

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

**Revision 1.0  
April 2021**

## Release History

Version	Comments	Issue Date
1.0	First release	April 2021

**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 © 2021, USB Promoter Group (Apple Inc., HP Inc., Intel Corporation, Microsoft Corporation, Renesas Corporation, STMicroelectronics, and Texas Instruments).

# Table of Contents

<b>1. INTRODUCTION</b>	<b>7</b>
1.1. SCOPE	7
1.2. REFERENCE DOCUMENTS	7
1.3. OVERVIEW	7
<b>2. INITIALIZATION</b>	<b>9</b>
2.1. RING 0 INITIALIZATION	9
2.2. CONNECTION MANAGER TRANSMIT FLOW	10
2.3. CONNECTION MANAGER RECEIVE FLOW	10
2.4. HOST INTERFACE RESET	11
2.5. DROM ACCESS	11
2.6. PCIe MEMORY BAR ACCESS	11
<b>3. ROUTER CONFIGURATION</b>	<b>11</b>
3.1. CONNECTION EVENT DETECTION	11
3.1.1. USB4 CONNECTION EVENTS	11
3.1.2. DISPLAY CONNECTION EVENTS	12
3.2. ROUTER INITIALIZATION	13
3.2.1. ROUTER ENUMERATION AND CONFIGURATION	13
3.2.2. ADAPTER ENUMERATION	15
3.2.3. ENABLING PROTOCOL TUNNELING	16
3.3. IDENTIFYING LOOPS AND INTER-DOMAIN LINKS	16
3.4. DETECTING A REMOVAL EVENT	19
3.4.1. LANE ADAPTER	19
3.4.2. DP ADAPTER	20
<b>4. CONTROL PACKETS</b>	<b>20</b>
4.1. ROUTING	20
<b>5. PATH CONFIGURATION</b>	<b>21</b>
5.1. FLOW CONTROL	22
5.1.1. FLOW CONTROL DISABLED	23
5.1.2. DEDICATED FLOW CONTROL	23
5.1.3. SHARED FLOW CONTROL	23
5.1.4. BUFFER ALLOCATION	23
5.2. PATH SETUP AND TEARDOWN	25
5.2.1. PATH SETUP	25
5.2.2. PATH TEARDOWN	25
5.3. PCIe PATH SETUP AND TEARDOWN	26
5.3.1. PCIe PATH SETUP	26
5.3.2. PCIe PATH TEARDOWN	28

<b>5.4. DISPLAYPORT PATH SETUP AND TEARDOWN</b>	<b>28</b>
5.4.1. CONNECTION MANAGER DISCOVERY	28
5.4.2. PAIRING DP ADAPTERS	29
5.4.3. DISPLAY PORT PATH SETUP AND TEARDOWN	29
<b>5.5. USB3 PATH SETUP AND TEARDOWN</b>	<b>33</b>
5.5.1. USB3 PATH SETUP	33
5.5.2. USB3 PATH TEARDOWN	35
5.5.3. USB3 ADAPTER DISCONNECT	35
<b>5.6. INTER-DOMAIN PATH SETUP AND TEARDOWN</b>	<b>35</b>
5.6.1. TRANSMIT RING SETUP	37
5.6.2. RECEIVE RING SETUP	37
<b>6. BANDWIDTH MANAGEMENT</b>	<b>38</b>
<b>6.1. BANDWIDTH ALLOCATION</b>	<b>38</b>
6.1.1. BANDWIDTH CALCULATIONS	38
6.1.2. CONFIGURATION MECHANISM	40
6.1.3. GUARANTEED BANDWIDTH	40
6.1.4. BANDWIDTH MANAGEMENT EVENTS	41
<b>6.2. BANDWIDTH CLIENTS</b>	<b>42</b>
6.2.1. DISPLAYPORT TUNNELING	42
6.2.2. USB3 TUNNELING	43
6.2.3. PCIE TUNNELING	43
6.2.4. HOST-TO-HOST TUNNELING	43
<b>7. USB4 PORT MANAGEMENT</b>	<b>45</b>
<b>7.1. LANE BONDING</b>	<b>45</b>
<b>7.2. DOWNSTREAM PORT RESET AND CHANGE OF LINK PARAMETERS</b>	<b>46</b>
<b>7.3. TIME SYNCHRONIZATION</b>	<b>46</b>
7.3.1. TMU MODES	46
7.3.2. ROUTER ACCURACY LEVEL REQUIREMENTS	47
7.3.3. TMU CONFIGURATION	47
<b>7.4. LANE DISABLE AND ENABLE</b>	<b>51</b>
7.4.1. INTRA-DOMAIN LINKS	51
7.4.2. INTER-DOMAIN LINKS	52
<b>7.5. LINK ERRORS AND NOTIFICATIONS</b>	<b>52</b>
7.5.1. NOTIFICATION ENABLE	53
7.5.2. NOTIFICATION ACKNOWLEDGEMENT	53
<b>8. POWER MANAGEMENT</b>	<b>54</b>
<b>8.1. LOW-POWER STATES</b>	<b>54</b>
<b>8.2. SLEEP AND WAKE</b>	<b>55</b>
8.2.1. ENTRY TO SLEEP STATE	55
8.2.2. EXIT FROM SLEEP STATE	57
8.2.3. BEHAVIOR WITH A NEIGHBOR DOMAIN IN SLEEP STATE	57

<b>9. OPERATIONS</b>	<b>59</b>
<b>9.1. ROUTER OPERATIONS</b>	<b>59</b>
<b>9.2. PORT OPERATIONS</b>	<b>60</b>

## **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 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 1.0 with Errata and ECN through October 15, 2020 (USB4 Specification)

Universal Serial Bus (USB4™) Inter-Domain Service Protocol, Version 1.0, to be published (USB4 Inter-Domain 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:

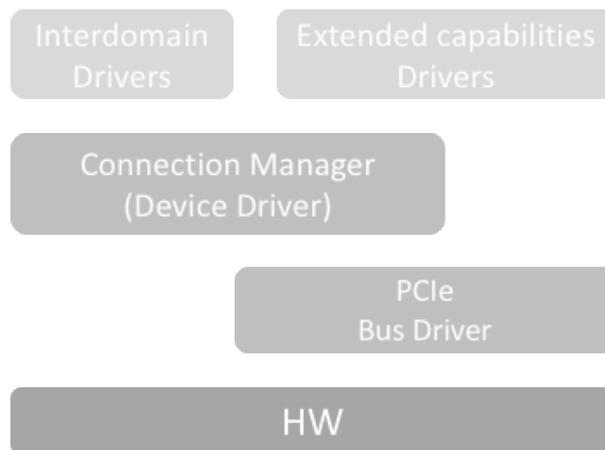
- Enumeration of Routers
- Setup and tear-down of Paths
- Initialization of a hot plugged Router and removal of a hot unplugged Router
- Configuration of QoS behavior including flow control and bandwidth allocation

A Connection Manager can also 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. Initialization

### 2.1. Ring 0 Initialization

To properly operate and enable Connection Manager functionality, the host interface driver needs to initialize 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 uses Transmit Ring 0 to inject Control Packets into a Domain. A Connection Manager uses Receive Ring 0 to receive Control Packets from a Domain.

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 the 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 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 is responsible for configuring 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 when it has data to transmit using Transmit Ring 0:

1. Create a Transmit Descriptor with the following fields:
  - The *Address* field points to the Data Buffer in Host Memory that holds the 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* and *Consumer Index* fields 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 is responsible for configuring 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 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 that need to be processed.

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

2. For each Receive Descriptor that needs 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 processed either free or recycle the data buffers as appropriate, increment the Consumer Index accordingly and write the result (modulo *Ring Size*) to the *Producer Index* and *Consumer Index* fields 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 can use 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 a single DW.

### **3. Router Configuration**

#### **3.1. Connection Event Detection**

A Router can detect and report 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 (include Host Router).
- The Router is hot plugged into the DFP of another Router in the Domain.

- The Router completes a Downstream Port Reset.
- The UFP of the Router is disabled, then re-enabled.
- The Router exits sleep state.

When a USB4 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 continues to resend the Hot Plug Event Packet until the Connection Manager responds with a Hot Plug Acknowledgment Notification Packets. If the Connection Manager receives multiple Hot Plug Event Packets for a hot plug event, it needs to send at least one Notification Packet for the event.

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

- A Hot Plug Event Packet.
- 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 Manger can conclude that another Router is connected downstream.
- An implementation-specific method (for Host Router).

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.

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

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

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.
2. The Connection Manager enumerates the Router by assigning it a Topology ID.
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.
4. The Connection Manager creates and configures one or more Paths through the Router.

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

A Connection Manager needs to be able to read from and write to the Configuration Space of a Router in order to enumerate it. Before accessing the Configuration Space of a newly connected Router, the Connection Manager needs to "unlock" 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 4.
  - 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 Rev 1.0 or higher, then the Router supports USB4.
  - Else, the Router does not support USB4.
4. The Connection Manager sends a single Write Request to the Router that writes the following information to 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
  - `Valid` bit = 1b

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

During 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 will clear 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.

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 can keep a list of Router UUIDs so that on subsequent connection events, the Connection Manager can compare the UUID value of a connected Router to the list to determine whether or not to enumerate the Router.

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
- `Wake Enable` bits

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

After configuring a Router, the Connection Manager polls 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.

### 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 starting from Adapter number 1 up to Adapter number {Max Adapter}. The Read Requests target `ADP_CS_0` to `ADP_CS_5` of 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 bonded to Lane 0 or be disabled before the Connection Manager can setup any Paths through the USB4 Port. The Connection Manager can initiate bonding for Lane 0 and Lane 1 as described in Section 7.1 or Lane disabled as described in Section 7.4.

If the Router has a USB3 Upstream Adapter, then it supports USB3 Tunneling. The Connection Manager makes a note that the Adapter is the USB3 Upstream Adapter. The Connection Manager can setup a Path for USB3 Tunneling as described in Section 5.5. Note that each USB3 Adapter only supports one USB3 Path.

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

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

If the Router has a PCIe Upstream Adapter, then PCIe tunneling is supported by the Router. The Connection Manager makes a note that the Adapter is the PCIe Upstream 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.

If the Router has one or more PCIe Downstream Adapters, the Connection Manager maps each PCIe Downstream 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 get 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 can read all the Configuration Spaces by following the Linked List as described in section 8.2.2 in the USB4 Specification. The

Connection Manager ignores 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 Router initialization is to enable Tunneled Protocols. By default, USB3 Tunneling and PCIe Tunneling are disabled in a Device Router.

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

- USB3 Tunneling is enabled in the Router that is connected to the Upstream Facing Port of the Device Router.
- The Device Router has a USB3 Upstream Adapter.
- Bandwidth Management requirements allows USB3 tunnel.

The Connection Manager sets the `ROUTER_CS_5.USB3 Tunneling On` bit to 1b to enable USB3 Tunneling. The Connection Manager sets the `ROUTER_CS_5.USB3 Tunneling On` bit to 0b to disable USB3 Tunneling.

The Connection Manager can optionally enable 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 a PCIe Upstream 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.

After enabling USB3 Tunneling and/or PCIe tunneling, the Connection Manager executes the following steps:

1. Set the `ROUTER_CS_5.Configuration Valid` bit to 1b.
2. 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. 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`. When the Connection Manager receives a Notification Packet with an *Event Code* = `ERR_ENUM`, 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 can identify a DFP-to-DFP connection and determine 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 Router receiving the Read Request responds with a Notification Packet with an *Event Code* = ERR\_ENUM.
4. The Connection Manager sends an Inter-Domain UUID Request Packet as defined in the USB4 Inter-Domain 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 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.

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**

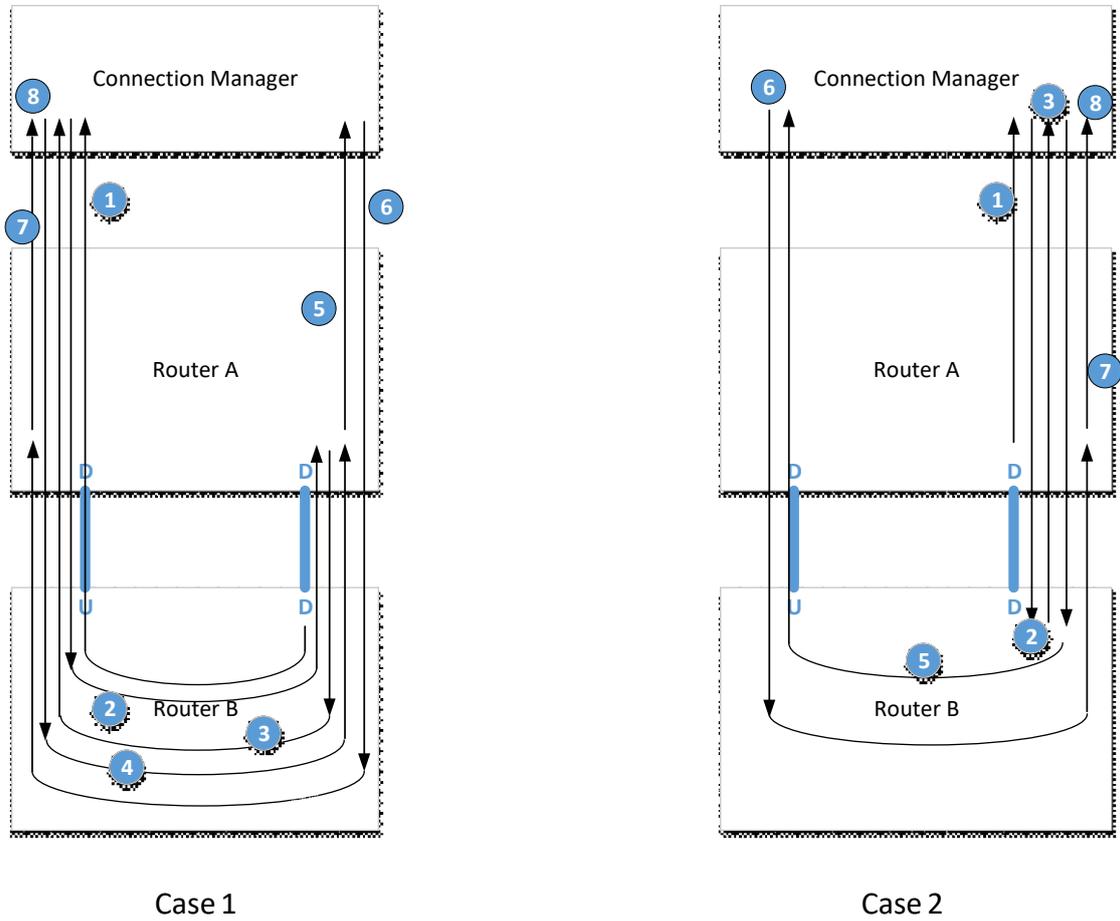
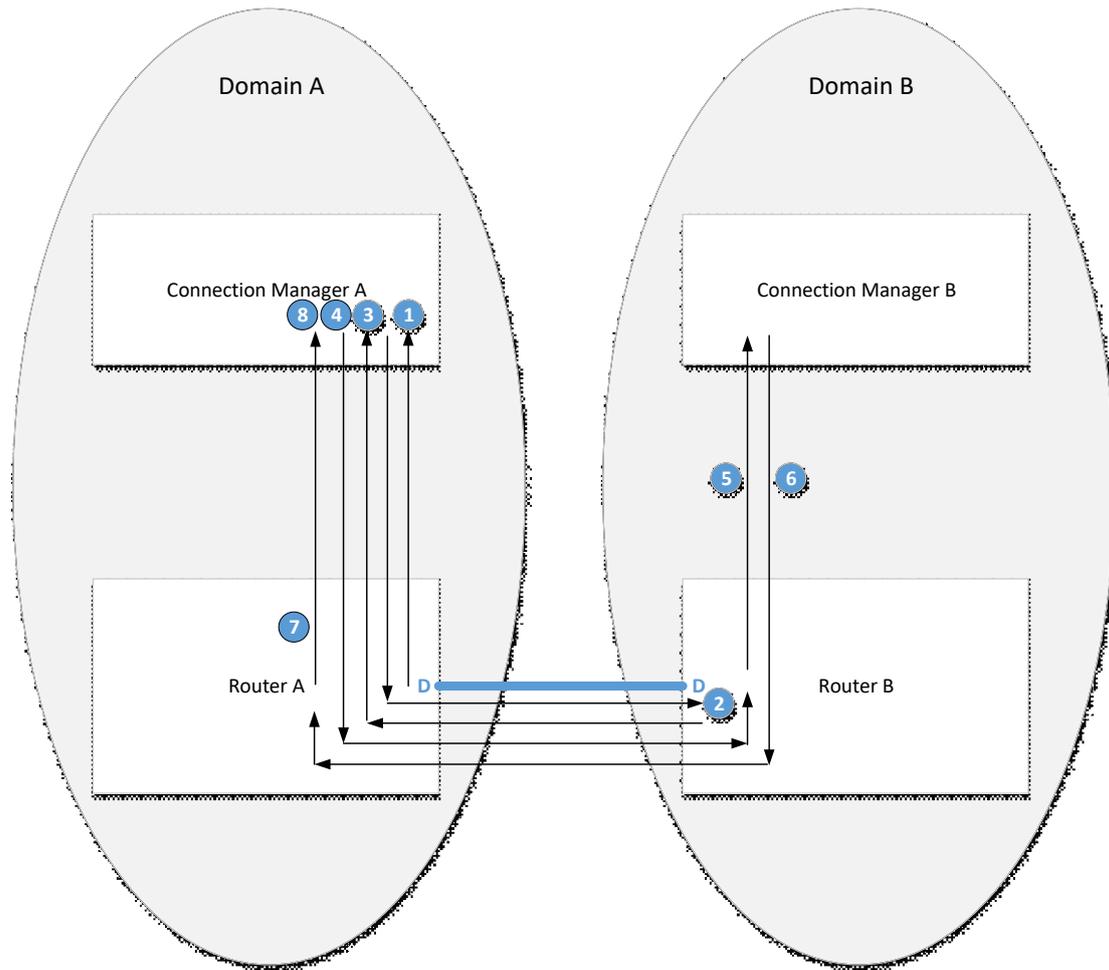


Figure 3-2 illustrates the steps above when there is in inter-Domain Link.

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



If the Connection Manager determines that there is an inter-Domain Link, it needs to set the `PORT_CS_19.USB4 Port is inter-Domain` bit to 1b.

### 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 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 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 Manger can conclude that another Router is connected downstream.
- An implementation-specific method.

After a removal event, the Connection Manager tears down all non-zero Paths 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 5.4.3.5.

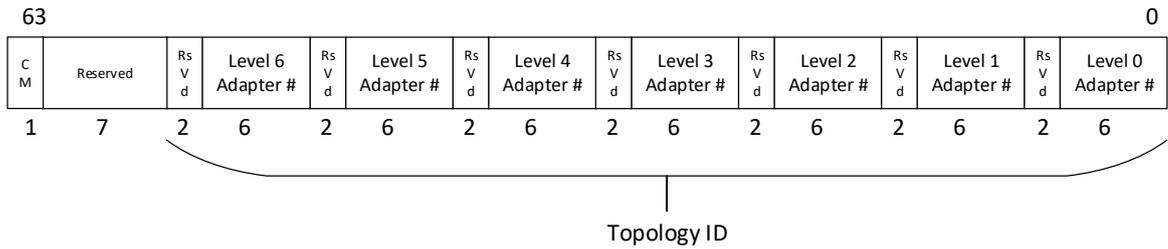
## **4. Control Packets**

### **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. The *CM* bit in the Route String is set to Control Packets flowing upstream are forwarded over the upstream Port of each Router and do not rely on the Route String for routing.

### 5. Path Configuration

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

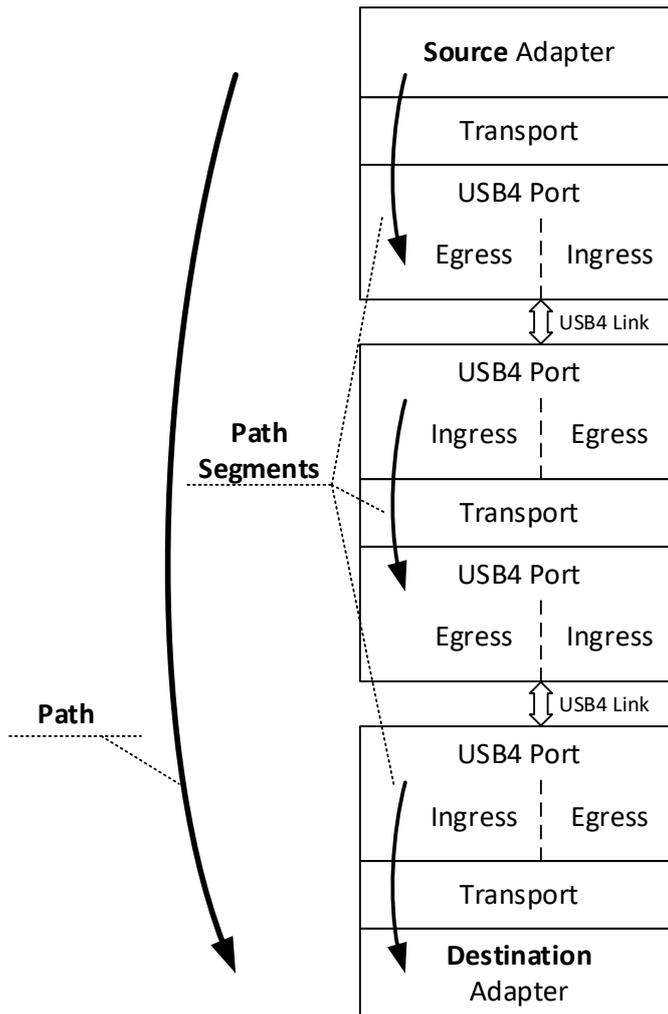
- Section 5.1 only applies to Path Setup in a USB4 Adapter.
- Section 5.2 applies to all Paths, regardless of the protocol that is tunneled.
- Section 5.3 describes the additional steps required for a Path that tunnels PCIe traffic.
- Section 5.4 describes the additional steps required for a Path that tunnels DisplayPort traffic.
- Section 5.5 describes the additional steps required for a Path that tunnels USB3 traffic.
- Section 5.6 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 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 WRR (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 needs to ensure 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 needs to ensure 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 credits that an Ingress Adapter can 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 credits in the *Non-Flow Controlled Buffers* field when Paths with flow control disabled are configured and decrease the number of credits 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 only uses the Flow Control Disabled scheme for the Main-Link Path of DisplayPort Tunneling.

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

When a Path with Dedicated Flow Control is setup, the Connection Manager allocates a number of credits 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 credits for the Dedicated Flow Control buffers are automatically deallocated from the Path and can be reallocated to a different buffer.

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

A Connection Manager sets the number of credits that an Ingress Adapter can 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 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.

A Connection Manager does not change the *Link Credits Allocated* field in a Protocol Adapter.

#### **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 Lane Adapter 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 a Buffer Allocation Request Router operation.
4. Determine the maximum number of DP Streams supported through the Lane Adapter, then calculate the number of credits needed for DP Tunneling.

The maximum number of supported DP Streams is potentially limited by the number of implemented credits, the number of DP IN Adapters, and the number of implemented Paths at the Lane Adapter.

The limitation due to number of credits can be calculated as follows:

$$\#DP \text{ Streams} = \text{Round down}((TB - CP \text{ Credits} - baMaxUSB3 - baMaxPCIe)/(baMinDPaux + baMinDPmain))$$

where:

TB = Total Buffers (see Step 1)

CP Credits = Control Path Credits (see Step 2)

*Note: The Connection Manager guarantees baMaxUSB3 for USB3 Tunneling to ensure correct scheduling of Isochronous traffic.*

*Note: If there are not enough DP Streams available, the Connection Manager may be able to allocate more credits for DP Tunneling by reducing the number of credits that are allocated to PCIe Tunneling. Note that a PCIe Tunneling Path requires at least 6 credits.*

$$DP \text{ Credits} = \#DP \text{ Streams} * (baMinDPaux + baMinDPmain)$$

where #DP Streams is the total number of current or future DP streams through that Lane Adapter.

5. Calculate the number of remaining buffers as follows:

$$\text{Remaining Buffers} = TB - (CP \text{ Credits} + DP \text{ Credits})$$

where:

TB = Total Buffers (see Step 1)

CP Credits = Control Path Credits (see Step 2)

DP Credits = DP Tunneling Credits (see Step 4)

6. Allocate baMaxUSB3 Credits to USB3 Tunneling.

7. Allocate credits for PCIe Tunneling as follows:

$$PCIe \text{ Credits} = \text{Max}(6, \text{Min}(baMaxPCIe, \text{Remaining Buffers} - baMaxUSB3))$$

8. The Connection Manager policy for reserving Host-to-Host Tunneling credits is implementation specific. A Connection Manager may reserve credits when allocating credits to the other Tunneling Paths or it may choose to not reserve credits and only allocate any remaining credits when a Host-to-Host connection is established.

- The maximum number of credits that can be allocated for a Host-to-Host Tunneling Path is baMaxHI (of the Host Router).
- A Host-to-Host Tunneling Path requires at least one credit.
- Note that rebalancing credits requires Path teardown and setup.

## 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 sends a Read Request to read 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 sends a Write Request to write 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 Output Adapter field when issuing a Write Request to a Path Entry.

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*

When configuring a Path for a specific Protocol Tunneling, it will have at least one Path at each direction.

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 credit allocation, will be received and processed by the Ingress Adapter at the other Router.*

2. The Connection Manager sets the *Path Enable* bit to 1b in both Adapters Configuration Spaces, 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 will have at least one Path at each direction and all Paths must be torn down when the Tunneled Protocol is removed.

The Connection Manger 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 Adapters, Source and Destination Adapters.
2. For each Hop along the Path:
  - a. Send a Write Request to the Ingress Adapter to set the *PATH\_CS\_0.Valid* bit to 0b. The Connection Manager does not change any other fields in the Path Configuration Space at this time.
  - b. Send periodic Read Requests to poll the *PATH\_CS\_1.Pending Packets* bit until it is 0b, which indicates that all Transport Layer Packets belonging to the Path have been dequeued.

*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 the PCIe Paths between two Routers.

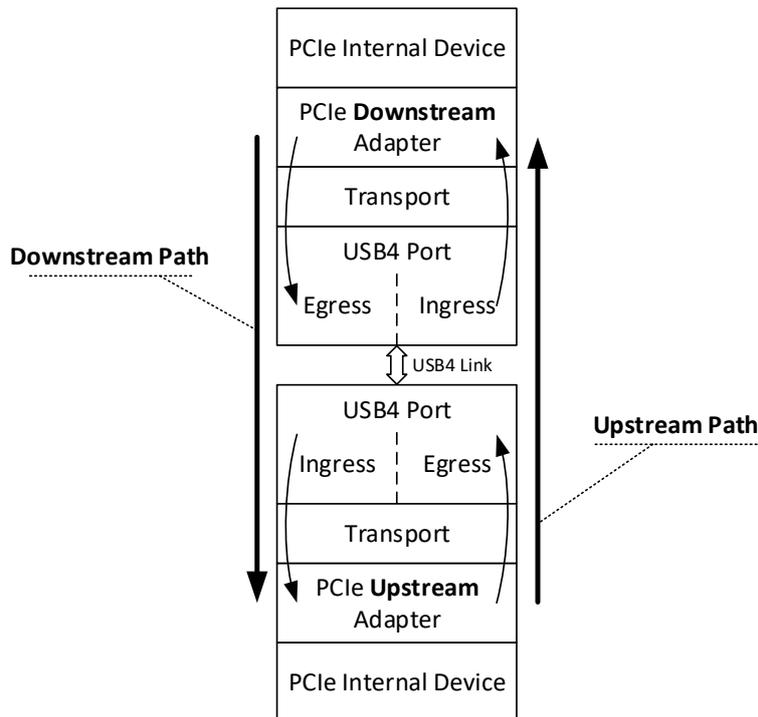
#### **5.3.1. PCIe Path Setup**

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 PCIe Downstream Adapter. The Path Attributes are defined in Table 5-1 (Path segment is from the PCIe Downstream 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 PCIe Upstream Adapter).
3. Set up the upstream Path.
  - Set Path Attributes in the Path Configuration Space of the PCIe Upstream Adapter. The Path Attributes are defined in Table 5-1 (Path segment is from the PCIe Upstream 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 PCIe Downstream Adapter).
4. Enable Transport Layer Packets on the PCIe Upstream Adapter by setting the `ADP_PCIE_CS_0.Path Enable` bit to 1b in the PCIe Upstream Adapter.
5. Enable Transport Layer Packets on the PCIe Downstream Adapter by setting the `ADP_PCIE_CS_0.Path Enable` bit to 1b in the PCIe Downstream Adapter.

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

Path Segment	Input HopID	Output HopID	Buffers Allocation: Dedicated	Priority	Weight	IFC	EFC	ISE	ESE
PCIe Port to USB4 Port	8		NA <sup>1</sup>	3	1	NA <sup>1</sup>	1	NA <sup>1</sup>	0
USB4 Port to PCIe Port		8	baMaxPCIe from the Buffer Allocation Request Operation <sup>2</sup>	3	1	1	0	0	0
1. A Connection Manager performs Read/Modify/Write to Path CS and does not change the <i>PATH_CS_0.Path Credits Allocated</i> , <i>PATH_CS_1.IFC Flag</i> and <i>PATH_CS_1.ISE Flag</i> fields at the PCIe Adapter 2. For more information regarding Buffer allocation, see Section 5.1.4									

**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 on the PCIe Downstream Adapter by setting the *ADP\_PCIE\_CS\_0.Path Enable* bit to 0b in the PCIe Downstream Adapter.
2. Disable Transport Layer Packets on the PCIe Upstream Adapter by setting the *ADP\_PCIE\_CS\_0.Path Enable* bit to 0b in the PCIe Upstream Adapter.
3. Teardown the downstream Path as defined in Section 5.2.2.
4. 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; 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; 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.

#### 5.4.2. 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 5.4.3.1.

#### 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 5.4.3.1.
2. Determine the relevant DP IN and DP OUT Adapter Capabilities as described in Section 5.4.3.2.
3. Configure the Adapters along the Path as described in Section 5.4.3.3.

Section 5.4.3.4 describes how a Path is configured.

##### 5.4.3.1. Available Bandwidth

A Connection Manager calculates the available USB4 bandwidth on each Link along the Path from the DP IN Adapter to the DP OUT Adapter according to Section 6. Then, based on the minimum available bandwidth, the Connection Manager concludes how much bandwidth can be allocated to the display stream. If needed, the Connection Manager can lower the maximal DP Link that would be established.

Table 5-2 shows the Asynchronous Bandwidth needed for each DP Link according to lane count and link rate.

**Table 5-2: DisplayPort Required Bandwidth (Gbps)**

	4 Lanes	2 Lanes	1 Lane
HBR3 (8.1Gbps)	26	13	6.5
HBR2 (5.4Gbps)	17.6	8.8	4.4
HBR (2.7Gbps)	8.8	4.4	2.2
RBR (1.62Gbps)	5.2	2.6	1.3

Note:  $Bandwidth = Lane\ Count * Link\ Rate * 8/10$

Note: Bandwidth calculation assumes max utilization over the DP Link and does not account for blanking and stuffing symbol removal

#### 5.4.3.2. 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.
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), 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.

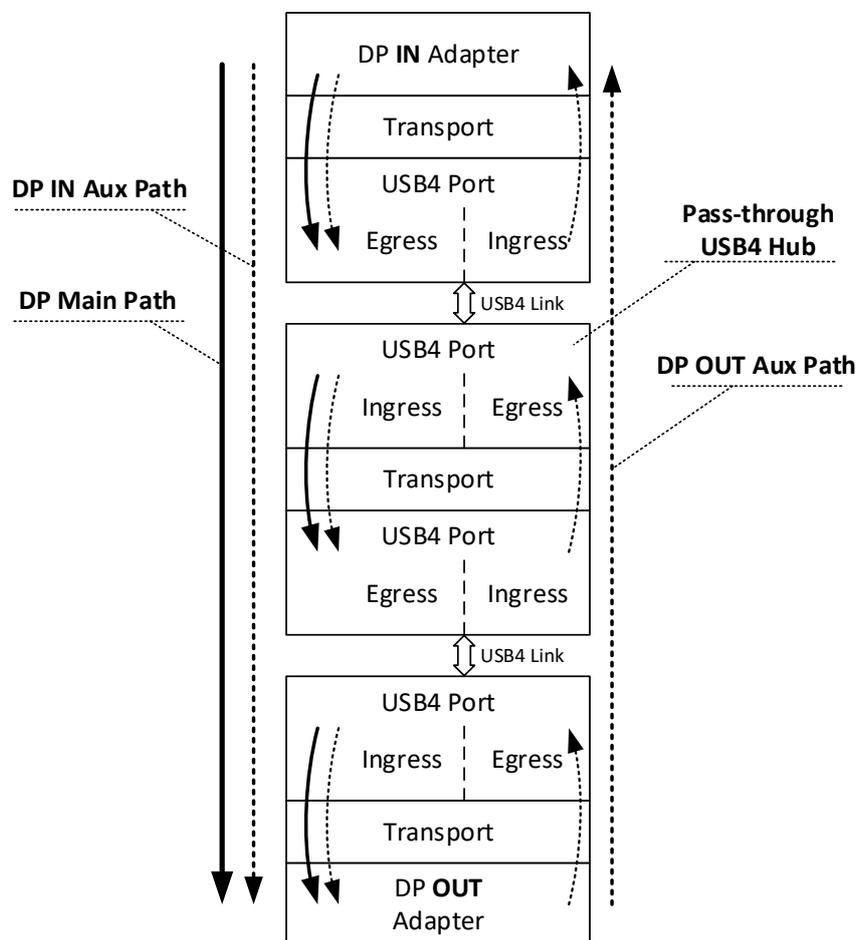
#### 5.4.3.3. Adapter Configuration

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 5-3 and is programmed to the Non-Flow Control field in the Adapter Configuration Space of the Lane Adapter.

#### 5.4.3.4. 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 5-3 shows an example of a DP Tunneling scheme with its basic constructs.

**Figure 5-3: 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 5-3 to the Path Configuration Space for 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 5-3 to the Path Configuration Space for 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 5-3 to the Path Configuration Space for the Ingress Lane Adapter.

## 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 5-3 to the Path Configuration Space for 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 5-3 to the Path Configuration Space for 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 5-3 to the Path Configuration Space of the Ingress Lane Adapter.

## 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 5-3 to the Path Configuration Space for 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 5-3 to the Path Configuration Space for 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 5-3 to the Path Configuration Space for the Lane Adapter.

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.
2. 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 OUT Adapter to enable the Paths.

**Table 5-3: 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>	2	1	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>	2	1	1	1	0	0
DP AUX: USB4 to DP Adapter		8		2	1	1	0	0	0

DP Main: DP IN Adapter to USB4	9		NA <sup>1</sup>	1	1	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>	1	1	0	0	0	0
DP Main: USB4 to DP OUT Adapter		9		1	1	0	0	0	0
1. A Connection Manager performs Read/Modify/Write to Path CS 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.									

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.

#### 5.4.3.5. 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 Path Setup and Teardown

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

#### 5.5.1. USB3 Path Setup

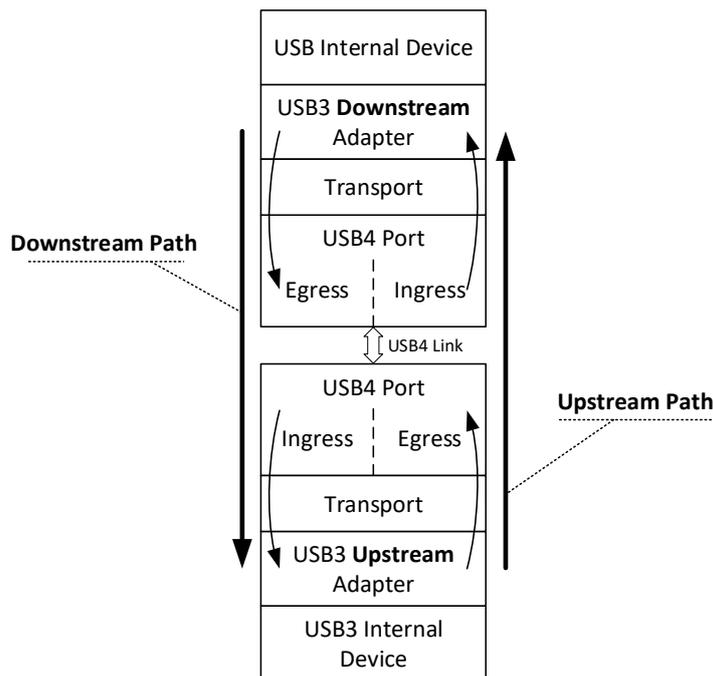
A USB3 Path is created between the USB3 Upstream Adapter and a USB3 Downstream Adapter of two physically connected Routers within the same Domain. A Connection

Manager does not set up a USB3 Path between two Routers unless they are directly connected to each other.

A Connection Manager can set up a Path for USB3 Tunneling when it first enumerates a Router that supports USB3 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 USB3 Tunneling. Figure 5-4 shows a USB3 Tunneling scheme with its basic constructs.

**Figure 5-4: USB3 Tunneling Setup**



A Connection Manager performs the following steps to establish a USB3 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 USB3 Downstream Adapter. The Path Attributes are defined in Table 5-4 (Path segment is from the USB3 Downstream Adapter to the USB4 Port).
  - Set Path Attributes in the Path Configuration Space of the Lane Adapter. The Path Attributes are defined in Table 5-4 (Path segment is from the USB4 Port to the USB3 Upstream 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 5-4 (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 5-4 (Path segment is from the USB4 Port to the USB3 Downstream Adapter).

4. Enable Transport Layer Packets on the USB3 Downstream Adapter by setting the *ADP\_USB3\_CS\_0.Path Enable* bit to 1b and the *ADP\_USB3\_CS\_0.Valid* bit to 1b in the USB3 Downstream Adapter. The *Valid* bit should not be set before the *Path Enable* bit is set.
5. Enable Transport Layer Packets on the USB3 Upstream Adapter by setting the *ADP\_USB3\_CS\_0.Path Enable* bit to 1b and the *ADP\_USB3\_CS\_0.Valid* bit to 1b in the USB3 Upstream Adapter. The *Valid* bit should not be set before the *Path Enable* bit is set.

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

Path Segment	Input HopID	Output HopID	Buffers Allocation: Dedicated Flow Control	Priority	Weight	IFC	EFC	ISE	ESE
USB3 Adapter to USB4 Port	8		NA	3	2	NA <sup>1</sup>	1	NA <sup>1</sup>	0
USB4 Port to USB3 Adapter		8	baMaxUSB3 from the Buffer Allocation Request Operation <sup>2</sup>	3	1	1	0	0	0
<ol style="list-style-type: none"> <li>1. A Connection Manager performs Read/Modify/Write to Path CS and does not change the <i>PATH_CS_0.Path Credits Allocated</i>, <i>PATH_CS_1.IFC Flag</i> and <i>PATH_CS_1.ISE Flag</i> fields at the USB3 Adapter.</li> <li>2. For more information regarding Buffer allocation, see Section 5.1.4.</li> </ol>									

### 5.5.2. USB3 Path Teardown

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

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

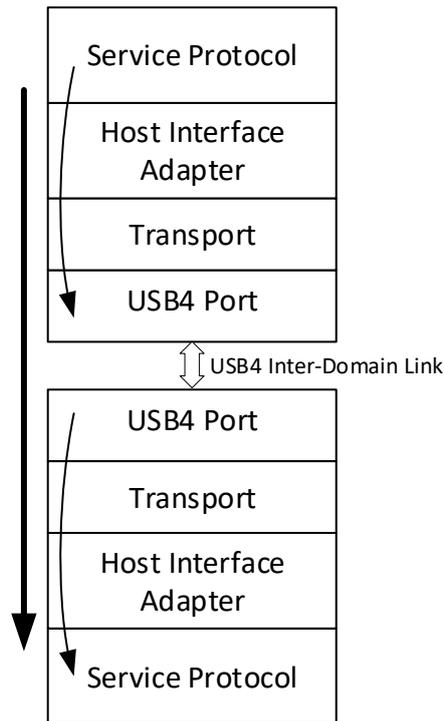
### 5.5.3. USB3 Adapter Disconnect

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

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

An Inter-Domain Tunnel is a set of Inter-Domain Paths that transport traffic between the USB4 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 Ethernet traffic is an example of a Service Protocol using the Inter-Domain Tunnel. Figure 5-5 shows an Inter-Domain Tunneling scheme between two hosts connected by an Inter-Domain Link.

**Figure 5-5: 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 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 supports Lane Bonding, then together with the peer Connection Manager set the Inter-Domain Link to operate in Lane Bonding mode.
5. To enable a Service Protocol, exchange protocol-specific Login Packets with the peer Connection Manager.
6. Configure and enable an inbound, an outbound Path, and the associated Descriptor Rings for the Service Protocol:
  - For an outbound Path, set the Path from a Transmit Descriptor Ring in the Host Router's Host Interface to the Inter-Domain Link, and enable the Transmit Descriptor Ring as defined in Section 5.6.1. The order of setting the Path and enabling the Descriptor Ring is vendor defined.
  - For an inbound Path, set the Path from the Inter-Domain Link to a Receive Descriptor Ring in the Host Router's Host Interface, and enable the Receive Descriptor Ring as defined in Section 5.6.2. The order of setting the Path and enabling the Descriptor Ring is vendor defined.

*Note: It is recommended that the Transmit Descriptor Ring is enabled before the Receive Descriptor Ring. Once the Receive Descriptor Ring is enabled, Credit Grant Packets can be transmitted to the peer Host Router.*

- Set the attributes for each Path:
  - Table 6-2 defines how to set the *Priority* and *Weight* fields.
  - 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.

#### **5.6.1. Transmit Ring Setup**

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.*

#### **5.6.2. Receive Ring Setup**

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.*

## 6. Bandwidth Management

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

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 the needs of the Tunneled Protocols with respect to USB4 bandwidth.

### 6.1. Bandwidth Allocation

#### 6.1.1. Bandwidth Calculations

This section describes how a Connection Manager calculates the maximum theoretical bandwidth and the available USB4 bandwidth.

### 6.1.1.1. Maximum 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 6-1 shows the raw bandwidth for each type of USB4 Link.

**Table 6-1: USB4 Raw Bandwidth**

USB4 Link	Raw Bandwidth (USB4)
Gen 2x1	10 Gbps
Gen 2x2	20 Gbps
Gen 3x1	20 Gbps
Gen 3x2	40 Gbps

*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 USB4 Link with the lowest raw bandwidth determines the maximum available USB4 bandwidth for the Path.

### 6.1.1.2. Available USB4 Bandwidth

A Connection Manager needs to calculate the available bandwidth at each Port before it sets up a Path. The Connection Manager uses Equation 1 to calculate the available downstream bandwidth for a Port. It uses Equation 2 to calculate the available upstream bandwidth for a Port.

#### Equation 1. Downstream Available Bandwidth

$$\text{Available BW} = [0.9 * (\text{Raw BW})] - \left[ \text{DP BW} + \text{USB3 BW} \left( \frac{\text{USB3 WRR} + \text{PCIe WRR}}{\text{USB3 WRR}} \right) \right]$$

where:

- Raw BW = the Raw USB4 Link Bandwidth for the Link
- DP BW = the sum of all DisplayPort Tunneling Main-Link Paths going through the Port in the downstream direction
- USB3 BW = value that the Connection Manager last wrote to the ADP\_USB3\_CS\_2.USB3 Allocated Downstream Bandwidth field in the USB3 Downstream Adapter of the Host Router
- USB3 WRR = weight assigned to the downstream USB3 Tunneling Path
- PCIe WRR = weight assigned to the downstream PCIe Tunneling Path. This value is 0 if PCIe Tunneling is not enabled

#### Equation 2. Upstream Available Bandwidth

$$\text{Available BW} = [0.9 * (\text{Raw BW})] - \left[ \text{DP BW} + \text{USB3 BW} \left( \frac{\text{USB3 WRR} + \text{PCIe WRR}}{\text{USB3 WRR}} \right) \right]$$

where:

- Raw BW = the Raw USB4 Link Bandwidth for the Link
- DP BW = the sum of all DisplayPort Tunneling Main-Link Paths going through the Port in the upstream direction

- USB3 BW = value that the Connection Manager last wrote to the USB3 ADP\_USB3\_CS\_2.Allocated Upstream Bandwidth field in the USB3 Downstream Adapter of the Host Router
- USB3 WRR = weight assigned to the upstream USB3 Tunneling Path
- PCIe WRR = weight assigned to the upstream PCIe Tunneling Path. This value is 0 if PCIe Tunneling is not enabled

### 6.1.2. Configuration Mechanism

A Router uses a 3-layer scheduling scheme to manage traffic. The Connection Manager uses this scheme to configure USB4 Link bandwidth. A Connection Manager sets the PATH\_CS\_1.Priority and PATH\_CS\_1.Weight fields to configure 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.

Table 6-2 shows the Priority and Weight values that are used for each type of Tunneled Protocol.

**Table 6-2: Tunneled Protocol Path Priority and Weights (Recommended)**

Path	Priority	Weight
DP Main-Link	1	1
DP AUX	2	1
USB3	3	2 (see Note 1)
PCIe	3	1 (see Note 1)
Host-to-Host	4	1
1. The 2:1 Ratio between USB3:PCIe Paths may vary.		

For a DP Tunneling Path, the Connection Manager can also configure the maximum Link Rate and Lane Count that the DP IN Adapter presents to DPTX.

For a USB3 Tunneling Path, the Connection Manager can control the bandwidth allocation for the Upstream and Downstream Paths using the Bandwidth Negotiation mechanism defined in the USB4 Specification.

### 6.1.3. Guaranteed Bandwidth

A Connection Manager can allocate guaranteed USB4 bandwidth to a Path. After it is allocated, guaranteed bandwidth can only be taken away if the Connection Manager can determine that the actual consumed bandwidth is less than the allocated bandwidth and there is a semaphore mechanism to prevent the Path from using additional USB4 bandwidth while it is being updated.

A Connection Manager allocates guaranteed bandwidth for a Path as follows:

1. Determine the maximum USB4 bandwidth that the priority group for the Path is allowed to consume.
  - Bandwidth can only be guaranteed for a priority group only if all higher priority groups have bounded bandwidth.
  - The maximum bandwidth that can be guaranteed for a priority group is calculated as follows:

$$Max\ BW = 0.9 * Raw\ BW - \sum Guarantees\ bandwidth\ for\ all\ higher\ priority\ groups$$

2. Set the priority group bandwidth so it does not exceed the maximum bandwidth.
3. Allocate a share of the priority group bandwidth to the Path:
  - Within the priority group, bandwidth is distributed by Weighted Round Robin (WRR).
  - The Weight of the Path determines what share of the bandwidth it gets.

#### **6.1.4. Bandwidth Management Events**

This section describes the events where a Connection Manager needs to manage the USB4 bandwidth. It also describes the high-level steps that a Connection Manager takes upon each event.

A Connection Manager must ensure that the *Available BW*, as defined in Equation 1, does not become negative as a result of creating a new Path.

##### **6.1.4.1. DisplayPort Path Setup**

When a new DP Tunneling Path is set up, the Connection Manager allocates bandwidth as follows:

- If (Minimum *Available USB4 BW* along the Path > Maximum supported DP BW):
  - Allocate maximum supported DP BW.
- Else:
  - Attempt to increase the Available USB4 BW by reducing the USB4 bandwidth allocated to USB3 Tunneling using the Bandwidth Negotiation method described in the USB4 Specification.
  - Configure the maximum Link Rate and Lane Count that the DP IN Adapter presents to DPTX so that the new DP Path does not require more USB4 bandwidth than the Available USB4 BW.
  - Return any unallocated USB4 BW to the USB3 Tunneling using the Bandwidth Negotiation method described in the USB4 Specification.

##### **6.1.4.2. DisplayPort Path Teardown**

After a DP Tunneling Path is torn down, the Connection Manager allocates bandwidth as follows:

- The USB4 bandwidth allocated to the DP Path is deallocated – The *Available USB4 BW* is increased by the amount that was allocated to that DP Path.
- Increase the USB4 bandwidth allocated to USB3 Tunneling using the Bandwidth Negotiation method described in the USB4 Specification.

##### **6.1.4.3. USB3 Tunneling Path Setup**

When a new USB3 Tunneling Path is set up, the Connection Manager allocates bandwidth as follows:

- When setting up a Path between the Host Router and a Device Router:
  - Read the *ADP\_CS\_USB3\_4.Maximum Supported Link Rate* field from the USB3 Downstream Adapter at the Host Router and the USB3 Upstream Adapter at the Device Router.

- Set the ADP\_USB3\_CS\_2.Allocated Upstream Bandwidth and ADP\_USB3\_CS\_2.Allocated Downstream Bandwidth fields to 90% of the lowest ADP\_CS\_USB3\_4.Maximum Supported Link Rate field.

*Note: After a USB3 link is established over a USB3 tunnel, the Allocated Upstream Bandwidth and Allocated Downstream Bandwidth can be decreased according to the ADP\_USB3\_CS\_4.Actual Link Rate field.*

- When setting up a Path between two Device Routers:
  - Verify that the ADP\_USB3\_CS\_2.Allocated Upstream Bandwidth and ADP\_USB3\_CS\_2.Allocated Downstream Bandwidth fields in the Host Router, will not violate the Available Upstream BW and Available Downstream BW equations for the new Path.  
  
For example, a Host Router is connected to USB4 Hub, and the Allocated Downstream Bandwidth is set to 15Gbps. If a new Device Router is connected to the Hub, establishing only Gen2x1 USB4 Link, then the Downstream Available Bandwidth on that USB4 Link will be smaller than the Allocated Downstream Bandwidth.
  - If Available USB4 BW can accommodate the Path, allocate USB4 bandwidth to the Path.
  - If Available USB4 BW cannot accommodate the new Path:
    - Attempt to increase the Available USB4 BW by reducing the USB4 bandwidth allocated to USB3 Tunneling using the Bandwidth Negotiation method described in the USB4 Specification:
      - If new Available USB4 BW can accommodate the Path, allocate USB4 bandwidth to the Path.
      - Else, the Connection Manager does not allocate USB4 bandwidth to the Path and does not enable the new Path.

#### 6.1.4.4. USB3 Tunneling Path Teardown

After a USB3 Tunneling Path is torn down, the Connection Manager allocates bandwidth as follows:

- If the Path is an upstream Path, the Allocated Upstream Bandwidth can be increased.
- If the Path is a downstream Path, the Allocated Downstream Bandwidth can be increased.

## 6.2. Bandwidth Clients

### 6.2.1. DisplayPort Tunneling

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

The DisplayPort Link 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.

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 Gbps
- The second DisplayPort Link is constrained by the Connection Manager, so it does not consume more than the remaining bandwidth:  
$$\text{HBR2x2Lanes} = 5.4\text{Gbps} * 2 * 0.8 = 8.64\text{ Gbps}$$

### **6.2.2. USB3 Tunneling**

USB3 Tunneling uses both Isochronous and Non-Isochronous bandwidth, depending on the connected USB 3.2 endpoints.

USB3 Tunneling bandwidth is dynamic – free USB4 bandwidth is allocated to USB3 Tunneling as it becomes available. USB3 Tunneling is allocated a bigger share of the USB4 bandwidth than PCIe Tunneling.

The Connection Manager and the internal host controller share a bandwidth negotiation mechanism, which is used to manage the allocation of the Isochronous bandwidth. The Connection Manager can allocate separate bandwidth for the downstream and upstream directions. The host controller dynamically reports the consumed bandwidth, allowing the Connection Manager to reclaim unused bandwidth when needed. See Chapter 9 of the USB4 Specification for more details.

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

### **6.2.3. PCIe Tunneling**

PCIe Tunneling is treated as Non-Isochronous and is not guaranteed bandwidth.

The Connection Manager does not directly allocate any bandwidth to a PCIe Path. Instead, the PCIe Path uses the bandwidth that is left over after the DP Tunneling and USB3 Tunneling Paths are configured. A PCIe Path and a USB3 Path are set to the same Priority group. Therefore, if the Connection Manager allocates bandwidth to the USB3 Path, it indirectly affects bandwidth allocation to the PCIe Path. The WRR ratio between the PCIe Path and the USB3 Path determine how much bandwidth is allocated to the PCIe Path.

### **6.2.4. 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. It is not guaranteed any bandwidth.

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 after USB4 Link Initialization:

- **Two Single-Lane Links** – Both Lane Adapters are in CL0 state. In this case, the Connection Manager can either bond the Lanes into one Dual-Lane Link (see Section 7.1) or it can disable the Lane 1 Adapter in order to operate with one single-Lane Link (see Section 7.4). The Connection Manager does not configure any Paths over the Port while the Port has two single-Lane Links.
- **One Single-Lane 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).

### 7.1. Lane Bonding

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*. If 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:

1. On both sides of the Link, read the `LANE_ADP_CS_1.Adapter State` field 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 of all Lane Adapters to x2.
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 with the *UPG* bit set to 1b for the Lane 1.
  - Polling the `LANE_ADP_CS_1.Negotiated Link Width` field in the Lane 0 Adapter until it indicates a USB4 Link width of x2.

When Lane Bonding is successful, a Dual-Lane 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. 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: 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 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, the Lanes of a USB4 Port will not be bonded. The Connection Manager needs to re-initiate Lane Bonding in order for the Link to operate as a Dual-Lane Link.*

## 7.3. 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.

### 7.3.1. 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	3

### 7.3.2. 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.*

### 7.3.3. TMU Configuration

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

- 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:
  1. Determine the accuracy level required by each Router in the Domain as described in 7.3.2.
  2. Determine which Router has the maximum accuracy level requirements.
  3. Select the TMU Mode that provides the maximum accuracy level.

The Connection Manager sets the TMU Mode by writing to the following fields:

- `TMU_RTR_CS_3.TSPacketInterval`
- `TMU_ADP_CS_3.EnableUniDirectionalMode`

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

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

TMU Mode	TSPacketInterval	EnableUniDirectionalMode
Off	0	0
LowRes	1000	1
HiFi-Uni	16	1
HiFi-Bi	16	0

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

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

TMU Mode	Freq Measurement Window	ErrorAvgConst, OffsetAvgConst, DelayAvgConst, FreqAvgConst
Off	NA	NA
LowRes	30	4
HiFi-Uni	800	8
HiFi-Bi	800	8

**7.3.3.1. 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 until either Lane Bonding completes (for a Dual-Lane Link) or Lane 1 is disabled (for a Single-Lane Link).

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 7.3.3.4.
3. Set the TMU Mode in the Router using the *TMU Enable* routine defined in Section 7.3.3.4.
4. Wait 50ms to allow the TMU to converge before Path Setup.

### 7.3.3.2. 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.

### 7.3.3.3. 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.

### 7.3.3.4. 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_ADP_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).
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 7.3.3.4.1.
  - b. If changing the TMU Mode to Off follow the TMU Disable routine in Section 7.3.3.4.2.
  - c. Else, follow the TMU Mode Switch routine in Section 7.3.3.4.3
5. Wait 50ms.
6. 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).
7. Set up the DP Paths that were torn down in step 1.

#### 7.3.3.4.1. 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 7-2 in the Downstream Router
2. 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 7-3 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.
3. Else:
  - a. Set the *TMU\_RTR\_CS\_3.TSPacketInterval* field according to Table 7-3 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 DPF.

#### 7.3.3.4.2. TMU Disable

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

1. If current TMU Mode is 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.
2. 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.

#### 7.3.3.4.3. TMU Mode Switch

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\_6.EnableUniDirectionalMode* bit in the DFP to indicate the new TMU Mode (0b for Bi-Directional, 1b for Uni-Directional).
2. Set the *TMU\_RTR\_CS\_3.TSPacketInterval* 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 in the UFP to indicate the new TMU Mode (0b for Bi-Directional, 1b for Uni-Directional).
  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.

#### 7.4. Lane Disable and Enable

The Connection Manager uses the *LANE\_ADP\_CS\_1.Lane Disable* bit to enable and disable a 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.

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

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 7.4.1 describes how a Connection Manager disables and enables a Lane on a Link that is not an inter-Domain Link. Section 7.4.2 describes how a Connection Manager disables and enables a Lane on an inter-Domain Link.

##### 7.4.1. Intra-Domain Links

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

1. If Port has a Dual-Lane Link, initiate a Downstream Facing Port Reset to transition to two Single-Lane Links and wait for Lane Initialization to complete.
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 this will cause the Router to send 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 will restart Lane Initialization on both Lane 1 and Lane 0.*

#### **7.4.2. Inter-Domain Links**

A Connection Manager disables Lane 1 in an Inter-Domain Link if either of the following are true:

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

The Connection Manager sets the `LANE_ADP_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.

#### **7.5. Link Errors and Notifications**

A Connection Manager can optionally configure a Router to send a Notification Packet for certain Link errors. Section 7.5.1 describes how a Connection Manager enables Notification Packets. Section 7.5.2 describes how a Connection Manager acknowledges a Notification Packet from a Router.

### 7.5.1. Notification Enable

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

1. Set the desired bits to 1b in the Adapter Configuration Space:
  - a. Set the corresponding bits to 1b in the *LANE\_ADP\_CS\_2.Logical Layer Errors Enable* field to enable Notification Packets for specific Link Layer errors.
  - 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.
2. 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.

### 7.5.2. Notification Acknowledgement

A Connection Manager needs to send a Notification Acknowledgment Packet 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

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

## 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.
- 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.
- 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.
- 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.
- 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 an inter-Domain Link.
- Do not enable any CLx states on a Link that uses an Active Cable that does not support CLx states.
- Enable CLx by setting the relevant bits in Lane 0 Adapter Space.

*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.*

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.

### 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 Port is Configured 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 Port is Configured 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 Port is Configured 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 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, 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. 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:
  - a. The Connection Manager continues to the next step when the *Sleep Ready* bits in all Routers in the Domain are 1b.
  - b. If, after 80ms, not all *Sleep Ready* bits are set to 1b, the Connection Manager stops polling and continues to the next step.
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 tells 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.
<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 Sections 0. It also need to perform Path setup as described in 5 to restore any Paths on the Domain that were present before the Domain entered sleep state.

### 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.

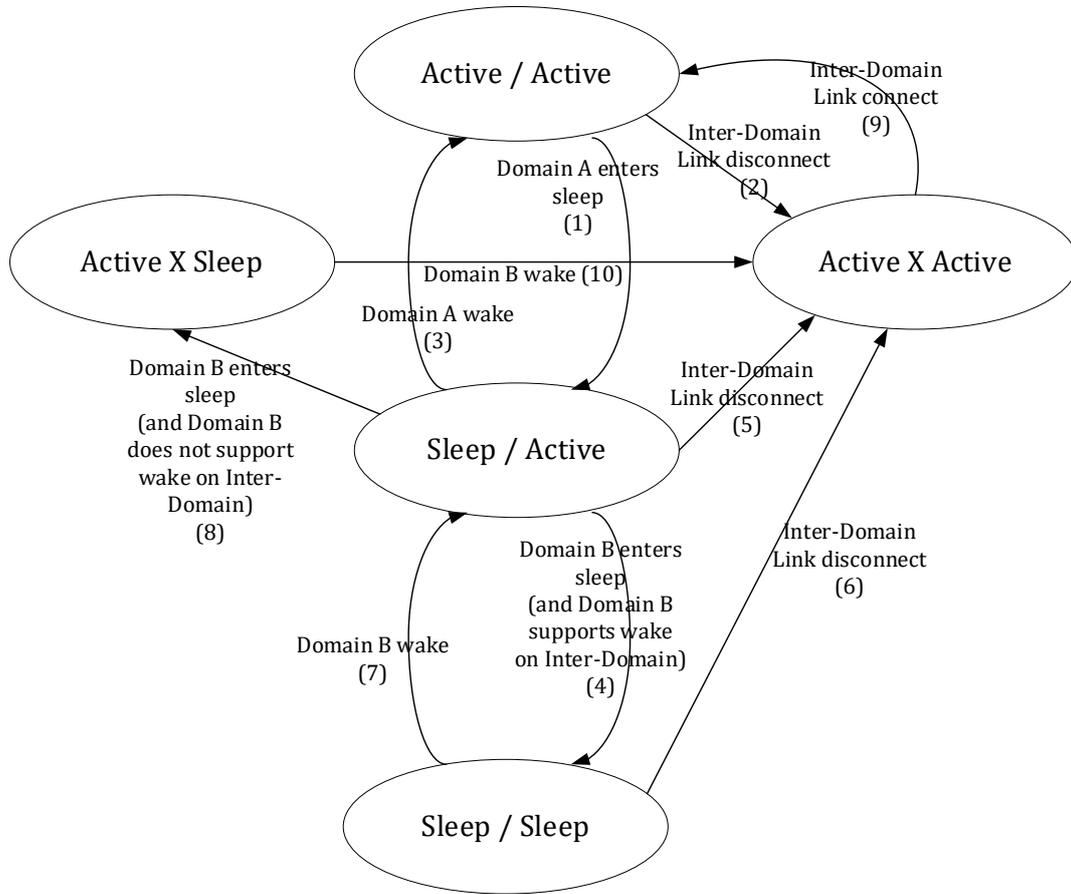
For convenience sake, the neighbor Domain is named Domain A in this section, while the Domain (and Connection Manager) whose behavior is documented is named Domain B in this section.

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>1</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

<sup>1</sup> In the USB4 Port Capability of the Router in Domain B facing the Inter-Domain Link

4	B sends an LT_LRoff transaction to A. A ignores the Transaction.	Connection Manager B brings Domain B into Sleep state.	N/A
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.

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.
2. If the Operation requires Metadata, write the metadata to ROUTER\_CS\_25.
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 will set the *Operation Valid* bit to 0b after it finishes processing the Operation. 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 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.

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 verifies 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.

When the *Pending* bit in SB Register Space is 1b, the Connection Manager does not write to registers `PORT_CS_27` through `PORST_CS_17` in Port Configuration Space. It also does not write to SB Register 18.