- Before enabling a CLx state, the Connection Manager verifies that both Routers support the relevant CLx state by reading the LANE\_ADAP\_CS\_0.CL1 Support, LANE\_ADAP\_CS\_0.CL2 Support, and/or LANE\_ADAP\_CS\_0.CL0s Support bits in both Routers.
- For a Gen 2/Gen3 Link if the Connection Manager sets the LANE\_ADAP\_CS\_1.CL2 *Enable* bit to 1b, it also sets the LANE\_ADAP\_CS\_1.CL1 *Enable* bit to 1b. For a Gen 4 Link if the Connection Manager sets the LANE\_ADAP\_CS\_1.CL1 *Enable* or LANE\_ADAP\_CS\_1.CL2 *Enable* bit to 1b, it also sets the LANE\_ADAP\_CS\_1.CL0s *Enable* bit to 1b.
- If the Connection Manager sets the LANE\_ADAP\_CS\_1.CL1 *Enable* bit to 1b, it also sets the LANE\_ADAP\_CS\_1.CL0s *Enable* bit to 1b.
- For Gen2 and Gen 3 the LANE\_ADAP\_CS\_1.CL2 *Enable* bit, the LANE\_ADAP\_CS\_1.CL1 *Enable* bit, and the LANE\_ADAP\_CS\_1.CL0s *Enable* bit must have the same values at both ends of a Link.
- For Gen 4 if the LANE\_ADAP\_CS\_1.CL2 *Enable* bit, the LANE\_ADAP\_CS\_1.CL1 *Enable* bit is set to 1b at one end of the Link, LANE\_ADAP\_CS\_1.CL0s *Enable* bit must be set to 1b on the other side of the Link. Connection Manager may enable CL2 without enabling CL1 or CL0s.
- Do not enable any CLx state in the Downstream Facing Ports of a Router if any CLx states are not enabled in the Upstream Facing Port of the Router.
- Do not enable any CLx states on Link that is used for Host-to-Host traffic. Set LANE\_ADAP\_CS\_1.CL2 *Enable*, LANE\_ADAP\_CS\_1.CL1 *Enable* and LANE\_ADAP\_CS\_1.CL0s *Enable* bit to 0b.
- If the TMU Mode is Enhanced Uni-Directional Mode, to enable CLx states wait at least 50ms after TMU enabling to allow TMU convergence.
- Do not enable any CLx states on a Link if PORT\_CS\_18. *Cable CLx Support* bit is set to 0b, value of 0b indicates the Cable does not support CLx states.
- Enable CLx by setting the relevant bits in Lane 0 Adapter Space.
- Do not set the CL0s Enable, CL1 Enable, or CL2 Enable bit to 1b in a USB4 Port that is part of an Inter-Domain Link.

*Note: When the TMU Mode on a Link is Enhanced Uni-Directional, the Link can enter CL1, CL2 or CL0s states.*

*Note: When the TMU Mode on a Link is HiFi Bi-Directional, the Link will not enter any CLx states since the Time Sync Handshakes are too frequent.*

*Note: When the TMU Mode on a Link is HiFi Uni-Directional, the Link will not enter CL1 or CL2 state since the Time Sync Handshakes are too frequent. The Link in the UFP-to-DFP direction may enter CL0s if traffic allows.*

*Note: When the Time Sync mode on a Link is LowRes Uni-directional the Link can enter CL1 or CL0s if traffic allows.*

Connection Manager may enable CLx after the Router is enumerated and configured, enabling CLx prior to the enumeration and configuration of the Router may take longer. Prior to enabling a CLx state, the Connection Manager needs to set the *PM Secondary* bit in the Lane 0 Adapters at both sides of the Link to resolve conflicts between concurrent requests to enter Low-Power states:

- The Connection Manager keeps the LANE\_ADAP\_CS\_1.*PM Secondary* bit set to 1b in the Upstream Facing Port.
- The Connection Manager sets the LANE\_ADAP\_CS\_1.*PM Secondary* bit to 0b in the Downstream Facing Port.

## 8.2. Sleep and Wake

A Connection Manager can transition its Domain into a sleep state. For example, when the host system transitions to a Low-Power state, the Connection Manager can transition the Routers in its Domain to the sleep state for greater power savings. Section 8.2.1 describes how a Connection Manager transitions a Domain to sleep state.

A Wake event brings the Domain out of sleep state. Section 8.2.2 describes how a Wake event causes the Domain to exit sleep state.

Section 8.2.3 describes how a Connection Manager responds when a neighbor Domain enters sleep state.

### 8.2.1. Entry to Sleep State

Before transitioning its Domain into sleep state, the Connection Manager sets the Wake Events that wake a Router from sleep as defined in Section 8.2.1.1. The Connection Manager also disables Low-Power states (CLx) before the transition to Sleep state. After setting the Wake Events, the Connection Manager transitions the Domain to sleep state as defined in Section 8.2.1.2.

#### 8.2.1.1. Wake Configurations

The Connection Manager enables Wake Events by setting one or more of the Wake Enable bits in Table 8-1 to 1b:

**Table 8-1: Wake Enable Bits**

Wake Enable Bit	Wake Event Configuration
<i>Enable Wake on Connect</i>	Can be set to 1b in any DFP that has the PORT_CS_19.USB4 <i>Port is Configured</i> bit set to 0b. Otherwise, set to 0b. Not Applicable to a UFP.
<i>Enable Wake on Disconnect</i>	Can be set to 1b in any DFP that has the PORT_CS_19.USB4 <i>Port is Configured</i> bit set to 1b. Otherwise, set to 0b. Not Applicable to a UFP.
<i>Enable Wake on USB4 Wake</i>	Can be set to 1b in any DFP that has the PORT_CS_19.USB4 <i>Port is Configured</i> bit set to 1b. Otherwise, set to 0b. Must be set to 1b in all UFP, so that the Host Router can wake the entire USB4 tree.
<i>Enable Wake on Inter-Domain</i>	Can be set to 1b in any DFP that is part of an Inter-Domain Link. Otherwise, set to 0b. Not applicable to a UFP.
<i>Enable Wake on PCIe</i>	Can be set to 1b in any Device Router. Actual wake from a PCIe wake event is determined by the native PCIe wake event.

<i>Enable Wake on USB3</i>	Can be set to 1b in any Router. Actual wake from a USB3 wake event is determined by the native USB3 wake event.
<i>Enable Wake on DP</i>	Can be set to 1b in any Router that supports DP tunneling.

If the Connection Manager enables one or more Wake Events in a Router, it does the following:

- Set the `PORT_CS_19.Enable Wake on USB4 Wake` bit to 1b in all Ports between that Router and the Host Router. This allows the Wake Event to propagate upstream and back to the Connection Manager.
- Set the `PORT_CS_19.USB4 Port is Configured` bit to 1b in both the UFP of the Router and the DFP that it is connected to.

*Note: If Wake Events are not enabled or on Inter-Domain Link, the Connection Manager may set the `PORT_CS_19.USB4 Port is Configured` bit to 1b or it may leave it at the default value of 0b.*

#### **8.2.1.2. Sleep Entry Sequence**

A Connection Manager does the following to initiate entry to sleep state:

1. If the Connection Manager initiated any transactions on the Sideband Channel, wait for those Transactions to complete.
2. Set the `ROUTER_CS_5.Enter Sleep` bit to 1b in the Router Configuration Space of all Routers in the Domain.

*Note: The internal USB SuperSpeed Plus hub or the internal USB peripheral device should be in U3 state prior to setting the Enter Sleep bit. The internal PCIe Switch or the internal PCIe Endpoint should be in D3 state prior to setting the Enter Sleep bit.*

*Note: The order of setting the Enter Sleep bit in the different Routers does not matter. The Connection Manager can set the Enter Sleep bit in different Routers in parallel and collect Sleep Ready indication from different Routers in parallel.*

3. (Version 1.0 Router) Poll the `ROUTER_CS_6.Sleep Ready` bit in the Router Configuration Space of each Router in the Domain. When the Sleep Ready bit is 1b, it indicates that a Router is ready for a Sleep Event

(Version 2.0 Router) Wait for reception of a Notification Packet from each Router with Event Code = `ROP_CMPLT` and the Event Info field set to 03h. Read the `ROUTER_CS_6.Sleep Ready` bit in the Router Configuration Space of each Router in the Domain to verify that a Router is ready for a Sleep Event. Alternatively, the Connection Manager may poll the `ROUTER_CS_6.Sleep Ready` bit until the bit is set to 1b:

- a. The Connection Manager continues to the next step when the *Sleep Ready* bits in all Routers in the Domain are 1b.

*Note: A Router sets the Sleep Ready bit to 1b within 50 milliseconds after the Enter Sleep bit is set to 1b<sup>1</sup>. It is recommended that the Connection Manager waits for 500 milliseconds before it times out.*

4. If the host system supports entry to sleep via assertion of PCIe PERST#, and if all Routers in the Domain support PCIe tunneling, the Connection Manager informs the host system that the Domain is ready to enter sleep state. The mechanism to inform the host system is implementation specific.
5. Else, the Connection Manager instructs the Host Router to send an LT\_LRoff Transaction on the Sideband Channel of each Downstream Facing Port. The Connection Manager then informs the host system that the Domain is ready to enter sleep state. The mechanism to inform the host system is implementation specific.

During the Sleep Entry Sequence, the Connection Manager does not do any of the following after setting the ROUTER\_CS\_5.Enter Sleep bit to 1b:

- Stop the transition to sleep state:
  - After setting the *Enter Sleep* bit to 1b, the Connection Manager must finish the sleep entry sequence defined in this section. When the sleep entry sequence is complete, the Connection Manager may then Wake the Domain to transition out of sleep state.
- Initiate a Transaction on the Sideband Channel:
  - The Connection Manager must wait until the Domain exits sleep state before it can initiate Transactions again.
- Write to the PORT\_CS\_19.USB4 Port is Configured bit.
- Write to the PORT\_CS\_19.USB4 Port is Inter-Domain bit.

#### **8.2.2. Exit from Sleep state**

A Domain exits sleep state after either of the following:

- A wake event in the Domain.
- A wake indication from the host system.

The Connection Manager can identify the location and cause of a wake event in the Domain by reading the bits in Table 8-2. The bits are read from each Router in the Domain to identify the source of the wake event.

**Table 8-2: Wake Status Bits**

Wake Enable Bit	Wake Event
<i>Wake on Connect Status</i>	Set to 1b after a wake event is generated by the Port as a result of a connect to the Port.
<i>Wake on Disconnect Status</i>	Set to 1b after a wake event is generated by the Port as a result of a disconnect from the Port.
<i>Wake on USB4 Wake Status</i>	Set to 1b after a wake event is generated by the Port as a result of a USB4 Wake Event.

<sup>1</sup> If any of the Router's Sideband Channels operate in TBT3-compatible mode, then the maximum value is 80ms.

<i>Wake on Inter-Domain Status</i>	Set to 1b after a wake event is generated by the Port as a result of an Inter-Domain Wake.
<i>Wake on PCIe Status</i>	Set to 1b when a PCIe Wake indication from a PCIe device connected to a PCIe Downstream Adapter causes the Router to exit from sleep
<i>Wake on USB3 Status</i>	Set to 1b when a USB3 Wake indication causes the Router to exit from sleep.
<i>Wake on DP Status</i>	Set to 1b when a DP connect or disconnect event

After exiting from sleep state, the Connection Manager needs to reconfigure each Router as described in chapter 3. It also performs Path setup as described in 0 to restore any Paths in the Domain that were present before the Domain entered sleep state.

After a Router exits from Sleep state, a Connection Manager may disable a USB3 Path that was enabled before Sleep state entry. To disable the Path, the Connection Manager sets the Path Enable bit in the USB3 Adapter Configuration Capability to 0b and sets the *Valid* bit in the USB3 Adapter Configuration Capability to 1b. If the Connection Manager does not disable a Path after exit from Sleep state, it shall set up the Path again as if it is a new Path.

### 8.2.3. Behavior with a Neighbor Domain in Sleep State

When the Routers in a neighbor Domain are in Sleep state, the Connection Manager needs to monitor the state of the Inter-Domain Link and behave as defined in this section.

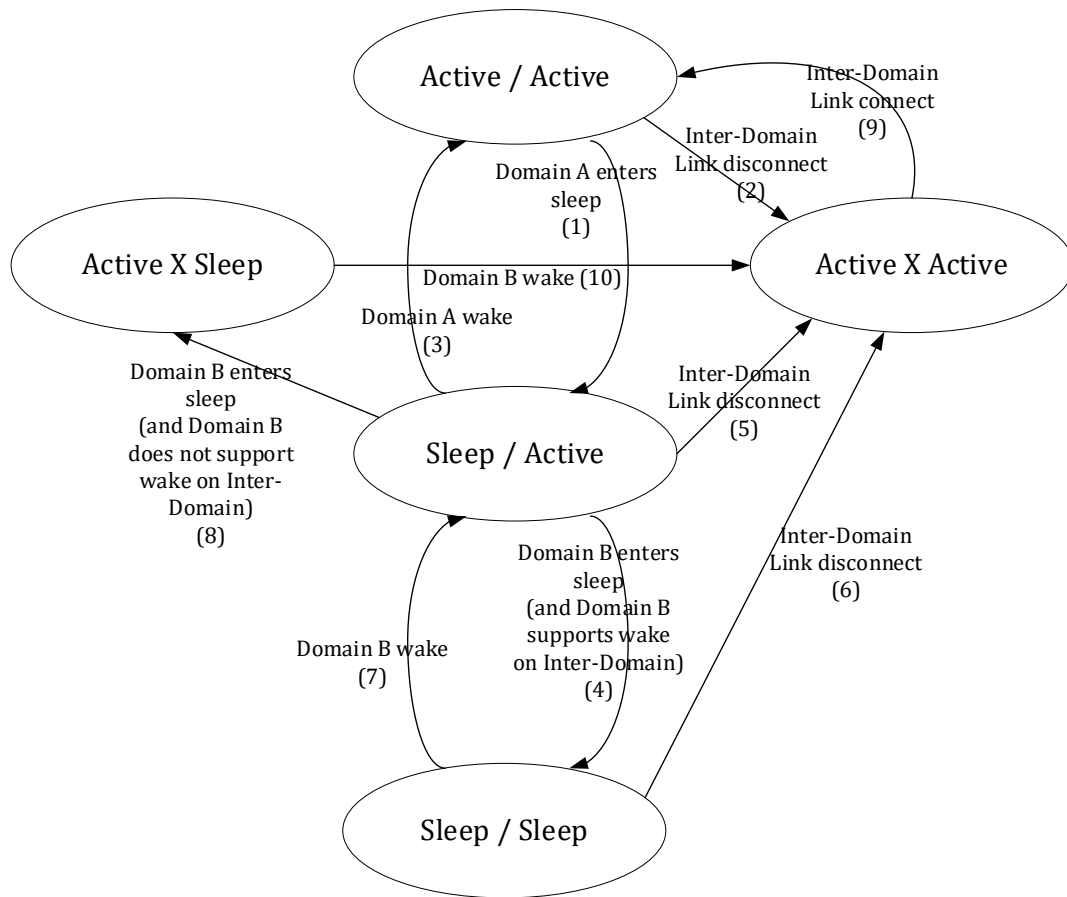
For convenience's sake, the neighbor Domain is named Domain A in this section, while the Domain (and Connection Manager) whose behavior is documented is named Domain B in this section.

When Domain A is in Sleep state with wake on Inter-Domain enabled, the Connection Manager in Domain B may wake Domain A by disconnecting the Inter-Domain Link. The Connection Manager in Domain B needs to therefore monitor the Sleep state in Domain A.

Figure 8-1 and Table 8-3 describe the behavior of the two Domains when entering Sleep state. Each state in the Figure denotes the state of the two Domains. For example, the state "Sleep / Active" refers to Domain A being in Sleep state, while Domain B is in Active (Enumerated) state. A sign of "X" in a state name (such as "Active X Sleep") indicates that the Inter-Domain Link is disconnected.

Each row in Table 8-3 corresponds to an arc in Figure 8-1. "A" refers to the Router in Domain A that faces the Inter-Domain Link. "B" refers to the Router in Domain B that faces the Inter-Domain Link. "Connection Manager A" refers to the Connection Manager in Domain A. "Connection Manager B" refers to the Connection Manager in Domain B.

**Figure 8-1: Monitoring Sleep State Status in a Neighbor Domain**



**Table 8-3: Monitoring Sleep state status in a neighbor Domain**

	Inter-Domain Link Behavior	Connection Manager Behavior in Domain B	Router Detected bit <sup>2</sup>
1	A sends an LT_LRoff Transaction to B. B disconnects Port (the Sideband Channel stays connected).	Connection Manager B receives a Hot Plug Event Packet with the <i>UPG</i> bit set to 1b. Connection Manager B reads the <i>Router Detected</i> bit to identify that Domain A is in Sleep state	1b
2	Link is disconnected. Not an Inter-Domain Link anymore.	Both Connection Managers receive a Hot Plug Event Packet with the <i>UPG</i> bit set to 1b. CM reads the <i>Router Detected</i> bit to identify that Link is disconnected.	0b
3	Link Initialization starts from Phase 3.	Both Connection Managers receive a Hot Plug Event Packet with the <i>UPG</i> bit set to 0b.	1b
4	B sends an LT_LRoff transaction to A. A ignores the Transaction.	Connection Manager B brings Domain B into Sleep state.	N/A

<sup>2</sup> In the USB4 Port Capability of the Router in Domain B facing the Inter-Domain Link

5	Link is disconnected. Not an Inter-Domain Link anymore.	Connection Manager B does not receive a Hot Plug Event Packet. Connection Manager B may identify that Link is disconnected by reading the <i>Router Detected</i> bit.	0b
6	Link is disconnected. Not an Inter-Domain Link anymore.	Connection Manager B receives a wake indication. Connection Manager B reads the <i>Router Detected</i> bit to identify that Link is disconnected.	0b
7	B starts Link Initialization. A ignores it.	Connection Manager B restores the Inter-Domain Link state. The <i>Router Detected</i> bit is set to 1b since the Sideband Channel is on.	1b
8	Link is disconnected. Not an Inter-Domain Link anymore.	On entry to Sleep state, Router B generates a disconnect of the Inter-Domain Link.	0b
9	Inter-Domain Link is established.	Connection Manager B runs the Inter-Domain Discovery Protocol.	1b
10	Link is disconnected.		0b

## 9. Operations

### 9.1. Router Operations

Router Operations are defined in Section 8.3.1 of the USB4 Specification. A Connection Manager initiates a Router Operation by writing to registers ROUTER\_CS\_9 through ROUTER\_CS\_26 in Router Configuration Space. A Connection Manager shall not issue a new Router Operation until the previous Router Operation has completed.

The Connection Manager does the following to initiate a Router Operation:

1. If the Operation requires data, write the data to ROUTER\_CS\_9 through ROUTER\_CS\_24. If the Operation does not require any data, then set ROUTER\_CS\_9 through ROUTER\_CS\_24 DWs to zeros.
2. If the Operation requires Metadata, write the metadata to ROUTER\_CS\_25. If the Operation does not require Metadata, then set ROUTER\_CS\_25 DW to zeros.
3. Write the Operation Opcode to bits 15:0 of ROUTER\_CS\_26.
4. Set the *Operation Valid* bit in ROUTER\_CS\_26 to 1b. The Connection Manager does not change the values of ROUTER\_CS\_9 through ROUTER\_CS\_26 while the *Operation Valid* bit is set to 1b.

A Router starts processing a Router Operation when *Operation Valid* bit is set to 1b. The Router sets the *Operation Valid* bit to 0b after it finishes processing the Operation. A Version 2.0 Router also sends a Notification Packet with *Event Code* = ROP\_CMPLT and the *Event Info* field set to 00h. On reception of the Notification Packet, the Connection Manager reads the *Operation Valid* bit to verify that the Router is done processing the Router Operation, then reads the results of the Router Operation. Alternatively, the Connection Manager polls the *Operation Valid* bit to see when the Router is done processing the Router Operation, then reads the results of the Router Operation.

A Connection Manager reads the results of a Router Operation as follows:

1. Read the *Operation Not Supported* bit:
  - a. If the *Operation Not Supported* bit is 1b, the Router does not support the Router Operation and any results are not valid.
  - b. If the *Operation Not Supported* bit is 0b, the Router supports the Router Operation and the results are valid.
2. Read the ROUTER\_CS\_26.*Status* field to see if the Router Operation finished successfully.
3. If the Completion Metadata is defined for the Router Operation, read ROUTER\_CS\_25.*Metadata*.
4. If the Completion Data is defined for the Router Operation, read ROUTER\_CS\_9.*Data* through ROUTER\_CS\_24.*Data*.

### 9.2. Port Operations

Router Operations are defined in Section 8.3.2 of the USB4 Specification. A Connection Manager uses registers 8 (Opcode), 9 (Metadata), and 18 (Data) in the Sideband (SB) Register Space of a USB4 Port to initiate a Port Operation. Section 4.1.1.3.2 in the USB4 Specification describes how a Connection Manager writes to SB Register Space.

A Connection Manager shall only issue a Port Operation to a USB4 Port after the previous Port Operation for that USB4 Port completes execution.

A Connection Manager does the following to issue a Port Operation:

1. Optionally write to the Metadata register in the SB Register Space of the target Port.
2. Optionally write to the Data register in the SB Register Space of the target Port.
3. Write to the Opcode register in the SB Register Space of the target Port.

A Connection Manager can only have one pending Transaction at a time. The Connection Manager needs to verify the completion of a write to the SB Register Space of the target Port before proceeding with the next write to SB Register Space.

*Notes:*

1. *If the write access is to the SB Register Space of a Router, then the Router completes the write access and sets the Pending bit to 0b within 50 milliseconds time after the Pending bit is set to 1b. It is recommended that the Connection Manager waits for 500 milliseconds before it times out.*
2. *If the write access is to the SB Register Space of a Re-timer, then the Router completes the write access and sets the Pending bit to 0b within 150 milliseconds time after the Pending bit is set to 1b. It is recommended that the Connection Manager waits for 1 second before it times out.*

After the Connection Manager writes to the Opcode register, the USB4 Port executes the Port Operation using the information in the Metadata and Data registers.

When a USB4 Port finishes executing a Port Operation, it updates the Metadata register with Completion Metadata (if any), and the Data register with Completion Data (if any). It also writes the completion status to the Opcode register as follows:

- If the Opcode register is set to a FourCC value of “!CMD” (444D4321h), the USB4 Port Operation is not supported by the Port.
- If the Opcode register is set to a FourCC value of “ERR ” (20525245h), the USB4 Port fails to execute a Port Operation or executed the Port Operation and failed.
- If the Opcode register is set to 0, the Port Operation completed successfully.

A Connection Manager should verify that a Port Operation completes successfully (Opcode Register is 0) before reading the Metadata and Data registers. If a Port Operation does not complete successfully, the values in the Metadata and Data registers are not valid. A Version 2.0 Router issues a Notification Packet with *Event Code* = POP\_CMPLT to indicate completion of the operation.

When the *Pending* bit in SB Register Space is 1b, the Connection Manager does not write to registers PORT\_CS\_1 through PORT\_CS\_17 in Port Configuration Space. It also does not write to SB Register 18.