

Universal Serial Bus 4 (USB4™) Specification

Apple Inc.
HP Inc.
Intel Corporation
Microsoft Corporation
Renesas Corporation
STMicroelectronics
Texas Instruments

Version 1.0 with Errata and ECN through May 19, 2021

May 2021

Release History

| Version | Comments | Issue Date |
|--|---|--------------|
| 1.0 | First release | August 2019 |
| 1.0 with Errata and ECN through May 4, 2020 | Includes errata and ECN through May 4, 2020 as part of the specification text. | June 2020 |
| 1.0 with Errata and ECN through October 15, 2020 | Includes errata and ECN through October 15, 2020 as part of the specification text. | October 2020 |
| 1.0 with Errata and ECN through May 19, 2021 | Includes errata and ECN through May 19, 2021 as part of the specification text. | May 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.

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 for further information

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

Acknowledgement of Technical Contribution

The authors of this specification would like to recognize the following people who participated in the USB4 Specification technical work group.

Apple Inc. – Promoter Company Employees

| | | | |
|-------------------|-----------------|------------------|--------------------|
| Majd Abu Tayeh | Nimrod Agmon | Lior Aloni | Brian Baek |
| Omer Bar-Lev | Moshe Benyamini | Gopu Bhaskar | Carlos Calderon |
| David Conroy | Bill Cornelius | Scott Deandrea | William Ferry |
| Amit Flanter | Itay Franko | Radia Gantous | Alex Gerber |
| Mark Goikhman | Nir Guetta | Yair Hershkovitz | Scott Jackson |
| Husam Khashiboun | Alan Kobayashi | Alexei Kosut | Christine Krause |
| Dmitri Krichevsky | Alex Lozovik | Rachel Menes | Shlomi Mor |
| Shlomi Museri | Mona Omari | Tal Ostro | Arie Peled |
| Collin Pieper | Idan Reller | Reese Schreiber | Oleg Schtofenmaher |
| Etan Shirron | Ori Moshe Stern | Jose Tierno | Maxim Tsudik |
| Anatoly Uskach | Jeff Wilcox | Dan Wilson | Ohad Zalcmán |

HP Inc. – Promoter Company Employees

| | | | |
|---------------|---------------|---------------|---------------|
| Roger Benson | Marcus Benzel | Alan Berkema | Kenneth Chan |
| Frank Chen | Phil Chen | Hosup Chung | Glen Dower |
| Mark Lessman | Nam Nguyen | Roger Pearson | Kenneth Smith |
| Chris Tabarez | | | |

Intel Corporation – Promoter Company Employees

| | | | |
|-------------------|------------------------------|---------------------|----------------------|
| Nausheen Ansari | Noam Arzy | Alexandre Audier | Binata Bhattacharyya |
| Huimin Chen | Hengju Cheng | Salauddin Choudhury | John Crouter |
| Maxim Dan | Jhuda Dayan | Yoni Dishon | Eran Galil |
| Saranya Gopal | Venkataramani Gopalakrishnan | Raul Gutierrez | Michael Gouzenfeld |
| Mickey Gutman | Benjamin Hacker | Yaniv Hayat | Uri Hermoni |
| Alon Horn | Abdul Ismail | Abhilash K V | Ziv Kabiry |
| Vijaykumar Kadgi | Vijay Kasturi | Sergey Khaykin | Lev Kolomiets |
| Vladislav Kopzon | Efraim Kugman | Edmond Lau | Uriel Lemberger |
| Yun Ling | Guobin Liu | Balaji Manoharan | Liran Manor |
| Uma Medepalli | Assaf Mevorach | Hezi Naaman | Ohad Navon |
| Naod Negussie | CheeLim Nge | Leonid Plaks | Duane Quiet |
| Rajaram Regupathy | Reuven Rozic | Oren Salomon | Zeeshan Sarwar |

| | | | |
|-------------------|-------------------|-------------------|-----------------|
| Brad Saunders | Leonid Shaposhnik | Ehud Shoor | Ari Sharon |
| Uri Soloveychik | Einat Surijan | Aviel Uzan | Karthi Vadivelu |
| Alex Vekker | Chen Vrubel | Stephanie Wallick | Tzewen Wang |
| Sarel Wechsler | Ady Weiss | Vitaly Zhivov | Gal Yedidia |
| Vladimir Yudovich | Aruni Nelson | | |

Microsoft Corporation – Promoter Company Employees

| | | | |
|---------------|----------------|----------------|-----------------|
| Randy Aull | Jim Belesiu | Martin Borve | Anthony Chen |
| Jesse Chen | Matt Chung | Aacer Daken | Rajib Dutta |
| Mark Friend | Philip Froese | David Hargrove | Robbie Harris |
| Kit Hui | Toby Nixon | Rahul Ramadas | Andrea Severson |
| Kiran Shastry | Nathan Sherman | Ji Sun | Shyamal Varma |

Renesas Corporation – Promoter Company Employees

| | | | |
|-------------|----------------|-------------------|------------------|
| Tam Do | Robert Dunstan | Philip Leung | Kiichi Muto |
| Ziba Nami | Hajime Nozaki | Raman Sargis | Yoshiyuki Tomoda |
| Starry Tsai | Jia Wei | Toshifumi Yamaoka | |

STMicroelectronics – Promoter Company Employees

| | | | |
|-----------------|-------------|------------|--|
| Nathalie Ballot | Joel Huloux | Gerard Mas | |
|-----------------|-------------|------------|--|

Texas Instruments – Promoter Company Employees

| | | | |
|-----------------------|-----------------|-----------------|-------------------|
| Mike Campbell | Anant Gole | Craig Greenburg | Michael Koltun IV |
| Sai Karthik Rajaraman | Anwar Sadat | Cory Stewart | Sue Vining |
| Deric Waters | Gregory Watkins | | |

Contributor Company Employees

| | | | |
|---------------------------------------|---------------|--------------|----------------|
| ACON, Advanced-Connectek, Inc. | Victory Chen | Conrad Choy | Vicky Chuang |
| | Jessica Feng | Sharon Hsiao | Wayne Wang |
| Advanced Micro Devices | Dennis Au | Nat Barbiero | Jason Chang |
| | Michael Comai | Walter Fry | Will Harris |
| | Jason Hawken | Jim Hunkins | Ling Kong |
| | Scott Ogle | Victor Salim | Joseph Scanlon |
| | Peter Teng | | |
| Allion Labs, Inc. | Howard Chang | Casper Lee | Brian Shih |
| Analogix Semiconductor, Inc. | Greg Stewart | Haijian Sui | Yueke Tang |
| | Ning Zhu | | |

| | | | |
|---|--------------------|----------------------|----------------------|
| Anritsu Corporation | Wataru Aoba | John Jerico Custodio | Kazuhiro Fujinuma |
| | Hiroshi Goto | Alessandro Messina | Tadanori Nishikobara |
| | Ryo Sunayama | Toshihiro Suzuki | Mitsuhiro Usuba |
| | Takeshi Wada | | |
| ASMedia Technology Inc. | Chang Chinyu | Chang Weiyun | Chen Chiahsin |
| | Chen Chunhung | Chuang Weber | Kuo Han Sung |
| | Lin Curtis | Lin ShuYu | Tseng PS |
| | Tseng YD | Wei Daniel | Wu ShengChung |
| Avery Design Systems, Inc. | Chris Browy | Chilai Huang | Zhihong Zeng |
| BitifEye Digital Test Solutions GmbH | Sebastian Muschala | Hermann Stehling | |
| Bizlink Technology, Inc. | Alex Chou | Morphy Hsieh | Kevin Tsai |
| Cadence Design Systems, Inc. | Marcin Behrendt | Jacek Duda | Shikha Gupta |
| | Gaurav Jain | Poonam Khatri | Yash Kothari |
| | Vinod Lakshman | Shivaji Magadum | Andy Mauffet-Smith |
| | Rohit Mishra | Uyen Nguyen | Raja Pounraj |
| | Thirumal Reddy | Anand RK | Anshul Shah |
| | Neelabh Singh | Ofer Steinberg | Mark Summers |
| | Claire Ying | Wasiq Zia | |
| Corigine, Inc. | Kevin Fan | Ali Khan | Xiao Xiao |
| Corning Optical Communications LLC | Mark Bradley | Wojciech Giziewicz | Ian McKay |
| | Jamie Silva | | |
| Cypress Semiconductor | Mark Fu | Naman Jain | Savan Javia |
| | Palani Subbiah | | |
| Dell Inc. | Mohammed Hijazi | Tom Lanzoni | Ken Nicholas |
| | Marcin Nowak | Scott Ogle | Adie Tan |
| | Lee Zaretsky | | |
| Diodes Incorporated | Qun Song | | |
| DisplayLink (UK) Ltd. | Pete Burgers | Dan Ellis | |
| DJI Technology Co., Ltd. | Steve Huang | | |
| Electronics Testing Center, Taiwan | Sophia Liu | | |

| | | | |
|---|---|---|--|
| Elka International Ltd. | Alvin Cheng Jui-Ming Yang | Chloe Hsieh | Roy Ting |
| Ellisys | Abel Astley Tim Wei | Mario Pasquali | Chuck Trefts |
| Etron Technology, Inc. | Andy Chen Chien-Cheng Kuo | Shihmin Hsu Jen Hong Larn | Bryan Huang |
| Foxconn / Hon Hai | Patrick Casher Fred Fons Christine Tran | Joe Chen Bob Hall A.J. Yang | Jason Chou Terry Little Jie Zheng |
| Fresco Logic Inc. | Tim Barilovits Jie Ni | Bob McVay Jeffrey Yang | Christopher Meyers |
| Genesys Logic, Inc. | Sean Chen Jerry Hu ChunYen Kuo Miller Lin Han Wu | Gerry Chou Perlman Hu Weddell Lee D.C. Lu Yihsun Wu | Thomas Hsieh Roy Huang Jimmy Lin Greg Tu |
| Google Inc. | Mark Hayter David Schneider | Benson Leung | Raj Mojumder |
| Granite River Labs | Nikhil Acharya Allen Chen Alan Chuang Tim Lin Rajaraman V | Yun Han Ang Cyan Chen Steven Lee Krishna Murthy Chin Hun Yaep | Sandy Chang Swee Guan Chua Caspar Lin Johnson Tan |
| Hotron Precision Electronic Ind. Corp. | Rosa Chen | Patrick Yeh | YF Zhang |
| I-PEX (Dai-ichi Seiko) | Alan Kinningham | Ro Richard | |
| Japan Aviation Electronics Industry Ltd. | Mark Saubert | Junichi Takeuchi | |
| JMicron Technology Corp. | Charon Chen | Mika Cheng | Kevin Liu |
| Kandou Bus SA | Brian Holden Andrew Stewart Atsushi Imaoka | Hitaish Sharma Mark Vennebarger Biing Lin Lem | David Stauffer Jit Lim |

| | | | |
|--|-----------------------|----------------------|-------------------|
| Keysight Technologies Inc. | Francis Liu | Roland Scherzinger | |
| L&T Technology Services | Binu Chinna Thankam | Sunil Kumar | Siddharth Pethe |
| | Badrinath Ramachandra | Arunkumar Selvam | Gayathri SN |
| LeCroy Corporation | Alan Blankman | Patrick Connally | Carl Damn |
| | David Fraticelli | Daniel H Jacobs | Farnoosh Jafary |
| | Tyler Joe | Carlo Mazzetti | Mike Micheletti |
| | Kathryn Morales | Jeff Sabuda | Joseph Schachner |
| | Chris Webb | | |
| Lenovo | Toshikazu Horino | Shinji Matsushima | Yuuki Matsuura |
| | Nozomu Nagata | Munefumi Nakata | Kazuya Shibayama |
| | Shunki Sugai | Chikara Takahashi | Masahiro Tokuno |
| | Kayanagi Tsuneo | | |
| LG Electronics Inc. | Do Kyun Kim | Yoon Jong Lee | Seung Yoo |
| Lintes Technology Co., Ltd. | Tammy Huang | Charles Kaun | RD Lintes |
| | Max Lo | CT Pien | Jin Yi Tu |
| | Jason Yang | | |
| Lotes Co., Ltd. | Regina Liu-Hwang | John Lynch | |
| Luxshare-ICT | Josue Casillo | CY Hsu | Antony Lin |
| | John Lin | Stone Lin | Scott Shuey |
| | Eric Wen | Pat Young | |
| Maxio Technology (Hangzhou) Ltd. | George Fang | | |
| MediaTek Inc. | Henry Chen | Alexyc Lin | Pochou Lin |
| | Chiachun Wang | | |
| MegaChips Corporation | Rahul Agarwal | Ramesh Dandapani | Satoru Kumashiro |
| | Ryuichi Mariizumi | Sireesha Vemulapalli | Nobu Yanagisawa |
| Mercedes-Benz Research & Development, North America, Inc. | Hans Wickler | | |
| Microchip Technology Inc. | Mark Bohm | Atish Ghosh | Fernando Gonzalez |
| | Mark Gordon | Richard Petrie | Brigham Steele |
| | Anthony Tarascio | Robert Zakowicz | |
| Molex LLC | Alan MacDougall | | |

| | | | |
|--|--------------------|-------------------|---------------------|
| MQP Electronics Ltd. | Sten Carlsen | Pat Crowe | |
| Newnex Technology Corp. | Sam Liu | | |
| NVIDIA | Jamie Aitken | Mark Overby | |
| NXP Semiconductors | Mahmoud El Sabbagh | Ken Jaramillo | Abhijeet Kulkarni |
| | Vijendra Kuroodi | Krishnan TN | |
| Oculus VR LLC | Marty Evans | Joaquin Fierro | Chao Hu |
| ON Semiconductor | Eduardo De Reza | Oscar Freitas | Christian Klein |
| | Amir Lahooti | | |
| Parade Technologies, Inc. | Jian Chen | Jimmy Chiu | Mark Qu |
| | Craig Wiley | Paul Xu | Kevin Yuan |
| | Alan Yuen | | |
| Phison Electronics Corp. | Jimmy Chen | Ko Hong Lipp | Sebastien Jean |
| | Stark Kuan | Thomas Lee | Anton Lin |
| | Winnie Lu | Wei Sui-Ning | James Tsai |
| | Michael Wu | Fu-Hua Yang | Chang Yuan-Cheng |
| Qualcomm, Inc | Tomer Ben Chen | Yiftach Benjamini | Richard Burrows |
| | Amit Gil | James Goel | Philip Hardy |
| | Raja Jagadeesan | Lalan Mishra | Dmitrii Vasilchenko |
| | Chris Wiesner | | |
| Realtek Semiconductor Corp. | Chung-Chun Chen | Jen Wen Chen | Shen Chen |
| | Jonathan Chou | Chang Ding | Yao Feng |
| | Bokai Huang | An-Ming Lee | Ray Lee |
| | Ryan Lin | Terry Lin | Luobin Wang |
| | Kay Yin | Chris Zeng | |
| Rohde & Schwarz GmbH & Co. KG | Johannes Ganzert | Randy White | |
| Samsung Electronics Co., Ltd. | Jaedeok Cha | KangSeok Cho | CheolYoon Chung |
| | Sangju Kim | Termi Kwon | Cheolho Lee |
| | Edward Lee | Jun Bum Lee | Chahoon Park |
| | Sunggeun Yoon | | |
| Seagate Technology LLC | Alvin Cox | Paul McParland | Michael Morgan |
| | Cuong Tran | | |

| | | | |
|--|--------------------|------------------------|--------------------------|
| Silicon Line GmbH | Ian Jackson | | |
| SiliConch Systems Private Limited | Kaustubh Kumar | Rakesh Polasa | Satish Anand Verkila |
| Softnautics LLP | Bhavesh Desai | Hetal Jariwala | Dipakkumar Modi |
| | Ishita Shah | Ujjwal Talati | |
| Spectra7 Microsystems Corp. | Alex Chow | James McGrath | |
| Specwerkz | Sydney Fernandes | Amanda Hosler | Diane Lenox |
| | Soren Petersen | | |
| STMicroelectronics | Nathalie Ballot | Joel Huloux | Gerard Mas |
| Sumitomo Electric Ind., Ltd., Optical Comm. R&D Lab | Sainer Siagian | Mitsuaki Tamura | |
| Synaptics Inc. | Jeff Lukanc | Mark Miller | Prashant Shamarao |
| Synopsys, Inc. | Prishkrit Abrol | Subramaniam Aravindhan | Jeanne Cai |
| | Jun Cao | Morten Christiansen | Scott Guo |
| | Eric Huang | Joseph Juan | Venkataraghavan Krishnan |
| | Jitendra Kushwaha | Behram Minwalla | Saleem Mohammad |
| | Rick Schmidt | Jasjeet Singh | Mahendra Singh |
| | John Stonick | Zongyao Wen | Fred Yu |
| Tektronix, Inc. | Madhusudan Acharya | Sourabh Das | Keyur Diwan |
| | Mark Guenther | Abhijeet Shinde | Gary Simontom |
| Thine Electronics, Inc. | Shuhei Yamamoto | | |
| Tyco Electronics Corp., a TE Connectivity Ltd. company | Simon Li | Jeff Mason | Jacky Mo |
| | Tommy Yu | Yuanbo Zhang | Tony Zhu |
| Varjo Technologies | Kai Inha | | |
| VIA Labs, Inc. | Wayne Tseng | | |
| VIA Technologies, Inc. | Benjamin Pan | Terrance Shih | Jay Tseng |
| | Fong-Jim Wang | | |
| Weltrend Semiconductor | Chao-Chee Ku | Jeng Cheng Liu | Wayne Lo |
| | Ho Wen Tsai | Eric Wu | Randolph Wu |
| | Simon Yeh | | |
| Western Digital | David Landsman | Larry McMillan | Rob Ryan |

Version 1.0
May 2021

- xi -

Universal Serial Bus 4
Specification

Wilder Technologies

Steve Bright

Zach Moore

Joe O'Brien

Majid Shayegh

Contents

| | | |
|---------|--|----|
| 1 | Introduction..... | 1 |
| 1.1 | Scope of the Document..... | 1 |
| 1.2 | USB Product Compliance..... | 1 |
| 1.3 | Document Organization..... | 1 |
| 1.4 | Design Goals..... | 1 |
| 1.5 | Related Documents..... | 1 |
| 1.6 | Conventions..... | 2 |
| 1.6.1 | Precedence..... | 2 |
| 1.6.2 | Keywords..... | 2 |
| 1.6.2.1 | Informative..... | 2 |
| 1.6.2.2 | May..... | 2 |
| 1.6.2.3 | N/A..... | 2 |
| 1.6.2.4 | Normative..... | 2 |
| 1.6.2.5 | Optional..... | 2 |
| 1.6.2.6 | Reserved..... | 2 |
| 1.6.2.7 | Shall..... | 2 |
| 1.6.2.8 | Should..... | 3 |
| 1.6.3 | Capitalization..... | 3 |
| 1.6.4 | Italic Text..... | 3 |
| 1.6.5 | Numbering..... | 3 |
| 1.6.6 | Bit, Byte, DW, and Symbol Conventions..... | 3 |
| 1.6.7 | Implementation Notes..... | 3 |
| 1.6.8 | Connection Manager Notes..... | 3 |
| 1.6.9 | Pseudocode..... | 3 |
| 1.6.10 | CRC Algorithms..... | 4 |
| 1.6.11 | FourCC..... | 4 |
| 1.7 | Reserved Values and Fields..... | 4 |
| 1.8 | Terms and Abbreviations..... | 5 |
| 2 | Architectural Overview..... | 11 |
| 2.1 | USB4 System Description..... | 11 |
| 2.1.1 | Architectural Constructs..... | 13 |
| 2.1.1.1 | Routers..... | 13 |
| 2.1.1.2 | Adapters..... | 13 |
| 2.1.1.3 | USB4 Ports and Links..... | 13 |
| 2.1.1.4 | USB4 Devices..... | 14 |
| 2.1.1.5 | USB4 Host..... | 15 |
| 2.1.1.6 | Re-timers..... | 16 |
| 2.1.1.7 | Connection Manager..... | 16 |
| 2.1.2 | USB4 Mechanical..... | 16 |
| 2.1.3 | USB4 Power..... | 16 |
| 2.1.4 | USB4 System Configuration..... | 16 |

| | | |
|----------|---|----|
| 2.1.5 | Thunderbolt™ 3 (TBT3) Compatibility Support | 16 |
| 2.1.6 | USB Type-C Alternate Mode Compatibility Support..... | 16 |
| 2.2 | USB4 Fabric Architecture | 16 |
| 2.2.1 | USB4 Functional Stack..... | 17 |
| 2.2.1.1 | Electrical Layer..... | 17 |
| 2.2.1.2 | Logical Layer | 18 |
| 2.2.1.3 | Transport Layer | 18 |
| 2.2.1.4 | Configuration Layer | 18 |
| 2.2.1.5 | Protocol Adapter Layer | 18 |
| 2.2.2 | USB4 Fabric Topology | 18 |
| 2.2.3 | Paths | 20 |
| 2.2.4 | Communication Constructs..... | 21 |
| 2.2.4.1 | USB4 Link..... | 21 |
| 2.2.4.2 | Sideband Channel..... | 22 |
| 2.2.5 | USB4 Host-to-Host Communications | 23 |
| 2.2.6 | Programming Model..... | 23 |
| 2.2.6.1 | Connection Manager | 23 |
| 2.2.6.2 | Configuration Spaces | 24 |
| 2.2.6.3 | Operations..... | 24 |
| 2.2.7 | Time Synchronization..... | 24 |
| 2.2.8 | USB4 Fabric Data Integrity | 24 |
| 2.2.9 | Global Life of a Router | 25 |
| 2.2.10 | Protocol Tunneling..... | 25 |
| 2.2.10.1 | USB3 Tunneling..... | 26 |
| 2.2.10.2 | Display Tunneling | 30 |
| 2.2.10.3 | PCIe Tunneling | 33 |
| 2.2.10.4 | Host Interface Adapter..... | 37 |
| 3 | Electrical Layer..... | 40 |
| 3.1 | Sideband Channel Electrical Specifications | 40 |
| 3.2 | USB4 Ecosystem..... | 41 |
| 3.2.1 | Insertion-Loss Considerations (Informative) | 41 |
| 3.2.2 | Coded Bit-Error-Ratio Considerations (Informative)..... | 42 |
| 3.3 | USB4 Electrical Compliance Methodology | 42 |
| 3.3.1 | System Compliance Test Point Definitions | 42 |
| 3.3.2 | AC Coupling Capacitors | 43 |
| 3.3.3 | Reference Clock-and-Data-Recovery (CDR) Function | 44 |
| 3.3.4 | Reference Equalization Function | 44 |
| 3.3.4.1 | Reference CTLE..... | 45 |
| 3.3.4.2 | Reference DFE | 47 |
| 3.3.5 | Time Domain Measurements..... | 47 |
| 3.3.6 | Compliance Boards..... | 47 |
| 3.3.6.1 | Compliance Plug Test Board | 47 |
| 3.3.6.2 | Compliance Receptacle Test Board..... | 47 |

| | | |
|---------|---|----|
| 3.4 | Router Assembly Transmitter Compliance | 47 |
| 3.4.1 | Transmitter Specifications Applied for All Speeds | 47 |
| 3.4.1.1 | Transmitter Frequency Variations during Link Training | 49 |
| 3.4.1.2 | Transmitter Differential Return Loss | 51 |
| 3.4.1.3 | Transmitter Common Mode Return Loss | 52 |
| 3.4.1.4 | Transmit Equalization | 52 |
| 3.4.2 | Transmitter Compliance Specifications for Gen 2 | 55 |
| 3.4.3 | Transmitter Compliance Specifications for Gen 3 Interconnects | 57 |
| 3.5 | Router Assembly Receiver Compliance | 58 |
| 3.5.1 | Receiver Specifications Applied for All Speeds | 58 |
| 3.5.1.1 | Receiver Differential Return Loss | 59 |
| 3.5.1.2 | Receiver Common Mode Return Loss | 60 |
| 3.5.2 | Receiver Uncoded BER Tolerance Testing | 61 |
| 3.5.3 | Receiver Multi Error-Bursts Testing | 63 |
| 3.6 | Captive Device Compliance | 64 |
| 3.6.1 | Captive Device Compliance Test Setup | 64 |
| 3.6.2 | Captive Device Transmitter Specifications | 65 |
| 3.6.2.1 | Conducted Energy in Wireless Bands | 65 |
| 3.6.2.2 | Transmitter Specifications | 66 |
| 3.6.2.3 | Transmitter Differential Return Loss | 69 |
| 3.6.2.4 | Transmitter Common Mode Return Loss | 69 |
| 3.6.2.5 | Transmit Equalization | 69 |
| 3.6.3 | Captive Device Receiver Specifications | 69 |
| 3.6.3.1 | Receiver Specifications Applied for All Speeds | 69 |
| 3.6.3.2 | Receiver Differential Return Loss | 70 |
| 3.6.3.3 | Receiver Common Mode Return Loss | 70 |
| 3.6.4 | Captive Device Receiver Uncoded BER Tolerance Testing | 70 |
| 3.6.5 | Captive Device Receiver Multi Error-Bursts Testing | 72 |
| 3.7 | Low Frequency Periodic Signaling (LFPS) | 73 |
| 3.7.1 | LFPS Signal Definition | 73 |
| 3.8 | Receiver Lane Margining (Testability) | 74 |
| 3.8.1 | Background | 74 |
| 3.8.1.1 | Software Margining Mode | 74 |
| 3.8.1.2 | Hardware Margining Mode | 75 |
| 3.8.2 | Receiver Voltage Margining and Timing Margining Requirements | 76 |
| 3.8.3 | Receiver Parameter Access | 78 |
| 4 | Logical Layer | 79 |
| 4.1 | Sideband Channel | 79 |
| 4.1.1 | Transactions | 80 |
| 4.1.1.1 | Symbols | 80 |
| 4.1.1.2 | Transaction Types | 80 |
| 4.1.1.3 | SB Register Space | 88 |
| 4.1.2 | Lane Initialization | 96 |

| | | |
|---------|---|-----|
| 4.1.2.1 | Phase 1 – Determination of Initial Conditions | 97 |
| 4.1.2.2 | Phase 2 – Router Detection..... | 99 |
| 4.1.2.3 | Phase 3 – Determination of USB4 Port Characteristics..... | 99 |
| 4.1.2.4 | Phase 4 – Lane Parameters Synchronization and Transmit Start | 100 |
| 4.1.2.5 | Phase 5 – Link Equalization..... | 100 |
| 4.2 | Logical Layer State Machine | 103 |
| 4.2.1 | Lane Adapter State Machine | 103 |
| 4.2.1.1 | Disabled..... | 104 |
| 4.2.1.2 | CLd..... | 104 |
| 4.2.1.3 | Training | 105 |
| 4.2.1.4 | CL0..... | 112 |
| 4.2.1.5 | Lane Bonding | 113 |
| 4.2.1.6 | Low Power (CL0s, CL1, and CL2) | 114 |
| 4.2.2 | USB4 Link Transitions | 130 |
| 4.2.2.1 | Transition from One Single-Lane Link to Two Single-Lane Links | 130 |
| 4.2.2.2 | Transition from Two Single-Lane Links to Dual-Lane Link | 131 |
| 4.2.2.3 | Transition from Dual-Lane Link to Two Single-Lane Links | 132 |
| 4.2.2.4 | Transition from Two Single-Lane Links to One Single-Lane Link | 132 |
| 4.2.3 | Logical Layer Link State..... | 133 |
| 4.3 | USB4 Link Encoding | 133 |
| 4.3.1 | Lane Distribution | 135 |
| 4.3.2 | Symbol Encoding..... | 136 |
| 4.3.2.1 | Symbol Encoding of Transport Layer Bytes..... | 136 |
| 4.3.3 | Ordered Sets | 137 |
| 4.3.4 | Bit Swap | 138 |
| 4.3.4.1 | Sync Bits | 138 |
| 4.3.4.2 | Data Symbol Payload | 138 |
| 4.3.4.3 | Ordered Set Symbol Payload..... | 139 |
| 4.3.5 | Scrambling | 140 |
| 4.3.6 | RS-FEC | 141 |
| 4.3.6.1 | RS-FEC Activation and Deactivation..... | 143 |
| 4.3.6.2 | Pre-Coding..... | 144 |
| 4.4 | USB4 Link Operation | 144 |
| 4.4.1 | Start of Data..... | 144 |
| 4.4.2 | Error Cases and Recovery | 144 |
| 4.4.3 | Clock Compensation and SKIP | 146 |
| 4.4.4 | Dual-Lane Skew | 146 |
| 4.4.5 | Disconnect..... | 147 |
| 4.4.5.1 | Upstream Facing Port Disconnect | 147 |
| 4.4.5.2 | Downstream Port Disconnect..... | 148 |
| 4.4.6 | Lane Adapter Disable and Enable | 151 |
| 4.4.6.1 | Disabled Adapter is the Upstream Adapter | 151 |

| | | |
|---------|---|-----|
| 4.4.6.2 | Disabled Adapter is not the Upstream Adapter | 152 |
| 4.4.7 | Time Sync Notification Ordered Set (TSNOS) | 154 |
| 4.5 | Sleep and Wake | 154 |
| 4.5.1 | Entry to Sleep | 154 |
| 4.5.2 | Behavior in Sleep State | 156 |
| 4.5.3 | Wake Events | 156 |
| 4.5.4 | Exit from Sleep..... | 157 |
| 4.5.4.1 | Upstream Facing Port Disconnect | 157 |
| 4.5.4.2 | Wake on USB4 Event | 157 |
| 4.6 | Timing Parameters | 158 |
| 5 | Transport Layer..... | 161 |
| 5.1 | Transport Layer Packets | 161 |
| 5.1.1 | Bit/Byte Conventions | 161 |
| 5.1.2 | Format | 161 |
| 5.1.2.1 | Header | 162 |
| 5.1.2.2 | Payload Padding | 163 |
| 5.1.2.3 | Error Correction Code (ECC)..... | 163 |
| 5.1.3 | Transport Layer Packets..... | 163 |
| 5.1.3.1 | Tunneled Packets | 164 |
| 5.1.3.2 | Control Packets | 164 |
| 5.1.3.3 | Link Management Packets..... | 164 |
| 5.1.4 | Effect of Link State on Transport Layer Packets | 167 |
| 5.1.5 | Minimum Headers Gap | 167 |
| 5.2 | Routing | 168 |
| 5.2.1 | Adapter Numbering Rules | 168 |
| 5.2.2 | HopID Rules | 169 |
| 5.2.3 | Routing Tables | 170 |
| 5.2.4 | Routing Rules | 171 |
| 5.2.4.1 | Control Packets | 171 |
| 5.2.4.2 | Link Management Packets..... | 171 |
| 5.2.4.3 | Tunneled Packets | 172 |
| 5.2.4.4 | Routing Example..... | 172 |
| 5.2.5 | Connectivity Rules | 173 |
| 5.3 | Quality of Service (QOS)..... | 174 |
| 5.3.1 | Packet Ordering..... | 174 |
| 5.3.2 | Flow Control..... | 174 |
| 5.3.2.1 | Ingress Adapter | 175 |
| 5.3.2.2 | Egress Adapter | 180 |
| 5.3.2.3 | Credit Counter Synchronization | 181 |
| 5.3.3 | Bandwidth Arbitration and Priority | 182 |
| 5.3.3.1 | Scheduling | 182 |
| 5.3.4 | Packet Forwarding Delay Jitter | 184 |
| 5.4 | Path Tear-down | 184 |
| 5.4.1 | Egress Adapter..... | 184 |

| | | |
|----------|---|-----|
| 5.4.2 | Ingress Adapter | 185 |
| 5.5 | Timing Parameters | 185 |
| 6 | Configuration Layer..... | 186 |
| 6.1 | Domain Topology | 186 |
| 6.2 | Router Addressing | 186 |
| 6.3 | Router States | 188 |
| 6.3.1 | Uninitialized Unplugged State | 189 |
| 6.3.2 | Uninitialized Plugged State..... | 189 |
| 6.3.3 | Sleep State..... | 189 |
| 6.3.4 | Enumerated State..... | 189 |
| 6.4 | Control Packet Protocol | 190 |
| 6.4.1 | Control Adapter | 190 |
| 6.4.2 | Control Packets..... | 190 |
| 6.4.2.1 | Bit/Byte Conventions..... | 190 |
| 6.4.2.2 | Format..... | 190 |
| 6.4.2.3 | Read Request..... | 191 |
| 6.4.2.4 | Read Response..... | 192 |
| 6.4.2.5 | Write Request | 194 |
| 6.4.2.6 | Write Response | 195 |
| 6.4.2.7 | Notification Packet | 196 |
| 6.4.2.8 | Notification Acknowledgement Packet | 197 |
| 6.4.2.9 | Hot Plug Event Packet..... | 198 |
| 6.4.2.10 | Inter-Domain Request..... | 199 |
| 6.4.2.11 | Inter-Domain Response..... | 200 |
| 6.4.3 | Control Packet Routing | 201 |
| 6.4.3.1 | Upstream-Bound Packets | 201 |
| 6.4.3.2 | Downstream-Bound Packets..... | 201 |
| 6.4.3.3 | Processing of Read and Write Requests..... | 203 |
| 6.4.4 | Control Packet Reliability..... | 204 |
| 6.5 | Notification Events | 205 |
| 6.6 | Notification Acknowledgement..... | 205 |
| 6.7 | Router Enumeration | 206 |
| 6.8 | Hot Plug and Hot Unplug Events | 208 |
| 6.8.1 | Router Hot Plug | 209 |
| 6.8.1.1 | Enumerated Routers..... | 209 |
| 6.8.1.2 | Uninitialized Routers..... | 209 |
| 6.8.1.3 | Hot Plugged Router | 209 |
| 6.8.2 | Router Hot Unplug..... | 209 |
| 6.8.2.1 | Hot Unplug on the Upstream Facing Port..... | 209 |
| 6.8.2.2 | Hot Unplug on a Downstream Facing Port..... | 210 |
| 6.9 | Downstream Facing Port Reset..... | 210 |
| 6.10 | Timing Parameters | 210 |
| 7 | Time Synchronization | 211 |

| | | |
|---------|--|-----|
| 7.1 | Time Synchronization Architecture..... | 211 |
| 7.1.1 | Synchronization Hierarchy | 211 |
| 7.1.1.1 | Intra-Domain Hierarchy | 211 |
| 7.1.1.2 | Inter-Domain Hierarchy | 212 |
| 7.1.2 | Time Sync Parameters | 212 |
| 7.1.2.1 | Local Time | 212 |
| 7.1.2.2 | Time Offset | 213 |
| 7.1.2.3 | Frequency Offset..... | 213 |
| 7.2 | Time Stamp Measurement | 214 |
| 7.2.1 | Asymmetry Corrections | 214 |
| 7.3 | Time Sync Protocol | 215 |
| 7.3.1 | Time Sync Handshake..... | 215 |
| 7.3.1.1 | Bi-Directional Time Sync Handshake..... | 216 |
| 7.3.1.2 | Uni-Directional Time Sync Handshake..... | 220 |
| 7.3.2 | Inter-Domain Time Sync | 222 |
| 7.3.3 | Packet Formats | 224 |
| 7.3.3.1 | Time Sync Notification Ordered Set Format | 224 |
| 7.3.3.2 | Follow-Up Packet Format..... | 224 |
| 7.3.3.3 | Inter-Domain Time Stamp Packet..... | 226 |
| 7.4 | Time Computations..... | 227 |
| 7.4.1 | Intra-Domain Equations | 229 |
| 7.4.2 | Inter-Domain Equations | 231 |
| 7.4.2.1 | Inter-Domain Time Stamp Computation | 232 |
| 7.4.2.2 | Inter-Domain Frequency Offset Computation..... | 232 |
| 7.4.2.3 | Inter-Domain Time Offset Computation | 233 |
| 7.4.2.4 | Inter-Domain Host Router Time Computation | 234 |
| 7.4.3 | Filtering | 235 |
| 7.5 | Time Synchronization Accuracy Requirements | 236 |
| 7.5.1 | Paired Measurement..... | 236 |
| 7.5.2 | Standalone Measurement | 236 |
| 7.5.3 | Measuring Method | 237 |
| 7.5.4 | Accuracy Parameters | 238 |
| 7.6 | Software Configuration | 239 |
| 7.6.1 | Intra-Domain Time Synchronization Setup | 239 |
| 7.6.2 | Inter-Domain Time Synchronization Setup | 239 |
| 7.6.3 | Post Time Mechanism | 239 |
| 7.6.4 | Time Disruption Bit..... | 240 |
| 8 | Configuration Spaces | 241 |
| 8.1 | Configuration Fields Access Types | 241 |
| 8.2 | Configuration Spaces..... | 241 |
| 8.2.1 | Router Configuration Space..... | 242 |
| 8.2.1.1 | Basic Configuration Registers..... | 244 |
| 8.2.1.2 | TMU Router Configuration Capability..... | 251 |
| 8.2.1.3 | Vendor Specific Capability (VSC) | 258 |

| | | |
|---------|---|-----|
| 8.2.1.4 | Vendor Specific Extended Capability (VSEC)..... | 259 |
| 8.2.2 | Adapter Configuration Space..... | 260 |
| 8.2.2.1 | Basic Configuration Registers..... | 262 |
| 8.2.2.2 | TMU Adapter Configuration Capability..... | 266 |
| 8.2.2.3 | Lane Adapter Configuration Capability..... | 268 |
| 8.2.2.4 | USB4 Port Capability..... | 272 |
| 8.2.2.5 | USB3 Adapter Configuration Capability..... | 278 |
| 8.2.2.6 | DP Adapter Configuration Capability..... | 280 |
| 8.2.2.7 | PCIe Adapter Configuration Capability..... | 292 |
| 8.2.3 | Path Configuration Space..... | 294 |
| 8.2.3.1 | Path 0 Entry..... | 294 |
| 8.2.3.2 | Lane Adapters..... | 295 |
| 8.2.3.3 | Protocol Adapters..... | 296 |
| 8.2.3.4 | Path Configuration Space Access..... | 298 |
| 8.2.4 | Counters Configuration Space..... | 300 |
| 8.3 | Operations..... | 301 |
| 8.3.1 | Router Operations..... | 301 |
| 8.3.1.1 | DP Tunneling Operations..... | 303 |
| 8.3.1.2 | NVM Operations..... | 305 |
| 8.3.1.3 | Router Discovery Operations..... | 311 |
| 8.3.1.4 | Port Control Operations..... | 317 |
| 8.3.2 | Port Operations..... | 318 |
| 8.3.2.1 | Compliance Port Operations..... | 320 |
| 8.3.2.2 | Service Port Operations..... | 329 |
| 8.3.2.3 | Receiver Lane Margining Port Operations..... | 330 |
| 9 | USB3 Tunneling..... | 340 |
| 9.1 | USB3 Adapter Layer..... | 341 |
| 9.1.1 | Encapsulation..... | 341 |
| 9.1.1.1 | LFPS Encapsulation..... | 342 |
| 9.1.1.2 | Ordered Set Encapsulation..... | 345 |
| 9.1.1.3 | Link Command Encapsulation..... | 347 |
| 9.1.1.4 | Idle Symbols..... | 347 |
| 9.1.1.5 | LMP Encapsulation..... | 347 |
| 9.1.1.6 | TP Encapsulation..... | 348 |
| 9.1.1.7 | ITP Encapsulation..... | 348 |
| 9.1.1.8 | Data Packet (DP) Encapsulation..... | 349 |
| 9.1.2 | Bandwidth Negotiation..... | 351 |
| 9.1.3 | Timing Parameters..... | 353 |
| 9.2 | Internal USB3 Device..... | 353 |
| 9.2.1 | Link Layer..... | 354 |
| 9.2.1.1 | Link Training and Status State Machine (LTSSM)..... | 354 |

| | | |
|----------|---|-----|
| 9.2.1.2 | Timers and Timeouts..... | 354 |
| 9.2.2 | USB3 Protocol Layer | 355 |
| 9.2.3 | Descriptors | 355 |
| 9.3 | Paths..... | 355 |
| 9.3.1 | Path Setup | 355 |
| 9.3.2 | Path Teardown..... | 355 |
| 10 | DisplayPort™ Tunneling..... | 357 |
| 10.1 | DP Adapter Protocol Stack | 357 |
| 10.1.1 | Transport Layer..... | 358 |
| 10.1.2 | Protocol Adapter Layer | 358 |
| 10.1.3 | DP Physical Layer | 358 |
| 10.2 | DP Adapter States | 358 |
| 10.2.1 | Reset | 359 |
| 10.2.2 | Present | 359 |
| 10.2.3 | Plugged..... | 360 |
| 10.2.4 | Paired | 360 |
| 10.3 | Interfaces | 360 |
| 10.3.1 | DisplayPort | 360 |
| 10.3.1.1 | LTTTPR Non-Transparent | 361 |
| 10.3.1.2 | Non-LTTTPR..... | 361 |
| 10.3.1.3 | LTTTPR Transparent..... | 362 |
| 10.3.2 | Programming Model..... | 362 |
| 10.3.2.1 | Adapter Configuration Space | 362 |
| 10.3.2.2 | Path Configuration Space..... | 362 |
| 10.3.3 | Hot Plug and Hot Removal Events | 363 |
| 10.3.3.1 | DP OUT Adapters..... | 363 |
| 10.3.3.2 | DP IN Adapters | 363 |
| 10.3.4 | DisplayPort Over USB4 Fabric | 365 |
| 10.3.4.1 | DisplayPort Data Packet Types | 365 |
| 10.3.4.2 | AUX Path Packet..... | 365 |
| 10.3.4.3 | Main-Link Path Packet Formats..... | 372 |
| 10.4 | System Flows | 372 |
| 10.4.1 | Connection Manager Discovery | 372 |
| 10.4.2 | Path Configuration..... | 373 |
| 10.4.2.1 | Setup | 373 |
| 10.4.2.2 | Tear-down..... | 375 |
| 10.4.3 | HPD Event Propagation..... | 376 |
| 10.4.3.1 | HPD Plug..... | 376 |
| 10.4.3.2 | HPD Unplug..... | 376 |
| 10.4.3.3 | IRQ | 376 |
| 10.4.3.4 | HPD Delay Requirements | 377 |
| 10.4.3.5 | Manual HPD Control | 377 |
| 10.4.4 | AUX Request and Response Handling | 377 |

| | | |
|-----------|--|-----|
| 10.4.4.1 | LTTTPR Non-Transparent Mode..... | 377 |
| 10.4.4.2 | Non-LTTTPR Mode..... | 379 |
| 10.4.4.3 | LTTTPR Transparent Mode..... | 382 |
| 10.4.4.4 | AUX Delay Requirements..... | 384 |
| 10.4.4.5 | Aggregated DisplayPort Capabilities..... | 384 |
| 10.4.4.6 | DPCD DP Tunneling over USB4..... | 385 |
| 10.4.5 | DP Adapters Init Flow..... | 386 |
| 10.4.5.1 | Multi-Function DP..... | 386 |
| 10.4.6 | Source Discovery..... | 386 |
| 10.4.6.1 | LTTTPR Recognition and Modes Change..... | 387 |
| 10.4.6.2 | DPRX Capabilities Read..... | 388 |
| 10.4.6.3 | Sink Count Read..... | 388 |
| 10.4.7 | Down-Spread Control..... | 388 |
| 10.4.8 | Stream Mode Set..... | 389 |
| 10.4.9 | DSC and FEC Enable..... | 389 |
| 10.4.10 | DP Link Training..... | 389 |
| 10.4.10.1 | LTTTPR..... | 389 |
| 10.4.10.2 | Non-LTTTPR and LTTTPR Transparent..... | 394 |
| 10.4.10.3 | Transition to High Speed Tunnel..... | 396 |
| 10.4.11 | Power States Set..... | 396 |
| 10.4.12 | DP Main-Link Disable..... | 397 |
| 10.4.13 | Link-Init..... | 397 |
| 10.4.14 | DP PHY Testability..... | 397 |
| 10.4.14.1 | DP IN Adapter PHY Layer Testing..... | 397 |
| 10.4.14.2 | DP OUT Adapter PHY Layer Testing..... | 398 |
| 10.5 | High Speed Tunneling..... | 398 |
| 10.5.1 | SST Tunneling..... | 399 |
| 10.5.1.1 | Video Data Packet..... | 400 |
| 10.5.1.2 | Main Stream Attribute Packet..... | 405 |
| 10.5.1.3 | Blank Start Packet..... | 406 |
| 10.5.1.4 | Secondary Data Packet..... | 408 |
| 10.5.1.5 | Fill Count..... | 412 |
| 10.5.2 | MST Tunneling..... | 415 |
| 10.5.2.1 | Sub-MTP TU..... | 415 |
| 10.5.2.2 | MTP to Sub-MTP TU Examples..... | 421 |
| 10.5.2.3 | MST Packet Format..... | 423 |
| 10.5.2.4 | MST Packets to DP MTP..... | 424 |
| 10.5.3 | FEC..... | 424 |
| 10.5.3.1 | SR Count..... | 424 |
| 10.5.3.2 | DP IN Adapter Requirements..... | 425 |
| 10.5.3.3 | DP OUT Adapter Requirements..... | 425 |
| 10.5.3.4 | FEC_DECODE Packet..... | 426 |

| | | |
|----------|--|-----|
| 10.5.4 | DP OUT Adapter Buffer | 426 |
| 10.5.4.1 | Buffer Operation..... | 427 |
| 10.5.4.2 | Accumulation Cycles..... | 427 |
| 10.5.5 | HDCP..... | 428 |
| 10.6 | DP Link Clock Sync..... | 428 |
| 10.6.1 | Synchronization Method..... | 429 |
| 10.6.1.1 | Events..... | 429 |
| 10.6.1.2 | Lifetime Counter..... | 429 |
| 10.6.1.3 | DP Clock Sync Packet..... | 431 |
| 10.6.2 | DP Adapter Requirements..... | 433 |
| 10.6.2.1 | DP IN Adapter Requirements | 433 |
| 10.6.2.2 | DP OUT Adapter Requirements | 433 |
| 10.7 | DP BW Allocation Mode..... | 434 |
| 10.7.1 | DP BW Allocation Mode Enablement | 434 |
| 10.7.2 | Interaction with DPTX..... | 435 |
| 10.7.2.1 | Estimated Bandwidth..... | 437 |
| 10.7.3 | Interaction with the Connection Manager | 438 |
| 10.8 | Timing Parameters | 440 |
| 11 | PCI Express Tunneling..... | 441 |
| 11.1 | PCIe Adapter Layer | 442 |
| 11.1.1 | Encapsulation..... | 442 |
| 11.1.1.1 | PCIe TLP and DLLP..... | 442 |
| 11.1.1.2 | PCIe Ordered Sets | 446 |
| 11.1.1.3 | Electrical Idle State..... | 448 |
| 11.1.1.4 | PERST | 448 |
| 11.1.2 | USB4 Hot-Plug..... | 449 |
| 11.2 | Internal PCIe Ports | 449 |
| 11.2.1 | PCIe Physical Layer Logical Sub-block..... | 449 |
| 11.2.1.1 | Encoding..... | 449 |
| 11.2.1.2 | Link Training and Status State Machine (LTSSM) | 449 |
| 11.2.1.3 | ASPM L1 Entry..... | 450 |
| 11.2.1.4 | Clock Tolerance Compensation..... | 450 |
| 11.2.1.5 | Compliance Mode..... | 450 |
| 11.2.1.6 | Clock Power Management..... | 450 |
| 11.2.1.7 | L2 State | 450 |
| 11.2.2 | PCIe Data Link Layer..... | 450 |
| 11.2.3 | PCIe Transaction Layer | 450 |
| 11.2.4 | PCIe Link Timers (Informative) | 451 |
| 11.2.5 | Precision Time Measurement (PTM) Mechanism | 452 |
| 11.2.5.1 | Parameter Generator..... | 454 |
| 11.2.5.2 | Parameter Consumer..... | 454 |
| 11.2.5.3 | PTM Calculations..... | 455 |
| 11.2.6 | Timing Parameters | 456 |

| | | |
|----------|---|-----|
| 11.3 | Paths..... | 457 |
| 11.3.1 | Path Set-Up | 457 |
| 11.3.2 | Path Tear-Down..... | 457 |
| 12 | Host Interface..... | 459 |
| 12.1 | Descriptor Ring Mode | 460 |
| 12.1.1 | DW, Byte, and Bit Order | 460 |
| 12.1.2 | Raw Mode | 460 |
| 12.1.3 | Frame Mode | 460 |
| 12.2 | End-to-End (E2E) Flow Control..... | 462 |
| 12.2.1 | E2E Flow Control Packets | 462 |
| 12.2.1.1 | E2E Credit Grant Packet..... | 462 |
| 12.2.1.2 | E2E Credit Sync Packet..... | 464 |
| 12.2.2 | Flow Control Rules | 464 |
| 12.2.2.1 | Credit Update | 464 |
| 12.2.2.2 | Credit Counter Synchronization..... | 464 |
| 12.2.2.3 | Transmitting Host Interface Rules | 465 |
| 12.2.2.4 | Receiving Host Interface Rules..... | 466 |
| 12.3 | Transmit Interface | 467 |
| 12.3.1 | Transmit Descriptor Structure | 467 |
| 12.3.2 | Transmit Flow..... | 468 |
| 12.3.2.1 | Frame Mode..... | 468 |
| 12.3.2.2 | Raw Mode..... | 469 |
| 12.4 | Receive Interface..... | 470 |
| 12.4.1 | Receive Descriptor Structure | 470 |
| 12.4.2 | Receive Flow..... | 472 |
| 12.4.2.1 | Frame Mode..... | 472 |
| 12.4.2.2 | Raw Mode..... | 473 |
| 12.5 | Interrupts | 474 |
| 12.5.1 | Interrupt Causes..... | 474 |
| 12.5.2 | Interrupt Masks | 474 |
| 12.5.3 | Interrupt Vectors | 474 |
| 12.5.4 | Interrupt Moderation | 474 |
| 12.6 | Programming Interface | 475 |
| 12.6.1 | Access Types..... | 476 |
| 12.6.2 | Registers Summary..... | 476 |
| 12.6.3 | Registers Description | 477 |
| 12.6.3.1 | Host Interface Control..... | 477 |
| 12.6.3.2 | Transmit Descriptor Rings..... | 479 |
| 12.6.3.3 | Receive Descriptor Rings..... | 481 |
| 12.6.3.4 | Interrupts..... | 484 |
| 12.7 | Timing Parameters | 489 |
| 13 | Interoperability with Thunderbolt™ 3 (TBT3) Systems | 490 |
| 13.1 | Electrical Layer | 490 |
| 13.2 | Logical Layer..... | 491 |

| | | |
|----------|--|-----|
| 13.2.1 | Sideband Channel | 491 |
| 13.2.1.1 | Bidirectional Re-timer | 491 |
| 13.2.1.2 | Transactions | 491 |
| 13.2.1.3 | SB Register Space | 493 |
| 13.2.1.4 | Lane Initialization | 494 |
| 13.2.2 | Logical Layer State Machine | 499 |
| 13.2.2.1 | CLd State..... | 499 |
| 13.2.2.2 | TS1 and TS2 Ordered Sets..... | 499 |
| 13.2.2.3 | Low Power (CL0s, CL1, and CL2) | 500 |
| 13.2.3 | USB4 Link Operation..... | 500 |
| 13.2.3.1 | USB4 Link Transitions | 500 |
| 13.2.3.2 | Pre-Coding..... | 500 |
| 13.2.4 | Sleep and Wake..... | 500 |
| 13.2.4.1 | Entry to Sleep..... | 500 |
| 13.2.4.2 | Behavior in Sleep State | 501 |
| 13.2.4.3 | Wake Events | 501 |
| 13.2.4.4 | Exit from Sleep | 501 |
| 13.2.5 | Timing Parameters | 501 |
| 13.3 | Transport Layer | 502 |
| 13.3.1 | Adapter Numbering Rules | 502 |
| 13.3.2 | Maximum HopID..... | 502 |
| 13.3.3 | Connectivity Rules | 502 |
| 13.3.4 | Buffer Allocation | 502 |
| 13.4 | Configuration Layer | 503 |
| 13.4.1 | Router Enumeration..... | 503 |
| 13.4.2 | Notification Packet | 503 |
| 13.4.3 | Bit Banging Interface | 503 |
| 13.4.4 | Control Packet Routing | 504 |
| 13.4.4.1 | Downstream-Bound Packets..... | 504 |
| 13.4.4.2 | Uninitialized Router Flow | 505 |
| 13.5 | Time Synchronization | 505 |
| 13.6 | Configuration Spaces..... | 505 |
| 13.6.1 | Router Configuration Space | 506 |
| 13.6.1.1 | Vendor Specific 1 Capability..... | 506 |
| 13.6.1.2 | Vendor Specific 3 Capability..... | 510 |
| 13.6.1.3 | Vendor Specific 4 Capability..... | 513 |
| 13.6.1.4 | Vendor Specific Extended 6 Capability..... | 514 |
| 13.6.2 | Adapter Configuration Space..... | 521 |
| 13.6.2.1 | Basic Attributes..... | 521 |
| 13.6.2.2 | USB4 Port Capability..... | 522 |
| 13.7 | PCI Express Tunneling | 522 |
| 13.7.1 | PCIe Power Management | 522 |
| 13.7.1.1 | L1..... | 522 |

| | | |
|----------|--|-----|
| 13.7.1.2 | L2..... | 523 |
| 13.8 | DisplayPort Tunneling..... | 523 |
| 13.8.1 | AUX Handling..... | 523 |
| 13.8.1.1 | DP IN Adapter Requirements..... | 523 |
| 13.8.1.2 | DP OUT Adapter Requirements..... | 523 |
| 13.8.2 | IRQ Handling..... | 524 |
| 13.8.3 | Connection Manager Discovery..... | 524 |
| 13.8.3.1 | TBT3 Connection Manager..... | 524 |
| 13.8.3.2 | TBT3 Router Discovery..... | 524 |
| 13.8.4 | Sink Count Read..... | 525 |
| 13.8.5 | Power States Set..... | 525 |
| 13.8.6 | DisplayPort Link Training..... | 525 |
| 13.8.6.1 | DP IN Adapter Requirements..... | 526 |
| 13.8.6.2 | DP OUT Adapter Requirements..... | 527 |
| 13.9 | USB3 Functionality..... | 528 |
| 13.10 | Host-to-Host Tunneling..... | 529 |
| A | Verification of CRC, Scrambling, and FEC Calculations..... | 530 |
| A.1 | Transport Layer Packet HEC..... | 530 |
| A.2 | Control Packet CRC..... | 530 |
| A.3 | Sideband Channel AT Transaction CRC..... | 530 |
| A.4 | Scrambler..... | 531 |
| A.5 | Logical Layer RS-FEC..... | 532 |
| A.6 | USB3 Tunneling CRC..... | 537 |
| A.7 | Host Interface Frame CRC..... | 538 |
| A.8 | ECC Examples..... | 541 |
| B | Summary of Transport Layer Packets..... | 543 |
| C | Examples of Link Power Management Flows..... | 544 |
| C.1 | Entry to Low Power States..... | 544 |
| C.1.1 | Successful Entry to CL2 State..... | 544 |
| C.1.2 | Successful Entry to CLOs State..... | 544 |
| C.1.3 | Rejection to Enter CL2 State..... | 545 |
| C.1.4 | Concurrent Requests to Enter Low Power State..... | 546 |
| C.1.5 | CL2_REQ Ordered Sets are Not Received..... | 546 |
| C.1.6 | CL2_REQ Ordered Sets are Partially Received..... | 547 |
| C.1.7 | Error in CL2_ACK Ordered Sets..... | 548 |
| C.1.8 | Error in CL_OFF Ordered Sets..... | 549 |
| C.2 | Exit from Low Power States..... | 550 |
| C.2.1 | Example: Exit from CLOs State..... | 550 |
| C.2.2 | Example: Exit from CL2 (or CL1) State..... | 552 |
| D | Serial Time Link Protocol (STLP)..... | 555 |
| D.1 | Time Synchronization..... | 555 |
| D.2 | Serial Time Link Packet Format..... | 555 |
| D.3 | TMU_CLK_OUT and TMU_CLK_IN..... | 558 |
| E | Ingress Buffer Space..... | 559 |
| E.1 | Target Bandwidth Buffer Calculation..... | 559 |

| | | |
|-------|---|-----|
| E.1.1 | Example for USB3 Tunneling Ingress Buffer Calculation | 559 |
| E.2 | Ingress Buffers Calculation for DP Main Path | 560 |

Figures

| | | |
|--------------|---|----|
| Figure 2-1. | USB4/USB3.2 Dual Bus System Architecture | 12 |
| Figure 2-2. | Single-Lane USB4 Link | 13 |
| Figure 2-3. | Dual-Lane USB4 Link | 14 |
| Figure 2-4. | Example of a USB4-Based Dock | 15 |
| Figure 2-5. | USB4 Functional Stack Layers | 17 |
| Figure 2-6. | USB4 Port (Lane Adapter), Protocol Adapter and Control Adapter across Functional Layers | 17 |
| Figure 2-7. | Example USB4 Physical Topology (No Loop) and Spanning Tree | 19 |
| Figure 2-8. | Example USB4 Physical Topology (with Loop) and Spanning Tree | 19 |
| Figure 2-9. | Paths across a USB4 Fabric | 20 |
| Figure 2-10. | USB4 Communication by Functional Layer | 21 |
| Figure 2-11. | Example Control Packet Traversing Several Routers..... | 22 |
| Figure 2-12. | Example USB4 Host-to-Host Connections | 23 |
| Figure 2-13. | Example of a USB4 Host with USB3 Tunneling Highlighted..... | 26 |
| Figure 2-14. | Example of a USB4 Hub with USB3 Tunneling Highlighted..... | 27 |
| Figure 2-15. | Example of a USB4 Peripheral Device with USB3 Tunneling Highlighted | 27 |
| Figure 2-16. | Protocol Stack for USB3 Tunneling..... | 28 |
| Figure 2-17. | Example of a USB4 Fabric with USB3 Tunneling..... | 29 |
| Figure 2-18. | Protocol Stacks along a USB3 Tunnel | 30 |
| Figure 2-19. | Example Topology for DisplayPort Tunneling | 30 |
| Figure 2-20. | DP IN and OUT Protocol Adapters in LTTTPR Non-Transparent and LTTTPR Transparent Modes..... | 31 |
| Figure 2-21. | DP IN and OUT Protocol Adapters in Non-LTTTPR Mode..... | 32 |
| Figure 2-22. | Protocol Stacks along a DisplayPort Tunneled Path | 33 |
| Figure 2-23. | Example Structure of a USB4 Host with PCIe Tunneling Highlighted..... | 34 |
| Figure 2-24. | Example USB4 Hub with PCIe Tunneling Highlighted | 34 |
| Figure 2-25. | Example USB4 Device with PCIe Tunneling Highlighted | 35 |
| Figure 2-26. | Protocol Stack for PCIe Tunneling | 35 |
| Figure 2-27. | Example of a USB4 Fabric with PCIe Tunneling..... | 36 |
| Figure 2-28. | Protocol Stacks along a PCIe Tunnel | 37 |
| Figure 2-29. | Protocol Stacks along a Path between Hosts..... | 38 |
| Figure 2-30. | Descriptor Ring and Data Buffers..... | 39 |
| Figure 3-1. | Combined Forward-Error-Correction and Pre-Coding Scheme..... | 42 |
| Figure 3-2. | Compliance Points Definition..... | 43 |
| Figure 3-3. | Examples for AC-Coupling Capacitor Placement..... | 43 |
| Figure 3-4. | Jitter Transfer Function | 44 |
| Figure 3-5. | Reference Receiver Equalization..... | 45 |
| Figure 3-6. | Frequency Response of Gen 2 Reference CTLE..... | 46 |
| Figure 3-7. | Frequency Response of Gen 3 Reference CTLE..... | 46 |
| Figure 3-8. | Router Assembly Transmitter Frequency Variation During Training..... | 50 |
| Figure 3-9. | Example Transmitter Frequency During Steady-State..... | 51 |
| Figure 3-10. | TX Differential Return Loss Mask..... | 51 |
| Figure 3-11. | TX Common-Mode Return Loss Mask | 52 |
| Figure 3-12. | Transmitter Equalizer Structure | 53 |
| Figure 3-13. | Transmitter Equalization Frequency Response for Gen 2 Systems | 54 |
| Figure 3-14. | Transmitter Equalization Frequency Response for Gen 3 Systems | 55 |
| Figure 3-15. | TX Mask Notations | 57 |
| Figure 3-16. | RX Differential Return-Loss Mask..... | 60 |
| Figure 3-17. | RX Common Mode Return-Loss Mask | 60 |
| Figure 3-18. | Receiver Tolerance Test Topologies | 61 |
| Figure 3-19. | Receiver Tolerance Test Setups | 62 |
| Figure 3-20. | Captive Device Compliance Test Setup | 65 |
| Figure 3-21. | Captive Device Receiver Test Setup | 71 |
| Figure 3-22. | Signaling During Power Management State Exit | 73 |

| | |
|--|-----|
| Figure 3-23. Software Margining Mode Example | 75 |
| Figure 3-24. Hardware Margining Flow..... | 76 |
| Figure 3-25. RX Margining Range Requirements..... | 77 |
| Figure 3-26. Optional RX Margining Range Capabilities..... | 78 |
| Figure 4-1. Cable Topologies (Informative) | 80 |
| Figure 4-2. Symbol and Bit Order on Sideband Channel | 81 |
| Figure 4-3. Propagation of a Broadcast RT Transaction..... | 84 |
| Figure 4-4. Sideband Channel Receive Transaction State Machine | 87 |
| Figure 4-5. Overview of Lane Initialization | 97 |
| Figure 4-6. Example of Lane Reversal..... | 98 |
| Figure 4-7. Progression of Link Equalization | 101 |
| Figure 4-8. The Lane Adapter State Machine..... | 103 |
| Figure 4-9. Training Sub-State Machine..... | 105 |
| Figure 4-10. Lane Bonding Sub-State Machine | 113 |
| Figure 4-11. Structure of a CL_WAKE1.X Ordered Set Symbol | 117 |
| Figure 4-12. Packet Flow in the Logical Layer | 134 |
| Figure 4-13. Byte Transmission Order on Lanes | 135 |
| Figure 4-14. Byte Ordering of Transport Layer Packets to the Logical Layer | 136 |
| Figure 4-15. Byte Ordering of Idle Packets to the Logical Layer..... | 136 |
| Figure 4-16. Symbol Encoding of Data Symbols..... | 137 |
| Figure 4-17. Symbol Encoding of Ordered Set Symbols | 138 |
| Figure 4-18. Bit and Byte Ordering on the Wire – Data Symbol Payload..... | 139 |
| Figure 4-19. Bit and Byte Ordering on the Wire – Ordered Set Symbol Payload | 140 |
| Figure 4-20. RS-FEC Data Structures..... | 143 |
| Figure 4-21. Lane Disable of the Upstream Adapter | 152 |
| Figure 4-22. Lane Disable Flow | 153 |
| Figure 5-1. Convention for Transport Layer Diagrams..... | 161 |
| Figure 5-2. Transport Layer Packet Format..... | 162 |
| Figure 5-3. Idle Packet Contents..... | 164 |
| Figure 5-4. Credit Grant Packet Format | 165 |
| Figure 5-5. Path Credit Sync Packet Format | 166 |
| Figure 5-6. Shared Buffers Credit Sync Packet Format..... | 166 |
| Figure 5-7. Two Concurrent Data Symbols Example | 168 |
| Figure 5-8. Routing Table | 171 |
| Figure 5-9. Routing Example..... | 173 |
| Figure 5-10. Example of Connectivity for USB3 Adapters | 174 |
| Figure 5-11. Egress Adapter Scheduler | 183 |
| Figure 6-1. Example of TopologyID Assignment..... | 187 |
| Figure 6-2. Host Router State Machine..... | 188 |
| Figure 6-3. Device Router State Machine | 188 |
| Figure 6-4 Control Packet Format | 190 |
| Figure 6-5. Route String Format..... | 191 |
| Figure 6-6. Read Request..... | 192 |
| Figure 6-7. Read Response..... | 194 |
| Figure 6-8. Write Request | 195 |
| Figure 6-9. Write Response | 196 |
| Figure 6-10. Notification Packet..... | 197 |
| Figure 6-11. Notification Acknowledgement Packet..... | 198 |
| Figure 6-12. Hot Plug Event Packet..... | 199 |
| Figure 6-13. Inter-Domain Request | 200 |
| Figure 6-14. Inter-Domain Response | 201 |
| Figure 6-15. Example of Control Packet Routing Between Domains | 202 |
| Figure 7-1. Time Synchronization Hierarchy within a Domain (Informative) | 212 |
| Figure 7-2. Local Time Counter Format | 213 |
| Figure 7-3. <i>TimeOffsetFromHR</i> Register Format | 213 |
| Figure 7-4. <i>FreqOffsetFromHR</i> Register Format | 213 |
| Figure 7-5. <i>Time Measurement Model for 64/66b Encoding</i> | 214 |
| Figure 7-6. Bi-Directional Time Sync Handshake..... | 216 |
| Figure 7-7. UFP State Machine for Bi-Directional Time Sync Handshake (Recommended) | 218 |

| | |
|--|-----|
| Figure 7-8. DFP State Machine for Bi-Directional Time Sync Handshake (Recommended) | 219 |
| Figure 7-9. Uni-Directional Time Sync Handshake | 220 |
| Figure 7-10. DFP State Machine for Uni-Directional Time Sync Handshake (Recommended) | 221 |
| Figure 7-11. UFP State Machine for Uni-Directional Time Sync Handshake (Recommended) | 222 |
| Figure 7-12. Inter-Domain Time Sync Protocol (Informative) | 224 |
| Figure 7-13. Follow-Up Packet Format | 225 |
| Figure 7-14. Inter-Domain Time Stamp Packet Format | 227 |
| Figure 7-15. Inter-Domain Topology (Informative) | 229 |
| Figure 7-16. Filter Attenuation | 235 |
| Figure 7-17. Dynamic Noise Types | 236 |
| Figure 7-18. Standalone Measurement Points | 237 |
| Figure 7-19. Time Events | 238 |
| Figure 7-20. Measuring Method | 238 |
| Figure 8-1. Structure of the Router Configuration Space | 243 |
| Figure 8-2. UUID Format | 251 |
| Figure 8-3. Structure of the TMU Router Configuration Capability | 252 |
| Figure 8-4. Structure of a Vendor Specific Capability | 259 |
| Figure 8-5. Structure of a Vendor Specific Extended Capability | 259 |
| Figure 8-6. Structure of the Adapter Configuration Space | 261 |
| Figure 8-7. Basic Configuration Registers of the Adapter Configuration Space | 262 |
| Figure 8-8. Structure of the TMU Adapter Configuration Capability | 266 |
| Figure 8-9. Structure of the Lane Adapter Configuration Capability | 269 |
| Figure 8-10. Structure of USB4 Port Capability | 272 |
| Figure 8-11. Structure of USB3 Adapter Configuration Capability | 278 |
| Figure 8-12. Structure of DP IN Adapter Configuration Capability | 281 |
| Figure 8-13. Structure of DP OUT Adapter Configuration Capability | 287 |
| Figure 8-14. Structure of PCIe Adapter Configuration Capability | 292 |
| Figure 8-15. Structure of Path 0 Entry Configuration Space | 294 |
| Figure 8-16. Structure of Path Entry 'n' in Path Configuration Space at Lane Adapter | 295 |
| Figure 8-17. Structure of Path Entry 'n' in Path Configuration Space of a Protocol Adapter | 296 |
| Figure 8-18. Configuration of a Path | 300 |
| Figure 8-19. Structure of the Counters Configuration Space | 300 |
| Figure 8-20. Get Capabilities Operation Data Response for Capability Index 0 | 314 |
| Figure 9-1. LFPS Tunneled Packet Format | 343 |
| Figure 9-2. Ordered Set Tunneled Packet Format | 346 |
| Figure 9-3. Link Command Tunneled Packet Format | 347 |
| Figure 9-4. Tunneled ITP Packet Format | 348 |
| Figure 9-5. Structure of an Unsegmented USB3 Data Packet | 349 |
| Figure 9-6. Segmentation of a USB3 Data Packet | 350 |
| Figure 9-7. Bandwidth Negotiation by the Internal Host Controller | 352 |
| Figure 9-8. Bandwidth Negotiation by the Connection Manager | 353 |
| Figure 10-1. DP Adapter Protocol Stack Layers | 358 |
| Figure 10-2. DP Adapter State Machine | 359 |
| Figure 10-3. DP Adapter Path Directions | 362 |
| Figure 10-4. DP Stream Resource Mapping Examples | 364 |
| Figure 10-5. AUX Channel Framing | 366 |
| Figure 10-6. AUX Packet Format | 366 |
| Figure 10-7. AUX Packet Example | 367 |
| Figure 10-8. HPD Packet Format | 368 |
| Figure 10-9. SET_CONFIG Packet Format | 368 |
| Figure 10-10. ACK Packet Format | 372 |
| Figure 10-11. Power On to HPD Sequence | 373 |
| Figure 10-12. Target AUX Transaction Flow | 378 |
| Figure 10-13. Snoop AUX Transaction Flow | 379 |
| Figure 10-14. DP IN Adapter AUX Handling State Machine | 381 |
| Figure 10-15. AUX Timing | 384 |
| Figure 10-16. Example DP Source Discovery Sequence | 387 |
| Figure 10-17. DP Link Training – LTTTPR_CR_DONE | 391 |
| Figure 10-18. DP Link Training – LTTTPR – EQ Phase | 392 |

| | |
|--|-----|
| Figure 10-19. DP Link Training – DPRX – CR_DONE Phase..... | 393 |
| Figure 10-20. DP Link Training – DPRX – EQ Phase..... | 394 |
| Figure 10-21. Main-Link SST Stream to Tunneled Packets | 399 |
| Figure 10-22. TU Set Packing for a 4-Lane Main-Link..... | 400 |
| Figure 10-23. TU Set Packing for a 2-Lane Main-Link..... | 401 |
| Figure 10-24. TU Set Packing for a 1-Lane Main-Link..... | 402 |
| Figure 10-25. EOC Symbol Packing Example..... | 403 |
| Figure 10-26. TU Set Header Format..... | 403 |
| Figure 10-27. Video Data Packet Format..... | 405 |
| Figure 10-28. MSA Header Format..... | 405 |
| Figure 10-29. MSA Packet Format..... | 406 |
| Figure 10-30. Blank Start Header Format..... | 407 |
| Figure 10-31. Blank Start Packet Format | 408 |
| Figure 10-32. Secondary TU Header Format..... | 409 |
| Figure 10-33. Tunneled Secondary Data Path Format | 411 |
| Figure 10-34. Secondary Data to Secondary TUs Examples | 412 |
| Figure 10-35. Non-Secondary Data Packet Fill Count Examples | 414 |
| Figure 10-36. Secondary Data Packet Fill Count Examples..... | 415 |
| Figure 10-37. Sub-MTP TU Structures..... | 416 |
| Figure 10-38. Sub-MTP TU Header Format..... | 416 |
| Figure 10-39. Sub-MTP TU 4-Lane Mapping..... | 420 |
| Figure 10-40. Sub-MTP TU 2-Lane Mapping..... | 420 |
| Figure 10-41. Sub-MTP TU 1-Lane Mapping..... | 421 |
| Figure 10-42. Unallocated Sequence, 1-Lane | 421 |
| Figure 10-43. Shifting SR, 1-Lane..... | 422 |
| Figure 10-44. ACT Sequence, 1-Lane..... | 422 |
| Figure 10-45. SF and VCPF Sequence 4-Lane | 423 |
| Figure 10-46. MST Packet Format | 424 |
| Figure 10-47. FEC_DECODE Packet Format | 426 |
| Figure 10-48. FEC Command Format..... | 426 |
| Figure 10-49. Active Video to Blanking..... | 427 |
| Figure 10-50. Adjust PLL Event Occurrence..... | 429 |
| Figure 10-51. Lifetime Counter Format | 430 |
| Figure 10-52. Filtered Lifetime Counter Logic Concept..... | 431 |
| Figure 10-53. DP Clock Sync Packet Format | 432 |
| Figure 10-54. DP Clock Sync Packet Example | 433 |
| Figure 10-55: DP IN Adapter Interaction with DPTX During DP BW Allocation | 437 |
| Figure 10-56: DP BW Allocation Interaction with Connection Manager..... | 439 |
| Figure 11-1. Tunneled PCIe TLP..... | 443 |
| Figure 11-2. Tunneled PTM Example..... | 444 |
| Figure 11-3. Tunneled PCIe DLLP..... | 445 |
| Figure 11-4. PCIe DLLP and TLP Tunneled Packet Payload..... | 446 |
| Figure 11-5: Example of PTM Relationships..... | 453 |
| Figure 11-6: PTM ResponseD Message | 454 |
| Figure 11-7: TMU to PTM Parameters Illustration | 456 |
| Figure 12-1. Segmentation of a Frame..... | 461 |
| Figure 12-2. Example of Forwarding an E2E Credit Grant Packet..... | 463 |
| Figure 12-3. E2E Credit Grant / Sync Packet Format..... | 463 |
| Figure 12-4. Transmit Descriptor Structure | 467 |
| Figure 12-5. Receive Descriptor Structure (Posted by Host) | 470 |
| Figure 12-6. Receive Descriptor Structure (Posted by Host Interface Adapter Layer) | 471 |
| Figure 12-7. Interrupt Moderation..... | 475 |
| Figure 12-8. Structure of the Interrupt Status Registers..... | 484 |
| Figure 12-9. Structure of the Interrupt Vector Allocation Registers (IVAR) | 487 |
| Figure 12-10. Structure of the Receive Ring Vacancy Control Register..... | 488 |
| Figure 13-1. Bidirectional Re-timer Topology | 491 |
| Figure 13-2. Bounce Mechanism | 493 |
| Figure 13-3. Structure of the Vendor Specific 1 Capability | 506 |
| Figure 13-4. Structure of the Vendor Specific 3 Capability | 511 |

| | |
|---|-----|
| Figure 13-5. Structure of the Vendor Specific 4 Capability | 513 |
| Figure 13-6. Structure of the Vendor Specific Extended 6 Capability | 514 |
| Figure 13-7. Example Vendor Specific Extended 6 Capability | 514 |
| Figure 13-8. Structure of the Common Region | 515 |
| Figure 13-9. Structure of a USB4 Port Region | 517 |
| Figure 13-10. DP IN Adapter Link Training State Machine | 526 |
| Figure 13-11. DP OUT Adapter Link Training State Machine..... | 527 |
| Figure 13-12. Example of a USB4-Based Dock with an Internal Host Controller | 528 |
| Figure A-1. Examples of Transport Layer Packet HEC Calculation | 530 |
| Figure A-2. Examples of USB3 Tunneling Calculations | 538 |
| Figure A-3. Example of a Credit Grant Record | 541 |
| Figure A-4. Example of an HPD Packet Payload | 541 |
| Figure A-5. Example of a SET_CONFIG Packet Payload..... | 541 |
| Figure A-6. Example of TU Set Header..... | 542 |
| Figure A-7. Example of a Sub-MTP TU Header..... | 542 |
| Figure A-8. Example of an E2E Credit Sync Packet Payload | 542 |
| Figure C-1. Successful Entry to CL2 State..... | 544 |
| Figure C-2. Successful Entry to CL0s State..... | 545 |
| Figure C-3. Failure to Enter CL2 State..... | 545 |
| Figure C-4. Concurrent Requests to Enter CL2 State..... | 546 |
| Figure C-5. Error in CL2_REQ Ordered Sets | 547 |
| Figure C-6. CL2_REQ Ordered Sets are Partially Received | 547 |
| Figure C-7. Errors in CL2_REQ Reception and CL_NACK Response..... | 548 |
| Figure C-8. Error in CL2_ACK Ordered Sets..... | 549 |
| Figure C-9. Error in CL_OFF Ordered Sets..... | 549 |
| Figure C-10. CL0s Exit..... | 551 |
| Figure C-11. CL2 (or CL1) Exit..... | 553 |
| Figure D-1. Pulse Width Modulation | 555 |
| Figure D-2. Serial Time Link Packet Structure..... | 556 |
| Figure D-3. Serial Time Link Packet Format | 556 |
| Figure D-4. TMU_CLK_OUT and TMU_CLK_IN Parameters | 557 |
| Figure D-5. Definition of TCO _{JTR} | 558 |

Tables

| | |
|---|----|
| Table 1-1. Rsvd Value and Field Handling | 4 |
| Table 3-1. SBTX and SBRX Specifications..... | 41 |
| Table 3-2. Electrical Compliance Test Points | 43 |
| Table 3-3. Transmitter Specifications Applied for All Speeds (at TP2) | 48 |
| Table 3-4. Transmitter Frequency Variation Limits During Link Training Before Obtaining Steady-State..... | 50 |
| Table 3-5. Transmit Equalization Presets..... | 53 |
| Table 3-6. Gen 2 Transmitter Specifications at TP2 | 55 |
| Table 3-7. Gen 2 Transmitter Specifications at TP3 | 56 |
| Table 3-8. Gen 3 Transmitter Specifications at TP2 | 57 |
| Table 3-9. Gen 3 Transmitter Specifications at TP3 | 58 |
| Table 3-10. Common Receiver Specifications at TP3'..... | 58 |
| Table 3-11. Stressed Signal for Gen 2 Receiver Compliance Testing | 62 |
| Table 3-12. Stressed Signal for Gen 3 Receiver Compliance Testing | 62 |
| Table 3-13. Wireless Band Conducted Limits (at TP3) | 65 |
| Table 3-14. Captive Device Transmitter Specifications at TP3 Applied for All Speeds | 66 |
| Table 3-15. Captive Device Transmitter Specifications at TP3 for Gen 2 Systems | 67 |
| Table 3-16. Captive Device Transmitter Specifications at TP3 for Gen 3 Systems | 68 |
| Table 3-17. Common Receiver Specifications at TP2..... | 69 |
| Table 3-18. Stressed Receiver Conditions for Gen 2 Captive Device Compliance Testing (at TP2)..... | 71 |
| Table 3-19. Stressed Receiver Conditions for Gen 3 Captive Device Compliance Testing (at TP2)..... | 71 |
| Table 3-20. LFPS Electrical Specifications | 73 |
| Table 3-21. RX Margining Voltage and Timing Requirements..... | 76 |
| Table 3-22. Optional RX Margining Voltage Capabilities..... | 78 |
| Table 4-1. LT Transaction Format..... | 81 |

| | |
|--|-----|
| Table 4-2. LSE Symbol..... | 81 |
| Table 4-3. AT Transaction Format..... | 82 |
| Table 4-4. STX Symbol for an AT Transaction..... | 82 |
| Table 4-5. Broadcast RT Transaction Format | 83 |
| Table 4-6. STX Symbol for a Broadcast RT Transaction | 83 |
| Table 4-7. Contents of Byte 2 in a Broadcast RT Transaction | 83 |
| Table 4-8. Contents of Byte 3 in a Broadcast RT Transaction | 84 |
| Table 4-9. Addressed RT Transaction Format | 84 |
| Table 4-10. STX Symbol for an Addressed RT Transaction..... | 85 |
| Table 4-11. Sideband Channel Receive Transaction State Machine | 88 |
| Table 4-12. AT/RT Command Data Symbols | 89 |
| Table 4-13. AT/RT Response Data Symbols | 89 |
| Table 4-14. Processing of a Received AT/RT Command | 89 |
| Table 4-15. SB Registers | 91 |
| Table 4-16. SB Register Fields Access Types | 92 |
| Table 4-17. SB Register Fields | 92 |
| Table 4-18. Lane Attributes | 99 |
| Table 4-19. Transmitter Behavior in Training Sub-states..... | 106 |
| Table 4-20. Training Sub-State Machine Transitions..... | 106 |
| Table 4-21. SLOS1 (64b/66b Encoding)..... | 108 |
| Table 4-22. SLOS2 (64b/66b Encoding)..... | 109 |
| Table 4-23. SLOS1 (128b/132b Encoding)..... | 110 |
| Table 4-24. SLOS2 (128b/132b Encoding)..... | 111 |
| Table 4-25. TS1 and TS2 Ordered Sets | 111 |
| Table 4-26. Transmitter Behavior in Bonding Sub-States..... | 113 |
| Table 4-27. Lane Bonding Sub-State Machine Transitions..... | 113 |
| Table 4-28. CL2_REQ Ordered Set | 114 |
| Table 4-29. CL1_REQ Ordered Set | 115 |
| Table 4-30. CL2_ACK Ordered Set..... | 115 |
| Table 4-31. CL1_ACK Ordered Set..... | 115 |
| Table 4-32. CL0s_ACK Ordered Set | 115 |
| Table 4-33. CL_NACK Ordered Set..... | 115 |
| Table 4-34. CL_OFF Ordered Set..... | 116 |
| Table 4-35. Ordered Set Structure..... | 137 |
| Table 4-36. Scrambling Rules | 140 |
| Table 4-37. START_RS_FEC Bit Sequence..... | 144 |
| Table 4-38. Error Cases and Impact on Logical Layer | 145 |
| Table 4-39. SKIP Ordered Set..... | 146 |
| Table 4-40. De-Skew Ordered Set..... | 147 |
| Table 4-41. TSN Ordered Set..... | 154 |
| Table 4-42. Router State Retained During Sleep..... | 156 |
| Table 4-43. Wake Events | 157 |
| Table 4-44. Logical Layer Timing Parameters | 158 |
| Table 5-1. Transport Layer Packet Header Format | 162 |
| Table 5-2. Credit Grant Packet Header | 164 |
| Table 5-3. Credit Grant Record Format..... | 165 |
| Table 5-4. Path Credit Sync Packet Header | 165 |
| Table 5-5. Path Credit Sync Packet Payload | 166 |
| Table 5-6. Shared Buffers Credit Sync Packet Header | 166 |
| Table 5-7. Shared Buffers Credit Sync Packet Payload..... | 166 |
| Table 5-8. Transport Layer Behavior per Link State | 167 |
| Table 5-9. Minimum Transport Layer Header Gap Requirements | 167 |
| Table 5-10. Ingress Adapter Flow Control Schemes | 175 |
| Table 5-11. Buffer Allocation Parameters..... | 176 |
| Table 5-12. Egress Adapter Flow Control Schemes | 180 |
| Table 5-13. Transport Layer Timing Parameters | 185 |
| Table 6-1. Control Packet Payload..... | 190 |
| Table 6-2. Content of a Read Request | 191 |
| Table 6-3. Content of a Read Response | 193 |

| | |
|---|-----|
| Table 6-4. Content of a Write Request..... | 194 |
| Table 6-5. Content of a Write Response..... | 196 |
| Table 6-6. Content of a Notification Packet..... | 197 |
| Table 6-7. Content of a Notification Acknowledgement Packet..... | 197 |
| Table 6-8. Content of a Hot Plug Event Packet..... | 198 |
| Table 6-9. Content of an Inter-Domain Request..... | 199 |
| Table 6-10. Content of an Inter-Domain Response..... | 200 |
| Table 6-11. Notification Events..... | 205 |
| Table 6-12. Configuration Layer Timing Parameters..... | 210 |
| Table 7-1. Bidirectional UFP Timeout Values..... | 217 |
| Table 7-2. Bidirectional DFP Timeout Values..... | 218 |
| Table 7-3. Follow-Up Packet Payload..... | 225 |
| Table 7-4. Inter-Domain Time Stamp Packet Payload..... | 227 |
| Table 7-5. Definition of Variables..... | 228 |
| Table 7-6. Index Notation..... | 228 |
| Table 7-7. Time Synchronization Accuracy Parameters..... | 239 |
| Table 8-1. Configuration Register Fields Access Types..... | 241 |
| Table 8-2. List of Router Configuration Capabilities..... | 243 |
| Table 8-3. Router Configuration Space Basic Attributes..... | 244 |
| Table 8-4. TMU Router Configuration Capability Fields..... | 253 |
| Table 8-5. Locked Registers Groups..... | 258 |
| Table 8-6. Vendor Specific Capability Fields..... | 259 |
| Table 8-7. Vendor Specific Extended Capability Fields..... | 260 |
| Table 8-8. List of Adapter Configuration Capabilities..... | 261 |
| Table 8-9. Adapter Configuration Space Basic Attributes..... | 262 |
| Table 8-10. Adapter Types..... | 266 |
| Table 8-11. TMU Adapter Configuration Capability Fields..... | 267 |
| Table 8-12. Contents of the Lane Adapter Configuration Capability..... | 269 |
| Table 8-13. USB4 Port Capability Fields..... | 273 |
| Table 8-14. USB3 Adapter Configuration Capability Fields..... | 278 |
| Table 8-15. DP IN Adapter Configuration Capability Fields..... | 281 |
| Table 8-16. DP OUT Adapter Configuration Capability Fields..... | 287 |
| Table 8-17. PCIe Adapter Configuration Capability Fields..... | 292 |
| Table 8-18. Contents of Path 0 Entry..... | 294 |
| Table 8-19. Contents of Path Entry in Path Configuration Space at Lane Adapter..... | 295 |
| Table 8-20. Contents of Path Entry in Path Configuration Space of a Protocol Adapter..... | 296 |
| Table 8-21. Counter Set Fields..... | 301 |
| Table 8-22. List of Router Operations..... | 302 |
| Table 8-23. Query DP Resource Availability Operation Metadata..... | 303 |
| Table 8-24. Query DP Resource Availability Completion Metadata and Status..... | 304 |
| Table 8-25. Allocate DP Resource Operation Metadata..... | 304 |
| Table 8-26. Allocate DP Resource Completion Metadata and Status..... | 304 |
| Table 8-27. De-Allocate DP Resource Operation Metadata..... | 305 |
| Table 8-28. De-Allocate DP Resource Completion Metadata and Status..... | 305 |
| Table 8-29. NVM Set Offset Operation Metadata..... | 306 |
| Table 8-30. NVM Set Offset Completion Metadata and Status..... | 306 |
| Table 8-31. NVM Write Operation Data..... | 307 |
| Table 8-32. NVM Write Completion Status..... | 307 |
| Table 8-33. NVM Authenticate Write Completion Status..... | 308 |
| Table 8-34. NVM Read Operation Metadata..... | 308 |
| Table 8-35. NVM Read Router Completion Metadata..... | 309 |
| Table 8-36. NVM Read Router Completion Data..... | 309 |
| Table 8-37. DROM Read Router Operation Metadata..... | 309 |
| Table 8-38. DROM Read Router Completion Metadata and Status..... | 310 |
| Table 8-39. DROM Read Router Completion Data..... | 310 |
| Table 8-40. Get NVM Sector Size Completion Metadata and Status..... | 311 |
| Table 8-41. Get PCIe Downstream Entry Mapping Completion Metadata and Status..... | 311 |
| Table 8-42. Get PCIe Downstream Entry Mapping Completion Data..... | 312 |
| Table 8-43. Get Capabilities Operation Metadata..... | 312 |

| | |
|--|-----|
| Table 8-44. Get Capabilities Operation Completion Metadata and Status..... | 313 |
| Table 8-45. List of Capabilities..... | 314 |
| Table 8-46. Set Capabilities Operation Metadata..... | 314 |
| Table 8-47. List of Capabilities..... | 315 |
| Table 8-48. Set Capabilities Operation Completion Status..... | 315 |
| Table 8-49. Buffer Allocation Request Router Completion Status and Metadata..... | 316 |
| Table 8-50. Buffer Allocation Request Router Completion Data DW Structure..... | 316 |
| Table 8-51. Get Container-ID Router Completion Status..... | 317 |
| Table 8-52. Get Container-ID Router Completion Data DW Structure..... | 317 |
| Table 8-53. Block Sideband Port Operation Completion Status..... | 317 |
| Table 8-54. Unblock Sideband Port Operation Completion Status..... | 318 |
| Table 8-55. List of Port Operations..... | 319 |
| Table 8-56. SET_TX_COMPLIANCE Operation Metadata..... | 323 |
| Table 8-57. SET_RX_COMPLIANCE Operation Metadata..... | 324 |
| Table 8-58. START_BER_TEST Operation Metadata..... | 325 |
| Table 8-59. END_BER_TEST Operation Metadata..... | 326 |
| Table 8-60. END_BER_TEST Completion Data..... | 326 |
| Table 8-61. END_BURST_TEST Operation Metadata..... | 327 |
| Table 8-62. END_BURST_TEST Completion Data..... | 327 |
| Table 8-63. READ_BURST_TEST Operation Metadata..... | 328 |
| Table 8-64. READ_BURST_TEST Completion Data..... | 328 |
| Table 8-65. ENTER_EI_TEST Operation Metadata..... | 329 |
| Table 8-66. ROUTER_OFFLINE_MODE Operation Metadata..... | 329 |
| Table 8-67. READ_LANE_MARGIN_CAP Completion Data..... | 330 |
| Table 8-68. RUN_HW_LANE_MARGINING Operation Metadata..... | 332 |
| Table 8-69. Contents Selection for RUN_HW_LANE_MARGINING Completion Data..... | 333 |
| Table 8-70. RUN_HW_LANE_MARGINING Completion Data..... | 334 |
| Table 8-71. RUN_SW_LANE_MARGINING Operation Metadata..... | 337 |
| Table 8-72. RUN_SW_LANE_MARGINING Completion Data..... | 338 |
| Table 8-73. READ_SW_MARGIN_ERR Completion Metadata..... | 338 |
| Table 9-1. PDF Values for USB3 Tunneling Packets..... | 342 |
| Table 9-2. LFPS Tunneled Packet Payload..... | 342 |
| Table 9-3. Ordered Set Tunneled Packet Payload..... | 346 |
| Table 9-4. USB3 Adapter Timing Parameters..... | 353 |
| Table 9-5. USB3 Timers and Timeout Values..... | 354 |
| Table 10-1. DisplayPort Modes Of Operation Over DisplayPort Tunneling..... | 361 |
| Table 10-2. Recommended Path Parameters..... | 362 |
| Table 10-3. DP Stream Resource Allocation Commands..... | 364 |
| Table 10-4. AUX Path Tunneled Packet Types..... | 365 |
| Table 10-5. Main-Link Path Tunneled Packet Types..... | 365 |
| Table 10-6. SET_CONFIG Message..... | 370 |
| Table 10-7. DisplayPort Required Bandwidth (Gbps)..... | 375 |
| Table 10-8. HPD Event Propagation Delay Requirement..... | 377 |
| Table 10-9. DPCD Internal Addresses..... | 380 |
| Table 10-10. DP IN Adapter AUX Handling State Machine..... | 381 |
| Table 10-11. AUX Delay Requirements..... | 384 |
| Table 10-12. Aggregated DisplayPort Capabilities..... | 384 |
| Table 10-13. DP Adapter Operation Mode Transitions..... | 387 |
| Table 10-14. Blank Start Control Link Symbols Mapping..... | 408 |
| Table 10-15. Fill Count Prev_Factor..... | 413 |
| Table 10-16. Slot Zero Sub-MTP TU Header Types..... | 416 |
| Table 10-17. Non-Slot Zero Sub-MTP TU Header Types..... | 417 |
| Table 10-18. Slot Zero Sub-MTP TU Packet Rules..... | 417 |
| Table 10-19. Non- Zero Slot Sub-MTP TU Packet Rules..... | 418 |
| Table 10-20. K-Code Index Nibble in Parameter Byte..... | 419 |
| Table 10-21. FLC Calculation Examples..... | 431 |
| Table 10-22. DPCD Bandwidth Allocation Registers..... | 435 |
| Table 10-23. DP IN Adapter Configuration Space Mapping..... | 436 |
| Table 10-24. DP Adapter Timing Parameters..... | 440 |

| | |
|--|-----|
| Table 11-1. PDF Values for PCIe Tunneled Packets | 442 |
| Table 11-2. TLP Pre-Header | 444 |
| Table 11-3. TS Ordered Sets | 447 |
| Table 11-4. Electrical Idle Ordered Sets | 447 |
| Table 11-5. PCIe Link Timer Ranges | 452 |
| Table 11-6. PCIe Adapter Timing Parameters | 457 |
| Table 12-1. Frame Mode Tunneled Packet Format | 461 |
| Table 12-2. E2E Credit Grant Packet Header | 463 |
| Table 12-3. E2E Credit Grant Packet Payload | 463 |
| Table 12-4. E2E Credit Sync Packet Header | 464 |
| Table 12-5. E2E Credit Sync Packet Payload | 464 |
| Table 12-6. Transmit Descriptor Contents | 467 |
| Table 12-7. Receive Descriptor Contents (Posted by Host) | 470 |
| Table 12-8. Receive Descriptor Contents (Posted by Host Interface Adapter Layer) | 471 |
| Table 12-9. Access Types | 476 |
| Table 12-10. Summary of Memory BAR Registers | 476 |
| Table 12-11. Host Interface Capabilities Register | 477 |
| Table 12-12. Host Interface Reset Register | 478 |
| Table 12-13. Host Interface Control Register | 478 |
| Table 12-14. Host Interface CL1 Enable | 478 |
| Table 12-15. Host Interface CL2 Enable | 478 |
| Table 12-16. Base Address Low Register | 479 |
| Table 12-17. Base Address High Register | 479 |
| Table 12-18. Producer and Consumer Indexes Register | 479 |
| Table 12-19. Ring Size Register | 480 |
| Table 12-20. Ring Control Register | 481 |
| Table 12-21. Base Address Low Register | 481 |
| Table 12-22. Base Address High Register | 481 |
| Table 12-23. Producer and Consumer Indexes Register | 482 |
| Table 12-24. Ring Size Register | 482 |
| Table 12-25. Ring Control Register | 483 |
| Table 12-26. PDF Bit Masks Register | 484 |
| Table 12-27. Interrupt Status | 485 |
| Table 12-28. Interrupt Status Clear | 485 |
| Table 12-29. Interrupt Status Set | 485 |
| Table 12-30. Interrupt Mask | 486 |
| Table 12-31. Interrupt Mask Clear | 486 |
| Table 12-32. Interrupt Mask Set | 486 |
| Table 12-33. Interrupt Throttling Rate (ITR) | 486 |
| Table 12-34. Interrupt Vector Allocation (IVAR) | 487 |
| Table 12-35. Receive Ring Vacancy Control | 488 |
| Table 12-36. Receive Ring Vacancy Status | 488 |
| Table 12-37. Host Interface Timing Parameters | 489 |
| Table 13-1. Thunderbolt 3 Parameters | 490 |
| Table 13-2. TBT3 LT Transaction Types | 491 |
| Table 13-3. STX Symbol | 492 |
| Table 13-4. Contents of Byte 2 in a Broadcast RT Transaction | 493 |
| Table 13-5. SB Registers | 494 |
| Table 13-6. SB Registers Fields | 494 |
| Table 13-7. Lane Attributes | 494 |
| Table 13-8. TS1 and TS2 Ordered Set Structure | 500 |
| Table 13-9. Router State Retained During Sleep | 501 |
| Table 13-10. Logical Layer Timing Parameters | 502 |
| Table 13-11. Buffer Allocation by TBT3 Connection Manager | 503 |
| Table 13-12. Configuration Register Fields Access Types | 505 |
| Table 13-13. List of TBT3-Compatible Router Configuration Capabilities | 506 |
| Table 13-14. Vendor Specific 1 Capability Fields | 507 |
| Table 13-15. Vendor Specific 3 Capability Fields | 511 |
| Table 13-16. Vendor Specific 4 Capability Fields | 513 |

| | |
|---|-----|
| Table 13-17. Common Region Fields..... | 515 |
| Table 13-18. USB4 Port Region Fields | 517 |
| Table 13-19. Adapter Configuration Space Basic Attributes | 521 |
| Table 13-20. USB4 Port Capability Fields | 522 |
| Table 13-21. DP IN Adapter Link Training State Machine Transition Table | 526 |
| Table 13-22. DP OUT Adapter Link Training State Machine Transition Table..... | 527 |
| Table A-1. Examples of Control Packet CRC Calculation | 530 |
| Table A-2. Example of a Read Command | 530 |
| Table A-3. Example of a Write Command..... | 531 |
| Table A-4. Examples of Scrambler Computations | 531 |
| Table A-5. Example 1 – RS-FEC Block | 532 |
| Table A-6. Example 2 – RS-FEC Block | 533 |
| Table A-7. Example 3 – RS-FEC Block | 534 |
| Table A-8. Example 4 – RS-FEC Block | 536 |
| Table B-1. Transport Layer Packet Summary | 543 |
| Table D-1. Serial Time Link Packet Fields..... | 556 |
| Table D-2. TMU_CLK_OUT and TMU_CLK_IN Specifications | 558 |

1 Introduction

1.1 Scope of the Document

The specification is primarily targeted at peripheral developers and platform/adaptor developers, but provides valuable information for platform operating system/BIOS/device driver, adaptor independent hardware vendors/independent software vendors, and system OEMs. This specification can be used for developing new products and associated software.

1.2 USB Product Compliance

Adopters of the USB4™ specification have signed the USB4 Adopters Agreement, which provides them access to a royalty-free reasonable and nondiscriminatory (RAND) license from the Promoters and other Adopters to certain intellectual property contained in products that are compliant with the USB4 specification. Adopters can demonstrate compliance with the specification through the testing program as defined by the USB Implementers Forum (USB-IF). Products that demonstrate compliance with the specification will be granted certain rights to use the USB-IF logos as defined in the logo license.

1.3 Document Organization

Chapters 1 and 2 provide an overview for all readers, while Chapters 3 through 13 contain detailed technical information defining USB4.

1.4 Design Goals

USB 3.1 and USB 3.2 were evolutionary steps to increase bandwidth. The goal for USB4 remains the same with the added goal of helping to converge the USB Type-C® connector ecosystem and minimize end-user confusion. Several key design areas to meet this goal are listed below:

- Offer display, data, and load/store functionality over a single USB Type-C connector.
- Retain compatibility with existing ecosystem of USB and Thunderbolt™ products.
- Define Port Capabilities for predictable and consistent user experience.
- Provide increased host flexibility to configure bandwidth, power management, and other performance-related parameters for system needs.

1.5 Related Documents

Universal Serial Bus 3.2 Specification, Revision 1.0, September 22, 2017 (USB 3.2 Specification)

USB Type-C® Cable and Connector Specification, Release 2.0 (USB Type-C Specification)

USB 3.0 Jitter Budgeting white paper (USB Jitter Paper)

Universal Serial Bus Power Delivery Specification, Release 3.0, Version 2.0, August 2019 (USB PD Specification)

PCI Express® Base Specification, Revision 4, Version 1, September 27, 2017 (PCIe Specification)

VESA DisplayPort™ Standard, Revision 1.2a, May 2012 (DisplayPort 1.2a Specification)

VESA DisplayPort™ Standard, Revision 1.4a, April 19, 2018 (DisplayPort 1.4a Specification)

VESA DisplayPort™ 1.4a PHY Layer Compliance Test Specification, Revision 1.0, 27 July, 2018 (DisplayPort 1.4a PHY CTS)

VESA DisplayPort™ Alt Mode on USB Type-C Standard, Revision 1.0b, November 03, 2017 (DisplayPort Alt Mode Specification)

eXtensible Host Controller Interface for Universal Serial Bus, Revision 1.1 (xHCI Specification)

USB4 Connection Manager (CM) Guide, Revision 1.0, [to be published] – (Connection Manager Guide)

USB4 Re-Timer Specification, [to be published] – (USB4 Re-Timer Specification)

USB4 Device ROM (DROM) Specification, Revision 1.0, [to be published] – (USB4 DROM Specification)

USB4 Inter-Domain Specification, Revision 1.0, [to be published] – (USB4 Inter-Domain Specification)

HDCP on DisplayPort™ Specification, Revision 2.3, January 22, 2019 (HDCP Specification)

1.6 Conventions

1.6.1 Precedence

If there is a conflict between text, figures, and tables, the precedence shall be tables, figures, and then text.

Note: The text, figures, and tables in this specification all contain necessary information for building an implementation and need to be read together to get a full understanding of USB4 technology.

1.6.2 Keywords

The following keywords differentiate between the levels of requirements and options.

1.6.2.1 Informative

Informative is a keyword that describes information with this specification that intends to discuss and clarify requirements and features as opposed to mandating them.

1.6.2.2 May

May is a keyword that indicates a choice with no implied preference.

1.6.2.3 N/A

N/A is a keyword that indicates that a field or value is not applicable and has no defined value and shall not be checked or used by the recipient.

1.6.2.4 Normative

Normative is a keyword that describes features that are mandated by this specification.

1.6.2.5 Optional

Optional is a keyword that describes features not mandated by this specification. However, if an optional feature is implemented, the feature shall be implemented as defined by this specification (optional normative).

1.6.2.6 Reserved

Reserved is a keyword indicating reserved bits, bytes, words, fields, and code values that are set-aside for future standardization. The use and interpretation of these may be specified by future extensions to this specification and, unless otherwise stated, shall not be utilized or adapted by vendor implementation. A reserved bit, byte, word or field shall be set to zero by the sender and shall be ignored by the receiver. Reserved field values shall not be sent by the sender and, if received, shall be ignored by the receiver.

1.6.2.7 Shall

Shall is a keyword indicating a mandatory (normative) requirement. Designers are mandated to implement all such requirements to ensure interoperability with other compliant devices.

1.6.2.8 **Should**

Should is a keyword indicating flexibility of choice with a preferred alternative. Equivalent to the phrase “it is recommended that”.

1.6.3 **Capitalization**

Some terms are capitalized to distinguish their definition in the context of this specification from their common English meaning. Words not capitalized have their common English meaning.

1.6.4 **Italic Text**

Italic text is used to identify variable names, register field and packet field names, or reference document titles.

1.6.5 **Numbering**

Numbers that are immediately followed by a lowercase “b” (e.g. 01b) are binary values. Numbers that are immediately followed by an uppercase “B” are byte values. Numbers that are immediately followed by a lowercase “h” (e.g. 3Ah) are hexadecimal values. Numbers not immediately followed by either a “b”, “B”, or “h” are decimal values.

1.6.6 **Bit, Byte, DW, and Symbol Conventions**

A bit, byte, DW, or Symbol residing in location n within an array is denoted as bit(n), byte(n), DW(n), or Symbol(n).

A sequence of bits, bytes, DWs, or Symbols residing in locations n to m (inclusive) within an array is denoted as bit[m:n], byte[m:n], DW[m:n], or Symbol[m:n].

1.6.7 **Implementation Notes**

Implementation Notes are not a normative part of this specification. They are included for clarification and illustration only. Implementation Notes within this specification are marked as such and set apart from the other text.

1.6.8 **Connection Manager Notes**

Connection Manager Notes provide requirements and recommendations for building a Connection Manager. Connection Manager Notes within this specification are marked as such and set apart from the other text.

Router behavior is undefined if a Connection Manager does not follow the requirements in a Connection Manager Note.

See the Connection Manager Guide for an example Connection Manager implementation.

1.6.9 **Pseudocode**

Throughout this specification, pseudocode is used to illustrate operating principles. Comments are demarcated by double forward slashes (“//”). The pseudocode conventions include:

```
if/else conditions;
    if (condition)
        // true operations
    else
        // false operations

for loops:
    for (conditions)
        // operations
```

1.6.10 CRC Algorithms

The CRC algorithms used in the specification are defined according to the following values:

- **WIDTH:** This is the width of the algorithm expressed in bits. This is one less than the width of the POLY.
- **POLY:** This is a binary value that should be specified in a hexadecimal number. The top bit of the polynomial should be omitted. For example, if the polynomial is 10110, the POLY equals 06. An important aspect of this parameter is that it represents the un-reflected polynomial; the bottom bit of this parameter is always the LSB of the divisor during the division regardless of whether the algorithm being modeled is reflected.
- **INIT:** This parameter specifies the initial value of the register when the algorithm starts. This is the value that is to be assigned to the register in the direct table algorithm. In the table algorithm, we may think of the register always commencing with the value zero, and this value being XORed into the register after the Nth bit iteration. This parameter should be specified as a hexadecimal number.
- **REFIN:** This is a Boolean parameter. If it is FALSE, inputs bytes are processed with bit 7 being treated as the most significant bit (MSB) and bit 0 being treated as the least significant bit. If this parameter is TRUE, each byte is reflected before being processed.
- **REFOUT:** This is a Boolean parameter. If it is set to FALSE, the final value in the register is fed into the XOROUT stage directly, otherwise, if this parameter is TRUE, the final register value is reflected first.
- **XOROUT:** This is a W-bit value that should be specified as a hexadecimal number. It is XORed to the final register value (after the REFOUT) stage before the value is returned as the official checksum.

For more information, see *A Painless Guide to CRC Error Detection Algorithms*, by Ross N. Williams, Version 3, August 19, 1993.

1.6.11 FourCC

A FourCC is a sequence of four bytes used to represent ASCII strings. It is limited to ASCII printable characters (one byte per character), with space characters reserved for padding shorter sequences.

For example, the FourCC string “ABC” is represented by a hexadecimal value of 20434241h, where the rightmost byte (41h) represents the first ASCII character (“A”).

1.7 Reserved Values and Fields

Unless otherwise specified, fields and values marked “Rsvd” shall be handled as described in Table 1-1.

Table 1-1. Rsvd Value and Field Handling

| Type | Handling |
|--------------------|--|
| Packet Values | A transmitter shall not use a value in this specification that is marked as “Rsvd”. The target of a packet shall ignore a packet that has any of its defined fields set to an Rsvd value and proceed as if the packet was never received. |
| Transaction Values | A transmitter shall not use a value in this specification that is marked as “Rsvd”. The target of a Transaction shall ignore a Transaction that has any of its defined fields set to an Rsvd value and proceed as if the Transaction was never received. |
| Ordered Set Values | A transmitter shall not use a value in this specification that is marked as “Rsvd”. The target of an Ordered Set shall ignore an Ordered Set that has any of its defined fields set to an Rsvd value and proceed as if the Ordered Set was never received. |
| Register Values | The Connection Manager shall not write a register with a value that is marked as “Rsvd”. Writing a register with a value that is marked as “Rsvd” results in undefined behavior. |

| Type | Handling |
|----------------------------------|---|
| Packet Fields | A transmitter shall set a field that is marked "Rsvd" to zero. A receiver shall ignore any fields that are marked "Rsvd". |
| Transaction Fields | A transmitter shall set a field that is marked "Rsvd" to zero. A receiver shall ignore any fields that are marked "Rsvd". |
| Register Fields | See Section 8.1. |
| Host Interface Descriptor Fields | A Host shall only write a zero value to a field that is marked "Rsvd". A Router may write any value in a field marked "Rsvd". See Section 12.6.1. |

1.8 Terms and Abbreviations

This section lists and defines the terms and abbreviations used throughout this specification. Note that terms and abbreviations not defined here, use their generally accepted or dictionary meaning.

| Term/Abbreviation | Description |
|-------------------------------|--|
| Active Cable | A cable with additional electronics to condition the data path signals. |
| Adapter | An addressable Router interface that includes additional functionality based on type. There are three types of Adapters: Lane Adapter, Protocol Adapter, and Control Adapter. |
| Adapter Configuration Space | The Configuration Space for Lane Adapters and Protocol Adapters. Control Adapters do not have an Adapter Configuration Space. |
| AT Transaction | A type of Transaction that is used by a Router to read from or write to the SB Register Space of a Link Partner. |
| Big Endian | A method of storing data that places the most significant byte of multiple-byte values at a lower storage address. For example, a 16-bit integer stored in Big Endian format places the least significant byte at the higher address and the most significant byte at the lower address. See also Little Endian. |
| Bit Error Rate (BER) | The ratio between the number of incorrect bits received to the total number of received bits. |
| Byte | A data element that is 8 bits in size. |
| Cable Re-timer | A Re-timer that is part of an Active Cable. |
| Captive Device | A Captive Device is device that is plugged directly to a USB Type-C Receptacle. A Captive Device can include a cable at its front-end as an integral part, providing access to a single test point at the connector. |
| CLx | Refers to any of CL0s, CL1, and/or CL2 Adapter states. |
| Configuration Layer | The protocol stack layer responsible for configuration tasks and handling Control Packets. |
| Configuration Space | An address space used by the Connection Manager to access a set of configuration registers. The USB4 architecture defines four Configuration Spaces: Router Configuration Space, Adapter Configuration Space, Path Configuration Space, and Counters Configuration Space. |
| Connection Manager (CM) | The software-based configuration entity that is responsible for managing a Domain. |
| Control Adapter | The Adapter within a Router that handles Control Packets and Router configuration tasks. The Control Adapter implements a Transport Layer and a Control Layer. The Control Adapter is assigned Adapter Number = 0. |
| Control Packet | A type of Transport Layer Packet that is used for configuration-related communication between a Connection Manager and a Router. |
| Counters Configuration Space | An optional Configuration Space that collects statistics within an Adapter. |
| Credit | The unit used to track receive buffer space. |
| Cyclic Redundancy Check (CRC) | A check performed on data to see if an error has occurred in transmitting, reading, or writing the data. The result of a CRC is typically stored or transmitted with the checked data. The stored or transmitted result is compared to a CRC calculated from the data to determine if an error has occurred. |

| Term/Abbreviation | Description |
|------------------------------|---|
| Crosstalk | A situation arising from one transmitted signal inadvertently causing an unwanted effect on another signal. Minimizing or (optimally) preventing cross-talk is a major consideration in maintaining signal integrity. |
| Data Dependent Jitter | A specific class of timing jitter. It is a form of deterministic jitter which is correlated with the sequence of bits in the data stream. It is also a form of Inter-Symbol Interference. |
| dBm | An abbreviation for the power ratio in decibels (dB) of the measured power referenced to one milliwatt (mW). |
| dequeue | The process of removing a packet from a buffer so that the buffer is free to be used for other packets. |
| device | A logical or physical entity that performs one or more functions. The entity described depends on the context of the reference. For example, "device" can refer to a single component such as a Router or a Re-timer, or it can refer to a collection of components that perform a particular function such as a USB4 Device. |
| Device Router | The Router in a USB4 hub or USB4 peripheral device. A Device Router has an Upstream Facing Port and optionally one or more Downstream Facing Ports. |
| De-Scrambling | Restoring a scrambled bit stream to its original state. See also Scrambling. |
| De-Skew | The procedure of aligning Symbols received across multiple Lanes of a Dual-Lane Link. |
| Destination Adapter | A Protocol Adapter that is the final recipient of Tunneled Packets routed through the USB4 Fabric. Note that a Protocol Adapter may be a Source Adapter on one Path and a Destination Adapter on another. |
| Domain | A collection of interconnected Routers managed by a single Connection Manager. |
| Doubleword, DW | 4 bytes. |
| Down-spread | Spreading a clock frequency downward from a peak frequency. As compared to "center-spread" which evenly splits the frequency around the center frequency. |
| Downstream | The direction of data flow away from the Connection Manager. |
| Downstream Facing Port (DFP) | A USB4 Port that is not an Upstream Facing Port. <i>Note: The DFP is not necessarily the power provider. Power roles are determined by USB PD negotiation as defined in the USB PD Specification.</i> |
| DP | DisplayPort™. |
| DPRX | DisplayPort receiver. |
| DPTX | DisplayPort transmitter. |
| Dual-Lane Link | A USB4 Link that uses two bonded Lanes to transmit and receive data. |
| Egress Adapter | An Adapter that transmits outgoing traffic. |
| Electrical Layer | The protocol stack layer that directly interacts with the communication medium between two Routers. The Electrical Layer decouples data transmission electrical specifications from the Logical Layer. |
| EMI | Electromagnetic interference. |
| FEC | Forward Error Correction. |
| Gen 2 | Refers to speeds of 10 Gbps (USB4) and/or 10.3125 Gbps (TBT3-Compatibility Mode). |
| Gen 3 | Refers to speeds of 20 Gbps (USB4) and/or 20.625 Gbps (TBT3-Compatibility Mode). |
| Host Router Clock | The clock that is the ultimate source of time for clock synchronization, either within a single Domain or between multiple interconnected Domains. |
| Host Router Time | The time as captured from the Host Router Clock. |
| HopID | A number assigned to a Transport Layer Packet to identify its Path in the context of a Link. |

| Term/Abbreviation | Description |
|------------------------|---|
| Host | The host computer system that interfaces with a Host Router. A Host includes the host hardware platform (CPU, bus, etc.) and the operating system in use. |
| Host Interface | The interface between a Host and a Host Router. |
| Host Router | The Router in a USB4 host. A Host Router has one or more Downstream Facing Ports. It does not have an Upstream Facing Port (the Host Interface Adapter serves as the Upstream Adapter). |
| Ingress Adapter | An Adapter that receives incoming traffic. |
| Inter-Domain Link | A USB4 Link that connects two Routers belonging to different Domains. |
| Lane | The dual simplex high speed differential signaling pair that provides communication between two interconnected Routers. The signaling rate at which a Lane operates defines the speed of communication for that Lane. |
| Lane 0 | The Lane associated with the Lane 0 Adapter. |
| Lane 1 | The Lane associated with the Lane 1 Adapter. |
| Lane Adapter | The Adapter that interfaces to a Lane. A USB4 Port contains two Lane Adapters: a Lane 0 Adapter and a Lane 1 Adapter. A Lane Adapter implements an Electrical Layer, a Logical Layer, and a Transport Layer. |
| Lane 0 Adapter | The lowest numbered Adapter in a USB4 Port. For example, in a USB4 Port containing Adapters 3 and 4, Adapter 3 is the Lane 0 Adapter. |
| Lane 1 Adapter | The highest numbered Adapter in a USB4 Port. For example, in a USB4 Port containing Adapters 3 and 4, Adapter 4 is the Lane 1 Adapter. |
| Lane Bonding | The process of turning two Single-Lane Links into a Dual-Lane Link. |
| Link | See USB4 Link. |
| Link Management Packet | A type of Transport Layer Packet that is used to communicate Link-specific information between two Link Partners. |
| Link Partner | The relationship between the two USB4 Ports on either end of a Link. For example, given two USB4 Ports (Port A and Port B) that are connected by a Link, Port B is the Link Partner of Port A and vice versa. |
| Little Endian | Method of storing data that places the least significant byte of multiple-byte values at lower storage addresses. For example, a 16-bit integer stored in Little Endian format places the least significant byte at the lower address and the most significant byte at the next address. See also Big Endian. |
| Local Clock | The free-running clock within a Router. The Local Clock is used for time synchronization purposes. |
| Local Time | The time as captured from the Local Clock. |
| Logical Layer | The protocol stack layer below the Transport Layer and on top of the Electrical Layer. The Logical Layer performs Link management tasks and transmission/reception of Symbols over a Link. |
| LTPR | LT-tunable PHY Repeater. Defined by the DisplayPort 1.4a Specification. |
| LT Transaction | A Type of Transaction that is used by a Router during USB4 Link initialization and to signal a change in Adapter state. |
| MFD | Multi-Function DP. A DisplayPort link which operates as part of a Multi-Function configuration. |
| Multi-Function | Configuration on the USB Type-C connector that supports USB 3.2 SuperSpeed signaling at the same time as DP Standard signaling. |
| NVM | Non Volatile Memory. |
| On-Board Re-timer | A Re-timer that is located between a Router and a USB Type-C connector. |
| Operation | An action initiated by a Connection Manager using the interface described in Section 8.3. There are two types of Operations: Router Operations and Port Operations. <i>Note: When not capitalized, "operation" has its ordinary meaning.</i> |
| Passive Cable | A cable that does not incorporate any electronics to condition the data path signals. |

| Term/Abbreviation | Description |
|---------------------------------|--|
| Path | The one-way path that carries data either from a Source Adapter to a Destination Adapter or between a Connection Manager and a Control Adapter. A Path traverses one or more Routers, which route traffic along the Path as configured by the Connection Manager. |
| Path 0 | The Path that corresponds to HopID 0. Path 0 is used exclusively for Control Packets. |
| Path Configuration Space | The Configuration Space for a Path. Each Path has its own Path Configuration registers. |
| PCB | Printed Circuit Board. |
| PCIe | PCI Express. |
| PCIe Host Interface | A Host Interface that is run over a PCIe bus. |
| PLL | Phase Locked Loop. |
| Platform Integrated Host Router | Host Router silicon that is integrated into the platform, either as part of the platform SoC or as a discrete chip on the motherboard. |
| Port Operation | A type of Operation that targets individual USB4 Port functionality (e.g. compliance tests and receiver Lane margining tests). |
| PRBS | Pseudo-Random Bit Sequence. |
| Protocol Adapter | The Adapter that interfaces to a protocol-specific entity. A Protocol Adapter implements a Protocol Adapter Layer and a Transport Layer. |
| Protocol Adapter Layer | The protocol stack layer above the Transport Layer. The Protocol Adapter Layer encapsulates Tunneled Protocol traffic into Tunneled Packets and hands them off to the Transport Layer. It also receives packets from the Transport Layer and extracts the Tunneled Protocol traffic. |
| Quality of Service (QoS) | About or relating to the specific bandwidth and latency requirements for a particular protocol or use case. |
| Receiver | The receiver in a Lane Adapter. |
| Re-timer Index | A one-based index number that uniquely identifies a Re-timer on a Link between two Routers. The Re-timer Index corresponds to the distance between a Re-timer and a Router where Re-timer index = (number of Re-timers between the Router and Re-timer) + 1. A Re-timer has two Re-timer Indexes – one for each Router on the Link. See Figure 4-3. |
| Route String | A string that represents the route between the Host Router and another Router. |
| Router | A component that manages and routes traffic through a USB4 Fabric and ensures time synchronization within the USB4 Fabric. |
| Router Assembly | A Router and any Re-timers between the USB4 Ports of the Router and their USB Type-C connectors. |
| Router Configuration Space | The Configuration Space for a Router. Each Router has a Router Configuration Space. |
| Router Operation | A type of Operation that targets Router-wide functionality (e.g. NVM update and reading Router capabilities). |
| RT Transaction | A type of Transaction that is used by a Router to communicate with a Re-timer. A Re-timer also uses RT Transactions to communicate with another Re-timer or Router. |
| Scrambling | The process of changing a bit stream in a pseudo-random way. See also De-Scrambling. |
| Sideband (SB) Channel | The connection between the USB4 Ports of two interconnected Routers that provides initial communication and setup functionality. |
| Sideband (SB) Registers | A set of registers that are used for Link setup and configuration over the Sideband Channel. A Router uses Transactions to read from and write to the Sideband Registers. |
| Single-Lane Link | A Link that uses only one Lane to transmit and receive data. |
| Skew | The presence of misalignment between Symbols across the Lanes of a Dual-Lane Link. |

| Term/Abbreviation | Description |
|------------------------------------|---|
| Source Adapter | A Protocol Adapter that is the source of Tunneled Packets routed through the USB4 Fabric. Note that a Protocol Adapter may be a Source Adapter on one Path and a Destination Adapter on another. |
| Spanning Tree | A loop-free Tree Topology used by a Connection Manager to manage a Domain. |
| Spread Spectrum Clock (SSC) | A method of varying the frequency of the clock over time to reduce EMI. |
| Standalone AIC Host Router | A Host Router that is added to the platform as an Add-In Card (AIC). |
| Symbol Time | The time it takes to transmit a 66b Symbol at a Gen 2 speed or a 132b Symbol at a Gen 3 speed. |
| TBT3 | Thunderbolt™ 3 |
| TBT3-Compatible | Able to interoperate with TBT3 products. |
| TBT3-Compatible Active Cable | A USB4 Active Cable that can interoperate with TBT3 and TBT3-Compatible Routers, Re-timers, and Connection Managers |
| TBT3-Compatible Connection Manager | A USB4 Connection Manager that can interoperate with TBT3 and TBT3-Compatible Routers, Re-timers, and Cables. |
| TBT3-Compatible Router | A Router that implements Chapter 13 of this specification. |
| Time Management Unit (TMU) | The functional block in each Router that is used to distribute and synchronize time throughout the USB4 Fabric. |
| TopologyID | A unique topological address representing the position of a Router within a Domain. |
| Transaction | The unit of communication across the Sideband Channel. A Transaction consists of a set of defined symbols. There are three types of Transactions: LT Transactions, AT Transactions, and RT Transactions. |
| Transport Layer | The protocol stack layer that is responsible for routing traffic through the USB4 Fabric and ensuring flow control and data integrity. |
| Transport Layer Packet (Packet) | A type of packet that either originates from or is routed by the Transport Layer. |
| Transmitter | Refers to a USB4 Transmitter. |
| Tree Topology | The type of network topology in which a central 'root' node (the top level of the hierarchy) is connected to one or more other nodes that are one level lower in the hierarchy with a point-to-point link. There may be more levels of hierarchy but each is connected to the level above via a point-to-point link. |
| Tunneled Packet | A type of Transport Layer Packet that is used to carry Tunneled Protocol traffic through a USB4 Fabric. |
| Tunneled Protocol | An I/O protocol tunneled through a USB4 Fabric. |
| TXFFE | Transmitter Feed Forward Equalization. |
| Unit Interval (UI) | Given a data stream of a repeating pattern of alternating 1 and 0 values, the Unit Interval is the value measured by averaging the time between voltage transitions, over a time interval long enough to make all intentional frequency modulation of the source clock negligible. |
| Upstream | The direction of data flow towards the Connection Manager. |
| Upstream Adapter | The Adapter that a Connection Manager uses to address a Router. |
| Upstream Facing Port (UFP) | The USB4 Port that faces the Connection Manager in the Spanning Tree topology of the Domain. |
| USB | Universal Serial Bus. |
| USB3 | USB 3.2 Specification. |
| USB4-Based Dock | A USB4-based dock product combines a USB4 hub (including at least one exposed USB Type-C downstream port) with additional capabilities that either exposes other connector types and/or includes other user-visible functions, e.g. storage, networking, etc. Example of functions that are not considered user-visible include firmware update or device authentication. |
| USB4 Fabric | One or more interconnected Domains. |

| Term/Abbreviation | Description |
|--------------------------|---|
| USB4 Link | The logical connection between the USB4 Ports of two interconnected Routers. A USB4 Link may include either one or two Lanes. Unless otherwise noted, the term "Link" without any preceding descriptor refers to a USB4 Link. |
| USB4 Port | A Router interface consisting of two Lane Adapters and a Sideband Channel. Protocol Adapters and Control Adapters are never part of a USB4 Port. |
| Word | 2 bytes. |

2 Architectural Overview

This chapter presents an overview of Universal Serial Bus 4 (USB4™) architecture and key concepts. USB4 is similar to earlier versions of USB in that it is a cable bus supporting data exchange between a host computer and a wide range of simultaneously accessible peripherals. However, USB4 also allows a host computer to setup data exchange between compatible peripherals. The attached peripherals share bandwidth as configured by the host computer. The bus allows peripherals to be attached, configured, used, and detached while the host and other peripherals are in operation.

When configured over a USB Type-C® connector interface, USB4 functionally replaces USB 3.2 while retaining USB 2.0 bus operating in parallel. Enhanced SuperSpeed USB, as defined in USB 3.2, remains the fundamental architecture for USB data transfer on a USB4 Fabric. The difference with USB4 versus USB 3.2 is that USB4 is a connection-oriented, tunneling architecture designed to combine multiple protocols onto a single physical interface, so that the total speed and performance of the USB4 Fabric can be dynamically shared. USB4 allows for USB data transfers to operate in parallel with other independent protocols specific to display, load/store and host-to-host interfaces. Additionally, USB4 extends performance beyond the 20 Gbps (Gen 2 x 2) of USB 3.2 to 40 Gbps (Gen 3 x 2) over the same dual-lane, dual-simplex architecture.

This specification introduces the concept of protocol tunneling to USB bus architecture. Besides tunneling Enhanced SuperSpeed USB (USB3), display tunneling based on DisplayPort (DP) protocol and load/store tunneling based on PCI Express (PCIe) are defined. These protocol tunnels operate independently over the USB4 transport and physical layers. Additionally, USB4 allocates packets for bus configuration and management, and packets can be allocated specifically for host-to-host data connections.


2.1 USB4 System Description

Figure 2-1 illustrates the dual bus architecture of USB 3.2 as augmented by USB4. As architected, backward compatibility is supported with minimum interoperability starting at USB 2.0, working up through USB 3.2, and finally up to USB4 based on the highest common bus level supported across the interconnected components.

Protocol tunnels are interfaced via Protocol Adapters specific to each protocol. For USB and PCIe protocols, native USB hubs and PCIe switches are required to handle protocol-related packet routing and buffering. For display tunneling, no intermediate DP-specific logic is required with display tunnels being established as an end-to-end link.

Time synchronization across a USB4 Fabric uses distributed time management units (TMUs) associated with each Router.

Figure 2-1. USB4/USB3.2 Dual Bus System Architecture



2.1.1 Architectural Constructs

The following summarizes basic constructs in the USB4 architecture.

2.1.1.1 Routers

The Router is a fundamental building block of the USB4 architecture. A Router maps Tunneled Protocol traffic to USB4 packets and routes packets through the USB4 Fabric. A Router also distributes and synchronizes time throughout the USB4 Fabric via its Time Management Unit (TMU).

A Router is discovered and configured by a Connection Manager located within the USB4 host. The concept of a Router is not to be confused with a USB4 hub. The Router includes a flat point-to-point, configurable switch necessary to create the internal paths between Adapters. One Router typically exists within each instance of a USB4 host, hub or device.

There are two types of Routers: Host Routers and Device Routers.

2.1.1.2 Adapters

Each Router contains up to 64 Adapters. Adapters provide an interface between a Router and an external entity. There are three types of Adapters: Protocol Adapters, Lane Adapters, and Control Adapters.

2.1.1.2.1 Protocol Adapter

A Protocol Adapter is used to translate between a supported native protocol and a USB4 tunnel. There are four types of Protocol Adapters: USB3 Adapters, DisplayPort (DP) Adapters, PCIe Adapters, and Host Interface (HI) Adapters.

2.1.1.2.2 Lane Adapter

A Lane Adapter provides an interface for a Lane. A USB4 Port has one Lane Adapter per Lane.

2.1.1.2.3 Control Adapter

A Router contains one Control Adapter. The Control Adapter is a logical Adapter and does not have a physical layer. The Control Adapter is the final consumer of Control Packets that are targeted to the Router. The Control Adapter also generates Control Packets, which are sent to the Connection Manager.

2.1.1.3 USB4 Ports and Links

A USB4 Port is the entity that provides the USB4 functional interface that resides on each end of a USB4 Link. It consists of the transmit and receive Lanes of the USB4 data bus along with a two-wire Sideband (SB) Channel (SBTX/SBRX). USB4 Ports operate as either a Single-Lane Link or Dual-Lane Link. When operating as a Single-Lane Link, Lane 1 of the USB4 Port is disabled. When operating as a Dual-Lane Link, Lanes 0 and 1 are logically bonded together to provide a single data channel. Figure 2-2 shows a Single-Lane Link. Figure 2-3 shows a Dual-Lane Link.

Figure 2-2. Single-Lane USB4 Link





Figure 2-3. Dual-Lane USB4 Link

The primary communication channel of the USB4 Fabric is over the USB4 Link that interconnects two USB4 Ports. The USB4 Link transports packets for both Tunneled Protocol traffic and bus management traffic between Routers. The Sideband Channel of a USB4 Port is used to initialize and manage the USB4 Link between the USB4 Ports.

For a USB4-enabled USB Type-C port, the complete interface includes a USB4 Port, a USB 2.0 data bus, and the USB Type-C Configuration Channel (CC) along with power/ground (VBUS, VCONN and GND).

2.1.1.4 USB4 Devices

2.1.1.4.1 USB4 Peripheral Device

A USB4 peripheral device has a single Upstream Facing Port. It does not have any Downstream Facing Ports.

A USB4 peripheral device contains a Device Router and can also optionally contain one or more of the following:

- An Enhanced SuperSpeed Function.
- A USB 2.0 Function.
- A PCIe Function.
- A DisplayPort Source or Sink.

A USB4 peripheral device supports 20G USB4 operation (Gen2x2) and optionally 40G USB4 operation (Gen3x2).

2.1.1.4.2 USB4 Hub

At a high level, a USB4 hub is functionally similar to a USB 3.2 hub – it consists of one Upstream Facing Port and one or more Downstream Facing Ports. USB4 hub functionally operates as a tree-like structure for enabling one or more Downstream Facing Ports to be served by one Upstream Facing Port, typically for the purpose of port expansion.

In addition to the USB4 specific hub functionality, USB 3.2 and USB 2.0 hub functionality is supported such that Downstream Facing Ports of a USB4 hub can support backward-compatibility with USB 3.2 and USB 2.0 devices.

A USB4 hub contains:

- A Device Router.
- An Enhanced SuperSpeed USB Hub.
- A PCIe Switch.
- A USB 2.0 Hub.


A USB4 hub supports 20G USB4 operation (Gen2x2) and 40G USB4 operation (Gen3x2).

A USB4 hub is required to support DisplayPort Alt Mode on all of its DFP. See the USB Type-C Specification for a full definition of the requirements and details regarding DisplayPort Alt Mode support on a DFP.

2.1.1.4.3 USB4-Based Dock

Figure 2-4 illustrates a USB4-Based Dock with two Downstream Facing Ports and peripheral device functions for each of the available protocols.

Figure 2-4. Example of a USB4-Based Dock



The requirements for a USB4 hub apply to a USB4-Based Dock.

2.1.1.5 USB4 Host

A USB4 host contains:

- A Host Router.
- A USB 2.0 Host.
- An Enhanced SuperSpeed Host.
- A DisplayPort Source.

A USB4 host can also optionally contain a PCIe controller.

A USB4 host supports 20G USB4 operation (Gen2x2) and optionally 40G USB4 operation (Gen3x2).

A USB4 host is required to support DisplayPort Alt Mode on all of its DFP. See the USB Type-C Specification for a full definition of the requirements and details regarding DisplayPort Alt Mode support on a DFP.

2.1.1.6 Re-timers

A USB4 product can contain up to two On-Board Re-timers per USB4 Port. USB4 Re-timers are described in the USB4 Re-Timer Specification.

2.1.1.7 Connection Manager

The Connection Manager is the software entity that is responsible for enumerating, configuring, and managing a Domain. The Connection Manager performs tasks such as Path setup/teardown, Hot plug/unplug, and bandwidth management. The Connection Manager is part of the USB4 host system.

2.1.2 USB4 Mechanical

The electro-mechanical specifications for USB cables and connector assemblies that support USB4 are documented by the USB Type-C Specification.

USB4 hubs and devices have an upstream connection. USB4 hosts and hubs have one or more downstream connections. For USB Type-C connectors, upstream and downstream behaviors are established using the configuration features of the USB Type-C functional architecture.

2.1.3 USB4 Power

Power for USB4 operation is established and managed as defined in the USB Type-C Specification and the USB PD Specification. Unlike USB 2.0 and USB 3.2, USB4 does not define its own VBUS-based power model.

2.1.4 USB4 System Configuration

USB4 connections between hosts, hubs and devices are established and managed as defined in the USB Type-C Specification. Once a USB4 connection is established, the USB4 host configures and manages USB4 operation using a software-based Connection Manager.

2.1.5 Thunderbolt™ 3 (TBT3) Compatibility Support

A USB4 host or USB4 peripheral device can optionally support interoperability with Thunderbolt 3 (TBT3) products.

A USB4 hub is required to support interoperability with Thunderbolt 3 products on all of its DFP. A USB4-Based Dock is required to support interoperability with Thunderbolt 3 products on its UFP in addition to all of its DFP.

When interoperating with a TBT3 product, Thunderbolt Alt Mode is established on the link between products. The USB Type-C Specification describes how a USB4 product negotiates and enters Thunderbolt Alt Mode.

Chapter 13 describes the additional requirements for a Router that supports TBT3-compatible interoperability.

2.1.6 USB Type-C Alternate Mode Compatibility Support

A USB4 product can optionally support interoperability with USB Type-C Alternate Modes as defined by the USB Type-C Specification.

2.2 USB4 Fabric Architecture

The USB4 Fabric is designed to meet the needs of multiple transport protocols. Its main features are:

- Signaling rates that support high throughput interconnects
 - 10 Gbps (for Gen 2) and 20 Gbps (for Gen 3).
 - Optional support for Thunderbolt 3-compatible rates of 10.3125 Gbps (for Gen 2) and 20.625 Gbps (for Gen 3).
- Hop-by-hop, credit-based flow control.

- Bandwidth management and prioritization.
- A programming model that allows a Connection Manager to initialize and manage a USB4 Domain in a manner that is transparent to the Tunneled Protocols and their software.
- A time synchronization protocol to synchronize real-time clocks across one or more USB4 Domains.
- Error detection, correction, and recovery.
- Link-level Power Management.

2.2.1 USB4 Functional Stack

Figure 2-5 shows the USB4 functional stack layers.

Figure 2-5. USB4 Functional Stack Layers





Figure 2-6 identifies which layers are implemented by USB4 Ports (Lane Adapter), Protocol Adapters and the Control Adapter.

Figure 2-6. USB4 Port (Lane Adapter), Protocol Adapter and Control Adapter across Functional Layers



2.2.1.1 Electrical Layer

The Electrical Layer defines electrical signaling characteristics of a USB4 Link including scrambling, encoding, jitter, and voltage.

2.2.1.2 Logical Layer

The Logical Layer establishes a USB4 Link between two Routers and provides services to transmit and receive streams of bytes between them.

The Logical Layer resides on top of the Electrical Layer and below the Transport Layer. It treats the traffic to and from the Transport Layer as a byte stream.

The services provided by the Logical Layer are:

- Establishment and maintenance of a USB4 Link with a Link Partner.
- Performance scalability via different speeds and widths.
- Error detection and recovery mechanisms.
- Operation with different media such as passive cable, active cable, and Re-timers.
- Support for mechanisms such as clocking compensation, data scrambling, and Forward-error-correcting codes.
- Power management.

A USB4 Link is assisted and managed by a companion Sideband Channel that:

- Configures parameters of the USB4 Link.
- Interacts with Re-timers (if present) and performs USB4 Link TxFFE handshake.
- Ensures a correct power down/wake up sequence of the USB4 Link transceivers and Re-timers.

2.2.1.3 Transport Layer

The Transport Layer forwards Tunneled Packets and Control Packets through the bus. It defines packet format, routing, Quality-of-Service (QoS) support, flow control, and time synchronization. The Transport Layer is where protocol MUXing is performed.

2.2.1.4 Configuration Layer

The Configuration Layer performs Router configuration tasks and handles incoming Control Packets. The Configuration Layer provides an addressing scheme for Control Packets within the Domain, processes Control Packets, and delivers a reliable transport mechanism for Control Packets.

Control Packets provide the Connection Manager with access to the Configuration Spaces of a Router.

2.2.1.5 Protocol Adapter Layer

The Protocol Adapter Layer performs mapping between Tunneled Protocol traffic and USB4 Transport Layer Packets. A Protocol Adapter Layer is defined by the type of Tunneled Protocol traffic it sends and receives. See Section 2.2.10 for more details of the different types of Protocol Adapter Layers.

2.2.2 USB4 Fabric Topology

The physical topology of a USB4 Fabric is an interconnected graph. The physical topology will typically resemble a tree with a USB4 host at the root and a set of USB4 hubs and/or devices connected downstream, in series and/or in parallel. However, loops can occur if there are multiple connections to a USB4 host. If the Connection Manager detects a loop in the physical topology of its Domain, it uses a subset of the interconnected graph in the form of a Spanning Tree, which removes any loops. If there are no loops in the physical topology, the Spanning Tree is the same as the physical tree. See the Connection Manager Guide for more information on how a Connection Manager detects and handles loops.

The Host Router is at the top of the Spanning Tree. The Spanning tree can have up to six levels, meaning that a Device Router can be up to 5 hops away from the Host Router along the Spanning Tree.

The Connection Manager accesses a Domain through the Host Router. Control Packets are only routed along Links that belong to the Spanning Tree. Tunneled Protocol traffic is not limited to the Spanning Tree.

Figure 2-7 shows an example of a Spanning Tree for a USB4 Fabric with no loops in its physical topology. Figure 2-8 shows an example of a Spanning Tree for a USB4 Fabric with a loop in its physical topology.

Figure 2-7. Example USB4 Physical Topology (No Loop) and Spanning Tree





Figure 2-8. Example USB4 Physical Topology (with Loop) and Spanning Tree




2.2.3 Paths

A Path represents a directional logical connection over the USB4 Fabric. A Path is defined either between two Protocol Adapters or between the Connection Manager and a Control Adapter.

A Path can be end-to-end (as is the case for Display and Host-to-Host Tunneling) or it can traverse a single USB4 Link (as is the case for USB3 tunneling and PCIe Tunneling). Figure 2-9 shows an example of a set of Paths that tunnel USB3 traffic between a USB4 host and a USB4 device.

Figure 2-9. Paths across a USB4 Fabric




A Tunnel is configured end-to-end before data transfer can take place.

Each packet carries an identifier (called a HopID) that identifies the Path in the context of a given Link. The HopID is configured by a Connection Manager and can vary across each Link in the Path. When a Router receives a packet, it uses the HopID to determine the packet's next HopID and which Egress Adapter to forward the packet to.

2.2.4 Communication Constructs

Figure 2-10 shows the various constructs that are used for communication over the USB4 Fabric.

Figure 2-10. USB4 Communication by Functional Layer



2.2.4.1 USB4 Link



The USB4 Link uses the packets and Ordered Sets described in this section.

2.2.4.1.1 Control Packets

Control Packets are used by a Connection Manager to configure and manage the Routers across the bus. They are also used by a Router to communicate with a Connection Manager. Control Packets are routed over the bus based on a Route String that identifies a Router's position in the Spanning Tree. When a Control Packet originates from the Connection Manager, the Route String identifies the Router that the packet is targeted to. When a Control Packet originates from a Router, the Router String identifies the Router that sent the packet.

Figure 2-11 shows an example of a Control Packet that traverses several Routers. The Control Adapter in a non-target Router forwards the packet to a USB4 Port. The Control Adapter of the target Router consumes the Control Packet.

Figure 2-11. Example Control Packet Traversing Several Routers



2.2.4.1.2 Tunneled Packets

Protocol traffic is encapsulated and tunneled over the USB4 Fabric in Tunneled Packets. Tunneled Packets traverse the USB4 Fabric along one or more Paths.

2.2.4.1.3 Link Management Packets

Link Management Packets are confined to a single USB4 Link. Link Management Packets originate in the Transport Layer of the Router at one end of the Link and terminate in the Transport Layer of the Router at the other end of the Link.

The following Link Management Packets are defined:

- Time Sync Packets – Used to synchronize the clocks of the Routers on the bus.
- Flow Control Packets – Used to prevent buffer overflow.
- Idle Packets – Ensure a steady byte stream is fed to the Logical Layer when no other Transport Layer Packets are being transmitted.

2.2.4.1.4 Ordered Sets

The Logical Layer uses Ordered Sets for tasks such as Symbol synchronization, Link training, and de-skew between Lanes. Ordered Sets are 66-bit Symbols (at Gen 2 speed) or 132-bit Symbols (at Gen 3 speed).

2.2.4.2 Sideband Channel

The Sideband Channel handles the following events:

- Lane Initialization.
- Connection or disconnect on a USB4 Port.
- Lane disable or enable.
- Entry or exit from Sleep state.


The Sideband Channel is a packet-based channel. Sideband Channel packets are called Transactions to distinguish them from USB4 Link packets.

Each USB4 Port implements a set of Link configuration registers called Sideband (SB) Register Space. Routers use Transactions to access the SB Register Space of another Router or Re-timer.

2.2.5 USB4 Host-to-Host Communications

Multiple Domains can be connected over a USB4 Fabric under certain conditions. Figure 2-12 illustrates two possible inter-Domain connections. An inter-Domain Link connects two Routers that reside in different Domains.

Figure 2-12. Example USB4 Host-to-Host Connections



The Connection Managers in two connected Domains communicate with each other using Host Interface Tunneled Packets. The most common use of Host-to-Host Tunneling is to allow two Connection Managers to exchange information through their respective Host Interfaces. For example, two USB4 hosts can exchange IP packets over USB4 using Host-to-Host Tunneling.

2.2.6 Programming Model

2.2.6.1 Connection Manager

A Connection Manager is the entity responsible for managing and configuring a Domain. The Connection Manager communicates with Routers in a Domain using Control Packets.

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

- Router initialization.
- Scanning and setting Adapter configurations.
- Setting and removing Paths.

- Configuration of QoS behavior in the USB4 Fabric including flow control and bandwidth allocation.
- Management across Domains.

2.2.6.2 Configuration Spaces

Routers are configured via configuration registers that reside in one of four Configuration Spaces:

- Router Configuration Space – Contains the Router-level configuration parameters. Each Router has a Router Configuration Space.
- Adapter Configuration Space – Contains capabilities, configuration, and error statistics for an Adapter. Protocol Adapters and Lane Adapters have an Adapter Configuration Space. Control Adapters do not have an Adapter Configuration Space.
- Path Configuration Space – Contains information used for Path setup. Protocol Adapters and Lane Adapters have a Path Configuration Space. Control Adapters do not have a Path Configuration Space. A Path Configuration Space includes an entry for each supported Path.
- Counters Configuration Space – Contains performance statistics for a set of selected Paths. Counters Configuration Space is optional. Protocol Adapters and Lane Adapters can optionally have a Counters Configuration Space. Control Adapters do not have a Counters Configuration Space.

2.2.6.3 Operations

A Router supports an interface, which allows a Connection Manager to initiate various Operations. There are two types of Operations: Router Operations and Port Operations.

Router Operations initiate Router-wide tasks such as NVM read/write and DisplayPort resource management. Router Operations are initiated by writing to Router Configuration Space.

Port Operations can be used to do the following:

- Initiate a read or write to the SB Register space of a Router, Link Partner, or a Re-timer in the Link between a Router and its Link Partner.
- Initiate Port-level tasks such as compliance tests and receiver Lane margining tests.

Port Operations are initiated by writing to the Port Capability in Adapter Configuration Space.

2.2.7 Time Synchronization

A USB4 Fabric provides time synchronization services between the Routers on the bus. Time synchronization is provided both within a single Domain and across multiple Domains.

Each Router tracks its time and frequency shift relative to a Host Router Clock on the bus. A Time Management Unit (TMU) in each Router is responsible for time synchronization operations within the Router.

2.2.8 USB4 Fabric Data Integrity

The USB4 Fabric defines the following data Integrity mechanisms to ensure that channel errors are detected and possibly corrected:

- Cyclic Redundancy Codes (CRC) are used to identify and correct packet errors. CRC are in the payload of AT Transactions and RT Transactions, in the header of Transport Layer Packets, in the payload of Control Packets, in the payload of Time Sync Packets, and in the payload of DisplayPort AUX Packets.
- Forward Error Correction (FEC) is used in the USB4 Link to reduce the BER of the channel.
- A Connection Manager may retransmit a Control Packet if a response is not received in time.

- A Router can retransmit an AT Transaction or an RT Transaction if a response is not received in time.
- Packets are checked for structure and contents.
- A Link-level flow control mechanism ensures that critical traffic is not lost due to buffer overflow. Link flow control operates independently of any flow control deployed by a Tunneled Protocol. Flow control deployed by a Tunneled Protocol is handled by the respective Tunneled Protocol stack.

2.2.9 Global Life of a Router

The following steps give a high level overview of what happens to a Router from the time it is hot-plugged into a bus until the time it is removed from the bus:

1. Router sets default values into the Configuration Space registers.
2. Router is hot-plugged into a Domain via its Upstream Facing Port.
3. Upstream Facing Port negotiates USB4 as described in the USB PD Specification.
4. Router participates in Lane Initialization to bring the USB4 Links into Active state (Section 4.1.2).
5. Router enables Control Packet routing and scheduling.
6. Connection Manager performs its first access to the Router, enumerating it.
7. Connection Manager configures the Router using Read Requests (Section 6.4.2.3) and Write Requests (Section 6.4.2.5). Router configuration includes setting up any Paths through the Router.
8. Router is ready to route and process Tunneled Protocol traffic.
9. If another Router is plugged into a Downstream Facing Port:
 - Downstream Facing Port negotiates USB4 as described in the USB PD Specification.
 - Router participates in Lane Initialization on the plugged Port.
 - After the USB4 Link is Active, the Router sends a Hot Plug Event Packet to the Connection Manager and waits for a Hot Plug Acknowledgment Packet (Section 6.8.1).
10. If a Downstream Router is unplugged:
 - Router discards any packets that would otherwise be routed to the unplugged Router.
 - Router sends a Hot Plug Event Packet to the Connection Manager and waits for a Hot Plug Acknowledgment Packet (Section 6.8.2).
 - Router inactivates Downstream Link (Section 4.4.5).
11. If the Router is disconnected from bus, it goes through Disconnect (Section 4.4.5).

2.2.10 Protocol Tunneling

A USB4 host supports USB3 Tunneling, DisplayPort Tunneling, and Host-to-Host Tunneling. A USB4 host can also optionally support PCIe Tunneling.

A USB4 hub supports USB3 Tunneling, DisplayPort Tunneling, PCIe Tunneling, and Host-to-Host Tunneling. There are multiple ways that a USB4 hub supports DisplayPort Tunneling:

- The USB4 hub acts as a “pass through” for DisplayPort Tunneling (i.e. the USB4 hub routes tunneled traffic directly between two of its USB4 Ports).
- The USB4 hub contains a DP OUT Adapter that receives Tunneled DisplayPort traffic from a USB4 Port and sends it to a DisplayPort Sink.

- The USB4 hub optionally contains a DP IN Adapter that sends DisplayPort traffic from a DisplayPort Source to a USB4 Port, which transmits Tunneled DisplayPort Traffic.

A USB4 hub acts as a “pass through” for Host-to-Host Tunneling. A USB4 hub does not contain a Host Interface Adapter.

A USB4 device optionally supports USB3 Tunneling, DisplayPort Tunneling, and/or PCIe Tunneling. A USB4 device Router does not support Host-to-Host Tunneling.

2.2.10.1 USB3 Tunneling

A USB4 host supports USB3 tunneling by one or more of the following methods:

- Incorporating an internal host controller (Enhanced SuperSpeed Host depicted in Figure 2-13).
- Implementing other means that meet the USB 3.2 Specification.

Figure 2-13 depicts a USB4 host that supports USB3 Tunneling.

Figure 2-13 Example of a USB4 Host with USB3 Tunneling Highlighted






Figure 2-14 depicts a USB4 hub. A USB4 hub contains an internal Enhanced SuperSpeed Hub. The internal Enhanced SuperSpeed Hub exposes one or more Downstream USB3 ports, which can be connected to a USB Endpoint or Downstream USB3 Protocol Adapter. The upstream port of the internal Enhanced SuperSpeed Hub interfaces with an Upstream USB3 Protocol Adapter that forwards packets to the Upstream Facing Port of the USB4 hub.

Figure 2-14. Example of a USB4 Hub with USB3 Tunneling Highlighted



A USB4 peripheral device can contain an internal Enhanced SuperSpeed Function. Figure 2-15 depicts a USB4 peripheral device with an internal USB peripheral device. The internal Enhanced SuperSpeed Function interfaces with the Router via a USB3 Protocol Adapter.

Figure 2-15. Example of a USB4 Peripheral Device with USB3 Tunneling Highlighted



A USB4 peripheral device can optionally implement multiple internal Enhanced SuperSpeed Functions and/or expose external downstream facing USB3 ports. In that case, an internal Enhanced SuperSpeed Hub can replace the Enhanced SuperSpeed Function shown in Figure 2-15 and a selection of functions would then connect to the downstream ports on that internal hub.

Figure 2-16 shows the USB3 Protocol stack as it applies to USB3 Protocol Adapter. The Link and Protocol Layers in the internal USB hub or the internal USB peripheral device implement a subset of the Link Layer defined in the USB 3.2 Specification.

Figure 2-16. Protocol Stack for USB3 Tunneling




Figure 2-17 shows an example flow that tunnels USB3 traffic over a USB4 Fabric.

Figure 2-17. Example of a USB4 Fabric with USB3 Tunneling





Figure 2-18 shows the protocol stacks traversed along the tunnel between the Enhanced SuperSpeed Host and the Enhanced SuperSpeed function in Figure 2-17.

Figure 2-18. Protocol Stacks along a USB3 Tunnel




2.2.10.2 Display Tunneling

The USB4 Display Tunneling protocol is based on the DisplayPort 1.4a Specification.

There are two types of DisplayPort™ Protocol Adapters: DP IN Protocol Adapters and DP OUT Protocol Adapters. A DP IN Protocol Adapter interfaces with a DisplayPort Source via a DP link. A DP OUT Protocol Adapter interfaces with a DisplayPort Sink via a DP link. A Router may contain one or more DP IN Protocol Adapters, one or more DP OUT Protocol Adapters, or a combination of DP IN and DP OUT Protocol Adapters. The DisplayPort tunnel over the USB4 Fabric has no directionality constraint other than starting at a DP IN Protocol Adapter and ending at a DP OUT Protocol Adapter.

Figure 2-19 shows the general topology of a system that tunnels DisplayPort traffic over the USB4 Fabric.


Figure 2-19. Example Topology for DisplayPort Tunneling



A DP Protocol Adapter can either operate in LTTTPR Non-Transparent mode, LTTTPR Transparent mode, or Non-LTTTPR mode. In LTTTPR Non-Transparent and LTTTPR Transparent modes, the DP OUT Protocol Adapter and DP IN Protocol Adapter behave as a single LTTTPR. In Non-LTTTPR mode, the DP OUT Protocol Adapter and DP IN Protocol Adapter behave such that DPRX appears to be directly attached to DPTX.


Figure 2-20 shows the system in Figure 2-19 from a DisplayPort perspective when the DP IN and DP OUT Protocol Adapters are in Non-Transparent and LTTTPR Transparent modes. Figure 2-21 shows the system in Figure 2-19 from a DisplayPort perspective when the DP IN and DP OUT Protocol Adapters are in Non-LTTTPR Mode.

Figure 2-20. DP IN and OUT Protocol Adapters in LTTTPR Non-Transparent and LTTTPR Transparent Modes



Note: A DP Tunnel appears as a single LTTTPR. While there may be up to 8 LTTTPRs between a DPTX and a DPRX, there can be no more than one DP Tunnel between a DPTX and a DPRX.

Figure 2-21. DP IN and OUT Protocol Adapters in Non-LTTPR Mode




A DP Protocol Adapter supports three Paths:

- An AUX Ingress Path for receiving AUX channel packets.
- An AUX Egress Path for sending AUX channel packets.
- A MAIN-Link Path for sending (in the case of a DP IN Protocol Adapter) or receiving (in the case of a DP OUT Protocol Adapter) Main-Link packets.

Figure 2-22 shows the protocol stacks traversed along a Path between the DP Source and DP Sink in Figure 2-19. The DisplayPort Link Layer, Physical Layer Logical Sub-Block, and Physical Layer Electrical Sub-Block are described in the DisplayPort 1.4a Specification.

Figure 2-22. Protocol Stacks along a DisplayPort Tunneled Path



A Router can incorporate an MST branch splitter in order to provide additional DisplayPort fanout. The methods for incorporating an MST branch splitter are outside the scope of this specification.

2.2.10.3 PCIe Tunneling


PCIe Tunneling is based on the PCI Express (PCIe) specification.

A USB4 host supports PCIe tunneling by any of the following methods:

- Incorporating an internal PCIe Switch.
- Connecting to a PCIe Root Complex via PCIe Root Ports.
- Implementing other means that meet the PCIe Specification.


Figure 2-23 depicts a USB4 host connected to PCIe Root Ports.

Figure 2-23. Example Structure of a USB4 Host with PCIe Tunneling Highlighted



A USB4 hub that tunnels PCIe traffic contains an internal PCIe Switch to forward PCIe traffic to its Downstream Facing Ports. Figure 2-24 depicts a USB4 hub highlighting the PCIe tunneling features. The internal Switch interfaces with the Router via PCIe Protocol Adapters. The PCIe Switch exposes one or more Downstream PCIe Ports, which can be connected to a PCIe Endpoint, PCIe Switch, or Downstream PCIe Protocol Adapter. The upstream port of the internal Switch interfaces with an Upstream PCIe Protocol Adapter that forwards packets to the Upstream Facing Port.

Figure 2-24. Example USB4 Hub with PCIe Tunneling Highlighted



A USB4 device may tunnel PCIe traffic by incorporating an internal PCIe Endpoint or an internal PCIe switch. Figure 2-25 depicts a USB4 device with an internal PCIe Endpoint.

Figure 2-25. Example USB4 Device with PCIe Tunneling Highlighted




Figure 2-26 shows the PCIe Protocol stack as it interfaces with a PCIe Protocol Adapter. The Transaction Layer, Data Link Layer, and Physical Layer Logical sub-block in the internal PCIe Switch or the internal PCIe Endpoint implement a subset of the PCIe Specification.

Figure 2-26. Protocol Stack for PCIe Tunneling




Figure 2-27 shows an example flow that tunnels PCIe traffic over a USB4 Fabric.

Figure 2-27. Example of a USB4 Fabric with PCIe Tunneling





Figure 2-28 shows the protocol stacks along the tunnel in Figure 2-27.

Figure 2-28. Protocol Stacks along a PCIe Tunnel




2.2.10.4 Host Interface Adapter

A Connection Manager interfaces with a Domain through the Host Interface Adapter in a Host Router. A Connection Manager sends Control Packets to and receives Control Packets from the Routers in its Domain via the Host Interface Adapter. A USB4 host can also communicate with another USB4 host via the Host Interface Adapter using USB4 Host-to-Host communication.

Figure 2-29 shows the protocol stacks traversed along a Path between two USB4 hosts.

Figure 2-29. Protocol Stacks along a Path between Hosts




Two Modes of data transfer can be selected for Host-to-Host Tunneling:

- **Raw Mode** – When sending data, a USB4 host posts the payload of a Tunneled Packet into a Host Memory buffer. The Host Interface Adapter Layer fetches the payload, encapsulates it in a Tunneled Packet, and forwards it to the Transport Layer within the Host Router. When receiving data, the Host Interface Adapter Layer posts the payload of received Tunneled Packets into Host Memory buffers.
- **Frame Mode** – When sending data, a USB4 host posts a Frame of up to 4096 bytes into a Host Memory buffer. The Host Interface Adapter Layer fetches the Frame, segments it, and encapsulates each segment into a separate Tunneled Packet. When receiving data, the Host Interface Adapter Layer assembles Tunneled Packets received from the USB4 Fabric into the original Frame and posts the Frame to a Host Memory buffer.

A USB4 host and a Host Interface Adapter interact via Descriptor Rings. Descriptor Rings reside in Host Memory. As depicted in Figure 2-30, a Descriptor Ring is a circular set of structures called Descriptors. Each Descriptor contains a reference to a Data Buffer in Host Memory and additional fields used to manage the transmit or receive flow.

A Host Interface Adapter Layer supports the same number of Transmit Descriptor Rings and Receive Descriptor Rings. Each Transmit Descriptor Ring is assigned to one egress Path. Each Receive Descriptor Ring is assigned to one ingress Path.


Figure 2-30. Descriptor Ring and Data Buffers



The Host Interface Adapter Layer employs an end-to-end (E2E), credit-based flow control mechanism to prevent overflow of a Receive Descriptor Ring. Flow control is managed individually for each Transmit Descriptor Ring together with the Receive Descriptor Ring that is the destination for packets sent from the Transmit Descriptor Ring.

This specification defines in detail a PCIe-based Host Interface Adapter Layer interface that uses PCI device headers, Memory Bars, and PCI interrupts. A Host Router that is not PCIe-based can implement a different interface that provides equivalent functionality.

3 Electrical Layer



This chapter describes the Electrical Layer specifications for a Router Assembly that supports the Gen 2 and Gen 3 modes of operation.

All Router Assemblies shall support Gen 2 speed (10 Gbps per Lane). A Router Assembly that is part of a USB4™ hub shall support Gen 3 speed (20 Gbps per Lane). A Router Assembly that is part of a USB4 host or USB4 device may optionally support Gen 3 speed (20 Gbps per Lane).

The Electrical Layer specifications detail the requirements and testing methodologies required for achieving reliable communication and obtaining interoperability of USB4 interconnects employing Passive or Active Cables. The Electrical Layer shall meet target Bit Error Ratio (BER) of 1E-12 or lower without applying Forward Error Correction. A USB4 Port shall use Spread-Spectrum-Clocking (SSC) to clock a USB4 Link. A USB4 Port employing dual transmitters shall use single clocking for both transmitters.

Note: The clock sources of two Link Partners are asynchronous with one another.

A USB4 Router Assembly is comprised of a Router and up to two Re-timers placed between the Router and the USB Type-C® connector. A USB4 Router Assembly may contain Linear Re-drivers between the Router and a Re-timer or between two Re-timers. If a Router Assembly contains Linear Re-drivers, a Re-timer shall be placed between the Linear Re-driver and the USB Type-C connector in order to meet the requirements in this section.

3.1 Sideband Channel Electrical Specifications

The Sideband Channel shall operate at 1 Mbps rate. A Router Assembly shall meet the Sideband Channel requirements defined in Table 3-1.

Table 3-1. SBTX and SBRX Specifications

| Symbol | Description | Min | Max | Units | Conditions |
|---|-------------------------------------|-------|-------|----------|--|
| SBTX _{VOH} | SBTX High Voltage | 2.4 | 3.47 | Volts | SBTXIOH = -600 μ A (set by 3.4 V/0.5 M Ω). See Note 1. |
| SBTX _{VOL} | SBTX Low Voltage | -0.05 | 0.4 | Volts | SBTXIOL = 600 μ A (set by 3.4 V/ 10 K Ω). See Note 1. |
| SBRX _{VIH} | SBRX High Voltage Detection | 2.0 | 3.72 | Volts | See Note 2, Note 6. |
| SBRX _{VIL} | SBRX Low Voltage Detection | -0.3 | 0.65 | Volts | See Note 7. |
| SBRX _{IHH} | SBRX High input current | -- | 25 | μ A | Vin = VDD. |
| SBRX _{ILL} | SBRX Low input current | -- | 0.4 | μ A | Vin = 0 V. |
| SBX _{TRTF} | SBTX/SBRX 10- 90% Rise/Fall time | 3.5 | -- | ns | Minimum to reduce crosstalk and EMI. See Note 3. |
| SBTX _{PULL_UP_RES} | SBTX pull-up resistor | 7.0K | 10.5K | Ω | See Note 4. |
| SBRX _{PULL_DOWN_RES} | SBRX pull-down resistor | 0.70M | 1.05M | Ω | See Note 5. |
| SBTX _{SOURCE_IMPEDANCE} | SBTX output impedance | 25 | 90 | Ω | |
| SBRX _{Cin} | SBRX Input Capacitance | -- | 8 | pF | |
| SBX _{UI} | UI duration | 980 | 1020 | ns | |
| Notes: | | | | | |
| <ol style="list-style-type: none"> 1. This parameter shall be verified in both transaction and steady state. The steady state condition shall be measured with a continuous high or low level. The transaction state condition shall be measured when sending SBX data. Over/undershoot shall be ignored. 2. A buffer may be used between the connector and the Router to meet these logic levels. When present, this buffer shall meet SBRXVIH and SBRXVIL as defined above. 3. Verify this parameter in transaction and not from power down to power up. The minimum is specified to control crosstalk and EMI. 4. A Router shall terminate the SBTX signal to 3.3 V nominal power. 5. A Router shall terminate the SBRX signal to GND. 6. Logical high maps to VIH. 7. Logical low maps to VIL. | | | | | |

3.2 USB4 Ecosystem

3.2.1 Insertion-Loss Considerations (Informative)

The insertion-loss of the physical media is a key factor for facilitating USB4 electrical compliance. It is recommended that a Router Assembly limit the total insertion-loss from the USB Type-C® receptacle to the USB4 transceiver as follows:

- The total insertion-loss for a Router Assembly supporting Gen 2 is less than or equal to 5.5 dB at 5 GHz, including the receptacle tongue, the PCB trace, the integrated circuit's package and die load.
- The total insertion-loss for a Router Assembly that supports Gen 3 is less than or equal to 7.5 dB at 10 GHz, including the receptacle tongue, the PCB trace, the integrated circuit's package and die load.


For a Captive Device that employs a passive attached cable, it is recommended that the device limit the total insertion-loss from the USB Type-C plug to the USB4 transceiver as follows:

- The total insertion-loss for a Captive Device supporting Gen 2 is less than or equal to 17.5 dB at 5 GHz, including the plug, the cable, the on-board connector, the PCB trace, the integrated circuit's package, and die load.
- The total insertion-loss for a Captive Device supporting Gen 3 is less than or equal to 15 dB at 10 GHz, including the plug, the cable, the on-board connector, the PCB trace, the integrated circuit's package, and die load.

3.2.2 Coded Bit-Error-Ratio Considerations (Informative)

The USB4 protocol employs a combination of Forward Error Correction (FEC) and Pre-Coding (PC) applied per Lane to provide enhanced data integrity. Figure 3-1 shows an example of this scheme.

Figure 3-1. Combined Forward-Error-Correction and Pre-Coding Scheme



The FEC scheme is based on Reed-Solomon RS(198,194) over GF(2⁸) code with two correctable errors per block. The Pre-Coder is designed for converting bursts of consecutive bit-errors into two errors at the beginning and end of the burst, supported by the RS-FEC. The combination of these two mechanisms provides high immunity for bursts of errors dominated by Decision Feedback Equalizers (DFE). Therefore, the coded BER performance is primarily impacted by the random bit-errors probability and not by long bursts of errors.

There is another effect that can cause more than two symbol errors in an RS-FEC block: If a multi-tap DFE is used in the receiver, when a burst of errors ends and the first DFE tap is fed by correct decision, errors still exist in the DFE pipeline, associated with the second tap and above. This increases the probability of having more bit errors until the pipeline is clean. As the first DFE tap is typically the most dominant tap, it is expected that the probability for additional errors is low, but still high as compared to the probability for random-errors, which correspond to the case where the DFE pipeline is clean. The probability that a burst will restart is dominated by the magnitude of the DFE taps with feedback delays of 2 UI or more. Therefore, any DFE taps after the first DFE tap need to be carefully controlled (or totally avoided) in order to minimize their impact on the coded performance.

Because direct measurement of the coded BER is not feasible due to the large measurement windows needed, an alternative indirect method is applied for validating the expected performance. Therefore, targets are specified for the random-errors probability and for the error burst restart probability, ensuring that the coded BER assumptions are met.

3.3 USB4 Electrical Compliance Methodology


3.3.1 System Compliance Test Point Definitions

All measurements shall be referenced to the electrical compliance test points in Table 3-2. Calibration shall be applied in cases where direct measurement is not feasible.

Table 3-2. Electrical Compliance Test Points

| Test Point | Description | Comments |
|------------|-----------------------------------|--|
| TP1 | Transmitter IC output | Not used for electrical testing. |
| TP2 | Transmitter port connector output | Measured at the plug side of the connector. |
| TP3 | Receiver port connector output | Measured at the receptacle side of the connector. All the measurements at this point shall be done while applying reference equalization function. |
| TP3' | Receiver port connector input | Measured at the plug side of the connector. |
| TP4 | Receiver IC input | Not used for electrical testing. |


Figure 3-2. Compliance Points Definition



3.3.2 AC Coupling Capacitors

All of the Lane 0 and Lane 1 electrical interfaces of a Router Assembly shall be AC-coupled. The SBTX and SBRX lines shall not be AC-coupled. All Lane 0 and Lane 1 transmit paths of a Router Assembly shall include AC-coupling capacitance between 135 nF and 265 nF. All Lane 0 and Lane 1 receive paths of a Router Assembly that are directly connected to a USB Type-C connector shall include AC-coupling capacitance between 300 nF and 363 nF, together with discharge resistors between 200 kΩ and 242 kΩ. AC-coupling capacitors (with discharge resistors) may be also placed at the receive paths of a Router Assembly that are not directly connected to USB Type-C connector.

Figure 3-3. Examples for AC-Coupling Capacitor Placement



3.3.3 Reference Clock-and-Data-Recovery (CDR) Function

All jitter and eye diagram measurements shall be performed while applying a reference clock-and-data-recovery (CDR) function. The reference CDR is modeled by a 2nd order PLL response (type II), which derives the following jitter transfer function, described in Laplace domain:

$$H_{jitter}(s) = \frac{s^2}{s^2 + 2 \cdot \zeta \cdot \omega_n \cdot s + \omega_n^2}$$

where


s is the frequency in Laplace domain

ζ is the damping factor

ω_n is the natural frequency of the system

The damping factor and natural frequency used for compliance testing shall be 0.94 and 2.2E7 rad/sec respectively, forming High-Pass-Filter (HPF) mask with 3 dB bandwidth at 5 MHz.

Figure 3-4. Jitter Transfer Function




3.3.4 Reference Equalization Function

All the measurements at TP3 compliance point shall be performed while applying a reference receiver equalization function. The reference receiver equalization function is comprised of parametric Continuous-Time-Linear-Equalizer (CTLE) and Decision-Feedback-Equalizer (DFE), as described in Section 3.3.4.1 and Section 3.3.4.2 respectively.

A measurement that is referenced to TP3 shall use equalization parameters that optimize the calculated eye-diagram.

Figure 3-5. Reference Receiver Equalization



3.3.4.1 Reference CTLE

The following equation describes the frequency response for the USB4 reference CTLE that shall be used for compliance testing:

$$H(s) = 1.41 \cdot \omega_{p2} \cdot \frac{s + \frac{A_{DC}}{1.41} \cdot \omega_{p1}}{(s + \omega_{p1}) \cdot (s + \omega_{p2})}$$

where

A_{DC} is the DC gain
 s is the frequency in Laplace domain

$$\omega_{p1} = \begin{cases} 2 \cdot \pi \cdot 1.5e^9 \frac{rad}{sec} & Gen 2 \\ 2 \cdot \pi \cdot 5e^9 \frac{rad}{sec} & Gen 3 \end{cases}$$

$$\omega_{p2} = \begin{cases} 2 \cdot \pi \cdot 5e^9 \frac{rad}{sec} & Gen 2 \\ 2 \cdot \pi \cdot 10e^9 \frac{rad}{sec} & Gen 3 \end{cases}$$

Ten different CTLE configurations shall be applied such that A_{DC} is one of $\{10^{\frac{-x}{20}} : x = 0, 1, \dots, 9 [dB]\}$.

Figure 3-6. Frequency Response of Gen 2 Reference CTLE





Figure 3-7. Frequency Response of Gen 3 Reference CTLE



3.3.4.2 Reference DFE

A 1-tap feedback filter is defined as part of the reference receiver equalizer used in the compliance testing. The DFE formula is described in the following equation:

$$Y_n = X_n - C_1 * \text{sign}(Y_{n-1})$$

where:

Y_n is the DFE output at time instant n

X_n is the DFE input (incoming signal after applying CTLE) at time instant n

C_1 is the DFE coefficient, which shall be limited to a range of 0 mV to 50 mV

3.3.5 Time Domain Measurements

Time domain measurements shall be performed using a real-time oscilloscope (or equivalent equipment) with an input bandwidth of 21 GHz \pm 1 GHz and a single-ended impedance of 50 Ω . The time domain measurement equipment shall support the USB4 reference CDR defined in Section 3.3.3 and the reference Equalization function defined in Section 3.3.4, which both shall be applied as required for the measurements.

3.3.6 Compliance Boards

3.3.6.1 Compliance Plug Test Board

A high quality USB Type-C plug-to-SMA/SMP test fixture shall be used to enable Router Assembly compliance testing. The fixture shall be comprised of a USB Type-C plug and a short paddle card that can be connected to a coaxial cable with a SMA/SMP connector at its end. The reference points TP2 and TP3' are defined such that the insertion-Loss from the connector pads to the compliance points is 0.5 dB \pm 0.25 dB at 5 GHz and 1 dB \pm 0.25 dB at 10 GHz, and the fixture's insertion-Loss shall be calibrated accordingly. Extra loss and distortion elements shall be compensated by physical and/or mathematical means.

The target single-ended impedance of the fixture shall be 42.5 Ω .

3.3.6.2 Compliance Receptacle Test Board

A high quality USB Type-C receptacle-to-SMA/SMP test fixture shall be used to enable Router Assembly testing. The fixture shall be comprised of a high quality USB Type-C receptacle and a short PCB trace that may be connected to coaxial cable with SMA/SMP connector at its end. The reference point TP3 is defined such that the insertion-loss from the connector pads to the compliance point is 0.5 dB \pm 0.25 dB at 5 GHz and 1 dB \pm 0.25 dB at 10 GHz. Extra loss and distortion elements shall be compensated by physical and/or mathematical means.

The target single-ended impedance of the fixture shall be 42.5 Ω .

3.4 Router Assembly Transmitter Compliance

Transmitter compliance testing for a Router Assembly is defined at two measurement point:

- The output of a compliance plug fixture at the TP2 reference point.
- The output of a compliance receptacle fixture at the TP3 reference point.

Compliance Plugs and Compliance Receptacles are defined in Section 3.3.6.1 and Section 3.3.6.2 respectively.

Unless otherwise specified, a transmitter shall drive PRBS31 pattern during compliance testing. All tests shall be performed with Spread-Spectrum-Clocking (SSC) enabled and while all neighboring transceivers are active.

3.4.1 Transmitter Specifications Applied for All Speeds

Table 3-3 defines the transmitter parameters that shall apply for both Gen 2 and Gen 3 modes of operation.

Table 3-3. Transmitter Specifications Applied for All Speeds (at TP2)

| Symbol | Description | Min | Max | Units | Conditions |
|-----------------------------|---|----------------------------|---------------------|--------|---|
| RL_DIFF | Differential Return Loss, 0.05-12GHz | -- | See Section 3.4.1.2 | dB | |
| RL_COMM | Common Mode Return Loss, 0.05 – 12GHz | -- | See Section 3.4.1.3 | dB | |
| TX_EQ | Transmitter Equalization Setting | -- | See Section 3.4.1.4 | | |
| SSC_DOWN_SPREAD_RANGE | Dynamic range of SSC down-spreading during steady-state | 0.4 | 0.5 | % | See Note 3, Note 4, and Figure 3-9. |
| SSC_DOWN_SPREAD_RATE | SSC down-spreading modulation rate during steady-state | 30 | 33 | KHz | See Note 4 and Figure 3-9. |
| SSC_PHASE_DEVIATION | Phase jitter associated with the SSC modulation during steady-state | 2.5 | 22 | ns pp | See Note 1, Note 4, and Figure 3-9. |
| SSC_SLEW_RATE | SSC frequency slew rate (df/dt) during steady-state | -- | 1250 | ppm/μs | See Note 2, Note 4, and Figure 3-9. |
| TX_FREQ_VARIATIONS_TRAINING | TX frequency variation during Link training, before obtaining steady-state | -- | See Section 3.4.1.1 | ppm | See Note 4. |
| LANE_TO_LANE_SKEW | Skew between dual transmit signals of the same USB4 Port | -- | 26 | ns | See Note 5. |
| RISE_FALL_TIME | TX rise/fall time measured between 20-80% levels | 10 | -- | ps | Test pattern shall be SQ128 (see Table 8-56). |
| V_ELEC_IDLE | Peak voltage during transmit electrical idle (one-sided voltage opening of the differential signal) | -- | 20 | mV | See Note 6. |
| V_TX_DC_AC_CONN | Instantaneous DC+AC voltages at the connector side of the AC coupling capacitors | -0.5 (min1) -0.3 (min2) | 1.0 | V | See Note 7. |

Notes:

1. SSC phase deviation shall be extracted from the transmitted signal. During this test, the transmitter shall be configured to send PRBS31 pattern. The SSC phase deviation shall be extracted from the signal phase after applying a 2nd order low-pass filter with 3 dB point at 5 MHz.
2. The SSC slew rate shall be extracted from the transmitted signal over measurement intervals of 0.5 μ s. The SSC slew-rate shall be extracted from the signal phase after applying a 2nd order low-pass filter with 3 dB point at 5 MHz.
3. SSC_DOWN_SPREAD_RANGE specifies the required SSC modulation depth, represented by the difference of the maximum and minimum modulated frequencies, referenced to the Link speed.
4. Steady-state clocking is applied from the point that SLOS training pattern is sent by the transmitter.
5. Total Lane-to-Lane skew measured at TP2 including the skew introduced by the physical media, by the Router IC TX, and by up to 2 re-timers placed on the board. Informative Lane-to-Lane skew budget: Router IC TX pins: 8 ns, each Re-timer input-to-output: 8 ns, physical media mismatches: 2ns.
6. V_ELEC_IDLE shall be extracted after applying first order low-pass filter with 3 dB point at 1.25 GHz.
7. The absolute single-ended voltage seen by the receiver. This requirement applies to all link states and during power-on, and power-off. (min1, max) is measured with a 200 K Ω receiver load, and (min2, max) is measured with a 50 Ω receiver load. The ground offset between a DFP and UFP does not contribute to V_TX_DC_AC_CONN.

3.4.1.1 Transmitter Frequency Variations during Link Training

3.4.1.1.1 Background (Informative)

A USB4 link can include up to 6 Re-timers, which forward data from one end of the Link to the other end. During Link training, the transmitters at both ends of the Link send SLOS pattern. The SLOS pattern is clocked with SSC down spreading as specified in Table 3-3.

Re-timer transmitters on the Link are all enabled in parallel. Initially, the Re-timer transmitters do not forward the incoming data and just send CL_WAKE1.X Ordered Sets clocked at a local constant frequency (without SSC modulation). In the later stages of the Link training, the Re-timer transmitters sequentially switch to forwarding the incoming data at the incoming frequency. As soon as all of the Re-timer transmitters complete the switching process, steady-state is obtained and SLOS pattern clocked with SSC is forwarded from one end of the Link to the other.

If a Link does not include any Re-timers, the transmitter frequency modulation is at its steady-state from the beginning of the Link training period, as clock switching does not take place.

3.4.1.1.2 Transmitter Specifications Applied to All Speeds

Table 3-4 specifies the limits on the transmitter frequency variation during Link training of a Router Assembly that includes one or more Re-timers:

Table 3-4. Transmitter Frequency Variation Limits During Link Training Before Obtaining Steady-State

| Symbol | Description | Min | Max | Units | Conditions |
|---------------------|---|------|------|-------|-------------|
| INIT_FREQ_VARIATION | Initial non-modulated transmit frequency applied while sending CL_WAKE1.x pattern | -300 | 300 | ppm | See Note 1. |
| DELTA_FREQ_200ns | Frequency variation during Link training over 200ns measurement windows | -- | 1400 | ppm | See Note 1. |
| DELTA_FREQ_1000ns | Frequency variation during Link training over 1µs measurement windows | -- | 2200 | ppm | See Note 1. |
| FREQ_OVERSHOOT | Maximum transient frequency offset from the Link baseline rate, including the clock source accuracy, dynamic clock switching effects and frequency variations induced by low frequency jitter | | 1400 | ppm | See Note 1. |

Notes:

1. Measurement shall be performed over the transmitted signal. The signal phase shall be extracted while applying a 2nd order low-pass filter with 3dB point at 5MHz.
2. INIT_FREQ_VARIATION corresponds to the transmitter average frequency offset from the Link baseline rate (10.0 Gbps in Gen 2, 20.0 Gbps in Gen3), without including low frequency jitter variations, which shall be filtered out by averaging the extracted frequency variation waveform over a window of at least 30µs.

. Figure 3-8 shows the transmit frequency variation during training for a Router Assembly. The transmit frequency variation following the clock-switching event shall be measured over time intervals of 200ns and 1µs.

Figure 3-8. Router Assembly Transmitter Frequency Variation During Training





Figure 3-9. Example Transmitter Frequency During Steady-State




3.4.1.2 Transmitter Differential Return Loss

Transmitter differential return-loss measurements shall be referenced to a single-ended impedance of 42.5 Ω. When measured at TP2, the differential mode return loss shall not exceed the limits given in the following equation:

$$SDD22(f) = \begin{cases} -8.5 & 0.05 < f_{GHz} \leq 3 \\ -3.5 + 8.3 \cdot \log_{10}\left(\frac{f_{GHz}}{12}\right) & 3 < f_{GHz} \leq 12 \end{cases}$$

Figure 3-10. TX Differential Return Loss Mask




3.4.1.3 Transmitter Common Mode Return Loss

Transmitter common-mode return-loss measurements shall be referenced to a single-ended impedance of 42.5 Ω. When measured at TP2, the common-mode return loss shall not exceed the limits given in the following equation:

$$SCC22(f) = \begin{cases} -6 & 0.05 < f_{GHz} \leq 2.5 \\ -3 & 2.5 < f_{GHz} \leq 12 \end{cases}$$

Figure 3-11. TX Common-Mode Return Loss Mask



3.4.1.4 Transmit Equalization

A Router Assembly shall support coefficient-based equalization at its transmitter output. The equalizer's structure is based on a 3-tap UI-spaced finite-impulse-response (FIR) filter as shown in Figure 3-12. The transmitted level corresponding to the nth symbol shall be generated as follows:

$$tx_out_n = \sum_{k=-1}^1 data_in_{n-k} \cdot C_k$$

where:

- tx_out_n is the transmitted level at time instant n
- data_in_{n-k} is the data symbol at time instant n-k (may be +1 or -1)
- C_k is the kth coefficient of the FIR filter

Figure 3-12. Transmitter Equalizer Structure




Table 3-5 lists the pre-shoot and de-emphasis transmitter equalization presets that a transmitter shall support. It also includes the corresponding informative coefficients values provided as a reference. Preset configurations 0-14 represent operation mode with full-swing transmitter output, while configuration 15 represent low swing mode. When configuration 15 is selected, the transmitter's output swing shall be attenuated by 3.5 ± 1 dB compared to its full-swing operation. The required tolerance of the pre-shoot and de-emphasis specifications is ± 1 dB.

For each preset, the pre-shoot and de-emphasis shall be measured at TP2 using SQ128 pattern (see Table 8-56). The pre-shoot shall be calculated as the ratio of the steady-state voltage obtained when configuring the transmitter equalizer such that the pre-cursor tap's magnitude is moved to the main tap, divided by the steady-state voltage with the standard transmit filter configuration. The de-emphasis shall be calculated as the ratio of the steady-state voltage obtained with the standard filter configuration, divided by the steady-state voltage when configuring the transmitter equalizer such that the post-cursor tap's magnitude is moved to the main tap.

Table 3-5. Transmit Equalization Presets

| Preset Number | Pre-shoot [dB] | De-emphasis [dB] | Informative Filter Coefficients | | |
|---------------|----------------|------------------|---------------------------------|----------------|----------------|
| | | | C ₋₁ | C ₀ | C ₁ |
| 0 | 0 | 0 | 0 | 1 | 0 |
| 1 | 0 | -1.9 | 0 | 0.90 | -0.10 |
| 2 | 0 | -3.6 | 0 | 0.83 | -0.17 |
| 3 | 0 | -5.0 | 0 | 0.78 | -0.22 |
| 4 | 0 | -8.4 | 0 | 0.69 | -0.31 |
| 5 | 0.9 | 0 | -0.05 | 0.95 | 0 |
| 6 | 1.1 | -1.9 | -0.05 | 0.86 | -0.09 |
| 7 | 1.4 | -3.8 | -0.05 | 0.79 | -0.16 |
| 8 | 1.7 | -5.8 | -0.05 | 0.73 | -0.22 |
| 9 | 2.1 | -8.0 | -0.05 | 0.68 | -0.27 |
| 10 | 1.7 | 0 | -0.09 | 0.91 | 0 |

| | | | | | |
|----|-----|------|-------|------|-------|
| 11 | 2.2 | -2.2 | -0.09 | 0.82 | -0.09 |
| 12 | 2.5 | -3.6 | -0.09 | 0.77 | -0.14 |
| 13 | 3.4 | -6.7 | -0.09 | 0.69 | -0.22 |
| 14 | 3.6 | 0 | -0.17 | 0.83 | 0 |
| 15 | 1.7 | -1.7 | -0.05 | 0.55 | -0.05 |

Notes:

1. The coefficients are normalized such that $|C_{-1}| + C_0 + |C_1|$ corresponds to full output swing. Preset configuration 15 represents operation mode with lower transmitter swing.
2. Preshoot and de-emphasis are calculated as following:

$$Preshoot = 20 \cdot \log_{10} \left(\frac{-C_{-1} + C_0 + C_1}{C_{-1} + C_0 + C_1} \right) \quad De - emphasis = 20 \cdot \log_{10} \left(\frac{C_{-1} + C_0 + C_1}{C_{-1} + C_0 - C_1} \right)$$

Figure 3-13 and Figure 3-14 depict the frequency responses of the different transmit equalization presets for Gen 2 and Gen 3 modes of operation, respectively.

Figure 3-13. Transmitter Equalization Frequency Response for Gen 2 Systems





Figure 3-14. Transmitter Equalization Frequency Response for Gen 3 Systems



3.4.2 Transmitter Compliance Specifications for Gen 2

A transmitter operating in Gen 2 mode shall meet the specifications in Table 3-3, Table 3-6, and Table 3-7.

Table 3-6. Gen 2 Transmitter Specifications at TP2

| Symbol | Description | Min | Max | Units | Comments |
|--------|---|-------|--------|-------|--|
| UI | Minimum Unit Interval | 99.97 | 100.03 | ps | The minimum UI value corresponds to the Link baseline speed of 10.0 Gbps with an uncertainty range of -300 ppm to 300 ppm. See Note 4. |
| AC_CM | TX AC Common Mode voltage | -- | 100 | mV pp | |
| TJ | Total Jitter | -- | 0.38 | UI pp | See Note 2 and Note 3. |
| UJ | Sum of uncorrelated DJ and RJ components (all jitter components except for DDJ) | -- | 0.31 | UI pp | See Note 2. |
| DDJ | Data-Dependent Jitter | -- | 0.15 | UI pp | See Note 5. |
| UDJ | Deterministic jitter that is uncorrelated to the transmitted data | -- | 0.17 | UI pp | |
| UDJ_LF | Low Frequency Uncorrelated Deterministic Jitter | -- | 0.04 | UI pp | See Note 6. |
| DCD | Even-odd jitter associated with Duty-Cycle-Distortion | -- | 0.03 | UI pp | |
| Y1 | TX eye inner height (one-sided voltage opening of the differential signal) | 140 | -- | mV | Measured for 1E6 UI. See Note 1, Note 2, and Figure 3-15. |
| Y2 | TX eye outer height (one-sided voltage opening of the differential signal) | -- | 650 | mV | Measured for 1E6 UI. See Note 1, Note 2, and Figure 3-15. |

Notes:


1. TX voltage is differential.
2. Measured while applying the reference CDR described in Section 3.3.3. Note that the measured jitter includes residual SSC jitter passing the reference CDR.
3. TJ is defined as the sum of all “deterministic” components plus 14.7 times the RJ RMS (the transmitter RJ RMS multiplier corresponds to the target BER with some margin on top).
4. UI shall be calculated dynamically using a uniform moving average filter with window size of 3000 symbols.
5. The transmit equalization shall be set such that the data dependent jitter is minimized.
6. UDJ_LF is the uncorrelated deterministic jitter measured after applying a 2nd order Low-Pass-Filter with 3 dB cut-off at 0.5 MHz on the measured jitter. This filter needs to be applied on top of the reference CDR function described in Section 3.3.3.

Table 3-7. Gen 2 Transmitter Specifications at TP3

| Symbol | Description | Min | Max | Units | Comments |
|--------|---|-----|------|-------|---|
| TJ | Total Jitter | -- | 0.60 | UI pp | See Note 2, Note 3. |
| UJ | Sum of uncorrelated DJ and RJ components (all jitter components except for DDJ) | -- | 0.31 | UI pp | Note 2. |
| UDJ | Deterministic jitter that is uncorrelated to the transmitted data | -- | 0.17 | UI pp | |
| X1 | TX eye horizontal deviation | -- | 0.24 | UI | Measured for 1E6 UI. See Note 2, Note 4, and Figure 3-15. |
| Y1 | TX eye inner height (one-sided voltage opening of the differential signal) | 47 | -- | mV | Measured for 1E6 UI. See Note 1, Note 2, and Figure 3-15. |
| Y2 | TX eye outer height (one-sided voltage opening of the differential signal) | -- | 650 | mV | Measured for 1E6 UI. See Note 1, Note 2, and Figure 3-15. |

Notes:

1. TX voltage is differential.
2. Measured while applying the reference CDR described in Section 3.3.3 and the reference equalizer defined in Section 3.3.4. Note that the measured jitter includes residual SSC jitter passing the reference CDR.
3. TJ is defined as the sum of all “deterministic” components plus 14.7 times the RJ RMS (the transmitter RJ RMS multiplier corresponds to the target BER with some margin on top).
4. X1 specification is informative but shall be assumed as a valid reference if direct TJ measurement cannot be reliably performed.

Figure 3-15. TX Mask Notations

3.4.3 Transmitter Compliance Specifications for Gen 3 Interconnects

A transmitter operating in Gen 3 mode shall meet the specifications in Table 3-3, Table 3-8, and Table 3-9.

Table 3-8. Gen 3 Transmitter Specifications at TP2

| Symbol | Description | Min | Max | Units | Comments |
|--------|---|--------|--------|-------|--|
| UI | Minimum Unit Interval | 49.985 | 50.015 | ps | The minimum UI value corresponds to the Link baseline speed of 20.0 Gbps with an uncertainty range of -300 ppm to 300 ppm. See Note 4. |
| AC_CM | TX AC Common Mode voltage | -- | 100 | mV pp | |
| TJ | Total Jitter | -- | 0.46 | UI pp | See Note 2, Note 3. |
| UJ | Sum of uncorrelated DJ and RJ components (all jitter components except for DDJ) | -- | 0.31 | UI pp | See Note 2. |
| DDJ | Data-Dependent Jitter | -- | 0.21 | UI pp | See Note 5. |
| UDJ | Deterministic jitter that is uncorrelated to the transmitted data | -- | 0.17 | UI pp | |
| UDJ_LF | Low Frequency Uncorrelated Deterministic Jitter | -- | 0.07 | UI pp | See Note 6. |
| DCD | Even-odd jitter associated with Duty-Cycle-Distortion | -- | 0.03 | UI pp | |
| Y1 | TX eye inner height (one-sided voltage opening of the differential signal) | 120 | -- | mV | Measured for 1E6 UI. See Note 1, Note 2, and Figure 3-15. |
| Y2 | TX eye outer height (one-sided voltage opening of the differential signal) | -- | 650 | mV | Measured for 1E6 UI. See Note 1, Note 2, and Figure 3-15. |

| Symbol | Description | Min | Max | Units | Comments |
|--|-------------|-----|-----|-------|----------|
| Notes: | | | | | |
| 1. TX voltage is differential. | | | | | |
| 2. Measured while applying the reference CDR described in Section 3.3.3. Note that the measured jitter includes residual SSC jitter passing the reference CDR. | | | | | |
| 3. TJ is defined as the sum of all "deterministic" components plus 14.7 times the RJ RMS (the transmitter RJ RMS multiplier corresponds to the target BER with some margin on top). | | | | | |
| 4. UI shall be calculated dynamically using a uniform moving average filter with window size of 6000 symbols. | | | | | |
| 5. The transmit equalization shall be set such that the data dependent jitter is minimized. | | | | | |
| 6. UDJ_LF is the uncorrelated deterministic jitter measured after applying a 2 nd order Low-Pass-Filter with 3 dB cut-off at 0.5 MHz on the measured jitter. This filter needs to be applied on top of the reference CDR function described in Section 3.3.3. | | | | | |

Table 3-9. Gen 3 Transmitter Specifications at TP3

| Symbol | Description | Min | Max | Units | Comments |
|---|---|-----|------|-------|---|
| TJ | Total Jitter | -- | 0.60 | UI pp | See Note 2, Note 3. |
| UJ | Sum of uncorrelated DJ and RJ components (all jitter components except for DDJ) | -- | 0.31 | UI pp | Note 2. |
| UDJ | Deterministic jitter that is uncorrelated to the transmitted data | -- | 0.17 | UI pp | |
| X1 | TX eye horizontal deviation | -- | 0.23 | UI | Measured for 1E6 UI. See Note 2, Note 4, and Figure 3-15. |
| Y1 | TX eye inner height (one-sided voltage opening of the differential signal) | 49 | -- | mV | Measured for 1E6 UI. See Note 1, Note 2, and Figure 3-15. |
| Y2 | TX eye outer height (one-sided voltage opening of the differential signal) | -- | 650 | mV | Measured for 1E6 UI. See Note 1, Note 2, and Figure 3-15. |
| Notes: | | | | | |
| 1. TX voltage is differential. | | | | | |
| 2. Measured while applying the reference CDR described in Section 3.3.3 and the reference equalizer defined in Section 3.3.4. Note that the measured jitter includes residual SSC jitter passing the reference CDR. | | | | | |
| 3. TJ is defined as the sum of all "deterministic" components plus 14.7 times the RJ RMS (the transmitter RJ RMS multiplier corresponds to the target BER with some margin on top). | | | | | |
| 4. X1 specification is informative but shall be assumed as a valid reference if direct TJ measurement cannot be reliably performed. | | | | | |

3.5 Router Assembly Receiver Compliance**3.5.1 Receiver Specifications Applied for All Speeds**

Table 3-10 defines the receiver parameters that shall apply for both Gen 2 and Gen 3 modes of operation.

Table 3-10. Common Receiver Specifications at TP3'

| Symbol | Parameters | Min | Max | Units | Comments |
|---------|--|-----|---------------------|-------|----------|
| RL_DIFF | Differential Return Loss, 0.05 – 12GHz | -- | See Section 3.5.1.1 | dB | |
| RL_COMM | Common Mode Return Loss, 0.05 – 12GHz | -- | See Section 3.5.1.2 | dB | |


| | | | | | |
|------------------------|---|------|------|-----|--------------------|
| LANE_TO_LANE_SKEW | Skew between dual incoming signals of the same USB4 Port | -- | 44 | ns | See Note 1. |
| RX_INIT_FREQ_VARIATION | Initial non-modulated signal frequency applied during training before obtaining steady-state operation | -300 | 300 | ppm | See Notes 2, 3, 4. |
| RX_DELTA_FREQ_200ns | Incoming signal's frequency variation during Link training over 200ns measurement windows | -- | 1400 | ppm | See Notes 2, 3. |
| RX_DELTA_FREQ_1000ns | Incoming signal's frequency variation during Link training over 1µs measurement windows | -- | 2200 | ppm | See Notes 2, 3. |
| RX_FREQ_OVERSHOOT | Incoming signal's maximum frequency offset from the Link baseline rate during training | -- | 1600 | ppm | See Notes 2, 3. |
| V_MAX | Peak Voltage | -- | 650 | mV | |
| | <p>Notes:</p> <ol style="list-style-type: none"> LANE_TO_LANE_SKEW specifies the maximum skew at the connector. On top of the skew measured at TP3', the following informative budget is assumed between the connector and the Router RX IC: each Re-timer input-to-output skew: 8ns, physical media mismatches: 2 ns. The clocking configuration during the different stages of the Link initialization is described in Section 3.4.1.1 and Figure 3-8. Extracted while applying a 2nd order low-pass filter with 3dB point at 5MHz over the signal phase. RX_INIT_FREQ_VARIATION corresponds to the incoming nominal frequency offset from the Link baseline rate (10.0 Gbps in Gen 2, 20.0 Gbps in Gen3), without including low frequency jitter induced variations. | | | | |

3.5.1.1 Receiver Differential Return Loss

Receiver differential return-loss measurements shall be referenced to a single-ended impedance of 42.5 Ω. When measured at TP3', the differential mode return loss shall not exceed the limits given in the following equation:

$$SDD11(f) = \begin{cases} -8.5 & 0.05 < f_{GHz} \leq 3 \\ -3.5 + 8.3 \cdot \log_{10} \left(\frac{f_{GHz}}{12} \right) & 3 < f_{GHz} \leq 12 \end{cases}$$

Figure 3-16. RX Differential Return-Loss Mask




3.5.1.2 Receiver Common Mode Return Loss

Receiver common-mode return-loss measurements shall be referenced to a single-ended impedance of 42.5 Ω. When measured at TP3', the common-mode return loss shall not exceed the limits given in the following equation:

$$SCC11(f) = \begin{cases} -6 & 0.05 < f_{GHz} \leq 2.5 \\ -3 & 2.5 < f_{GHz} \leq 12 \end{cases}$$

Figure 3-17. RX Common Mode Return-Loss Mask



3.5.2 Receiver Uncoded BER Tolerance Testing

A receiver shall operate at BER of $1E-12$ or lower with neither Forward Error Correction nor Pre-Coding applied when a stressed signal is driven at its input. Tolerance testing shall be performed with down-spreading of the clock enabled and while all USB4 Ports are active.

There are two test setups for evaluating the receiver tolerance:

- “Case 1”, which addresses installations with low Insertion-Loss.
- “Case 2”, which addresses installations with maximum Insertion-Loss. For this setup, the following end-to-end insertion loss from the pattern-generator output to TP3 calibration point is assumed:
 - Gen 2: 18dB at 5GHz (including 2m USB-C™ passive cable as part of the setup).
 - Gen 3: 16dB at 10GHz (including 0.8m USB-C passive cable as part of the setup).

Figure 3-18, Figure 3-19, Table 3-11 and Table 3-12 describe the test topologies and setups for Case 1 and Case 2.

Figure 3-18. Receiver Tolerance Test Topologies




Figure 3-19. Receiver Tolerance Test Setups




Table 3-11. Stressed Signal for Gen 2 Receiver Compliance Testing

| Test Case | Inner Eye Voltage [mV peak] | Random Jitter (RJ) [UI peak-to-peak] | Periodic Jitter [UI peak-to-peak] | Total Jitter [UI peak-to-peak] |
|-------------|-----------------------------|--------------------------------------|-----------------------------------|--------------------------------|
| 1 (at TP3') | 350 | 0.14 | 0.17 | 0.35 |
| 2 (at TP3) | 47 | 0.14 | 0.17 | 0.60 |

Notes:

- The Total jitter and Random jitter shall be referenced to 1E-12 statistics. The Inner Eye Voltage shall be measured over 1E6 UI.

Table 3-12. Stressed Signal for Gen 3 Receiver Compliance Testing

| Test Case | Inner Eye Voltage [mV peak] | Random Jitter (RJ) [UI peak-to-peak] | Periodic Jitter [UI peak-to-peak] | Total Jitter [UI peak-to-peak] |
|-------------|-----------------------------|--------------------------------------|-----------------------------------|--------------------------------|
| 1 (at TP3') | 350 | 0.14 | 0.17 | 0.38 |
| 2 (at TP3) | 49 | 0.14 | 0.17 | 0.60 |

Notes:

- The Total jitter and Random jitter shall be referenced to 1E-12 statistics. The Inner Eye Voltage shall be measured over 1E6 UI.

A receiver shall be tested by injecting several different periodic jitter components, one at a time. The testing shall include sinusoidal jitter frequencies of 1 MHz, 2 MHz, 10 MHz, 50 MHz, and 100 MHz. In all cases, the incoming signal shall include SSC modulation on top of the sinusoidal jitter component at the range of 5600ppm. PRBS31 pattern shall be used for receiver compliance testing, however, calibration of the stressed signal source may be performed with a periodic pattern shorter than PRBS31. AC common-mode noise shall be added at the pattern-generator output to ensure worst-case transmitter characteristics. The total common-mode noise shall be

100 mV peak-to-peak at the pattern-generator output, while the added noise profile shall be sine-wave at a frequency not smaller than 400 MHz. The periodic jitter, random jitter and the common-mode noise values shall be calibrated for Case 1 setup at TP3'. Case 2 shall follow with the same pattern-generator settings for these components.

All specified jitter values shall be calibrated while applying the reference CDR defined in Section 3.3.3. The parameters specified at TP3 shall be calibrated after applying the reference equalizer defined in Section 3.3.4.

A receiver may configure its Link Partner's TX equalizer during the USB4 Link establishment. The pattern generator shall support tunable 3-tap FIR at its output, which may be adjusted during the test by the receiver under test through out-of-band software channel. Case 1 shall be calibrated first and Case 2 shall follow with the same TX FIR configuration as an initial setting.

3.5.3 Receiver Multi Error-Bursts Testing

When a receiver employs DFE with more than one tap, it shall take steps to limit the probability that a burst of errors is restarted right after the reception of one or more correct bits (see Section 3.2.2). The receiver multi error-burst probability shall be characterized as follows:

- Definitions:
 - N is a parameter that defines the observation window for burst restart. It also defines the interval between error bursts. N shall be at least 32 bits.
 - An Error Capture is an observation window that starts with the detection of a bit error which is preceded by at least N consecutive bits without errors.
 - A Burst Restart event is an Error Capture that contains an error burst of one or more consecutive errors followed by one or more correct bits and then by one or more errors within the observation window.
- Initialize the receiver test setup with "Case 2" configuration used for testing the uncoded BER with periodic jitter component of 100 MHz (see Section 3.5.2). PRBS31 test pattern shall be used and neither Forward Error Correction nor Pre-Coding shall be applied.
- After initialization, the periodic jitter magnitude shall be increased to the point where uncoded BER of 1E-8 is observed.
- The receiver under test shall trigger on random bit-errors and capture errors that follow. This shall be done using the method described above for Error Captures. At least N consecutive bits shall be examined for errors starting from the initial trigger.
- The probability for obtaining Burst Restart events shall not exceed 5E-7 (i.e. one error burst restart per 2 million error captures on average).
- Error Capture events and Burst Restart events shall be counted and reported out as detailed in Section 8.3.2.1.3 and Section 8.3.2.1.4. The Burst Restarts Count shall not increment more than once in an Error Capture event.

The following pseudocode example describes how to update the Error Capture Count and the Burst Restarts Count:

```
Start: Wait for N consecutive bits without errors

Wait for a bit with error

Start an observation window of N bits

Increment Error Capture Count by 1

If a transition from a bit without errors to a bit with error is
detected within the observation window, increment Burst Restarts
Count by 1

Wait for end of observation window

Go to Start
```

The following is an example Error Capture (N=32):

No burst restart (Error Captures count shall be incremented by 1, Burst Restarts count shall not change):

```
captured_data[31:0]=0000000000000000000000001111111111
```

Burst restart (Error Captures count shall be incremented by 1, Burst Restarts count shall be incremented by 1):

```
captured_data[31:0]=000000000000000000000000111001111111
```

where '1' represents a bit error and '0' represents a correct bit, as expected from "exclusive or" (XOR) operation between the received bits and the synchronized reference PRBS31 pattern. captured_data[0] corresponds to the initial error event trigger.

Note: A burst of errors contains one or more consecutive bit errors.


3.6 Captive Device Compliance

A Captive Device is a Router Assembly. However, because a Captive Device has a cable permanently attached, it requires some adjustments in how compliance testing is performed.

3.6.1 Captive Device Compliance Test Setup

The compliance test requires connecting a Compliance Receptacle to the Captive Device's connector, as illustrated in Figure 3-20. Details of the Compliance Receptacle can be found in Section 3.3.6.2.

Figure 3-20. Captive Device Compliance Test Setup



3.6.2 Captive Device Transmitter Specifications

Transmitter compliance testing for a Captive Device is defined at the output of a Compliance Receptacle fixture referenced to TP3.

Unless otherwise specified, a Captive Device's Transmitter shall transmit PRBS31 pattern throughout the compliance testing. Compliance testing shall be performed with Spread-Spectrum-Clocking (SSC) enabled, and while the neighbor interfaces are all active.

3.6.2.1 Conducted Energy in Wireless Bands

In order to avoid interference to wireless systems, a Captive Device shall conform to the limits on conducted energy of the USB4 differential pair exiting the device as set forth in Table 3-13. The limits are specified both in the differential and the common mode domains.

Table 3-13. Wireless Band Conducted Limits (at TP3)

| Symbol | Description | Max | Units | Comments |
|------------------|--|-----|-------|--|
| PWR_diff_1-14 | Differential Power in each of the Wi-Fi bands 1-14, which are centered around the following frequencies (in MHz): 2412, 2417, 2422, 2427, 2432, 2437, 2442, 2447, 2452, 2457, 2462, 2467, 2472, 2484 | -17 | dBm | Measured in 18MHz window around the center frequency |
| PWR_comm_1-14 | Common mode Power in each of the Wi-Fi bands 1-14, which are centered around the following frequencies (in MHz): 2412, 2417, 2422, 2427, 2432, 2437, 2442, 2447, 2452, 2457, 2462, 2467, 2472, 2484 | -47 | dBm | Measured in 18MHz window around the center frequency |
| PWR_diff_36-64 | Differential Power in each of the Wi-Fi bands 36-64, which are centered around the following frequencies (in MHz): 5180, 5190, 5200, 5210, 5220, 5230, 5240, 5260, 5280, 5300, 5320 | -22 | dBm | Measured in 18MHz window around the center frequency |
| PWR_comm_36-64 | Common mode Power in each of the Wi-Fi bands 36-64, which are centered around the following frequencies (in MHz): 5180, 5190, 5200, 5210, 5220, 5230, 5240, 5260, 5280, 5300, 5320 | -43 | dBm | Measured in 18MHz window around the center frequency |
| PWR_diff_100-140 | Differential Power in each of the Wi-Fi bands 100-140, which are centered around the following frequencies (in MHz): 5500, 5520, 5540, 5560, 5580, 5600, 5620, 5640, 5660, 5680, 5700 | -23 | dBm | Measured in 18MHz window around the center frequency |

| Symbol | Description | Max | Units | Comments |
|------------------|--|-----|-------|---|
| PWR_comm_100-140 | Common mode Power in each of the Wi-Fi bands 100-140, which are centered around the following frequencies (in MHz): 5500, 5520, 5540, 5560, 5580, 5600, 5620, 5640, 5660, 5680, 5700 | -44 | dBm | Measured in 18MHz window around the center frequency. |
| PWR_diff_149-165 | Differential Power in each of the Wi-Fi bands 149-165, which are centered around the following frequencies (in MHz): 5745, 5765, 5785, 5805, 5825 | -23 | dBm | Measured in 18MHz window around the center frequency. |
| PWR_comm_149-165 | Common mode Power in each of the Wi-Fi bands 149-165, which are centered around the following frequencies (in MHz): 5745, 5765, 5785, 5805, 5825 | -45 | dBm | Measured in 18MHz window around the center frequency. |

3.6.2.2 Transmitter Specifications

Table 3-14, Table 3-15, and Table 3-16 define the parameters for Captive Device transmitters operating in Gen 2 and Gen 3 modes.

Table 3-14. Captive Device Transmitter Specifications at TP3 Applied for All Speeds

| Symbol | Description | Min | Max | Units | Comments |
|-----------------------------|---|-----|---------------------|--------------|---|
| RL_DIFF | Differential Return Loss, 0.05-12 GHz | -- | See Section 3.6.2.3 | dB | |
| RL_COMM | Common Mode Return Loss, 0.05-12 GHz | -- | See Section 3.6.2.4 | dB | |
| TX_EQ | Transmitter Equalization Settings | -- | See Section 3.4.1.3 | | |
| SSC_DOWN_SPREAD_RANGE | Dynamic range of SSC down-spreading during steady-state | 0.4 | 0.5 | % | See Note 3, Note 4, and Figure 3-9. |
| SSC_DOWN_SPREAD_RATE | SSC down-spreading modulation rate during steady-state | 30 | 33 | KHz | See Note 4 and Figure 3-9. |
| SSC_PHASE_DEVIATION | Phase jitter associated with the SSC modulation during steady-state | 2.5 | 22 | ns pp | See Note 1, Note 4 and Figure 3-9. |
| SSC_SLEW_RATE | SSC frequency slew rate (df/dt) during steady state | -- | 1250 | ppm/ μ s | See Note 2, Note 4 and Figure 3-9. |
| TX_FREQ_VARIATIONS_TRAINING | TX frequency variations during Link training, before obtaining steady-state | -- | See Section 3.4.1.1 | ppm | See Note 4. |
| LANE_TO_LANE_SKEW | Skew between dual transmit signals of the same USB4 Port | -- | 44 | ns | See Note 5. |
| RISE_FALL_TIME | TX rise/fall time measured between 20-80% levels | 10 | | ps | Test pattern shall be SQ128 (see Table 8-56). |
| AC_CM | Output AC Common Mode Voltage | -- | 200 | mV pp | |

| Symbol | Description | Min | Max | Units | Comments |
|-----------------|---|----------------------------|-----|-------|-------------|
| V_ELEC_IDLE | Peak voltage during transmit electrical idle (one-sided voltage opening of the differential signal) | -- | 20 | mV | See Note 6. |
| V_TX_DC_AC_CONN | Instantaneous DC+AC voltages at the connector side of the AC coupling capacitors | -0.5 (min1) -0.3 (min2) | 1.0 | V | See Note 7. |

Notes:

1. The SSC phase deviation shall be extracted from the transmitted signal. During this test, the transmitter shall be configured to send PRBS31 pattern. The SSC phase deviation shall be extracted from the signal phase after applying a 2nd order low-pass filter with 3 dB point at 5 MHz.
2. The SSC slew rate shall be extracted from the transmitted signal over measurement intervals of 0.5 μ s. The SSC phase slew-rate shall be extracted from the signal phase after applying a 2nd order low-pass filter with 3 dB point at 5 MHz.
3. SSC_DOWN_SPREAD_RANGE specifies the required SSC modulation depth, represented by the difference of the maximum and minimum modulated frequencies, referenced to the Link speed.
4. Steady-state clocking is applied from the point that SLOS training pattern is sent by the transmitter.
5. Total Lane-to-Lane skew measured at TP3 including the skew introduced by the physical media, by the Router IC TX, and by up to 4 re-timers placed on the board and inside the cable plugs. Informative Lane-to-Lane skew budget: Router IC TX pins: 8 ns, each Re-timer input-to-output: 8 ns, physical media mismatches: 4 ns.
6. V_ELEC_IDLE shall be extracted after applying first order low-pass filter with 3 dB point at 1.25 GHz.
7. The absolute single-ended voltage seen by the receiver. This requirement applies to all link states and during power-on, and power-off. (min1, max) is measured with a 200 K Ω receiver load, and (min2, max) is measured with a 50 Ω receiver load. The ground offset between a DFP and UFP does not contribute to V_TX_DC_AC_CONN.

Table 3-15. Captive Device Transmitter Specifications at TP3 for Gen 2 Systems

| Symbol | Description | Min | Max | Units | Comments |
|--------|---|-------|--------|-------|--|
| UI | Minimum Unit Interval | 99.97 | 100.03 | ps | The minimum UI value corresponds to the Link baseline speed of 10.0 Gbps with an uncertainty range of -300 ppm to 300 ppm. See Note 5. |
| TJ | Total Jitter | -- | 0.60 | UI pp | See Note 2, Note 3. |
| UJ | Sum of uncorrelated DJ and RJ components (all jitter components except for DDJ) | -- | 0.31 | UI pp | See Note 2. |
| UDJ | Deterministic jitter that is uncorrelated to the transmitted data | -- | 0.17 | UI pp | |
| UDJ_LF | Low Frequency Uncorrelated Deterministic Jitter | -- | 0.04 | UI pp | See Note 6. |
| DCD | Even-odd jitter associated with Duty-Cycle-Distortion | -- | 0.03 | UI pp | |
| X1 | TX eye horizontal deviation | -- | 0.24 | UI | Measured for 1E6 UI. See Note 2, Note 4, and Figure 3-15. |
| Y1 | TX eye inner height (one-sided voltage opening of the differential signal) | 47 | -- | mV | Measured for 1E6 UI. See Note 1, Note 2, and Figure 3-15. |

| Symbol | Description | Min | Max | Units | Comments |
|---|--|-----|-----|-------|--------------------------------------|
| Y2 | TX eye outer height (one-sided voltage opening of the differential signal) | -- | 650 | mV | See Note 1, Note 2, and Figure 3-15. |
| Notes: | | | | | |
| <ol style="list-style-type: none"> TX voltage is differential. Measured while applying the reference CDR described in Section 3.3.3 and the reference equalizer defined in Section 3.3.4. Note that the measured jitter includes residual SSC jitter passing the reference CDR. TJ is defined as the sum of all "deterministic" components plus 14.7 times the RJ RMS (the transmitter RJ RMS multiplier corresponds to the target BER with some margin on top). X1 specification is informative but shall be assumed as a valid reference if direct TJ measurement cannot be reliably performed. UI shall be calculated dynamically using a uniform moving average filter with window size of 3000 symbols. UDJ_LF is the uncorrelated deterministic jitter measured after applying a 2nd order Low-Pass-Filter with 3 dB cut-off at 0.5 MHz on the measured jitter. This filter needs to be applied on top of the reference CDR function described in Section 3.3.3. | | | | | |

Table 3-16. Captive Device Transmitter Specifications at TP3 for Gen 3 Systems

| Symbol | Description | Min | Max | Units | Comments |
|--------|---|--------|--------|-------|--|
| UI | Minimum Unit Interval | 49.985 | 50.015 | ps | The minimum UI value corresponds to the Link baseline speed of 20.0 Gbps with an uncertainty range of -300 ppm to 300 ppm. See Note 5. |
| TJ | Total Jitter | -- | 0.60 | UI pp | See Note 2, Note 3. |
| UJ | Sum of uncorrelated DJ and RJ components (all jitter components except for DDJ) | -- | 0.31 | UI pp | See Note 2. |
| UDJ | Deterministic jitter that is uncorrelated to the transmitted data | -- | 0.17 | UI pp | |
| UDJ_LF | Low Frequency Uncorrelated Deterministic Jitter | -- | 0.07 | UI pp | See Note 6. |
| DCD | Even-odd jitter associated with Duty-Cycle-Distortion | -- | 0.03 | UI pp | |
| X1 | TX eye horizontal deviation | -- | 0.23 | UI | Measured for 1E6 UI. See Note 2, Note 4, and Figure 3-15. |
| Y1 | TX eye inner height (one-sided voltage opening of the differential signal) | 49 | -- | mV | Measured for 1E6 UI. See Note 1, Note 2, and Figure 3-15. |
| Y2 | TX eye outer height (one-sided voltage opening of the differential signal) | -- | 650 | mV | See Note 1, Note 2, and Figure 3-15. |

Notes:

1. TX voltage is differential.
2. Measured while applying the reference CDR described in Section 3.3.3 and the reference equalizer defined in Section 3.3.4. Note that the measured jitter includes residual SSC jitter passing the reference CDR.
3. TJ is defined as the sum of all “deterministic” components plus 14.7 times the RJ RMS (the transmitter RJ RMS multiplier corresponds to the target BER with some margin on top).
4. X1 specification is informative but shall be assumed as a valid reference if direct TJ measurement cannot be reliably performed.
5. UI shall be calculated dynamically using a uniform moving average filter with window size of 6000 symbols.
6. UDJ_LF is the uncorrelated deterministic jitter measured after applying a 2nd order Low-Pass-Filter with 3 dB cut-off at 0.5 MHz on the measured jitter. This filter needs to be applied on top of the reference CDR function described in Section 3.3.3.

3.6.2.3 Transmitter Differential Return Loss

Transmitter differential return-loss measurements shall be referenced to a single-ended impedance of 42.5 Ω. When measured at TP3, the differential mode return loss shall not exceed the limits given in the following equation:

$$SDD22(f) = \begin{cases} -8.5 & 0.05 < f_{GHz} \leq 3 \\ -3.5 + 8.3 \cdot \log_{10} \left(\frac{f_{GHz}}{12} \right) & 3 < f_{GHz} \leq 12 \end{cases} \text{ [dB]}$$

3.6.2.4 Transmitter Common Mode Return Loss

Transmitter common-mode return-loss measurements shall be referenced to a single-ended impedance of 42.5 Ω. When measured at TP3, the common-mode return loss shall not exceed the limits given in the following equation:

$$SCC22(f) = \begin{cases} -4 & 0.05 < f_{GHz} \leq 2.5 \\ -2 & 2.5 < f_{GHz} \leq 12 \end{cases}$$

3.6.2.5 Transmit Equalization

A Captive Device shall implement tunable 3-tap finite-impulse-response (FIR) equalization at its output. The transmit equalization shall support preset configurations with different de-emphasis and pre-shoot settings as specified in Table 3-5.

3.6.3 Captive Device Receiver Specifications**3.6.3.1 Receiver Specifications Applied for All Speeds**

Table 3-17 defines parameters for receivers that shall apply for both Gen 2 and Gen 3 modes of operation.

Table 3-17. Common Receiver Specifications at TP2

| Symbol | Description | Min | Max | Units | Comments |
|-------------------|--|-----|---------------------|-------|-------------|
| RL_DIFF | Differential Return Loss, 0.05-12GHz | -- | See Section 3.6.3.2 | dB | |
| RL_COMM | Common Mode Return Loss, 0.05-12GHz | -- | See Section 3.6.3.3 | dB | |
| LANE_TO_LANE_SKEW | Skew between dual incoming signals of the same USB4 Port | -- | 26 | ns | See Note 1. |

| Symbol | Description | Min | Max | Units | Comments |
|--|--|------|------|-------|--------------------|
| RX_INIT_FREQ_VARIATION | Initial non-modulated signal frequency applied during training before obtaining steady-state operation | -300 | 300 | ppm | See Notes 2, 3, 4. |
| RX_DELTA_FREQ_200ns | Incoming signal's frequency variation during Link training over 200ns measurement windows | | 1400 | ppm | See Notes 2, 3. |
| RX_DELTA_FREQ_1000ns | Incoming signal's frequency variation during Link training over 1µs measurement windows | | 2200 | ppm | See Notes 2, 3. |
| RX_FREQ_OVERSHOOT | Incoming signal's maximum frequency offset from the Link baseline rate during training | | 1400 | ppm | See Notes 2, 3. |
| V_MAX | Peak voltage | -- | 650 | mV | |
| Notes: | | | | | |
| <ol style="list-style-type: none"> LANE_TO_LANE_SKEW specifies the maximum skew at the connector. On top of the skew measured at TP2, the following informative budget is assumed between the connector and the Router RX IC: each Re-timer input-to-output skew: 8ns, physical media mismatches: 4 ns. The clocking configuration during the different stages of the Link initialization is described in Section 3.4.1.1 and Figure 3-8. Extracted while applying a 2nd order low-pass filter with 3dB point at 5MHz over the signal phase. RX_INIT_FREQ_VARIATION corresponds to the incoming nominal frequency offset from the Link baseline rate (10.0 Gbps in Gen 2, 20.0 Gbps in Gen3), without including low frequency jitter induced variations. | | | | | |

3.6.3.2 Receiver Differential Return Loss

Receiver differential return-loss measurements shall be referenced to a single-ended impedance of 42.5 Ω. When measured at TP2, the differential mode return loss shall not exceed the limits given in the following equation:

$$SDD11(f) = \begin{cases} -8.5 & 0.05 < f_{GHz} \leq 3 \\ -3.5 + 8.3 \cdot \log_{10}\left(\frac{f_{GHz}}{12}\right) & 3 < f_{GHz} \leq 12 \end{cases} \text{ [dB]}$$

3.6.3.3 Receiver Common Mode Return Loss

Receiver common-mode return-loss measurements shall be referenced to a single-ended impedance of 42.5 Ω. When measured at TP2, the common-mode return loss shall not exceed the limits given in the following equation:

$$SCC11(f) = \begin{cases} -4 & 0.05 < f_{GHz} \leq 2.5 \\ -2 & 2.5 < f_{GHz} \leq 12 \end{cases}$$

3.6.4 Captive Device Receiver Uncoded BER Tolerance Testing

The ability of a Captive Device Receiver to tolerate the worst-case incoming signal is examined using a stressed receiver test.

A Captive Device shall be able to reliably receive the stressed input signals specified in Table 3-18 and Table 3-19 (referenced to TP2), and operate with Bit-Error-Ratio of 1E-12 or lower, with neither Forward Error Correction nor Pre-Coding applied. A Captive Device receiver shall be tested by injecting several different periodic jitter components, one at a time. The test shall

include sinusoidal jitter frequencies of 1 MHz, 2 MHz, 10 MHz, 50 MHz, and 100 MHz. In all cases, the incoming signal shall include SSC modulation on top of the sinusoidal jitter component at the range of 5600ppm. PRBS31 pattern shall be used for Captive Device compliance testing, however calibration of the stressed signal source may be performed with a periodic pattern shorter than PRBS31. AC common-mode noise shall be added at the pattern-generator output to ensure worst-case transmitter characteristics. The total common-mode noise shall be 100 mV peak-to-peak at TP2, while the added noise profile shall be sine-wave at a frequency not smaller than 400 MHz. All the specified jitter values shall be calibrated while applying the reference CDR defined in Section 3.3.3.

A Captive Device receiver may configure its Link Partner's TX equalizer during Lane Initialization. The pattern generator shall support tunable 3-tap FIR at its output, which may be adjusted during the test by the receiver under test through out-of-band software channel.

The Captive Device receiver test setup is described in Figure 3-21.


Table 3-18. Stressed Receiver Conditions for Gen 2 Captive Device Compliance Testing (at TP2)

| Inner Eye Voltage [mV peak] | Data Dependent Jitter [UI peak-to-peak] | Random Jitter (RJ) [UI peak-to-peak] | Periodic Jitter [UI peak-to-peak] | Total Jitter [UI peak-to-peak] |
|---|---|--------------------------------------|-----------------------------------|--------------------------------|
| 140 | 0.10 | 0.14 | 0.17 | 0.41 |
| Notes: | | | | |
| 1. The Total jitter and Random jitter shall be referenced to 1E-12 statistics. The Inner Eye Voltage shall be measured over 1E6 UI. | | | | |

Table 3-19. Stressed Receiver Conditions for Gen 3 Captive Device Compliance Testing (at TP2)

| Inner Eye Voltage [mV peak] | Data Dependent Jitter [UI peak-to-peak] | Random Jitter (RJ) [UI peak-to-peak] | Periodic Jitter [UI peak-to-peak] | Total Jitter [UI peak-to-peak] |
|---|---|--------------------------------------|-----------------------------------|--------------------------------|
| 120 | 0.15 | 0.14 | 0.17 | 0.46 |
| Notes: | | | | |
| 1. The Total jitter and Random jitter shall be referenced to 1E-12 statistics. The Inner Eye Voltage shall be measured over 1E6 UI. | | | | |

Figure 3-21. Captive Device Receiver Test Setup



3.6.5 Captive Device Receiver Multi Error-Bursts Testing

When a Captive Device's receiver employs DFE with more than one tap, it shall take steps to limit the probability that a burst of errors is restarted right after the reception of one or more correct bits (see Section 3.2.2). The receiver multi error-burst probability shall be characterized as follows:

- Definitions:
 - N is a parameter that defines the observation window for burst restart. It also defines the interval between error bursts. N shall be at least 32 bits.
 - An Error Capture is an observation window that starts with the detection of a bit error which is preceded by at least N consecutive bits without errors.
 - A Burst Restart event is an Error Capture that contains an error burst of one or more consecutive errors followed by one or more correct bits and then by one or more errors within the observation window.
- Initialize the receiver test setup as configured for testing the uncoded BER with periodic jitter component of 100 MHz (see Section 3.6.4). PRBS31 test pattern shall be used and neither Forward Error Correction nor Pre-Coding shall be applied.
- After initialization, the periodic jitter magnitude shall be increased to the point where uncoded BER of 1E-8 is observed.
- The receiver under test shall trigger on random bit-errors and capture errors that follow. This shall be done using the method described above for Error Captures. At least N consecutive bits shall be examined for errors starting from the initial trigger.
- The probability for obtaining Burst Restart events shall not exceed 5E-7 (i.e. one error burst restart per 2 million error captures on average).
- Error capture events and error-burst restart events shall be counted and reported out as detailed in Section 8.3.2.1.3 and Section 8.3.2.1.4. The Burst Restarts Count shall not increment more than once in an Error Capture event.

The following pseudo-code example describes how to update the Error Capture Count and the Burst Restarts Count:

```
Start: Wait for N consecutive bits without errors

Wait for a bit with error

Start an observation window of N bits

Increment Error Capture Count by 1

If a transition from a bit without errors to a bit with error is
detected within the observation window, increment Burst Restarts
Count by 1

Wait for end of observation window

Go to Start
```

The following is an example analysis:

No burst restart (Error Captures count shall be incremented by 1, Burst Restarts count shall not change):

captured_data[31:0]=0000000000000000000000001111111111

Burst restart (Error Captures count shall be incremented by 1, Burst Restarts count shall be incremented by 1):

captured_data[31:0]=00000000000000000000000011100111111111

where '1' represents a bit error and '0' represents a correct bit, as expected from "exclusive or" (XOR) operation between the received bits and the synchronized reference PRBS31 pattern. captured_data[0] corresponds to the initial error event trigger.

Note: A burst of errors contains one or more consecutive bit errors.

3.7 Low Frequency Periodic Signaling (LFPS)

Low frequency periodic signaling (LFPS) is used for in-band communication between two Link Partners when exiting power management Link states.


3.7.1 LFPS Signal Definition

The following table defines the LFPS electrical specification.

Table 3-20. LFPS Electrical Specifications

| Symbol | Description | Min | Max | Units | Comments |
|---------------------|---|-----|------|----------|--|
| tPeriod | Period of LFPS cycle (clock pattern) | 20 | 80 | ns | |
| tPreData | Period of time in which electrical-idle shall be set after the LFPS sequence | 80 | 120 | ns | See Figure 3-22. |
| V_CM_AC_LFPS | Common mode noise | -- | 100 | mV pk-pk | |
| V_TX_DIFF_PP_LFPS | LFPS peak-to-peak differential amplitude | 800 | 1200 | mV pk-pk | |
| V_LFPS_RX_DETECT_TH | LFPS RX detect threshold | 100 | 300 | mV pk-pk | Signal shall be rejected below minimum level and shall be accepted above maximum level. |
| tRiseFall | Rise/Fall time of LFPS signal, measured from 20% to 80% of the signal dynamic range | -- | 4 | ns | |
| LFPS_DUTY_CYCLE | Duty-Cycle of LFPS signal | 45 | 55 | % | Applies for all the LFPS cycles including the first and the last, which should not be cut in the middle. |

Figure 3-22. Signaling During Power Management State Exit





IMPLEMENTATION NOTE

Typical implementations for transitioning from LFPS to USB3 Data might not be suitable for USB4. These implementations could have an uncertainty of up to LFPS $t_{Period\ max}$ (80ns) for the actual duration of LFPS signaling. This uncertainty would not allow generation of $t_{PreData}$ (electrical idle) with a duration of 80ns – 120ns. It is implementation dependent how to transition from LFPS via electrical idle to USB4 Data as shown in Figure 3-22.

3.8 Receiver Lane Margining (Testability)

3.8.1 Background

All USB4 Ports shall support receiver Lane margining as defined in this section. Receiver Lane margining provides a means to assess end-to-end Link performance, enables the validation of Links that use Re-timers or device down topologies (i.e. do not have a connector that can be used for compliance testing), and enables a standard method for margining systems in production or in the field. Receiver Lane margining is performed while the Link is in the CL0 state. This allows derivation of the Link electrical performance that would be experienced by the end user.

Receiver margin values and error counts described in this section are not designed to determine compliance for any component. Rather, Receiver Lane Margining standardizes methods already in use by many component and system manufacturers to determine a subjective measure of link robustness.


A USB4 Port shall support either the software (SW) margining mode defined in Section 3.8.1.1 or the hardware (HW) margining mode defined in Section 3.8.1.2. A USB4 Port may optionally support both SW and HW margining modes.

Section 8.3.2.3 defines the Operations that are used to obtain receiver margining capabilities, configure margining parameters, run a margining test, and get margining results.

3.8.1.1 Software Margining Mode

The SW margining mode enables application software executing on the system under test to obtain the margin information using a test flow that is implemented in software. In this mode, software can perform sequential Operations of `RUN_SW_LANE_MARGINING` and `READ_SW_MARGIN_ERR` to implement the margining flow. A receiver that supports the SW mode and time margining shall be able to support a voltage and time offset concurrently, such that the voltage margin at each time offset can be obtained. A simplified example of a SW margining flow is depicted in Figure 3-23.

Figure 3-23. Software Margining Mode Example



¹ See Table 8-67 for READ_LANE_MARGIN_CAP Completion Data

² See Table 8-71 for RUN_SW_LANE_MARGINING Operation Metadata

³ See Table 8-73 for READ_SW_MARGIN_ERR Completion Metadata


A Router shall implement a 4-bit error counter for each Lane when performing a non-destructive margining test in SW mode. A bit error counter shall increment when it detects a bit error on its Lane and shall stop when it reaches its maximum value. A Router shall allow system software to reset the bit error counters at any time.

A Router does not need to implement bit error counters if it only supports HW margining mode.

3.8.1.2 Hardware Margining Mode

HW margining mode implements the margin test flow in hardware or firmware. Application software executing on the system under test configures the test parameters and initiates a margining test by writing to configuration and command registers. The implementation of the margining flow is performed by hardware or firmware and software obtains the margin results through the reading of status and result registers. This flow is depicted in Figure 3-24.

Figure 3-24. Hardware Margining Flow



¹ See Table 8-67 for READ_LANE_MARGIN_CAP Completion Data

² See Table 8-68 for RUN_HW_LANE_MARGINING Operation Metadata

³ See Table 8-70 for RUN_HW_LANE_MARGINING Completion Metadata

3.8.2 Receiver Voltage Margining and Timing Margining Requirements

All receivers shall support voltage margining whereby a receiver sampler is offset from the nominal sampling position in the voltage (vertical) dimension. Independent voltage margining in the positive (high) and negative (low) directions is optional but recommended. A USB4 Port shall be capable of performing voltage margining for each Lane independently. Voltage margining shall be non-destructive (i.e. not introduce actual bit errors on the Link). Voltage margining can be implemented using a monitor sampler that is offset from the nominal sample position while the data sampler remains at the nominal position. A receiver shall use the range and step size in Table 3-21 for voltage margining.


Support for timing margining, whereby a receiver sampler is offset from the nominal sampling position in the timing (horizontal) dimension, is optional. Independent margining in the positive (right) and negative (left) directions is optional but recommended. A USB4 Port that supports timing margining shall be capable of performing timing margining for each Lane independently. The timing margining may be destructive (i.e. cause actual bit errors in the Link). Timing margining can be implemented using a jitter injection circuit to inject jitter onto the Rx sampling clock to offset the data sampler from the nominal position. A receiver shall use the range and step size in Table 3-21 for timing margining.

Table 3-21. RX Margining Voltage and Timing Requirements

| Parameter | Min | Max | Units |
|---|---------|---------|-------|
| Voltage Margin Range | +/- 74 | +/- 200 | mV |
| Voltage Margining Step Size | Note 1 | 3 | mV |
| Voltage Margining Steps (per direction) | Note 2 | 127 | |
| Timing Margining Range | +/- 0.2 | +/- 0.5 | UI |
| Timing Margining Step Size | Note 3 | 0.03 | UI |

| Parameter | Min | Max | Units |
|---|--------|-----|-------|
| Timing Margining Steps (per direction) | Note 4 | 31 | |
| Notes: | | | |
| <ol style="list-style-type: none"> 1. The minimum Voltage Margining Step Size is bounded by $\text{Min}(\text{Voltage Margin Range})/\text{Max}(\text{Voltage Margining Steps})$. 2. The minimum Voltage Margining Steps (per direction) is bounded by $\text{ceiling}(\text{Min}(\text{Voltage Margin Range})/\text{Max}(\text{Voltage Margin Step Size}))$. 3. The minimum Timing Margining Step Size is bounded by $\text{Min}(\text{Timing Margining Range})/\text{Max}(\text{Timing Margining Steps})$. 4. The minimum Timing Margining Steps (per direction) is bounded by $\text{ceiling}(\text{Min}(\text{Timing Margin Range})/\text{Max}(\text{Timing Margin Step Size}))$. | | | |

Figure 3-25. RX Margining Range Requirements




The purpose of timing margining is to measure the actual timing margin seen by the receiver when operating in a fully functional mode. Therefore, when performing timing margining, it is recommended that all RX adaptive circuitry be enabled and operating in a closed loop manner.

In addition to supporting the mandatory voltage requirements displayed in Table 3-21, a receiver may optionally support the Rx voltage margining capabilities defined in Table 3-22.

Table 3-22. Optional RX Margining Voltage Capabilities

| Parameter | Min | Max | Units |
|---|--------|---------|-------|
| Voltage Margin Range | +/- 74 | +/- 400 | mV |
| Voltage Margining Steps (per direction) | 25 | 127 | |


Figure 3-26. Optional RX Margining Range Capabilities



3.8.3 Receiver Parameter Access

It is recommended that all receiver training (e.g. equalization) parameters be accessible to system software to enable a system integrator with greater visibility into the operation of the receiver. This allows a system integrator to observe the stability and performance of the link and provides an aid for debug. The receiver training parameters should be accessible from the vendor specific address range within the SB Registers Space. The number and type of parameter values is implementation specific. The system integrator needs to work directly with the receiver IP vendor to understand how to interpret the stored receiver parameters.

4 Logical Layer




4.1 Sideband Channel

Sideband Channel transactions are sent over the SBTX wire and received over the SBRX wire defined in Section 3.1.

Figure 4-1 depicts several examples of SBTX and SBRX connectivity between two Routers using On-Board Re-timers, and Passive or Active Cables.

A Router may support a maximum of two On-Board Re-timers between it and the cable.

Figure 4-1. Cable Topologies (Informative)



4.1.1 Transactions

4.1.1.1 Symbols

A Sideband Channel shall encode all transmitted Symbols using the 10-bit Start/Stop encoding scheme as follows:


- A Start bit (logical 0b).
- Eight bits of payload.
- A Stop bit (logical 1b).

4.1.1.2 Transaction Types

Transactions are carried over the Sideband Channel in one of three transaction types:

- Link Type (LT).
- Administrative Type (AT).
- Re-timer Type (RT).

The symbols within a transaction shall be sent in ascending order. The bits within a symbol shall be sent in the order from bit 0 to bit 7. Figure 4-2 illustrates this ordering.

Figure 4-2. Symbol and Bit Order on Sideband Channel

An inter-Symbol gap of any duration may be inserted between Symbols of the same transaction.

4.1.1.2.1 LT Transactions

LT Transactions are used during Lane Initialization. They are also used to signal a change in Adapter state due to events such as a Lane disconnect or transition to a Low Power state. An LT Transaction shall consist of the three Symbols described in Table 4-1.

Table 4-1. LT Transaction Format

| Byte | Payload Value | Description |
|------|---------------------------|--|
| 0 | FEh | Data Link Escape (DLE) Symbol - indicates the beginning of a transaction. |
| 1 | See Table 4-2 | Lane State Event (LSE) - identifies a change in originator's state or in the Adapter state. |
| 2 | Bitwise complement of LSE | Complement LSE (CLSE) |

Table 4-2. LSE Symbol

| Bits | Name | Function |
|-------|------------------|---|
| [3:0] | <i>LSESymbol</i> | This field defines the LT Transaction type. Undefined values are Rsvd. 0000b - LT_Fall (Lane Disable event) 0010b - LT_Resume (USB4™ Port started USB4 transmission) 0011b - LT_LRoff (disconnect or System Sleep state) |
| [4] | <i>Rsvd</i> | Reserved |
| [5] | <i>LSELane</i> | Indicates the Lane affected by this transaction. Shall be 0b (for Lane 0) or 1b (for Lane 1). Shall be set to 0b when issuing an LT_LRoff Transaction. |
| [7:6] | <i>StartLT</i> | Identifies that this is an LT Transaction. Shall be set to 10b. |

The recipient of an LT Transaction shall verify that the CLSE Symbol payload is a bitwise complement of the LSE Symbol payload. An LT Transaction that fails this check shall be dropped and no further action shall be taken on its behalf.

4.1.1.2.2 AT Transactions

There are two types of AT Transactions:

- AT Commands – used to read from or write to SB Register space of a Router.
- AT Responses – used to respond to an AT Command.

An AT Transaction shall consist of the Symbols in Table 4-3.

Table 4-3. AT Transaction Format

| Byte | Payload Value | Description |
|---------|-------------------------------|---|
| 0 | FEh | Data Link Escape (DLE) Symbol – indicates the beginning of a transaction. |
| 1 | See Table 4-4 | Start Transaction (STX) Symbol – defines the operation of the transaction. |
| [n+1:2] | See Section 4.1.1.3 | Data Symbols – the number of Data Symbols is not explicitly defined and is inferred from the end-of-transaction delimiters (DLE-ETX). The number of Data Symbols (n) shall not exceed 66. |
| n+2 | Low-order byte of 16-bit CRC | Low CRC (LCRC) Symbol |
| n+3 | High-order byte of 16-bit CRC | High CRC (HCRC) Symbol |
| n+4 | FEh | Data Link Escape (DLE) Symbol |
| n+5 | 40h | End of Transaction (ETX) Symbol |

Table 4-4. STX Symbol for an AT Transaction

| Bits | Name | Function |
|-------|---------------------|--|
| [0] | <i>CmdNotResp</i> | Identifies whether the AT Transaction is an AT Command or an AT Response. Shall be set to 0b for an AT Response or 1b for an AT Command. |
| [1] | <i>ReturnBounce</i> | Shall be set to 0b. |
| [2] | <i>Recipient</i> | Identifies the intended final recipient of the transaction. Shall be set to 1b. |
| [3] | <i>Bounce</i> | Shall be set to 0b. |
| [4] | <i>Responder</i> | Shall be set to 0b. |
| [5] | <i>Rsvd</i> | Reserved. |
| [7:6] | <i>StartAT</i> | Identifies that this is an AT Transaction. Shall be set to 00b. |

A Router that receives an AT Command with the *Recipient* bit set to 1b shall respond with an AT Response.

Note: Only a Router issues AT Transactions. A Re-timer does not issue AT Transactions. See the USB4 Re-Timer Specification for more information.

4.1.1.2.3 RT Transactions

There are two types of RT Transactions:

- RT Command – used by a Router or a Re-timer to communicate with another Router or Re-timer. Can be Broadcast (Section 4.1.1.2.3.1) or Addressed (Section 4.1.1.2.3.2).
- RT Response – used to respond to an RT Command (Addressed RT Transactions only).

4.1.1.2.3.1 Broadcast RT Transactions

A Router sends Broadcast RT Transactions to enumerate the Re-timers along a Link and to propagate Link attributes to Re-timers along a Link. A Broadcast RT Transaction does not cause an RT Response.

A Broadcast RT Transaction shall have the format shown in Table 4-5.

Table 4-5. Broadcast RT Transaction Format

| Byte | Payload Value | Description |
|------|-------------------------------|--|
| 0 | FEh | Data Link Escape (DLE) Symbol – indicates the beginning of a transaction. |
| 1 | See Table 4-6 | Start Transaction (STX) Symbol – defines the operation of the transaction. |
| 3:2 | See Table 4-7 and Table 4-8 | Link Parameters – contains a list of the Link parameters selected during Lane Initialization. |
| 4 | Low-order byte of 16-bit CRC | Low CRC (LCRC) Symbol |
| 5 | High-order byte of 16-bit CRC | High CRC (HCRC) Symbol |
| 6 | FEh | Data Link Escape (DLE) Symbol |
| 7 | 40h | End of Transaction (ETX) Symbol |

Table 4-6. STX Symbol for a Broadcast RT Transaction

| Bits | Name | Function |
|-------|-------------------|---|
| [7:6] | <i>StartRT</i> | Identifies that this is an RT Transaction. Shall be set to 01b. |
| [5] | <i>Broadcast</i> | Identifies that this is a Broadcast RT Transaction. Shall be set to 1b. |
| [4:1] | <i>Index</i> | Index of the transmitting Router or Re-timer. Shall be set to 0. |
| [0] | <i>CmdNotResp</i> | Identifies whether the RT Transaction is an RT Command or an RT Response. Shall be set to 1b. |

Table 4-7. Contents of Byte 2 in a Broadcast RT Transaction


| Bits | Name | Function |
|-------|-----------------------------|---|
| [7:5] | <i>Rsvd</i> | Reserved |
| [4] | <i>TBT3-CompatibleSpeed</i> | Set to 0b. |
| [3] | <i>SSCAwaysOn</i> | Set to 0b when the Router sending the Transaction exits from a CLx state with SSC off. Set to 1b when the Router sending the Transaction always operates with SSC on, including during exit from CLx state. It is recommended that a Router set <i>SSCAwaysOn</i> to 1b. <i>Note: A Re-timer sets this field to 0b before forwarding a Broadcast RT Transaction.</i> |
| [2] | <i>RS_FEC</i> | Set to 0b when RS-FEC is not enabled on the Link. Set to 1b when RS-FEC is enabled on the Link. |
| [1] | <i>Rsvd</i> | Reserved |
| [0] | <i>USB4</i> | Set to 1b. |

Table 4-8. Contents of Byte 3 in a Broadcast RT Transaction

| Bits | Name | Function |
|-------|---------------------|---|
| [7:4] | <i>SelectedGen</i> | Indicates the Gen speed selected for operation. 0000b - Reserved 0001b - Gen 2 speed 0010b - Gen 3 speed Else - Reserved. |
| [3:2] | <i>Rsvd</i> | Reserved. |
| [1] | <i>Lane1Enabled</i> | Shall equal the value of the <i>Enabling Decision (Lane 1)</i> bit in the Link Configuration register of the SB Register Space (see Section 4.1.1.3). |
| [0] | <i>Lane0Enabled</i> | Shall equal the value of the <i>Enabling Decision (Lane 0)</i> bit in the Link Configuration register of the SB Register Space (see Section 4.1.1.3). |

Figure 4-3 depicts an example of how a Broadcast RT Transaction is used to assign indexes to a Link with six Re-timers. In the example, each Re-timer that receives the Broadcast RT Transaction increments the index by one and stores the resulting Re-timer Index locally as its index in the direction of the Broadcast RT Transaction. The Re-timer then sends the Broadcast RT Transaction on its other USB4 Port.

Figure 4-3. Propagation of a Broadcast RT Transaction



Note: Only a Router issues Broadcast RT Transactions. A Re-timer does not issue Broadcast RT Transactions. See the USB4 Re-Timer Specification for more information.

4.1.1.2.3.2 Addressed RT Transactions

A Router uses an Addressed RT Transaction to access the SB Register Space of a Re-timer on the Link. The Router sets the *Index* field in the STX symbol of the Addressed RT Transaction to the Re-timer Index of the target Re-timer.

A Router can also use an Addressed RT Transaction to access the SB Register Space of an adjacent Re-timer or Router during TxFFE negotiation. A Router sets the *Index* field in the STX symbol of the Addressed RT Transaction to 0 to target an adjacent Router or Re-timer.

An Addressed RT Transaction shall have the format shown in Table 4-9.

Table 4-9. Addressed RT Transaction Format

| Byte | Payload Value | Description |
|---------|------------------------------|---|
| 0 | FEh | Data Link Escape (DLE) Symbol - indicates the beginning of a transaction. |
| 1 | See Table 4-10 | Start Transaction (STX) Symbol - defines the operation of the transaction. |
| [n+1:2] | See Section 4.1.1.3 | Data Symbols - a number (n) of data symbols. The number of Data Symbols shall not exceed 66. <i>Note: The number of Data Symbols is not explicitly defined and is inferred from the end-of-transaction delimiters (DLE-ETX).</i> |
| n+2 | Low-order byte of 16-bit CRC | Low CRC (LCRC) Symbol |

| Byte | Payload Value | Description |
|------|-------------------------------|--|
| n+3 | High-order byte of 16-bit CRC | High CRC (HCRC) Symbol |
| n+4 | FEh | Data Link Escape (DLE) Symbol |
| n+5 | 40h | End of Transaction (ETX) Symbol |

Table 4-10. STX Symbol for an Addressed RT Transaction

| Bits | Name | Function |
|-------|-------------------|---|
| [7:6] | <i>StartRT</i> | Identifies that this is an RT Transaction. Shall be set to 01b. |
| [5] | <i>Broadcast</i> | Shall be set to 0b. |
| [4:1] | <i>Index</i> | <p>Addressed RT Command: Shall be set to 0 if the target of the Transaction is the first Router or Re-timer that receives the Transaction. Else, shall be set to the Re-timer Index of the Re-timer that is the target of the Command.</p> <p>Addressed RT Response: Shall be set to the value of the <i>Index</i> field in the corresponding Addressed RT Command.</p> |
| [0] | <i>CmdNotResp</i> | Identifies whether the RT Transaction is an RT Command or an RT Response. Shall be set to 0b for an RT Response or 1b for an RT Command. |

A Router that receives an Addressed RT Command with the *Index* field set to 0 shall respond with an Addressed RT Response. A Router shall not respond to Addressed RT Commands with a non-zero *Index* field.

4.1.1.2.4 AT and RT Transaction Rules

A transmitter may abort an AT Transaction or an RT Transaction after sending the first DLE Symbol (in order to send a higher priority LT Transaction), but shall not abort an AT Transaction or an RT Transaction after the STX Symbol is sent. When a receiver receives two or more leading DLE symbols it shall discard the extra leading DLE symbols and process the received LT Transaction as if only one leading DLE symbol was received.

To avoid identifying a pattern of FEh-40h inside the payload as an end-of-transaction delimiter, if any Data Symbol or a CRC Symbol in an AT Transaction or an RT Transaction contains the same payload as a DLE Symbol, the transmitter of the AT Transaction or RT Transaction shall insert a Symbol with payload of FEh in front of that Data Symbol. The recipient of an AT Transaction or an RT Transaction shall strip all duplicating FEh Symbols that immediately precede a Data Symbol or a CRC Symbol. See also Figure 4-4.

Each AT Transaction or RT Transaction shall include a 16-bit CRC. The 16-bit CRC is broken up into two bytes, which make up the Low CRC and High CRC Symbols described in Table 4-3, Table 4-5, and Table 4-9. The CRC protects the STX and Data Symbols only. Thus, only the STX and Data Symbols shall be used in CRC calculation. FEh Symbols inserted by the transmitter into the Data Symbols of the transaction are not CRC protected.

The CRC shall be calculated in increasing Symbol order, starting with the STX Symbol. Within each Symbol, CRC shall be calculated from bit[7] to bit[0]. The CRC shall be calculated using the following rules:

- Width: 16
- Poly: 8005h
- Init: FFFFh
- RefIn: True
- RefOut: True
- XorOut: 0000h

Appendix A provides an example of a CRC calculation.

4.1.1.2.5 AT and RT Command Rules

A Router shall process AT Commands and AT Responses arriving from the Link Partner or Re-timer in the order received. For example, a Router shall respond to a received AT Command before it processes an AT Response that arrived after the AT Command.

A Router shall process Addressed RT Commands and Addressed RT Responses arriving from a Re-timer in the order received. For example, a Router shall respond to a received Addressed RT Command before it processes an Addressed RT Response that arrived after the Addressed RT Command.

Note: A Router may send Commands and Responses in any order so long as timing constraints are met. Thus, when retransmitting a Transaction, a Router may send a Response before reissuing a Command, or it may reissue a Command, then send a Response.

A Router shall not send an AT Command or Addressed RT Command while it is waiting for a response for either a previously sent AT Command or a previously sent Addressed RT Command.

4.1.1.2.5.1 AT Commands

The recipient of an AT Command shall send an AT Response within `tCmdResponse` of receiving the AT Command. It is recommended that the recipient of an AT Command respond with an AT Response as soon as possible.

If a Router sends an AT Command, then receives at least two AT Commands from the target of the outstanding AT Command within `tATTimeout`, it shall stop waiting for an AT Response and shall immediately reissue the outstanding AT Command. An example of when this can occur, is when one Router powers up before the other Router.

Otherwise, a Router shall wait `tATTimeout` for an AT Response and may timeout the outstanding AT Command if an AT Response is not received within `tATTimeout`. When a timeout occurs, the originator may either resend the AT Command, issue another AT Command, or handle the timeout in an implementation-specific manner.

4.1.1.2.5.2 Addressed RT Commands

The recipient of an Addressed RT Command shall send an Addressed RT Response within `tCmdResponse` of receiving the Addressed RT Command. It is recommended that the recipient of an Addressed RT Command respond with an Addressed RT Response as soon as possible.

A Router shall wait `tRTTimeout` for an Addressed RT Response and may timeout the outstanding Addressed RT Command if an Addressed RT Response is not received within `tRTTimeout`. When a timeout occurs, the originator may either resend the Addressed RT Command, issue another Addressed RT Command, or handle the timeout in an implementation-specific manner.

4.1.1.2.6 Receiver Decoding of LT, AT, and RT Transactions

Figure 4-4 and Table 4-11 depict the receive state-machine to decode LT, AT, and RT Transactions over the Sideband Channel. For each entry in the Table, if Condition is met, then Action is taken and the state-machine transitions to Next State.

A Router shall ignore and discard an AT Transaction or an Addressed RT Transaction if any of the following are true:

- The CRC in the Transaction is invalid.
- The Transaction has no data and no *CRC* field.

Any sequence of Symbols not handled as an LT Transaction, an AT Transaction, or an RT Transaction shall be discarded. An AT Response or an RT Response shall not be sent in response to such a sequence.

Figure 4-4. Sideband Channel Receive Transaction State Machine




Table 4-11. Sideband Channel Receive Transaction State Machine

| Initial State | Next State | Condition | Action |
|--|------------|--|--|
| Idle | DLE1 | Received a DLE Symbol | -- |
| Idle | Idle | Received other Symbol | -- |
| DLE1 | LT | Received LSE Symbol. | Store LSE Symbol. |
| DLE1 | DLE1 | Received a DLE Symbol. | -- |
| DLE1 | AT/RT | Received an STX Symbol. | Initialize <i>Count</i> to 1. Store STX Symbol. Initialize CRC machine and update with STX Symbol. |
| DLE1 | Idle | Received other Symbol. | -- |
| LT | DLE1 | Received a DLE Symbol. | Discard stored LSE Symbol. |
| LT | Idle | Received CLSE Symbol. | Process LT Transaction. See Section 4.1.1.2.1. |
| LT | Idle | Received other Symbol. | Discard stored LSE Symbol. |
| AT/RT | DLE2 | Received a DLE Symbol. | -- |
| AT/RT | AT/RT | Received other Symbol AND Count < Max Transaction Size. See Note 1. | Store Symbol. Update CRC. Increment Count. |
| AT/RT | Idle | Received other Symbol AND Count >= Max Transaction Size. See Note 1. | Discard stored Symbols. |
| DLE2 | AT/RT | Received DLE Symbol AND Count < Max Transaction Size. See Note 1. | Store Symbol. Update CRC. Increment Count. |
| DLE2 | AT/RT | Received an STX Symbol not equal to 40h. | Initialize <i>Count</i> to 1. Store STX Symbol. Initialize CRC machine and update with STX Symbol. |
| DLE2 | Idle | Received an ETX Symbol. | Process AT Transaction or RT Transaction. |
| DLE2 | Idle | Received other Symbol. | Discard stored symbols. |
| Notes: | | | |
| 1. Max Transaction Size is 69 Symbols for an AT Transaction or Addressed RT Transaction or 5 Symbols for a Broadcast RT Transaction. | | | |

4.1.1.3 SB Register Space

A Router uses AT Transactions to access the SB Register Space of its Link Partner. A Router uses Addressed RT Transaction to access the SB Register Space of a Re-timer on the Link. Section 4.1.1.3.1 describes how a Router accesses SB Register Space.

A Connection Manager can also initiate an access to the SB Register Space of a Router or Re-timer. Section 4.1.1.3.2 describes how a Connection Manager accesses SB Register Space.

A Router shall maintain one SB Register Space per USB4 Port. Section 4.1.1.3.3 describes the format and contents of SB Register Space.

4.1.1.3.1 Router Access

A Router uses an AT Transaction with an AT Command to access the SB Register Space of its Link Partner. A Router uses an Addressed RT Transaction with an AT Command to access the SB Register Space of a Re-timer on the Link. An AT Command or RT Command shall consist of the Symbols described in Table 4-12.

Table 4-12. AT/RT Command Data Symbols

| Symbol | Bits | Name | Function |
|--------|------|--------------|--|
| 0 | 7:0 | <i>REG</i> | 8-bit address of the register being read from or written to. |
| 1 | 6:0 | <i>LEN</i> | Number of bytes to read/write. Shall not be greater than 64. |
| 1 | 7 | <i>WnR</i> | 0b for a Read Command. 1b for a Write Command. |
| m:2 | 7:0 | COMMAND_DATA | For a Write Command: A sequence of up to 64 bytes to be written. Register Contents shall appear least significant byte first. For a Read Command: There is no COMMAND_DATA. |

When a Router receives an AT Command it sends an AT Response. When a Re-timer receives an RT Command, it sends an RT Response. An AT Response or RT Response shall consist of the Symbols described in Table 4-13.

Table 4-13. AT/RT Response Data Symbols

| Symbol | Bits | Name | Function |
|--------|------|---------------|---|
| 0 | 7:0 | <i>REG</i> | 8-bit address of the register being read to or written from. |
| 1 | 6:0 | <i>LEN</i> | See Table 4-14. |
| 1 | 7 | <i>WnR</i> | 0b for a Read Response. 1b for a Write Response. |
| m:2 | 7:0 | RESPONSE_DATA | For a Write Response: One of the Result Codes defined in Table 4-14. For a Read Response: The sequence of bytes returned from target. Each Symbol contains one byte as payload. Register contents shall appear low-ordered byte first. |

When a Router receives an AT Command or an RT Command, it shall process the AT or RT Command according to Table 4-14 and in the order listed.

Table 4-14. Processing of a Received AT/RT Command

| STX Symbol Case | Action Taken | <i>LEN</i> Field in Response | <i>RESPONSE_DATA</i> field in Response |
|---|--|--|--|
| Operation with a single Data Symbol. | Send Response with no <i>RESPONSE_DATA</i> . | 0 | None. |
| Read Operation cases: | | | |
| Read operation to an unsupported vendor defined register or to an undefined register. | Send Response with no <i>RESPONSE_DATA</i> . | 0 | None. |
| Read operation with more than two Data Symbols. | Send Response with no <i>RESPONSE_DATA</i> . | 0 | None. |
| Read operation of more bytes than the target register length. | Send Response with a truncated read of the entire register contents. | Set to the size in bytes of the register being accessed. | Register being accessed. |
| Read operation of fewer bytes than the target register length. | Send Response with the number of bytes requested. | | Bytes requested. |
| All other read operations. | Perform the operation. Send Response. | | Register being accessed. |

| STX Symbol Case | Action Taken | LEN Field in Response | RESPONSE_DATA field in Response |
|--|---|---|---------------------------------|
| Write Operation cases: | | | |
| Write operation to an unsupported vendor defined register or to an undefined register. | Do not perform the write operation. Send Response. | 0 | Result Code = 01h (ERROR). |
| Write operation with <i>LEN</i> field that does not match the Transaction size. | | | |
| Write operation to a register with all Read-only (RO) fields. | | | |
| Write operation of more bytes than the target register length. | | | |
| Write operation of less bytes than the target register length. | Perform the write operation only on the requested bytes. Send Response. | Size in bytes of the register being accessed. | Result Code = 00h (SUCCESS). |
| All other write operations. | Perform the write operation. Send Response. | Size in bytes of the register being accessed. | Result Code = 00h (SUCCESS). |

4.1.1.3.2 Connection Manager Access

A Connection Manager initiates a read from or write to SB Register Space via the USB4 Port Capability in Adapter Configuration Space. The Connection Manager writes to the *Target*, *Address*, *WnR*, *LEN*, *Index*, and *Data* fields then sets the *Pending* bit to 1b, which causes the Router to issue a SB Register access. When the SB Register access completes, the Router writes back to the *LEN*, *No Response*, *Result Code*, and *Data* fields and sets the *Pending* bit to 0b.

When the *Pending* bit in a USB4 Port Capability is set to 1b, a Router shall:

- If the *Target* field in the USB4 Port Capability is set to 000b, the Router shall issue an access to the SB Register Space of the Port.
 - The Router shall access the register identified in the *Address* field in the USB4 Port Capability.
 - The Router shall access the number of bytes indicated in the *Length* field.
 - If the *WnR* field is set to 0b, the Router shall read from the SB Register Space.
 - If the *WnR* field is set to 1b, the Router shall write to the SB Register Space. The Router shall write the contents from the *Data* DWs in the USB4 Port Capability. Register contents shall be written least significant byte first.
- If the *Target* field in the USB4 Port Capability is set to 001b, the Router shall send an AT Command on the Sideband Channel of the Port to access the SB Register Space of the Link Partner. The AT Command shall have the following contents:
 - The *REG* field shall be set to the contents of the *Address* field in the USB4 Port Capability.
 - The *LEN* field shall be set to the contents of the *Length* field in the USB4 Port Capability.
 - The *WnR* field shall be set to the contents of the *WnR* field in the USB4 Port Capability.

- If the *WnR* field is set to 1b, the *COMMAND_DATA* Symbols shall contain the contents from the *Data* DWs in the USB4 Port Capability. Register contents shall appear least significant byte first.
- If the *Target* field in the USB4 Port Capability is set to 010b, the Router shall send an Addressed RT Command on the Sideband Channel of the Port to access the SB Register Space of a Re-timer on the Link. The Addressed RT Command shall have the following contents:
 - The *Index* field shall be set to the contents of the Re-timer Index in the USB4 Port Capability.
 - The *REG* field shall be set to the contents of the *Address* field in the USB4 Port Capability.
 - The *LEN* field shall be set to the contents of the *Length* field in the USB4 Port Capability.
 - The *WnR* field shall be set to the contents of the *WnR* field in the USB4 Port Capability.
 - If the *WnR* field is set to 1b, the *COMMAND_DATA* Symbols shall contain the contents from the *Data* DWs in the USB4 Port Capability. Register contents shall appear least significant byte first.

A Router shall process an AT Command as described in Table 4-14. When a response to a local access is ready or when a Response Transaction is received, the Router shall update the USB4 Port Capability as follows:

- The *LEN* field in the AT Response, the RT Response, or the local access is copied to the *Length* field in the USB4 Port Capability.
- The *No Response* bit is updated.
- The *Result Code* bit in the USB4 Port Capability is updated.
- For a read, the *Data* DWs in the USB4 Port Capability are updated.

The Router shall then set the *Pending* bit in the USB4 Port Capability to 0b.

4.1.1.3.3 SB Register Definitions

Table 4-15 lists the registers in the SB Register Space of a Router. See the USB4 Re-Timer Specification for the structure of the SB Register Space of a Re-timer.

Table 4-15. SB Registers

| Register | Size (Bytes) | Name | Description |
|----------|--------------|---------------------------|---|
| 0 | 4 | <i>Vendor ID</i> | Identifies the manufacturer of the Router silicon. |
| 1 | 4 | <i>Product ID</i> | Identifies the type of the Router. |
| 2 to 7 | N/A | -- | Undefined. |
| 8 | 4 | <i>Opcode</i> | A Port Operation Opcode in FourCC format. |
| 9 | 4 | <i>Metadata</i> | Metadata written or read with a Port Operation. |
| 10 to 11 | N/A | -- | Undefined. |
| 12 | 3 | <i>Link Configuration</i> | Defines Link Configuration for a USB4 Port. |
| 13 | 4 | <i>TxFE</i> | Used to exchange the TxFE parameters of USB4 Link transmitters. |
| 14 | N/A | -- | Undefined. |

| Register | Size (Bytes) | Name | Description |
|------------|-----------------|---------------------------------|--|
| 15 | 4 | <i>Sideband Channel Version</i> | A vendor defined version of the Sideband Channel implementation. |
| 16 to 17 | vendor specific | <i>vendor specific</i> | Vendor specific register. |
| 18 | 64 | <i>Data</i> | Data written or read with a Port Operation. |
| 19 to 127 | vendor specific | vendor specific | Vendor specific register. |
| 128 to 255 | N/A | -- | Undefined. |

Table 4-16 defines the Commands that are allowed for a particular SB Register field.

Table 4-16. SB Register Fields Access Types

| Access Type | Description |
|-------------|--|
| RO | Read Only. A Write operation to a field with this access type shall have no effect. A Read operation shall return a meaningful value. |
| RW | Read/Write. A field with this access type shall be capable of both Read operation and Write operation. The value read from this field shall reflect the last value written to it unless the field was reset in the interim. |
| Rsvd | Reserved. Reserved for future implementation. A Write operation to this field shall have no effect. |

The SB Register Space registers shall have the structure and fields described in Table 4-17. Registers not listed in Table 4-17 are undefined and shall not be used.

Table 4-17. SB Register Fields

| Register | Register Name | Byte | Bits | Field Name and Description | Type | Default Value |
|----------|---------------|------|------|---|------|----------------|
| 0 | Vendor ID | 0 | 7:0 | Vendor ID Low – Identifies the manufacturer of the Router silicon. Shall contain the same value as the lower byte of the <i>Vendor ID</i> field in Router Configuration Space. | RO | Vendor Defined |
| | | 1 | 7:0 | Vendor ID High – Identifies the manufacturer of the Router silicon. Shall contain the same value as the higher byte of the <i>Vendor ID</i> field in Router Configuration Space. | RO | Vendor Defined |
| | | 2 | 7:0 | Reserved | Rsvd | 0 |
| | | 3 | 7:0 | Reserved | Rsvd | 0 |
| 1 | Product ID | 0 | 7:0 | Product ID Low – Assigned by the manufacturer to identify the type of the Router. Shall contain the same value as the lower byte of the <i>Product ID</i> field in Router Configuration Space. | RO | Vendor Defined |
| | | 1 | 7:0 | Product ID High – Assigned by the manufacturer to identify the type of the Router. Shall contain the same value as the higher byte of the <i>Product ID</i> field in Router Configuration Space. | RO | Vendor Defined |
| | | 2 | 7:0 | Reserved | Rsvd | 0 |
| | | 3 | 7:0 | Reserved | Rsvd | 0 |
| 8 | Opcode | 0 | 7:0 | Opcode 0 Contains the first character of the Opcode. | RW | 0 |
| | | 1 | 7:0 | Opcode 1 Contains the second character of the Opcode. | RW | 0 |

| Register | Register Name | Byte | Bits | Field Name and Description | Type | Default Value |
|----------|--------------------|------|------|--|------|----------------|
| 8 | Opcode | 2 | 7:0 | Opcode 2 Contains the third character of the Opcode. | RW | 0 |
| | | 3 | 7:0 | Opcode 3 Contains the fourth character of the Opcode. | RW | 0 |
| 9 | Metadata | 3:0 | 7:0 | Metadata A Connection Manager optionally writes to this field when initiating a Port Operation. A Router shall write the Completion Metadata (if any) to this field after executing a Port Operation. See Section 8.3.2 for more information on Port Operations. | RW | 0 |
| 12 | Link Configuration | 0 | 0 | Enabling Decision (Lane 0) – Shall indicate whether or not the Lane 0 Adapter is enabled during Lane Initialization. 0b – Not enabled 1b – Enabled | RO | 0b |
| | | | 1 | Enabling Decision (Lane 1) – Shall indicate whether or not the Lane 1 Adapter is enabled during Lane Initialization. 0b – Not enabled 1b – Enabled | RO | 0b |
| | | | 7:2 | Reserved | Rsvd | 0 |
| | | 1 | 0 | Enabling Request (Lane 0) – Indicates whether the Router requests enabling for Lane 0. 0b – No request to enable 1b – Request to enable A Router shall set this bit to 0b when the <i>Lane Disable</i> bit in the Lane Adapter Configuration Capability for Lane 0 is 1b. | RO | Vendor Defined |
| | | | 1 | Enabling Request (Lane 1) – Indicates whether the Router requests enabling for Lane 1. 0b – No request to enable 1b – Request to enable A Lane 1 Adapter shall not request enabling unless the Lane 0 Adapter requests enabling. A USB4 Port shall only set this bit to 1b if all On-Board Re-timers connected between the Router and the cable support dual Lanes. The method of conveying the capabilities of an On-Board Re-timer to the Router is implementation specific. A Router shall set this bit to 0b when the <i>Lane Disable</i> bit in the Lane Adapter Configuration Capability for Lane 1 is 1b. | RO | Vendor Defined |
| | | | 3:2 | Reserved | Rsvd | 0 |
| | | | 4 | Bonding Support – Indicates whether Lane Bonding is supported for this USB4 Port. 0b – Lane Bonding is not supported 1b – Lane Bonding is supported | RO | Vendor Defined |

| Register | Register Name | Byte | Bits | Field Name and Description | Type | Default Value |
|----------|--------------------|------|-----------------|---|------|--|
| 12 | Link Configuration | 1 | 5 | <p>Gen 3 Support – Indicates whether Gen 3 speeds are supported on this USB4 Port.</p> <p>0b – Gen 3 speeds are not supported 1b – Gen 3 speeds are supported</p> <p>A USB4 Port shall only set this bit to 1b if all of the following are true:</p> <ul style="list-style-type: none"> The Port supports Gen 3 speeds. All On-Board Re-timers connected between the Port and the cable support Gen 3 speeds. <p><i>Note: The method of conveying the capabilities of an On-Board Re-timer to the Router is implementation specific.</i></p> <ul style="list-style-type: none"> The Port implements both a Lane 0 Adapter and a Lane 1 Adapter. The <i>Target Link Speed</i> field in the Lane Adapter Configuration Capability is 1100b. <p>Otherwise this bit shall be 0b.</p> | RO | 1b for a USB4 hub, else Vendor Defined |
| | | | 6 | <p>RS-FEC Request (Gen 2) – Indicates whether the Router requests enabling RS-FEC in Gen 2 Speeds for this USB4 Port.</p> <p>0b – No request to enable 1b – Request to enable</p> <p>A USB4 Port shall set this bit to the same value as the <i>Request RS-FEC Gen 2</i> bit in the USB4 Port Capability.</p> | RO | 1b |
| | | | 7 | <p>RS-FEC Request (Gen 3) – Indicates whether the Router requests enabling RS-FEC in Gen 3 Speeds for this USB4 Port.</p> <p>0b – No request to enable 1b – Request to enable</p> <p>A USB4 Port shall set this bit to the same value as the <i>Request RS-FEC Gen 3</i> bit in the USB4 Port Capability.</p> | RO | 1b |
| | | 2 | 0 | <p>USB4 Sideband Channel Support</p> <p>Indicates that this USB4 Port supports a Sideband Channel as defined in Chapter 4 of this specification. Shall be set to 1b.</p> | RO | 1b |
| | | | 1 | <p>TBT3-Compatible Speeds Support – Indicates whether TBT3-Compatible speeds are supported on this USB4 Port.</p> <p>0b – TBT3-Compatible speeds are not supported 1b – TBT3-Compatible speeds are supported</p> <p>A USB4 Port shall only set this bit to 1b if all On-Board Re-timers connected between the Router and the cable support TBT3-Compatible Mode speeds. The method of conveying the capabilities of an On-Board Re-timer to the Router is implementation specific.</p> | RO | Vendor Defined |
| | | 7:2 | Reserved | Rsvd | 0 | |

| Register | Register Name | Byte | Bits | Field Name and Description | Type | Default Value |
|----------|---------------|------|------|--|------|----------------|
| 13 | TxFFE | 0 | 3:0 | TxFFE Request (Lane 0) – Identifies one of 16 predefined TxFFE configurations requested by the receiver. | RO | Vendor Defined |
| | | | 4 | Rx Locked (Lane 0) – Indicates that the receiver completed TxFFE negotiation. 0b – TxFFE negotiation is not done 1b – TxFFE negotiation is done | RO | 0b |
| | | | 5 | Rx Active (Lane 0) – Indicates whether or not the receiver is active. 0b – Receiver is inactive 1b – Receiver is active | RO | 0b |
| | | | 6 | Clock Switch Done (Lane 0) Shall be set to the value of the <i>Rx Locked (Lane 0)</i> bit. | RO | 0b |
| | | | 7 | New Request (Lane 0) – Indicates whether or not the receiver is processing a TxFFE Request. 0b – Receiver finished processing the previous request 1b – Receiver is still processing the previous request | RO | 0b |
| | | 1 | 3:0 | TxFFE Request (Lane 1) – Identifies one of 16 predefined TxFFE configurations requested by the receiver. | RO | Vendor Defined |
| | | | 4 | Rx Locked (Lane 1) – Indicates that the receiver completed TxFFE negotiation. 0b – TxFFE negotiation is not done 1b – TxFFE negotiation is done | RO | 0b |
| | | | 5 | Rx Active (Lane 1) – Indicates whether or not the receiver is active. 0b – Receiver is inactive 1b – Receiver is active | RO | 0b |
| | | | 6 | Clock Switch Done (Lane 1) Shall be set to the value of the <i>Rx Locked (Lane 1)</i> bit. | RO | 0b |
| | | | 7 | New Request (Lane 1) – Indicates whether or not the receiver is processing a TxFFE Request. 0b – Receiver finished processing the previous request 1b – Receiver is still processing the previous request | RO | 0b |
| | | 2 | 3:0 | TxFFE Setting (Lane 0) – Index of the TxFFE configuration to be loaded to the transmitter. | RW | Vendor Defined |
| | | | 5:4 | Reserved | Rsvd | 0 |
| | | | 6 | Request Done (Lane 0) – Indicates whether the transmitter loaded the recent requested index of TxFFE configuration. 0b – Transmitter has not yet loaded the recent requested index of TxFFE configuration 1b – Transmitter has loaded the recent requested index of TxFFE configuration | RW | 0b |
| | | | 7 | Tx Active (Lane 0) – Indicates whether the transmitter is transmitting a valid signal. 0b – Transmitter is not transmitting a valid signal 1b – Transmitter is transmitting a valid signal | RW | 0b |

| Register | Register Name | Byte | Bits | Field Name and Description | Type | Default Value | |
|--|--------------------------|------|------|--|--|----------------|----------------|
| 13 | TxFFE | 3 | 3:0 | TxFFE Setting (Lane 1) – Index of the TxFFE configuration to be loaded to the transmitter. | RW | Vendor Defined | |
| | | | 5:4 | Reserved | Rsvd | 0 | |
| | | | 6 | Request Done (Lane 1) – Indicates whether the transmitter loaded the recent requested index of TxFFE configuration. 0b – Transmitter has not yet loaded the recent requested index of TxFFE configuration 1b – Transmitter has loaded the recent requested index of TxFFE configuration | RW | 0b | |
| | | | 7 | Tx Active (Lane 1) – Indicates whether the transmitter is transmitting a valid signal. 0b – Transmitter is not transmitting a valid signal 1b – Transmitter is transmitting a valid signal | RW | 0b | |
| 15 | Sideband Channel Version | 0 | 7:0 | Sub Version – The binary coded decimal digit of the sub version number of the Sideband Channel implementation. Contains a vendor defined value. | RO | Vendor Defined | |
| | | | 1 | 7:0 | Minor Version – The binary coded decimal digit of the minor version number of the Sideband Channel implementation. Contains a vendor defined value. | RO | Vendor Defined |
| | | | 2 | 7:0 | Major Version LSB – The least-significant binary coded decimal digit of the major version number of the Sideband Channel implementation. Contains a vendor defined value. | RO | Vendor Defined |
| | | | 3 | 7:0 | Major Version MSB – The most-significant binary coded decimal digit of the major version number. Contains a vendor defined value. | RO | Vendor Defined |
| 18 | Data | 63:0 | 7:0 | Data A Connection Manager optionally writes to this field when initiating a Port Operation. A Router shall write the Completion Data (if any) to this field after executing a Port Operation. See Section 8.3.2 for more information on Port Operations. | RW | 0 | |
| Notes: | | | | | | | |
| <ol style="list-style-type: none"> Byte 0 in Register 13 is the <i>Rx Status & TxFFE Request</i> byte for Lane 0. Byte 1 in Register 13 is the <i>Rx Status & TxFFE Request</i> byte for Lane 1. Byte 2 in Register 13 is the <i>Tx Status</i> byte for Lane 0. Byte 3 in Register 13 is the <i>Tx Status</i> byte for Lane 1. | | | | | | | |


4.1.2 Lane Initialization

This section defines how a Sideband Channel is used to initialize the Lane(s) that make up a USB4 Link. Lane Initialization is the process by which the Electrical Layer of a Lane goes from inactive to actively transmitting and receiving traffic.

The Sideband Channel shall initialize each Lane independently. A Sideband Channel is used to initialize all enabled Lanes in its USB4 Port. Initialization shall occur for all enabled USB4 Ports on a Router.

Lane Initialization consists of a series of phases, which are summarized in Figure 4-5 and described in the subsections below. Lane Initialization starts with the Adapter in CLd state and ends with the Adapter in Training.LOCK1 state. The transition from CLd state to Training state is defined in Section 4.2.1.2.3.

Figure 4-5. Overview of Lane Initialization



Lane Initialization is described from the point of view of a single Router and its Link Partner, which are called “Router A” and “Router B” respectively for distinction. Note that Lane Initialization is symmetrical, meaning that the steps defined for Router A are also followed by Router B in parallel (with Router A and Router B switching roles) for each phase of Lane Initialization. Furthermore, the steps for each phase are defined within the context of a single Lane and are repeated for each Lane shared between Router A and Router B.

4.1.2.1 Phase 1 – Determination of Initial Conditions

During phase 1, Router A discovers the following connection information:

- Whether or not the Link is USB4.
- Whether or not there is a reverse insertion at the USB Type-C® connector.
- Whether or not it is connected by an Active Cable that contains Re-timers.
 - If connected by an Active Cable that contains Re-timers, whether or not the cable is a TBT3 Active Cable (See Section 13.2.1.1).
 - Else, whether or not Cable supports Gen 3 speed.

A USB4 Port shall not continue on to Phase 2 until it has obtained the connection information described in this section. See the USB PD Specification and the USB Type-C Specification for how to determine the connection information.

A USB4 Port shall not proceed to Phase 2 if the Link is not USB4.

A USB4 Port shall drive SBTX to logic low by default.



IMPLEMENTATION NOTE

Each Router has a pull-up resistor that pulls SBTX up to logical high unless the Router actively drives SBTX low. Because of this, it is necessary to actively drive SBTX in Phase 1, so that the pull-up resistor doesn't cause the Sideband Channel to inadvertently read as logic high.

4.1.2.1.1 Lane Reversal

Lane mismatch is caused by reverse insertion of a USB Type-C connector. Lane mismatch results in a Link where:


- The SBTX wires for the USB4 Ports at each end of the Link are connected to each other.
- The SBRX wires for the USB4 Ports at each end of the Link are connected to each other.
- The Lane 0 Adapters at each end of the Link are connected to the Lane 1 Adapters at the other end of the Link.

When a Router detects a reverse insertion on a USB Type-C connector, it shall perform Lane Reversal in the USB4 Port that faces the reversed connector. Lane Reversal shall be performed during phase 1. Lane Reversal consists of the following:

- The Router shall swap the designation of Lane 0 and Lane 1 and shall associate the Lane 0 Adapter with the updated Lane 0 and the Lane 1 Adapter with the updated Lane 1.
- If there are no On-Board Re-timers between Router and the USB Type-C connector, the Router shall swap its SBTX and SBRX lines facing the connector.

Figure 4-6 show an example of a Link before and after Lane reversal. Note that in Figure 4-6 Lane mismatch occurs on both ends of the Link, therefore both ends of the Link perform Lane Reversal.

Figure 4-6. Example of Lane Reversal



4.1.2.1.2 Polarity Inversion

Polarity mismatch is caused by connecting a Lane between a Router and an On-Board Re-timer or a Router and a USB Type-C connector such that the positive half of a differential pair at one end is connected to the negative half of a differential pair at the other end.

A Router may be configured to perform polarity inversion during Phase 1 of Lane Initialization. The method to perform polarity inversion is outside the scope of this specification.

4.1.2.2 Phase 2 – Router Detection

The second phase of Lane Initialization consists of each Router detecting any connected Routers.

After completion of phase 1, a Host Router shall initiate Router detection by driving SBTX to logic high on all of its Downstream Facing Ports.

When a Device Router detects a logic high on SBRX of its Upstream Facing Port for $t_{ConnectRx}$ time, it shall drive SBTX to logic high on all of its USB4 Ports.

After a USB4 Port drives its SBTX to logic high and detects a logic high on its SBRX for $t_{ConnectRx}$ time, it shall set its *Router Detected* bit to 1b, then transition to Phase 3 of Lane Initialization.

Note: If a USB4 Port receives a Transaction while it is still in Phase 2, it may ignore or drop the Transaction.

4.1.2.3 Phase 3 – Determination of USB4 Port Characteristics

The third phase of Lane Initialization consists of each Router acquiring information about its Link Partner.

During phase 3, Router A shall read the Link Configuration register (register 12) of Router B using AT Transactions. Router A shall issue at least one more register read to Router B in order to avoid a $t_{ATTimeout}$ delay at the Link Partner (see Section 4.1.1.2.5). Router A may re-read Register 12 or it may read other SB Registers during this phase.

Router A shall then decide the following Lane attributes using the decision criteria in Table 4-18.

Table 4-18. Lane Attributes

| Attribute | Action |
|-----------|---|
| Enabling | If <i>Enabling Request</i> = 1b in the SB Register Space of the Adapters of both Router A and Router B, Router A shall: <ul style="list-style-type: none"> • Proceed to initialize the Lane. • Set the local <i>Enabling Decision</i> bit for the Adapter to 1b. Else, Router A shall: <ul style="list-style-type: none"> • Not initialize the Lane. The Adapter shall remain in CLd state. • Set the local <i>Enabling Decision</i> bit for the Adapter to 0b. |
| Dual-Lane | Router A shall set the <i>Bonding Enabled</i> bit in the USB4 Port Capability to 1b if all of the following are true: <ul style="list-style-type: none"> • The USB4 Ports of Router A and Router B both have the <i>Enabling Request</i> bit set to 1b for both Lanes of the USB4 Port. • The USB4 Ports of Router A and Router B both support Lane bonding (i.e. the <i>Bonding Support</i> bit is 1b in the SB Register Space of the Adapters on both sides of the Lane). Otherwise, Router A shall set the <i>Bonding Enabled</i> bit to 0b. |

| Attribute | Action |
|-----------|---|
| Speed | <p>Gen 2 vs. Gen 3 selection:</p> <p>Router A shall operate at Gen 3 Lane speed if all of the following are true:</p> <ul style="list-style-type: none"> • The Adapters on both sides of the Lane support Gen 3 (<i>Gen 3 Support</i> is set to 1b). • The cable over which Router A and Router B are communicating supports Gen 3. <p>Otherwise, Router A shall operate at Gen 2 Lane Speed.</p> <p>Router A shall set the <i>Current Link Speed</i> field in the Lane Adapter Configuration Capability to reflect whether it is operating at Gen 2 or Gen 3 Lane Speed.</p> |
| RS-FEC | <p>Gen 2 speed:</p> <p>Router A shall enable RS-FEC if both sides of the Link request it (i.e. the <i>RS-FEC Request (Gen 2)</i> bit is set to 1b in the SB Register Space of both the local USB4 Port and its Link Partner). Otherwise, RS-FEC shall not be enabled.</p> <p>Router A shall set the <i>RS-FEC Enabled (Gen 2)</i> bit in the USB4 Port Capability to reflect whether it is operating with RS-FEC.</p> <p>Gen 3 speed:</p> <p>Router A shall enable RS-FEC if both sides of the Link request it (i.e. the <i>RS-FEC Request (Gen 3)</i> bit is set to 1b in the SB Register Space of both the local USB4 Port and its Link Partner). Otherwise, RS-FEC shall not be enabled.</p> <p>Router A shall set the <i>RS-FEC Enabled (Gen 3)</i> bit in the USB4 Port Capability to reflect whether it is operating with RS-FEC.</p> |

Note: The values that are read from the Link Configuration registers in Phase 3 determine the Link characteristics for the rest of the Lane Initialization sequence. If the values in the Link Configuration registers change after they are read in Phase 3, the new values do not take effect until the next Lane Initialization.

4.1.2.4 Phase 4 – Lane Parameters Synchronization and Transmit Start

During Phase 4, a USB4 Port sets Lane speed and start transmission. The following sequence takes place on the Sideband Channel between Router A and Router B:

1. Router A shall send a Broadcast RT Transaction every $t_{\text{LaneParams}}$ with the parameter values in Table 4-18. Router A shall continue sending Broadcast RT Transactions until all of the following conditions are true, then continue to Step 2:
 - At least t_{LTPhase4} time has passed from completion of Phase 2.
 - Router A has sent Broadcast RT Transactions at least twice.
 - Router A has received a Broadcast RT Transaction from Router B.
2. Router A shall activate the transmitter on each enabled Lane at the selected speed and shall send SLOS1 as defined in Section 4.2.1.3.
3. After its transmitter is transmitting a valid signal, Router A shall then send an LT_Resume Transaction for each enabled Lane in the USB4 Port and shall set to 1b the *Tx Active* bit in the Tx Status byte of the TxFFE Register in the SB Register Space to indicate that it has started transmission on the target Lane. The *LSELane* field in the LT_Resume Transaction shall equal the Lane number associated with the transmitter.

Note: The behavior of the transmitter is undefined prior to setting the Tx Active bit to 1b.

4. Router A's USB4 Port goes to Phase 5.


4.1.2.5 Phase 5 – Link Equalization

The fifth phase of Lane Initialization consists of negotiating transmitter Feed-Forward Equalization (TxFFE) parameters between each Router or Re-timer and a Router or Re-timer adjacent to it. During TxFFE negotiation, the receiver cycles through one or more of the potential preset numbers defined in Table 3-5. The receiver examines its behavior for each

preset and selects one preset value to use. The receiver may follow any order while cycling through the preset numbers.

Figure 4-7 shows two Routers that share a Link that includes six Re-timers. In Figure 4-7, Router A's transmitters perform TxFFE negotiation with the receivers of On-Board Re-timer 1. The transmitters in On-Board Re-timer 6 perform TxFFE negotiation with the receivers of Router B.

Figure 4-7. Progression of Link Equalization



Router A's transmitter shall perform the transmitter flow in the symmetric equalization flow defined in Section 4.1.2.5.1.1. A Router shall use Addressed RT Transactions with the *Index* field set to 0 to access the SB Register Space of the adjacent component, which can be either an On-Board Re-timer, a Cable Re-timer, or Router B.

Router B's receiver shall perform the receiver flow in the symmetric equalization flow defined in Section 4.1.2.5.1.2. The Router shall use Addressed RT Transactions with the *Index* field set to 0b to access the SB Register Space of the adjacent component, which can be either an On-Board Re-timer, a Cable Re-timer, or Router A.

A Router shall set the *Clock Switch Done* bit to 1b in the SB Register space of the Adapters of a USB4 Port after all of the USB4 Port's receivers complete the equalization flow.

The equalization flow for Re-timers is defined in the USB4 Re-Timer Specification.

4.1.2.5.1 Symmetric TxFFE Negotiation

During symmetric TxFFE negotiation, the transmitter negotiates TxFFE parameters with a connected receiver. Both the transmitter and receiver can issue commands to access the SB Register Space of their negotiation partner.

4.1.2.5.1.1 Transmitter Flow

1. The transmitter shall start with the following default values in the *Tx Status* byte of the TxFFE register
 - *Tx Active* bit = 1b (see Section 4.1.2.4).
 - *Request Done* bit = 0b.
2. The transmitter shall read the *Rx Status & TxFFE Request* byte of the receiver. To do this, the transmitter sends a read Command that targets the receiver's TxFFE register.
3. On reception of a Response from the receiver, the transmitter shall do the following:
 - If *Rx Locked*=1, then negotiation is complete and no further TxFFE negotiation steps are taken.
 - Else, if *New Request*=0, the receiver has not provided a new request yet. The transmitter shall retry Step 2 within *tPollTxFFE* of receiving the Response.
 - Else, this is a new request to update TxFFE parameters. Continue on to Step 4.

4. The transmitter shall update its transmitter parameters based on the new parameters in the received Response. To do this, the transmitter loads one of 16 predefined TxFFE configurations that matches the *TxFFE Request* field in the Response. Table 3-5 contains the structure of the Transmit Equalization Presets.
 - The transmitter shall update its *Tx Status* byte with the following values:
 - *TxFFE Setting* = the index (from 16 possible values) loaded above to the TxFFE configuration configured at the transmitter.
 - *Request Done* = 1b.

It is recommended that the transmitter loads the requested TxFFE Preset and updates Tx Status byte within 1ms from reception of the new request to update TxFFE parameters.
5. The transmitter shall read the *Rx Status & TxFFE Request* byte of the receiver. To do this, the transmitter sends a read Command that targets the receiver's TxFFE register.
6. On reception of a Response from the receiver, the transmitter shall do the following:
 - If *New Request*=1, the receiver is still trying to lock on a previous request. The transmitter shall retry Step 5 within tPollTXFFE of receiving the Response.
 - Else, the transmitter shall set the *Request Done* bit in the *Tx Status* byte to 0b, and return to Step 2.

4.1.2.5.1.2 Receiver Flow

1. The receiver shall start with the following default values in the *Rx Status & TxFFE Request* byte of the TxFFE register:
 - *Rx Locked* = 0b.
 - *New Request* bit = 0b.
 - *Rx Active* bit = 0b.
2. The receiver shall read the transmitter's *Tx Status* byte. It does so by sending a read Command to the transmitter that targets its TxFFE register.
3. On reception of a Response from the transmitter, the receiver shall do the following:
 - If *Tx Active* = 1b, then enable the receiver, set the *Rx Active* bit to 1b, and continue on to Step 4.
 - Else, repeat Step 2 within tPollTXFFE of receiving the Response.
4. The receiver shall evaluate its receiver behavior and shall set the *Rx Locked* bit to 1b if equalization is complete.
5. The receiver shall do the following:
 - If *Rx Locked* = 1, then TXFFE negotiation is complete and no further negotiation steps are taken.
 - Else, the receiver shall:
 - Select a new set of TxFFE parameters and shall set the *TxFFE Request* field to the index of the selected set of TXFFE parameters.
 - Set the *New Request* bit to 1b.
 - Continue with the steps below.
6. The receiver shall read the transmitter's *Tx Status* byte. It does so by sending a read Command to the transmitter that targets its TxFFE register.
7. On reception of a Response from the transmitter, the receiver shall do the following:
 - If (*Tx Active* = 1b) AND (*Request Done* = 1b) AND (*TxFFE Setting* = value of *TxFFE request* in the local *Rx Status & TxFFE Request* byte), then continue on to Step 8.
 - Else, repeat Step 6 within tPollTXFFE of receiving the Response.

8. The receiver shall evaluate its receiver behavior and set the *Rx Locked* bit to 1b if equalization is complete.
9. The receiver shall set the *New Request* bit to 0b.
10. The receiver shall read the transmitter's *Tx Status* byte by sending a read Command to the transmitter that targets its TxFFE register.
11. On reception of a Response from the transmitter, the receiver shall do the following:
 - If (Tx Active = 1b) and (Request Done = 0b), then go to Step 5.
 - Else, repeat Step 10 within tPollTXFFE of receiving the Response.

4.2 Logical Layer State Machine


The behavior of the Logical Layer in an Adapter is described by the state machine defined in Section 4.2.1. Section 4.2.2 defines the behavior of a USB4 Link and how the Link is affected by Adapter state.

4.2.1 Lane Adapter State Machine

The state machine in Figure 4-8 describes the behavior of the Logical Layer in a Lane Adapter. A detailed description of the states and transitions between states follows.

- Disabled state – The Lane Adapter disables the Lane.
- CLd state – Lane Adapter transmitter and receiver are inactive.
- Training state – The Lane Adapter performs Symbol synchronization and transfer of Lane parameters.
- CL0 state – The Lane Adapter can transmit and receive Transport Layer Packets across the Lane.
- Lane Bonding state – bonds two Single-Lane Links into a Dual-Lane Link.
- CL0s, CL1, CL2 states – Low Power states.

Figure 4-8. The Lane Adapter State Machine



4.2.1.1 Disabled

4.2.1.1.1.1 Entry to State

An Adapter shall enter this state from Training state or CL0 state when the *Lane Disable* bit in the Lane Adapter Configuration Capability is set to 1b. See Section 4.4.6 for details.

A Lane Adapter shall set the *Plugged* bit to 0b upon transitioning to Disabled state.

4.2.1.1.1.2 Behavior in State

Lane common mode voltages do not need to be maintained while in this state.

4.2.1.1.1.3 Exit from State

An Adapter shall exit this state when the *Lane Disable* bit in the Lane Adapter Configuration Capability is set to 0b. See Section 4.4.6 for details.

Note: This is the only event on which the Lane Adapter exits the Disabled state.

A disabled Adapter shall stay in the Disabled state for a minimum of *tDisabled*. If the *Lane Disable* bit is set to 0b less than *tDisabled* after sending the LT_Fall Transaction, the Adapter shall not transition to the CLd state until *tDisabled* has elapsed.

4.2.1.2 CLd

4.2.1.2.1 Entry to State

A Lane Adapter shall enter this state on any of the following events:

- Router power-on.
- The USB4 Port is disconnected and *Lane Disable* bit in the Lane Adapter Configuration Capability is set to 0b (see Section 4.4.5).
- Router enters Sleep state (see Section 4.5.1).

In addition to the events listed above, a Lane Adapter that is not the Upstream Adapter shall enter this state on any of the following events:

- Adapter exits from the Disabled state.
- Adapter receives an LT_Fall Transaction.

A Lane Adapter shall set the *Plugged* bit to 0b upon transitioning to the CLd state.

4.2.1.2.2 Behavior in State

During this state, Lane common mode voltages do not need to be maintained. A Lane Adapter is in this state when Lane Initialization begins. Lane Initialization is described in Section 4.1.2.

A Lane Adapter that enters this state due to a disconnect shall perform Lane Initialization starting from Phase 1, except for the following cases:

- Downstream Facing Port Disconnect due to reception of LT_LRoff (see Section 4.4.5.2.2).

A Lane Adapter that enters this state from the Disabled State performs Lane Initialization after the Lane is enabled. The Lane Adapter shall start Lane Initialization from Phase 4. The USB4 Port shall maintain any state acquired in Phases 1 through 3 of previous Lane Initialization. See Section 4.4.6 for more information on Lane Disable and Enable.

A Lane Adapter that enters this state due to the Router entering Sleep state performs Lane Initialization after a Wake event. The Lane Adapter shall start Lane Initialization from Phase 2. It is recommended that the Adapter starts Lane Initialization with the last set of TxFFE parameters used prior to entry to sleep state. See Section 4.5 for more information on Router Sleep and Wake.

Note: The Link Partner may not necessarily start Phase 5 with the last set of TxFFE parameters used prior to entry to sleep state. Therefore, a receiver needs to verify the TxFFE parameters initiated at the Link Partner's transmitter. (see Section 4.5.4)

A Lane Adapter that enters this state due to Link training timeout shall perform Lane Initialization starting from Phase 1.

A Lane Adapter (that is not the Upstream Adapter) that enters this state due to reception of an LT_Fall Transaction starts Lane Initialization when it receives a Broadcast RT Transaction. The Lane Adapter shall start Lane Initialization from Phase 4. The USB4 Port shall maintain any state acquired in Phases 1 through 3 of previous Lane Initialization. The Lane 0 Adapter shall not start Lane Initialization until it receives a Broadcast RT Transaction with the *Lane0Enabled* bit set to 1b. The Lane 1 Adapter shall not start Lane Initialization until it receives a Broadcast RT Transaction with the *Lane1Enabled* bit set to 1b.

4.2.1.2.3 Exit from State

A Lane Adapter shall exit this state when its transmitter is transmitting (completion of Phase 4 of Lane Initialization, see Section 4.1.2.4) and its receiver is enabled.

After exiting the CLd state, a Lane Adapter shall transition to the Training.LOCK1 sub-state.

4.2.1.3 Training

4.2.1.3.1 Entry to State

A Lane Adapter shall enter this state on any of the following events:

- After exiting the CLd state.
- When recovering from a USB4 Link error.
- After exiting the CL2 or CL1 states.

4.2.1.3.2 Behavior in State

During this state Symbols are synchronized and Lane parameters are transferred between the two ends of the Lane. The transmitter and receiver are on while in this state.

A Lane Adapter shall follow the Training sub-state machine described in Figure 4-9 with the behavior described in Table 4-19 and the sub-state transitions described in Table 4-20. The sub-state transitions shall occur within tTrainingTransition time from receiving the last bit of the relevant Symbols.

Figure 4-9. Training Sub-State Machine




Table 4-19. Transmitter Behavior in Training Sub-states

| State | Transmitter Behavior |
|-------|--|
| LOCK1 | Transmitter shall send back-to-back SLOS1. When a Lane Adapter transitions to the Training state from the CL0 state, its transmitter may transmit up to 16 Symbol Times of random bits before sending the first SLOS1 that is not RS-FEC encoded. |
| LOCK2 | Transmitter shall send back-to-back SLOS2. |
| TS1 | Transmitter shall send back-to-back TS1 Ordered Sets. |
| TS2 | Transmitter shall send back-to-back TS2 Ordered Sets. |

Table 4-20. Training Sub-State Machine Transitions

| Transition | From State | To State | Conditions* |
|------------|------------|----------|---|
| 1 | LOCK1 | LOCK2 | <ul style="list-style-type: none"> Received 2 SLOS Symbols in a row. Sent at least 2 complete SLOS1. Receiver completed TxFFE negotiation (i.e. <i>Rx Locked</i> = 1b). |
| 2 | LOCK2 | TS1 | <ul style="list-style-type: none"> Received 2 SLOS2 Symbols in a row. Sent at least 2 complete SLOS2. |
| 3** | TS1 | TS2 | Gen 2: <ul style="list-style-type: none"> Received 2 TS1 Ordered Sets in a row. Sent at least 32 TS1 Ordered Sets. Gen 3: <ul style="list-style-type: none"> Received 2 TS1 Ordered Sets in a row. Sent at least 16 TS1 Ordered Sets. |
| 4 | TS1 | LOCK2 | <ul style="list-style-type: none"> Received 2 SLOS1 Symbols in a row. |
| 5 | TS2 | LOCK2 | <ul style="list-style-type: none"> Received 2 SLOS Symbols in a row. |
| 6 | TS2 | CL0 | Gen 2: <ul style="list-style-type: none"> Received 2 TS2 Ordered Sets in a row. Sent at least 16 TS2 Ordered Sets. Gen 3: <ul style="list-style-type: none"> Received 2 TS2 Ordered Sets in a row. Sent at least 8 TS2 Ordered Sets. |
| 7 | Any | CLd | <ul style="list-style-type: none"> Adapter remains in Training state for tTrainingAbort1 or tTrainingAbort2 time. |
| 8 | Any | LOCK1 | <ul style="list-style-type: none"> A timeout error (see Section 4.4.2). <i>Note: This transition is optional.</i> |

* All conditions must be met before transition takes place.

** Router may also optionally transition from TS1 state to TS2 state after receiving 2 TS2 Ordered sets and sending the required number of TS1 Ordered Sets.

4.2.1.3.3 Exit from State

A Lane Adapter that transitions from CLd state to Training state shall complete training and transition to the CL0 state within tTrainingAbort1 after entering the Training state. If the Adapter does not transition to CL0 state within tTrainingAbort1, the Router shall initiate a Disconnect by driving SBTX to a logical low state for tDisconnectTx.

A Lane Adapter that transitions from a state other than CLd to Training state shall complete training and transition to the CL0 state within tTrainingAbort2 after entering the Training state.

If the Adapter does not transition to CL0 state within $t_{\text{TrainingAbort2}}$, the Router shall initiate a Disconnect by driving SBTX to a logical low state for $t_{\text{DisconnectTx}}$.

A Hot Plug Event Packet shall be sent on transition to CL0 state if the Adapter is part of a Downstream Facing Port and it entered Training state from a CLd state.

A Lane Adapter shall set the *Plugged* bit to 1b when it transitions from the Training state to the CL0 state.

4.2.1.3.4 SLOS1 and SLOS2

The Symbol Lock Ordered Sets (SLOS) are used to establish bit and Symbol lock, both for initial Lane training and for re-establishing bit/Symbol lock on a previously trained USB4 Link. These Ordered Sets have a high number of transitions to facilitate bit lock and easily detected Symbol boundaries to facilitate Symbol lock. There are two Ordered Sets used for Symbol Lock, SLOS1 and SLOS2, each of which consists of a defined series of bit patterns. The SLOS1 bit pattern is generated using a known pseudo-random binary sequence (PRBS11). A PRBS11 of length 2047 with the following characteristics is used:

- PRBS11 polynomial: $G(x) = x^{11} + x^9 + 1$
- PRBS11 initial state: 400h

The bit pattern for SLOS1 is composed of <0b, 2047 bits of PRBS11>. The bit pattern for SLOS2 is composed of <1b, 2047 bits of logically inverted PRBS11>.

When operating in Gen 2 mode with RS-FEC encoding disabled, SLOS shall be encoded using 64b/66b encoding. When operating in Gen 2 mode with RS-FEC encoding enabled, SLOS that are not RS-FEC encoded shall be encoded using 128b/132b encoding.

When operating in Gen 3 mode, SLOS that are not RS-FEC encoded shall be encoded using 128b/132b encoding.

Section 4.3.6 defines the structure of SLOS that are RS-FEC encoded.

The SLOS1 and SLOS2 for 64b/66b encoding are transmitted as 32 66-bit Symbols, where the n th Symbol equals <b10b, n 'th 64b word of the bit pattern>. See Table 4-21 for SLOS1 with 64b/66b encoding and Table 4-22 for SLOS2 with 64b/66b encoding.

The SLOS1 and SLOS2 for 128b/132b encoding are transmitted as 16 132-bit Symbols, where the n th Symbol equals <b1010b, n 'th 128b word of the bit pattern>. See Table 4-23 for SLOS1 with 128b/132b encoding and Table 4-24 for the SLOS2 with 128b/132b encoding.

Note: A Symbol containing part of a SLOS is referred to as a "SLOS Symbol".

The SLOS1 and SLOS2 shall not be scrambled, and the scrambler shall not advance upon receive/transmit.

When transmitting SLOS1 or SLOS2 using 64b/66b encoding, a Router shall transmit all 32 SLOS Symbols in their entirety. When using 128b/132b encoding, a Router shall transmit all 16 SLOS Symbols in their entirety. A Router shall not transmit an incomplete SLOS.

Table 4-21. SLOS1 (64b/66b Encoding)

| Symbol number | Sync Bits | SLOS1 Contents (63:0) |
|---------------|-----------|---|
| 0 | 10 | 01000000001010000001000100001010 10100100000001101000001110010001 |
| 1 | 10 | 10111010111010100010100001010001 00100010101101010000110000100111 |
| 2 | 10 | 10010111001110010111101110010010 10111011000010101110010000101110 |
| 3 | 10 | 10010010100110110001111011101100 10101011110000001001100001011111 |
| 4 | 10 | 00100100011101101011010110001100 01110111101101010010110000110011 |
| 5 | 10 | 1001111101111000010100110010001 11111010110000100011100101011011 |
| 6 | 10 | 10000110101100111000111110110110 00101101110100110101001111000011 |
| 7 | 10 | 10011001101111111110100000001001 00000101101000100110010101111110 |
| 8 | 10 | 00010000110010100111110001110001 10110110111011011010101101100000 |
| 9 | 10 | 11011100011101011011010001101100 10111011110010101001110000011101 |
| 10 | 10 | 100011010111011100010101101000 00011001000011111010011000100111 |
| 11 | 10 | 110101110001000101101010011000 00011111000011000110011110111111 |
| 12 | 10 | 0010100001110001001101101011101 10001001011101011001010001111000 |
| 13 | 10 | 10110011010011111100111000011110 11001100101111111100100000011101 |
| 14 | 10 | 0000110100100111001101110111101 01010001000000101010000100000100 |
| 15 | 10 | 10100010110001010011101000111010 01011010011001100111111111110000 |
| 16 | 10 | 00000110000000111100000110011000 11111111011000000101110000100101 |
| 17 | 10 | 10010110011110011111001111000111 10011011001111101111100010100011 |
| 18 | 10 | 01000101110010100101110001100101 10111110011010001111100101100011 |
| 19 | 10 | 10011101101111010110100100011001 10101111111000100000110101000111 |
| 20 | 10 | 00001011011001001101111011110100 10100100110001101111101110100010 |
| 21 | 10 | 10100101000001100010001111010101 10010000011110100011001001011111 |
| 22 | 10 | 01100100010111101010010010000110 11010011101100111010111110100010 |
| 23 | 10 | 00100101010101100000000111000000 11011000011101110011010101111100 |
| 24 | 10 | 00010001100010101111010000100100 10010110110110011011011111101101 |
| 25 | 10 | 00001011001001001111011011100101 10101110011000101111110100100001 |
| 26 | 10 | 00110100101111001100100111111101 11000001010110001000011101010011 |
| 27 | 10 | 0100001111001001100111011111101 01000001000010001010010101000110 |
| 28 | 10 | 00001011110001001001101011011110 00110100110111001111010111100100 |
| 29 | 10 | 01001110101011101000001010010001 00011010101011100000001011000001 |
| 30 | 10 | 00111000101110110100101011001100 00111111100110000011111100011000 |
| 31 | 10 | 01101111001110100111101001110010 01110111011101010101010000000000 |

Table 4-22. SLOS2 (64b/66b Encoding)

| Symbol number | Sync Bits | SLOS2 Contents (63:0) |
|---------------|-----------|---|
| 0 | 10 | 101111111010111110111011110101 01011011111110010111110001101110 |
| 1 | 10 | 01000101000101011101011110101110 11011101010010101111001111011000 |
| 2 | 10 | 01101000110001101000010001101101 01000100111101010001101111010001 |
| 3 | 10 | 01101101011001001110000100010011 01010100001111110110011110100000 |
| 4 | 10 | 11011011100010010100101001110011 10001000010010101101001111001100 |
| 5 | 10 | 01100000010000111101011001101110 00000101001111011100011010100100 |
| 6 | 10 | 01111001010011000111000001001001 11010010001011001010110000111100 |
| 7 | 10 | 0110011001000000000101111110110 11111010010111011001101010000001 |
| 8 | 10 | 11101111001101011000001110001110 01001001000100100101010010011111 |
| 9 | 10 | 00100011100010100100101110010011 01000100001101010110001111100010 |
| 10 | 10 | 01110010100010001110101010010111 11100110111100000101100111011000 |
| 11 | 10 | 00101000111011101001010101100111 11100000111100111001100001000000 |
| 12 | 10 | 11010111100011101100100101000010 01110110100010100110101110000111 |
| 13 | 10 | 01001100101100000011000111100001 00110011010000000011011111100010 |
| 14 | 10 | 11110010110110001100100010000010 1010111011111101010111101111011 |
| 15 | 10 | 01011101001110101100010111000101 10100101100110011000000000001111 |
| 16 | 10 | 11111001111111000011111001100111 00000000100111111010001111011010 |
| 17 | 10 | 01101001100001100000110000111000 01100100110000010000011101011100 |
| 18 | 10 | 10111010001101011010001110011010 01000001100101110000011010011100 |
| 19 | 10 | 01100010010000101001011011100110 01010000000111011111001010111000 |
| 20 | 10 | 11110100100110110010000100001011 01011011001110010000010001011101 |
| 21 | 10 | 01011010111110011101110000101010 01101111100001011100110110100000 |
| 22 | 10 | 10011011101000010101101101111001 00101100010011000101000001011101 |
| 23 | 10 | 11011010101010011111111000111111 00100111100010001100101010000011 |
| 24 | 10 | 11101110011101010000101111011011 01101001001001100100100000010010 |
| 25 | 10 | 11110100110110110000100100011010 01010001100111010000001011011110 |
| 26 | 10 | 11001011010000110011011000000010 00111110101001110111100010101100 |
| 27 | 10 | 10111100001101100110001000000010 10111110111101110101101010111001 |
| 28 | 10 | 11110100001110110110010100100001 11001011001000110000101000011011 |
| 29 | 10 | 10110001010100010111110101101110 11100101010100011111110100111110 |
| 30 | 10 | 11000111010001001011010100110011 11000000011001111100000011100111 |
| 31 | 10 | 10010000110001011000010110001101 10001000100010101010101111111111 |

Table 4-23. SLOS1 (128b/132b Encoding)

| Symbol number | Sync Bits | SLOS1 Contents (127:0) |
|---------------|-----------|--|
| 0 | 1010 | 01000000001010000001000100001010 10100100000001101000001110010001 10111010111010100010100001010001 00100010101101010000110000100111 |
| 1 | 1010 | 10010111001110010111101110010010 10111011000010101110010000101110 10010010100110110001111011101100 10101011110000001001100001011111 |
| 2 | 1010 | 00100100011101101011010110001100 01110111101101010010110000110011 10011111101111000010100110010001 11111010110000100011100101011011 |
| 3 | 1010 | 10000110101100111000111110110110 00101101110100110101001111000011 10011001101111111110100000001001 00000101101000100110010101111110 |
| 4 | 1010 | 00010000110010100111110001110001 101101101110110110101101100000 11011100011101011011010001101100 10111011110010101001110000011101 |
| 5 | 1010 | 100011010111011100010101101000 00011001000011111010011000100111 110101110001000101101010011000 00011111000011000110011110111111 |
| 6 | 1010 | 00101000011100010011011010111101 10001001011101011001010001111000 10110011010011111100111000011110 1100110010111111110010000011101 |
| 7 | 1010 | 00001101001001110011011101111101 01010001000000101010000100000100 10100010110001010011101000111010 01011010011001100111111111110000 |
| 8 | 1010 | 00000110000000111100000110011000 11111111011000000101110000100101 10010110011110011111001111000111 10011011001111101111100010100011 |
| 9 | 1010 | 01000101110010100101110001100101 10111110011010001111100101100011 10011101101111010110100100011001 10101111111000100000110101000111 |
| 10 | 1010 | 00001011011001001101111011110100 10100100110001101111101110100010 10100101000001100010001111010101 10010000011110100011001001011111 |
| 11 | 1010 | 01100100010111101010010010000110 11010011101100111010111110100010 00100101010101100000000111000000 11011000011101110011010101111100 |
| 12 | 1010 | 00010001100010101111010000100100 10010110110110011011011111101101 00001011001001001111011011100101 10101110011000101111110100100001 |
| 13 | 1010 | 00110100101111001100100111111101 11000001010110001000011101010011 0100001111001001100111011111101 01000001000010001010010101000110 |
| 14 | 1010 | 00001011110001001001101011011110 00110100110111001111010111100100 01001110101011101000001010010001 00011010101011100000001011000001 |
| 15 | 1010 | 00111000101110110100101011001100 00111111100110000011111100011000 01101111001110100111101001110010 01110111011101010101010000000000 |

Table 4-24. SLOS2 (128b/132b Encoding)

| Symbol number | Sync Bits | SLOS2 Contents (127:0) |
|---------------|-----------|--|
| 0 | 1010 | 101111111101011111011101110101 01011011111110010111110001101110 0100010100010101110101110101110 11011101010010101111001111011000 |
| 1 | 1010 | 01101000110001101000010001101101 01000100111101010001101111010001 01101101011001001110000100010011 01010100001111110110011110100000 |
| 2 | 1010 | 11011011100010010100101001110011 10001000010010101101001111001100 01100000010000111101011001101110 00000101001111011100011010100100 |
| 3 | 1010 | 01111001010011000111000001001001 11010010001011001010110000111100 0110011001000000000101111110110 11111010010111011001101010000001 |
| 4 | 1010 | 11101111001101011000001110001110 01001001000100100101010010011111 00100011100010100100101110010011 01000100001101010110001111100010 |
| 5 | 1010 | 01110010100010001110101010010111 11100110111100000101100111011000 001010001110111010010101100111 11100000111100111001100001000000 |
| 6 | 1010 | 11010111100011101100100101000010 01110110100010100110101110000111 01001100101100000011000111100001 00110011010000000011011111100010 |
| 7 | 1010 | 11110010110110001100100010000010 10101110111111010101111011111011 01011101001110101100010111000101 10100101100110011000000000001111 |
| 8 | 1010 | 11111001111111000011111001100111 00000000100111111010001111011010 01101001100001100000110000111000 01100100110000010000011101011100 |
| 9 | 1010 | 10111010001101011010001110011010 01000001100101110000011010011100 01100010010000101001011011100110 01010000000111011111001010111000 |
| 10 | 1010 | 11110100100110110010000100001011 01011011001110010000010001011101 01011010111110011101110000101010 01101111100001011100110110100000 |
| 11 | 1010 | 10011011101000010101101101111001 00101100010011000101000001011101 11011010101010011111111000111111 00100111100010001100101010000011 |
| 12 | 1010 | 11101110011101010000101111011011 01101001001001100100100000010010 11110100110110110000100100011010 01010001100111010000001011011110 |
| 13 | 1010 | 11001011010000110011011000000010 00111110101001110111100010101100 10111100001101100110001000000010 101111101111011101011010111001 |
| 14 | 1010 | 11110100001110110110010100100001 11001011001000110000101000011011 10110001010100010111110101101110 11100101010100011111110100111110 |
| 15 | 1010 | 11000111010001001011010100110011 11000000011001111100000011100111 10010000110001011000010110001101 10001000100010101010101111111111 |

4.2.1.3.5 TS1 and TS2 Ordered Sets

A TS1 Ordered Set and a TS2 Ordered Set shall have the structure in Table 4-25.

Table 4-25. TS1 and TS2 Ordered Sets

| Bits | Name | Description |
|-------|----------------------------|---|
| 63:59 | <i>Reserved</i> | Reserved. Transmitter may send any value and receiver shall ignore. It is recommended that this field be zero. |
| 58:56 | <i>Lane Bonding Target</i> | Lane Bonding Target. Used to set the target Link width (see also Section 4.2.1.5). A transmitter shall set this field according to the value of the <i>Target Link Width</i> field of the Lane Adapter Configuration Capability 000b – Establish two Single-Lane Links 001b – Establish a Dual-Lane Link All other values are reserved and shall not be used. |
| 55:48 | <i>Lane Number</i> | Lane Number. Transmitter shall set this value to match the Lane number. 00h – Lane 0 01h – Lane 1 All other values are reserved and shall not be used. |

| Bits | Name | Description |
|-------|------------------------------|---|
| 47:32 | <i>Reserved</i> | Reserved. Transmitter may send any value and receiver shall ignore. It is recommended that this field be zero. |
| 31:29 | <i>Reserved</i> | Reserved. Transmitter shall write 0 and receiver shall ignore. |
| 28:26 | <i>Lane Bonding Target 2</i> | Lane Bonding Target 2. Transmitter shall set this value to match the <i>Lane Bonding Target</i> field. Receiver shall ignore this field. |
| 25:16 | <i>Reserved</i> | Reserved. Transmitter may send any value and receiver shall ignore. It is recommended that this field be zero. |
| 15:10 | <i>TSID</i> | TSID. Training Sequence ID 10 0110b - TS1 01 1001b - TS2 |
| 9:0 | <i>SCR</i> | SCR - Shall be set to 00 1111 0010b to indicate that Ordered Set contents are scrambled. |

4.2.1.4 CLO

4.2.1.4.1 Entry to State

A Lane Adapter shall enter this state upon any of the following events:

- Successful completion of Lane training.
- Successful completion of Lane Bonding.
- Exit from CLOs state.

4.2.1.4.2 Behavior in State

When a Lane Adapter is in the CLO state, the Transport Layer may utilize the USB4 Link for data transfer. The transmitter and receiver are on while in this state.

4.2.1.4.3 Exit from State

A Lane Adapter shall exit this state after one of the following occurs:

- Adapter Disable. A Lane Adapter that exits this state due to an Adapter disable shall transition to either the Disabled state or the CLd state as defined in Section 4.4.6.
- Adapter disconnect. A Lane Adapter that exits this state due to a disconnect shall transition to the CLd state (see Section 4.4.5).
- Reception of an LT_Fall Transaction. The Lane Adapter shall transition to the CLd state.
- Transition to Training state when either:
 - An error event occurs that transitions the Lane Adapter to the Training.LOCK1 sub-state. Section 4.4.2 describes potential error events and how they are handled.
 - Reception of any 2 SLOS Symbols in a row transitions the Lane Adapter to either the Training.LOCK1 sub-state or the Training.LOCK2 sub-state.
- Transition to CLOs, CL1, or CL2 states (see Section 4.2.1.6.2).
- Transition to Lane Bonding state when either:
 - The *Lane Bonding* bit in the Lane Adapter Configuration Capability Register of either Adapter in the USB4 Port is set to 1b.
 - 3 TS1 Ordered Sets are received in a row.

A Lane Adapter shall not exit this state to enter Lane Bonding state while it is sending a Transport Layer Packet. The Adapter shall complete sending the Packet before entering Lane Bonding state.

4.2.1.5 Lane Bonding

4.2.1.5.1 Entry to State

A Lane Adapter shall enter this state from CL0 state on any of the following events:

- The *Lane Bonding* bit in the Lane Adapter Configuration Capability Register of either Adapter of the USB4 Port is set to 1b.
- Three TS1 Ordered Sets are received in a row.

4.2.1.5.2 Behavior in State

Lane Bonding is the state that bonds two Single-Lanes Links into a Dual-Lane Link. See Section 4.2.2 for details on transition to Dual-Lane operation and to Single-Lane operation.

A Lane Adapter shall follow the Lane Bonding sub-state machine described in Figure 4-10 with the behavior described in Table 4-26 and the state transitions described in Table 4-27.

Figure 4-10. Lane Bonding Sub-State Machine




Table 4-26. Transmitter Behavior in Bonding Sub-States

| State | Transmitter Behavior |
|-------|---|
| TS1 | Transmitter shall send back-to-back TS1 Ordered Sets. |
| TS2 | Transmitter shall send back-to-back TS2 Ordered Sets. |

Table 4-27. Lane Bonding Sub-State Machine Transitions

| Transition | From State | To State | Conditions* |
|------------|------------|----------------|---|
| 1 | TS1 | TS2 | Gen 2: <ul style="list-style-type: none"> • Received 2 TS1 Ordered Sets in a row. • Sent at least 32 TS1 Ordered Sets. Gen 3: <ul style="list-style-type: none"> • Received 2 TS1 Ordered Sets in a row. • Sent at least 16 TS1 Ordered Sets. |
| 2 | TS1 | Training.LOCK2 | <ul style="list-style-type: none"> • Received 2 SLOS1 Symbols in a row. <i>Note: at Gen 3 speed, it is sufficient to match either the high 64 bits or the low 64 bits of a 128 bit Symbol when determining whether 2 SLOS1 Symbols are received in a row.</i> |

| Transition | From State | To State | Conditions* |
|------------|------------|----------------|--|
| 3 | TS2 | CL0 | Gen 2: <ul style="list-style-type: none"> Received 2 TS2 Ordered Sets in a row. Sent at least 16 TS2 Ordered Sets. Gen 3: <ul style="list-style-type: none"> Received 2 TS2 Ordered Sets in a row. Sent at least 8 TS2 Ordered Sets. |
| 4 | TS2 | Training.LOCK2 | <ul style="list-style-type: none"> Received 2 SLOS Symbols in a row. |
| 5 | TS1/TS2 | CLd | <ul style="list-style-type: none"> Did not transition to CL0 state after tBonding time. |

* All conditions must be met before transition takes place.

The transmitter and receiver are on while in this state.

4.2.1.5.3 Exit from State

A Lane Adapter shall exit this state as defined in Table 4-27.

A Lane Adapter that exits this state due to successful completion (i.e. Transition 3 in Table 4-27) transitions to the CL0 state. A Lane Adapter that transitions to CL0 state shall continue sending TS2 Ordered Sets until the other Lane Adapter in the USB4 Port enters CL0 state.

A Lane Adapter that exits this state due to receiving SLOS Symbols (i.e. Transitions 2 and 4 in Table 4-27) transitions to the Training.LOCK2 sub-state.

A Lane Adapter that exits this state due to bonding timeout (i.e. Transition 5 in Table 4-27) transitions to CLd state.

4.2.1.6 Low Power (CL0s, CL1, and CL2)

The CL0s, CL1, and CL2 Low Power states are used to reduce transmitter and receiver power when a Lane is idle. Support for the CLx Low Power states is optional.

When a Lane Adapter supports CLx states, it shall enter or reject a CLx state as described in Section 4.2.1.6.2. When a Lane Adapter does not support CLx states, it shall reject entry to CLx state as described in Section 4.2.1.6.2.

4.2.1.6.1 Ordered Sets

The following Ordered Sets are used to enter and exit a Low Power state:

- CL2_REQ Ordered Set is used to request entry to CL2 state. See Table 4-28.
- CL1_REQ Ordered Set is used to request entry to CL1 state. See Table 4-29.
- CL2_ACK Ordered Set is used to approve entry to CL2 state. See Table 4-30.
- CL1_ACK Ordered Set is used to approve entry to CL1 state. See Table 4-31.
- CL0s_ACK Ordered Set is used to approve entry to CL0s state. See Table 4-32.
- CL_NACK Ordered Set is used to reject entry to a Low Power state. See Table 4-33.
- CL_OFF Ordered Set is used to complete entry to Low Power state. See Table 4-34.
- CL_WAKE1.X and CL_WAKE2.X Ordered Sets are used to exit a Low Power state. See Section 4.2.1.6.4.

Table 4-28. CL2_REQ Ordered Set

| Bits | Value | Description |
|-------|---|-----------------------|
| 63:10 | 11 1011 0000 0100 1011 1011 0000 0100 1011 1011 0000 0100 1011 1000b | Ordered Set contents. |

| | | |
|-----|---------------|--|
| 9:0 | 00 1111 0010b | SCR – Shall be set to this value to indicate that the Ordered Set contents are scrambled. |
|-----|---------------|--|

Table 4-29. CL1_REQ Ordered Set

| Bits | Value | Description |
|-------|---|--|
| 63:10 | 11 1011 0000 0100 0111 1011 0000 0100 0111 1011 0000 0100 0111 1011b | Ordered Set contents. |
| 9:0 | 00 1111 0010b | SCR – Shall be set to this value to indicate that the Ordered Set contents are scrambled. |

Table 4-30. CL2_ACK Ordered Set

| Bits | Value | Description |
|-------|---|--|
| 63:10 | 10 1011 0000 0100 1010 1011 0000 0100 1010 1011 0000 0100 1011 1100b | Ordered Set contents. |
| 9:0 | 00 1111 0010b | SCR – Shall be set to this value to indicate that the Ordered Set contents are scrambled. |

Table 4-31. CL1_ACK Ordered Set

| Bits | Value | Description |
|-------|---|--|
| 63:10 | 10 1011 0000 0100 0110 1011 0000 0100 0110 1011 0000 0100 0110 1011b | Ordered Set contents. |
| 9:0 | 00 1111 0010b | SCR – Shall be set to this value to indicate that the Ordered Set contents are scrambled. |

Table 4-32. CL0s_ACK Ordered Set

| Bits | Value | Description |
|-------|---|--|
| 63:10 | 10 1011 0000 0100 0101 0110 1011 0000 0100 0101 0110 1011 0000 0100b | Ordered Set contents. |
| 9:0 | 00 1111 0010b | SCR – Shall be set to this value to indicate that the Ordered Set contents are scrambled. |

Table 4-33. CL_NACK Ordered Set

| Bits | Value | Description |
|-------|---|--|
| 63:10 | 11 1111 0000 0100 0111 1111 0000 0100 0111 1111 0000 0100 0100 0100b | Ordered Set contents. |
| 9:0 | 00 1111 0010b | SCR – Shall be set to this value to indicate that the Ordered Set contents are scrambled. |

Table 4-34. CL_OFF Ordered Set

| Bits | Value | Description |
|-------|---|--|
| 63:10 | 00 0011 1111 1100 0011 1111 1100 0011 1111 1100 0011 1111 1100 0100b | Ordered Set contents. |
| 9:0 | 00 1111 0010b | SCR – Shall be set to this value to indicate that the Ordered Set contents are scrambled. |

4.2.1.6.1.1 CL_WAKE1.X Ordered Sets

CL_WAKE1.X Ordered Sets are sent by Re-timers.

When operating in Gen 2 mode with RS-FEC encoding disabled, a CL_WAKE1.X Ordered Set Symbol has the structure of an SLOS1 Symbol with 64b/66b encoding and the following modifications:

- Bits [63:56] of an even numbered Symbol payload are CXh, where “X” is the hexadecimal Re-timer Index of the Re-timer that generated the Ordered Set.
- Bits [63:56] of an odd numbered Symbol payload are the logical inverse of CXh, where “X” is the hexadecimal Re-timer Index of the Re-timer that generated the Ordered Set.


When operating in Gen 2 mode with RS-FEC encoding enabled, a CL_WAKE1.X Ordered Set Symbol has the structure of an SLOS1 Symbol with 128b/132b Encoding and the following modifications (see Figure 4-11):

- Bits [127:120] of an even numbered Symbol payload are CXh, where “X” is the hexadecimal Re-timer Index of the Re-timer that generated the Ordered Set.
- Bits [63:56] of an even numbered Symbol payload are CXh, where “X” is the hexadecimal Re-timer Index of the Re-timer that generated the Ordered Set.
- Bits [127:120] of an odd numbered Symbol payload are the logical inverse of CXh, where “X” is the hexadecimal Re-timer Index of the Re-timer that generated the Ordered Set.
- Bits [63:56] of an odd numbered Symbol payload are the logical inverse of CXh, where “X” is the hexadecimal Re-timer Index of the Re-timer that generated the Ordered Set.


When operating in Gen 3 mode, a CL_WAKE1.X Ordered Set Symbol has the structure of an SLOS1 Symbol with 128b/132b Encoding and the following modifications (see Figure 4-11):

- Bits [127:120] of an even numbered Symbol payload are CXh, where “X” is the hexadecimal Re-timer Index of the Re-timer that generated the Ordered Set.
- Bits [127:120] of an odd numbered Symbol payload are the logical inverse of CXh, where “X” is the hexadecimal Re-timer Index of the Re-timer that generated the Ordered Set.


Figure 4-11. Structure of a CL_WAKE1.X Ordered Set Symbol



(a) CL_WAKE1.X Ordered Set Symbol in Gen 2 mode and RS-FEC disabled



(b) CL_WAKE1.X Ordered Set Symbol in Gen 2 mode and RS-FEC enabled



(c) CL_WAKE1.X Ordered Set Symbol in Gen 3 mode

A CL_WAKE1.X Ordered Set shall not be scrambled, and the scrambler shall not advance upon receive/transmit. Unless otherwise mentioned, a CL_WAKE1.X Ordered Set shall be transmitted in its entirety.

4.2.1.6.1.2 CL_WAKE2.X Ordered Sets

A CL_WAKE2.X Ordered Sets are sent by Routers.

When operating in Gen 2 mode with RS-FEC encoding disabled, a CL_WAKE2.X Ordered Set Symbol has the structure of an SLOS2 Symbol with 64/66b encoding and the following modifications:

- Bits [63:56] of an even numbered Symbol payload shall be CXh, where “X” is the hexadecimal index from the last CL_WAKE1.X Ordered Set received.
- Bits [63:56] of an odd numbered Symbol payload shall be the logical inverse of CXh, where “X” is the hexadecimal index from the last CL_WAKE1.X Ordered Set received.

When operating in Gen 2 mode and RS-FEC encoding is enabled, a CL_WAKE2.X Ordered Set Symbol has the structure of an SLOS2 Symbol with 128b/132b encoding and the following modifications:

- Bits [127:120] of an even numbered Symbol payload shall be CXh, where “X” is the hexadecimal index from the last CL_WAKE1.X Ordered Set received.
- Bits [63:56] of an even numbered Symbol payload shall be CXh, where “X” is the hexadecimal index from the last CL_WAKE1.X Ordered Set received.

- Bits [127:120] of an odd numbered Symbol payload shall be the logical inverse of CXh, where “X” is the hexadecimal index from the last CL_WAKE1.X Ordered Set received.
- Bits [63:56] of an odd numbered Symbol payload shall be the logical inverse of CXh, where “X” is the hexadecimal index from the last CL_WAKE1.X Ordered Set received.

When operating in Gen 3 mode, a CL_WAKE2.X Ordered Set Symbol has the structure of an SLOS2 Symbol with 128b/132b Encoding and the following modifications:

- Bits [127:120] of an even numbered Symbol payload shall be CXh, where “X” is the hexadecimal index from the last CL_WAKE1.X Ordered Set received.
- Bits [127:120] of an odd numbered Symbol payload shall be the logical inverse of CXh, where “X” is the hexadecimal index from the last CL_WAKE1.X Ordered Set received.

A CL_WAKE2.X Ordered Set shall not be scrambled, and the scrambler shall not advance upon receive/transmit. Unless otherwise mentioned, a CL_WAKE2.X Ordered Set shall be transmitted in its entirety.

4.2.1.6.2 Entry to State

The Lane Adapters in a USB4 Port either enter or reject entry to a CLx state as described in this section.

A USB4 Port uses a set of objections to prevent its Adapters from entering into a Low Power state. When a USB4 Port asserts an objection, its Lane Adapters cannot enter a Low Power state. When a USB4 Port does not assert an objection, its Lane Adapters may optionally enter a Low Power state. Objections are defined in Section 4.2.1.6.3.

For a Single Lane Link, only the Lane 0 Adapter executes the flow in this section. For a Dual-Lane Link, both the Lane 0 Adapter and the Lane 1 Adapter execute the flow. The USB4 Port that initiates the transition to a CLx state is referred to as the “Requesting Port.” The USB4 Port on the other side of the Link is referred to as the “Responding Port.”

- In the Requesting Port, an Adapter requests entry to a Low Power state by sending a request Ordered Set as follows. The request Ordered Set shall be sent back-to-back until a response Ordered Set is received from the Link Partner.
 - A Lane Adapter uses the CL2_REQ Ordered Set to request entry to the CL2 state. An Adapter shall send CL2_REQ Ordered Sets when its USB4 Port does not assert any objections to enter CL2 state.
 - A Lane Adapter uses the CL1_REQ Ordered Set to request entry to the CL1 state. An Adapter shall send CL1_REQ Ordered Sets when its USB4 Port asserts an objection to enter CL2 state but does not assert any objections to enter CL1 state.
 - If a Lane Adapter receives a CL1_REQ Ordered Set or a CL2_REQ Ordered Set from its Link Partner, it shall not request entry to a Low Power state until after transitioning back to CL0.
 - If the Requesting Port asserts an objection after the Lane Adapter has sent a request Ordered Set, the Lane Adapter shall ignore the objection until the Lane Adapter is either in a CLx state or receives a CL_NACK Ordered Set.

Note: If the Requesting Port receives a Transport Layer Packet during this time and flow control for the Packet is enabled, the Transport Layer Packet is queued and causes the Port to assert an objection. The Lane Adapters in the Port will process the objection after receiving a CL_NACK or transitioning to the CLx state. The Packet will be transmitted when the Lane Adapters are in CL0 state.

- In the Responding Port, a Lane Adapter responds to a request to enter a Low Power state by sending back-to-back response Ordered Sets as follows:
 - A Lane Adapter shall reject a request to enter a Low Power state when all of the following are true:
 - The Adapter has already sent a request to enter the same Low Power state.
 - The *PM Secondary* bit in the Lane 0 Adapter and/or the Lane 1 Adapter of the Responding Port is set to 0b.

The Lane Adapter shall send CL_NACK Ordered Sets for as long as it receives the request Ordered Set from the Link Partner. After the Adapter stops receiving request Ordered Sets, it may send request Ordered Sets.

- A Lane Adapter that receives a CL1_REQ Ordered Set after it has sent a CL2_REQ Ordered Set, shall accept the request by responding with CL1_ACK Ordered Sets. The Adapter shall stop sending CL2_REQ Ordered Sets.
- A Lane Adapter that receives a CL2_REQ Ordered Set after it has sent a CL1_REQ Ordered Set, shall not respond to the request and shall continue sending CL1_REQ Ordered Sets.
- Else, if the Responding Port does not assert an objection to enter CL2 state, it shall respond to CL2_REQ Ordered Sets with a CL2_ACK Ordered Set. The first CL2_ACK shall be sent within tCLxRequest after receiving the request. The CL2_ACK Ordered Set shall be sent 375 times.
- Else, if the Responding Port does not assert an objection to enter CL1 state, it shall respond to CL2_REQ or CL1_REQ Ordered Sets with a CL1_ACK Ordered Set. The first CL1_ACK shall be sent within tCLxRequest after receiving the request. The CL1_ACK Ordered Set shall be sent 375 times.
- Else, if the *CL0s Enable* bit is set to 1b in the Lane 0 Adapter of the Responding Port, a Lane Adapter shall respond to a request to enter a Low Power state with a CL0s_ACK Ordered Set. The first CL0s_ACK shall be sent within tCLxRequest after receiving the request. The CL0s_ACK Ordered Set shall be sent 16 times.

Note: Unless otherwise noted, the Connection Manager can change the value of the CL0s Enable bit at any time.

- Else, a Lane Adapter shall respond to a request to enter a Low Power state with CL_NACK Ordered Sets. The CL_NACK Ordered Sets shall be sent 16 times. The Adapter shall resume regular CL0 operation in the transmit direction once it stops sending the CL_NACK Ordered Sets.
 - If the Responding Port asserts an objection after the Lane Adapter has sent a CLx_ACK response Ordered Set, but before the transition to CLx state is complete, the Lane Adapter shall ignore the objection until it transitions to the CLx state. After transitioning to CLx state, the objection will cause the Adapter to exit the CLx state as defined in Section 4.2.1.6.5.
- In the Requesting Port, a Lane Adapter shall stop sending a request to enter a Low Power state when it receives a response Ordered Set from the Link Partner.
 - If the response is a CL2_ACK, a CL1_ACK, or a CL0s_ACK Ordered Set, the Lane Adapter shall send 375 CL_OFF Ordered Sets. The CL_OFF Ordered sets shall be sent back-to-back. The first CL_OFF Ordered Set shall be sent within tCLxResponse after detection of the response. If the response is a CL2_ACK or a CL1_ACK Ordered Set, the Adapter shall also shut down its receiver.
 - If the response is a CL_NACK Ordered Set, the Adapter shall not send another CL2_REQ Ordered Set or CL1_REQ Ordered Set for tCLxRetry after receiving the CL_NACK Ordered Set.

- If the response is a CL_NACK Ordered Set, all Lane Adapters in the Requesting Port shall resume regular CL0 operation.
- If the Requesting Port detects Link errors in the direction of the Link Partner before receiving a response Ordered Set from the Link Partner, it shall:
 - Stop sending the request to enter a Low Power state.
 - Transition its Lane Adapters to the Training.LOCK1 sub-state and send the first SLOS within tCLxResponse of detecting the Link error.
- In the Requesting Port, a Lane Adapter shall do the following after sending 375 CL_OFF Ordered Sets:
 1. Shut down its transmitter within tTxOff time. The Adapter may send additional CL_OFF Ordered Sets during the tTxOff period.
 - If RS-FEC is enabled, the transmitter may shut down before the current RS_FEC block ends.
 2. Transition state as follows:
 - If the response from the Link Partner was CL2_ACK, transition to CL2 state.
 - If the response from the Link Partner was CL1_ACK, transition to CL1 state.
 - If the response from the Link Partner was CL0s_ACK, transition to CL0s state.
 3. Enable exit from CLx state as follows:
 - If the response from the Link Partner was CL2_ACK or CL1_ACK, wait tEnterLFPS1 time after the state transition in Step 2), then enable transmission and detection of Low Frequency Periodic Signaling (LFPS) (see Section 4.2.1.6.4).
 - If the response from the Link Partner was CL0s_ACK, wait tEnterLFPS4 time after the state transition in Step 2), then enable transmission of Low Frequency Periodic Signaling (LFPS).
- In the Responding Port, a Lane Adapter shall shut down its receiver after receiving a CL_OFF Ordered Set. If the Adapter sent CL0s_ACK Ordered Sets, it shall also transition to the CL0s state. The Adapter shall then enable exit from the Low Power state as follows:
 - If the Adapter sent CL2_ACK or CL1_ACK, wait tEnterLFPS2 time after shutting down the receiver, then enable transmission and detection of Low Frequency Periodic Signaling (LFPS).
 - If the Adapter sent CL0s_ACK, wait tEnterLFPS5, then enable detection of Low Frequency Periodic Signaling (LFPS).
- In the Responding Port, a Lane Adapter shall do the following after the equivalent of 375 Symbol Times has passed since sending the first response Ordered Set:
 - If the Adapter sent CL2_ACK Ordered Sets, it shall shut down its transmitter and shall shut down its receiver if it has not done so already. It shall then transition to the CL2 state.
 - If the Adapter sent CL1_ACK Ordered Sets, it shall shut down its transmitter and shall shut down its receiver if it has not done so already. It shall then transition to the CL1 state.
 - If the Adapter sent CL0s_ACK Ordered Sets, it shall shut down its receiver if it has not done so already. It shall then transition to the CL0s state.

- Enable exit from CLx state as follows:
 - If the Adapter sent CL2_ACK or CL1_ACK Ordered Sets, wait $t_{EnterLFPS3}$ time, then enable transmission and detection of Low Frequency Periodic Signaling (LFPS).
 - If the Adapter sent CL0s_ACK Ordered Sets, wait 240 Symbol Times + 100ns, then enable detection of Low Frequency Periodic Signaling (LFPS).

A Lane Adapter may transition to Training.LOCK1 sub-state as a result of Logical Layer errors during the entry to Low Power state with the following exceptions:

- After sending the first CL2_ACK, CL1_ACK, or CL0s_ACK Ordered Set, a Lane Adapter shall not enter Training state as a result of Logical Layer errors in its receivers. It is recommended that the Adapter turn off RS-FEC immediately when it detects Logical Layer errors in order to detect SLOS.
- A Lane Adapter that is sending CL_OFF Ordered Sets shall complete the transition to CL2, CL1, or CL0s state. The Adapter may then initiate an exit from the Low Power state to the Training state.

If a Lane Adapter receives SLOS at any time during the entry flow, it shall abort the entry flow and transition to either the Training.LOCK1 sub-state or the Training.LOCK2 sub-state as defined in Section 4.2.1.4.3.

Entry time to a CLx state is calculated by the following equations:

Equation 4-1. Entry Time to CL0s State

$$t_{CL0sEntry} \cong t_{CLxRequest} + t_{CLxResponse} + 375 \text{ Symbol Time} + t_{TxOff} + t_{EnterLFPS4}$$

Equation 4-2. Entry Time to CL1/ CL2 State

$$t_{CLxEntry} \cong t_{CLxRequest} + t_{CLxResponse} + t_{EnterLFPS2}$$

Note: If there are Re-timers on the Link, the Re-timer latency can add up to 2 μ s to the calculated CLx entry time.

See Appendix C.1 for examples of CL2, CL1, and CL0s entry.



CONNECTION MANAGER NOTE

A Connection Manager shall not do any of the following:

- *Set the CL0s Enable, CL1 Enable, and/or CL2 Enable bit to 1b when USB4 Port has two Single-Lane Links that are not yet bonded.*
- *Set the CL0s Enable, CL1 Enable, and/or CL2 Enable bits to 1b when an Active Cable that does not support Low Power states is used.*
- *Set the CL0s Enable bit to 0b while the CL1 Enable bit is 1b.*
- *Set the CL1 Enable bit to 0b while the CL2 Enable bit is 1b.*
- *Set different values for the CL0s Enable bits at the two ends of a Link.*
- *Set different values for the CL1 Enable bits at the two ends of a Link.*
- *Set different values for the CL2 Enable bits at the two ends of a Link.*
- *Set either the CL0s Enable, CL1 Enable, or CL2 Enable bit to 1b in a USB4 Port operating with Bi-Directional Time Sync Handshakes.*

A Connection Manager may set the CL0s Enable, CL1 Enable, and CL2 Enable bits to 1b for a Lane 0 Adapter when the Lane 1 Adapter is in the Disabled state or CLd state.



CONNECTION MANAGER NOTE

A Connection Manager shall set the PM Secondary bit to 1b in a Lane Adapter that is part of the Upstream Facing Port. For a Downstream Facing Port, the Connection Manager shall set the PM Secondary bit as follows:

- *If the Downstream Facing Port connects to an Upstream Facing Port, set the PM Secondary bit to 0b in at least one of the Adapters in the Downstream Facing Port.*
- *If the Downstream Facing Port connects to another Downstream Facing Port:*
 - *In one Downstream Facing Port, set the PM Secondary bit to 0b in at least one Lane Adapter.*
 - *In the other Downstream Facing Port, keep the PM Secondary bit to 1b in both Lane Adapters.*

4.2.1.6.3 Objections

A USB4 Port shall assert an objection to enter CL2 state if:

- The CL2 Support bit in the Lane 0 Adapter is 0b.
- The CL2 Enable bit in the Lane 0 Adapter is 0b.

Note: Unless otherwise noted, the Connection Manager can change the value of the CL2 Enable bit at any time.

- There is a Transport Layer Packet to be sent over the USB4 Port.

Note: Link-level flow control credits do not need to be available in order to assert this objection.

- The Lane 0 Adapter is referenced in an *Egress Adapter* field of a PCIe Adapter's Routing Table and the PCIe Adapter is not in PCIe L1 state.
- The Lane 0 Adapter is referenced in an *Egress Adapter* field of the Upstream PCIe Adapter's Routing Table, and either the No-Snoop Latency value or the Snoop Latency value in the last LTR Message transmitted upstream is smaller than the sum of the CL2 entry and exit latency.
- The Lane 0 Adapter is referenced in an *Egress Adapter* field of a Downstream PCIe Adapter's Routing Table, and either the No-Snoop Latency value or the Snoop Latency value in the last LTR Message received by the Downstream PCIe Adapter is smaller than the sum of the CL2 entry and exit latency.
- The Lane 0 Adapter is referenced in an *Egress Adapter* field of a DP IN Adapter's Routing Table.
- The Lane 0 Adapter is referenced in an *Egress Adapter* field of a DP OUT Adapter's Routing Table and a Packet is issued from the DP OUT Adapter.
- The Lane 0 Adapter is referenced in an *Egress Adapter* field of a USB3 Adapter's Routing Table and the USB3 link between the USB3 Adapter and the internal USB3 device is not in U2 or U3 state.
- The Lane 0 Adapter is referenced in an *Egress Adapter* field of a USB3 Adapter's Routing Table, the USB3 link between the USB3 Adapter and the internal USB3 device is in U2 state, and CL2 entry is disabled in U2 state.

Note: Whether or not to enable CL2 entry when the USB3 link between the USB3 Adapter and the internal USB3 device is in U2 state is a device-specific decision. The method for managing CL2 entry in U2 state is outside the scope of this specification.

- The Lane 0 Adapter is referenced in an *Egress Adapter* field of a USB3 Adapter's Routing Table, the USB3 link between the USB3 Adapter and the internal USB3 device is in U3 state, and CL2 entry is disabled in U3 state.

Note: Whether or not to enable CL2 entry when the USB3 link between the USB3 Adapter and the internal USB3 device is in U3 state is a device-specific decision. The method for managing CL2 entry in U3 state is outside the scope of this specification.

- Entry to CL2 state would delay a pending Time Sync handshake. This objection shall be asserted until the Time Sync handshake is complete.
- Host Routers only:
 - The Lane 0 Adapter is referenced in an *Egress Adapter* field of a Host Interface Adapter's Routing Table, whose Path corresponds to a Transmit Descriptor Ring that disables CL2 entry.
 - The Lane 0 Adapter is referenced in an *Egress Adapter* field of a Host Interface Adapter's Routing Table and the Host Interface has a Packet to send over the Adapter.

Note: End-to-end flow control credits are not required to be available in order to assert this objection.

- Device Routers only:
 - One of its Ports is in the process of CL0s, CL1 or CL2 exit flow.

A USB4 Port shall assert an objection to enter CL1 state if:

- The *CL1 Support* bit in the Lane 0 Adapter is 0b.
- The *CL1 Enable* bit in the Lane 0 Adapter is 0b.

Note: Unless otherwise noted, the Connection Manager can change the value of the CL1 Enable bit at any time.

- There is a Transport Layer Packet to be sent over the USB4 Port.

Note: Link-level flow control credits do not need to be available in order to assert this objection.

- The Lane 0 Adapter is referenced in an *Egress Adapter* field of a PCIe Adapter's Routing Table and the PCIe Adapter is not in PCIe L1 state.
- The Lane 0 Adapter is referenced in an *Egress Adapter* field of the Upstream PCIe Adapter's Routing Table, and either the No-Snoop Latency value or the Snoop Latency value in the last LTR Message transmitted upstream is smaller than the sum of the CL1 entry and exit latency.
- The Lane 0 Adapter is referenced in an *Egress Adapter* field of a Downstream PCIe Adapter's Routing Table, and either the No-Snoop Latency value or the Snoop Latency value in the last LTR Message received by the Downstream PCIe Adapter is smaller than the sum of the CL1 entry and exit latency.
- The Lane 0 Adapter is referenced in an *Egress Adapter* field of a DP IN Adapter's Routing Table.
- The Lane 0 Adapter is referenced in an *Egress Adapter* field of a DP OUT Adapter's Routing Table and a Packet is issued from the DP OUT Adapter.
- The Lane 0 Adapter is referenced in an *Egress Adapter* field of a USB3 Adapter's Routing Table and the USB3 link between the USB3 Adapter and the internal USB3 device is not in U2 or U3 state.
- The Lane 0 Adapter is referenced in an *Egress Adapter* field of a USB3 Adapter's Routing Table, the USB3 link between the USB3 Adapter and the internal USB3 device is in U2 state, and CL1 entry is disabled in U2 state.

Note: Whether or not to enable CL1 entry when the USB3 link between the USB3 Adapter and the internal USB3 device in U2 state is a device-specific decision. The method of managing CL1 entry in U2 state is outside the scope of this specification.

- The Lane 0 Adapter is referenced in an *Egress Adapter* field of a USB3 Adapter's Routing Table, the USB3 link between the USB3 Adapter and the internal USB3 device is in U3 state, and CL1 entry is disabled in USB U3 state.

Note: Whether or not to enable CL1 entry when the USB3 link between the USB3 Adapter and the internal USB3 device in U3 state is a device-specific decision. The method of managing CL1 entry in U3 state is outside the scope of this specification.

- Entry to CL1 state would delay a pending Time Sync handshake. The objection shall be asserted until the Time Sync handshake is complete.
- Host Routers only:
 - The Lane 0 Adapter is referenced in an *Egress Adapter* field of a Host Interface Adapter's Routing Table, whose Path corresponds to a Transmit Descriptor Ring that disables CL1 entry.
 - The Lane 0 Adapter is referenced in an *Egress Adapter* field of a Host Interface Adapter's Routing Table and the Host Interface has a Packet to send over the Adapter.

Note: End-to-end flow control credits are not required to be available in order to assert this objection.

- Device Routers only:
 - One of its Ports is in the process of CL0s, CL1 or CL2 exit flow.

A Lane Adapter may also assert implementation-specific objections to CL1 and/or CL2 entry.

4.2.1.6.4 Behavior in State

While in CL2 state, the transmitter shall be in electrical idle. Lane common mode voltages shall be maintained.

While in CL1 state, the transmitter shall be in electrical idle. Lane common mode voltages shall be maintained.

While in CL0s state, the transmitter at the requesting USB4 Port shall be in electrical idle. Lane common mode voltages shall be maintained.

Receiver termination shall be maintained in CL0s, CL1, and CL2 states.

4.2.1.6.5 Exit from State

A Lane Adapter shall initiate transition out of CL2, CL1, or CL0s state when:

- An objection is set in the USB4 Port that would have prevented the Adapter from entering the Low Power state.
- The Adapter is in CL0s state and a CL2_REQ Ordered Set or a CL1_REQ Ordered Set is received from the Link Partner.
- The Adapter is in CL0s state and it detects Link errors that cause the Adapter to transition to Training state.
- The Adapter is in a CL1 state or a CL2 state and it is referenced in an *Egress Adapter* field of a Lane Adapter's Routing Table and the other Adapter's receiver is exiting from CL0s, CL1, or CL2 states.
- The Adapter is in CL0s state and it is referenced in an *Egress Adapter* field of a Lane Adapter's Routing Table and the other Adapter's receiver is exiting from CL0s, CL1, or CL2 states.

A Lane Adapter may also exit from CL0s, CL1, or CL2 states by means which are implementation specific.

While in CL2 or CL1 state, the USB4 Ports at either end of the Lane may initiate exit from the state. While in CL0s state, only the USB4 Port whose transmitter is in electrical idle may initiate exit from the state.

See Appendix C for examples of end-to-end flows describing the behavior of Adapters and Re-timers during CL2, CL1, or CL0s exit.

4.2.1.6.5.1 Exit flow from CL0s state

The USB4 Port initiating exit from CL0s state shall:

1. Send a Low Frequency Periodic Signaling (LFPS) burst on all Lanes for the duration of at least 16 LFPS cycles (see Section 3.7) and for no more than tLFPSDuration.
2. Return to Electrical Idle for tPreData (see Table 3-20).
3. Start transmitting SLOS1 on each Lane of the USB4 Port. Any received CL_WAKE Ordered Sets shall be ignored.
 - A USB4 Port may exit CL0s state with SSC enabled or disabled.

4. On detection of 2 back-to-back TS2 Ordered Sets, stop sending SLOS1 and send at least 16 TS2 Ordered Sets. The first TS2 Ordered Set shall be sent within $t_{\text{TrainingTransition}}$ after detection of the second TS2 Ordered Set. Before transmitting the first TS2 Ordered Set:
 - The scrambler shall load a new seed as defined in Section 4.3.5.
 - Activate RS-FEC as defined in Section 4.3.6.1.
 - Enable SSC if SSC is disabled.

Note: If the Adapter is in Training state (due to receiving Link errors while in CL0s state), then the Adapter proceeds with the Training state-machine rather than sending the TS2 Ordered Sets.

If the receiver did not detect 2 back-to-back TS2 Ordered Sets within $t_{\text{TrainingAbort2}}$ time after the transmitter started sending SLOS1 it shall initiate a Disconnect by driving SBTX to a logical low state for $t_{\text{DisconnectTx}}$.

5. Transition to CL0 state.
 - If the USB4 Port operated as a Dual-Lane Link prior to entry to CL0s state, the USB4 Port shall resume operation as a Dual-Lane Link independent of the setting of the TS2 Ordered Sets. A de-skew Ordered Set shall be sent as defined in Section 4.4.4. The scrambler shall load a new seed as defined in Section 4.3.5.
 - If the Router initiated exit from CL0s state due to receiving CL1_REQ or CL2_REQ Ordered Sets, then the Router shall respond to the request Ordered Sets. The Router shall not send any Transport Layer Packets before completing the CLx entry flow.

Upon detecting 2 LFPS cycles, a Lane Adapter in CL0s state shall:

1. Enable the receiver to start bit and symbol synchronization not earlier than $t_{\text{CLxIdleRx}}$ after the last LFPS cycle received. An Adapter shall complete Symbol lock within $t_{\text{WarmUpCL0s}}$ time from the reception of the first LFPS cycle.
2. On reception of 3 back-to-back CL_WAKE1.X Ordered Set Symbols by the Lane 0 Adapter, the transmitter shall transmit at least 8 CL_WAKE2.X Ordered Set Symbols on each enabled Lane of the USB4 Port. Note that Transport Layer traffic is momentarily interrupted to transmit the Symbols. The first CL_WAKE2.X Ordered Set Symbol shall be sent within $t_{\text{WakeResponse}}$ after receiving the last bit of the third CL_WAKE1.X Ordered Set Symbol.
 - If the receiver loses Symbol alignment lock after the transmitter sends the CL_WAKE2.X Ordered Set the receiver shall regain Symbol alignment lock within $t_{\text{SymbolLock}}$ time.
 - If RS-FEC is on in the transmitting direction, then the transmitted CL_WAKE2.X Ordered Set Symbols are RS-FEC encoded. If RS-FEC is off in the transmitting direction, then the transmitted CL_WAKE2.X Ordered Set Symbols are not RS-FEC encoded.
 - If 3 back-to-back CL_WAKE1.(X+1) Ordered Set Symbols or 3 back-to-back SLOS Symbols are not received within $t_{\text{CL0sSwitch}}$ time after receiving a CL_WAKE1.X Symbol, then the Adapter shall transition to the Training.LOCK1 sub-state.
 - If the Link is operating at Gen 2 speed, the Adapter may transmit a partial CL_WAKE2.X Ordered Set in order to send the required number of CL_WAKE2.X Ordered Set Symbols. Otherwise, the Wake Ordered Set shall be transmitted in its entirety.
 - If the Router initiated exit from CL0s state by sending CL1_REQ or CL2_REQ Ordered Sets, then the Router shall continue sending the Ordered Sets

instead of sending Transport Layer Packets. The Router shall not send any Transport Layer Packets after sending the first CL1_REQ or a CL2_REQ Ordered Set.

3. On detection of 3 back-to-back SLOS Symbols by all enabled Adapters of the USB4 Port, transmit 16 TS2 Ordered Sets in each enabled Lane of the USB4 Port.
 - If the Router initiated exit from CL0s state by sending CL1_REQ or CL2_REQ Ordered Sets, then the Router may send more than 16 TS2 Ordered Sets on each enabled Lane.
4. On detection of 2 back-to-back TS2 Ordered Sets, transition to CL0 state.
 - If the Router initiated exit from CL0s state by sending CL1_REQ or CL2_REQ Ordered Sets, then the Router shall continue to send the Ordered Sets. The Router shall not send any Transport Layer Packets before completing the CLx entry flow.
 - If the Adapter does not detect 2 back-to-back TS2 Ordered Sets in $t_{TS2Timeout}$ from transmitting TS2 Ordered Sets, the Lane Adapters in the Port shall enter the Training state.

Exit time from CL0s is calculated by the following equation:

Equation 4-3. CL0s Exit Time without Re-timers

$$t_{CL0sExit} \cong t_{WarmUpCL0s} + 3 \times t_{TrainingTransition}$$

Equation 4-4. CL0s Exit Time with Re-timers on Link

$$t_{CL0sExit} \cong t_{WarmUpCL0s} + N_{Retimers} \times (t_{WakeResponse} + t_{SymbolLock}) + (N_{Retimers} - 1) \times t_{SwitchNoSSC} + t_{SwitchSSC} + 3 \times t_{TrainingTransition}$$

Note: $t_{SwitchNoSSC}$ and $t_{SwitchSSC}$ are defined in the USB4 Re-timer specification.

Note: If there are Re-timers on the Link, the Re-timer latency can add up to $2\mu s$ to the calculated CLx exit time.

In order to limit the CL0s exit time to $245\mu s$, a Router shall comply with the following equation:

$$t_{WarmUpCL0s} + 6 \times t_{WakeResponse} + t_{TrainingTransition} < 80\mu s$$

4.2.1.6.5.2 Exit flow from CL1 or CL2 state (No Re-timers on the Link)

This section applies when there are no USB4 Re-timers on the Link.

The USB4 Port initiating exit from CL1 or CL2 state shall:

1. Send a Low Frequency Periodic Signaling (LFPS) burst on each Lane until the receiver detects LFPS. If the receiver did not detect LFPS after $t_{TrainingAbort2}$ time the Router shall initiate a Disconnect by driving SBTx to a logical low state for $t_{DisconnectTx}$.
2. Return to Electrical Idle for $t_{PreData}$.
3. Start transmitting SLOS1 on the Lane.

- A USB4 Port may exit CL2 or CL1 state with SSC enabled or disabled.
4. Enable the receiver to start bit and symbol synchronization not earlier than $t_{CLxIdleRx}$ after the last LFPS cycle received. A Lane Adapter shall complete Symbol lock within t_{RxLock} time.
 5. Transition the Lane Adapter to Training.LOCK1 sub-state.
 - On transition to the TS1 sub-state, the USB4 Port shall enable SSC if SSC is disabled.

Upon detecting 2 LFPS cycles, a Lane Adapter in CL1 or CL2 state shall:

1. Send a Low Frequency Periodic Signaling (LFPS) burst on the Lane for a duration of at least 3 LFPS cycles (see Section 3.7) and for no more than $t_{LFPSDuration}$. If the Lane Adapter is in CL1 state, the first LFPS shall be sent within $t_{WarmUpCL1}$ after receiving the first LFPS cycle. If the Lane Adapter is in CL2 state, the first LFPS shall be sent within $t_{WarmUpCL2}$ after receiving the first LFPS cycle.
2. Return to Electrical Idle for $t_{PreData}$.
3. Start transmitting SLOS1 on the Lane.
 - A USB4 Port may exit CL2 or CL1 state with SSC enabled or disabled.
4. Enable the receiver to start bit and symbol synchronization not earlier than $t_{CLxIdleRx}$ after the last LFPS cycle received. A Lane Adapter shall complete Symbol lock within t_{RxLock} time.
5. Transition to Training.LOCK1 sub-state.
 - On transition to the TS1 sub-state, the USB4 Port shall enable SSC if SSC is disabled.

Exit time from CL1 and CL2 states is calculated by the following equations:

Equation 4-5. CL1 Exit Time with no Re-timers

$$t_{CL1Exit} = t_{WarmUpCL1} + t_{LFPSDuration} + t_{RxLock} + 4 \times t_{TrainingTransition}$$

Equation 4-6. CL2 Exit Time with no Re-timers

$$t_{CL2Exit} = t_{WarmUpCL2} + t_{LFPSDuration} + t_{RxLock} + 4 \times t_{TrainingTransition}$$

4.2.1.6.5.3 Exit flow from CL1 or CL2 state (Re-timers on the Link)

This section applies when there is at least one USB4 Re-timer on the Link.

The USB4 Port initiating exit from CL1 or CL2 state shall:

1. Send a Low Frequency Periodic Signaling (LFPS) burst on each Lane until its receiver detects LFPS. If the receiver did not detect LFPS after $t_{TrainingAbort2}$ time the Router shall initiate a Disconnect by driving SBTX to a logical low state for $t_{DisconnectTx}$.
2. Return to Electrical Idle for $t_{PreData}$.
3. Start transmitting SLOS1 on the Lane.
 - A USB4 Port may exit CL2 or CL1 state with SSC enabled or disabled.

4. Enable the receiver to start bit and symbol synchronization not earlier than $t_{CLxIdleRx}$ after the last LFPS cycle received. A Lane Adapter shall complete Symbol lock within t_{RxLock} time.
5. Upon reception of 3 back-to-back $CL_WAKE1.X$ Ordered Set Symbols, start transmitting $CL_WAKE2.X$ Ordered Set Symbols on the Lane. The first $CL_WAKE2.X$ Ordered Set Symbol shall be sent within $t_{WakeResponse}$ after receiving the third $CL_WAKE1.X$ Ordered Set Symbol. The Adapter shall ignore any received $CL_WAKE2.Y$ (where Y is any value) Ordered Set Symbols interleaved with $CL_WAKE1.X$ Ordered Set Symbols when it determines the reception of back-to-back $CL_WAKE1.X$ Ordered Set Symbols. If the receiver loses Symbol alignment lock after the transmitter sends the $CL_WAKE2.X$ Ordered Set, the receiver shall regain Symbol alignment lock within $t_{SymbolLock}$ time.
6. Upon reception of 7 back-to-back $CL_WAKE2.X$ Ordered Set Symbols or 7 back-to-back SLOS Symbols, transition the Adapter to Training.LOCK1 sub-state within $t_{WakeResponse}$ time.
 - On transition to the TS1 sub-state, the USB4 Port shall enable SSC if SSC is disabled.

Upon detecting 2 LFPS cycles, a Lane Adapter in CL1 or CL2 state shall:

1. Send a Low Frequency Periodic Signaling (LFPS) burst on the Lane for a duration of at least 3 LFPS cycles (see Section 3.7) and for no more than $t_{LFPSDuration}$. If the Lane Adapter is in CL1 state, the first LFPS shall be sent within $t_{WarmUpCL1}$ after receiving the first LFPS cycle. If the Lane Adapter is in CL2 state, the first LFPS shall be sent within $t_{WarmUpCL2}$ after receiving the first LFPS cycle.
2. Return to Electrical Idle for $t_{PreData}$.
3. Start transmitting SLOS1 on the Lane.
 - A USB4 Port may exit CL2 or CL1 state with SSC enabled or disabled.
4. Enable the receiver to start bit and symbol synchronization not earlier than $t_{CLxIdleRx}$ after the last LFPS cycle received. An Adapter shall complete Symbol lock within t_{RxLock} time.
5. Upon reception of 3 back-to-back $CL_WAKE1.X$ Ordered Set Symbols, start transmitting $CL_WAKE2.X$ Ordered Set Symbols on the Lane. The first $CL_WAKE2.X$ Ordered Set Symbol shall be sent within $t_{WakeResponse}$ after receiving the third $CL_WAKE1.X$ Ordered Set Symbol. The Adapter shall ignore any received $CL_WAKE2.Y$ (where Y is any value) Ordered Set Symbols interleaved with $CL_WAKE1.X$ Ordered Set Symbols when it determines the reception of back-to-back $CL_WAKE1.X$ Ordered Set Symbols. If the receiver loses Symbol lock after sending the $CL_WAKE2.X$ Ordered Set it shall regain Symbol lock within $t_{SymbolLock}$ time.
6. Upon reception of 7 back-to-back $CL_WAKE2.X$ Ordered Set Symbols or 7 back-to-back SLOS Symbols, transition to Training.LOCK1 sub-state.
 - On transition to the TS1 sub-state, the USB4 Port shall enable SSC if SSC is disabled.

Exit time from CL1 and CL2 states is calculated by the following equations:

Equation 4-7. CL1 Exit Time with Re-timers on the Link

$$\begin{aligned}
 t_{CL1Exit} &= t_{WarmUpCL1} + t_{LFPSDuration} + (N_{Re-timers} + 1) \times (t_{WakeResponse} + t_{SymbolLock}) + t_{SwitchSSC} \\
 &+ t_{CLxLock} + (N_{Re-timers} - 1) \times t_{SwitchNoSSC} + 4 \times t_{TrainingTransition}
 \end{aligned}$$

Equation 4-8. CL2 Exit Time with Re-timers on the Link

$$t_{CL2Exit} = t_{WarmUpCL2} + t_{LFPSDuration} + (N_{Re-timers} + 1) \times (t_{WakeResponse} + t_{SymbolLock}) + t_{SwitchSSC} + t_{CLxLock} + (N_{Re-timers} - 1) \times t_{SwitchNoSSC} + 4 \times t_{TrainingTransition}$$

Note: $t_{SwitchNoSSC}$, $t_{SwitchSSC}$ and $t_{CLxLock}$ are defined in the USB4 Re-timer specification.

Note: If there are Re-timers on the Link, the Re-timer latency can add up to $2\mu s$ to the calculated CLx exit time.

4.2.2 USB4 Link Transitions

A USB4 Port shall support the following Link configurations:

- One Single-Lane Link.
 - Lane 0 Adapter is enabled and in CL0 state. Lane 1 Adapter is disabled.
 - *Negotiated Link Width* = x1.

Note: This is a fallback configuration that occurs if the Connection Manager disables Lane 1 after Lane Initialization or Lane Bonding fail.

- Two independent Single-Lane Links.
 - Both the Lane 0 Adapter and Lane 1 Adapter are enabled and in CL0 state.
 - Lane 0 and Lane 1 are not bonded.
 - *Negotiated Link Width* = x1.

Note: This is a transient configuration that occurs between successful completion of Lane Initialization and successful completion Lane Bonding.

- One Dual-Lane Link.
 - Both the Lane 0 Adapter and Lane 1 Adapter are enabled and in CL0 state.
 - Lane 0 and Lane 1 are successfully bonded.
 - *Negotiated Link Width* = x2.



CONNECTION MANAGER NOTE

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

4.2.2.1 Transition from One Single-Lane Link to Two Single-Lane Links

A USB4 Port shall transition from one Single-Lane Link to two Single-Lane Links when the Lane 1 Adapter is enabled. The transition from one Single-Lane to two Single-Lane Links is complete when the Lane 1 Adapter transitions to the CL0 state.



CONNECTION MANAGER NOTE

A Connection Manager shall not enable any Path other than Path 0 in a USB4 Port with two Single-Lane Links that have not yet been bonded.

4.2.2.2 Transition from Two Single-Lane Links to Dual-Lane Link

A Connection Manager sets the *Lane Bonding* bit to 1b in the Lane Adapter Configuration Capability of either Adapter in a USB4 Port to initiate bonding two Single-Lane Links into a Dual-Lane Link. This causes the Lane Adapters in the USB4 Port to transition to the Lane Bonding state (see Section 4.2.1.5.1).



CONNECTION MANAGER NOTE

A Connection Manager shall only set the Lane Bonding bit to 1 if the Target Link Width fields in the Ports at both side of the Link allow a Dual-Lane Link.

A USB4 Port shall transition its Lane Adapters to the Lane Bonding state when all of the following are true:

- The *Supported Link Widths* field in the Lane Adapter Configuration Capability Register of both Adapters is set to x2 support or more.
- The *Target Link Width* field in the Lane Adapter Configuration Capability Register of both Adapters is set to establish a Dual-Lane Link.

The Logical Layer shall transition to a Dual-Lane Link when the following conditions are met:

- Both Adapters have transitioned successfully to CL0 state within tBonding time after sending the first TS1 Ordered Set with *Lane Bonding Target* set to 001b.
- Link Partner has responded on both Lanes with TS1 and TS2 Ordered Sets that have the *Lane Bonding Target* field set to 001b.

Note: If the Lane Adapter transitions from Bonding state to Training state due to an error on the Link and still manages to transition to CL0 state within tBonding time from entry to Bonding state, it is considered a successful transition.

If Lane bonding is successful, then a Router shall:

- Set the *Adapter State* field in the Lane Adapter Configuration Capability of the Lane 0 Adapter to CL0.
- Set the *Negotiated Link Width* field in the Lane Adapter Configuration Capability of the Lane 0 Adapter to indicate a USB4 Link width of x2.
- Send a Hot Plug Event Packet with the *UPG* bit set to 1b for the Lane 1 Adapter in the Downstream Facing Port.

Note: After successful Lane bonding, the Lane 1 Adapter is “unplugged” and the Configuration Space of the Lane 0 Adapter is used to configure the Link.

If one of the Lane Adapters is not in CL0 tBonding time after entry to the Lane Bonding state, the Router shall initiate a Disconnect by driving SBTX to a logical low state for tDisconnectTx.



CONNECTION MANAGER NOTE

A Connection Manager can conclude that Lane bonding failed if, after initiating Lane Bonding, it receives a Hot Plug Event Packet with UPG=1 for Lane 0.

The Connection Manager decides what to do after a bonding failure. The Connection Manager can retry Lane bonding, configure a Single-lane Link, or take other measures based on the desired policy.

The Adapters of a Dual-Lane Link operate in CL0 state in tandem with the following dependencies:

- Any Ordered Set sent on the Link shall be sent simultaneously on both Lanes within the permitted transmit skew.
- When either Adapter of a Dual-Lane Link transitions to one of the Training sub-states, the other Adapter in the USB4 Port shall transition to the same Training sub-state.

4.2.2.2.1 Training a Dual-Lane Link

When an Adapter that is part of a Dual-Lane Link enters Training state, the other Adapter in the USB4 Port shall enter Training state as well. During the training state, Lane Bonding is re-initiated when the Adapters of a Dual-Lane Link enter Training state from CL0, CL1, or CL2 states. While attempting to bond the Lanes, the following values are sent during Training state in all TS1 and TS2 Ordered Sets:

- *Lane Bonding Target* is set to 001b.
- *Lane Bonding Target 2* is set to 001b.

The Logical Layer shall resume Dual-Lane Link operation if both Adapters meet the transition conditions in Step 6 of the Training state machine in Table 4-20 within $t_{\text{TrainingAbort2}}$ time. The Adapter that transitions to the CL0 state first shall send TS2 Ordered Sets until the other Adapter in the USB4 Port exits the training state.

If either Adapter does not meet the transition conditions in Step 6 of the Training state machine within $t_{\text{TrainingAbort2}}$ time, then the Adapters initiate a disconnect (see Section 4.2.1.3.3).

4.2.2.3 Transition from Dual-Lane Link to Two Single-Lane Links

When a USB4 Port that is operating with a Dual-Lane Link is reset, the USB4 Port transitions to two Single-Lane USB4 Links after the USB4 Port completes Lane Initialization of its Lanes (see Section 4.1.2 for Lane Initialization flow).



CONNECTION MANAGER NOTE

To transition a Link from a Dual-Lane Link to a Single Lane Link, a Connection Manager should reset the Link and wait for Lane Initialization to complete. After Lane Initialization is complete, the Connection Manager can disable the Lane 1 Adapter to transition to a Single-Lane Link.

4.2.2.4 Transition from Two Single-Lane Links to One Single-Lane Link

A USB4 Port shall transition from two Single-Lane Links to one Single-Lane Link when one of its Adapters transitions to the Disabled state.



CONNECTION MANAGER NOTE

A Connection Manager disables the Lane 1 Adapter in a USB4 Port to transition to a Single-Lane Link. The Connection Manager shall transition a Dual-Lane Link to Two Single-Lane Links (see Section 4.2.2.3), before disabling the Lane 1 Adapter.

A Connection Manager shall not disable the Lane 0 Adapter of a USB4 Port until after it has disabled the Lane 1 Adapter.

4.2.3 Logical Layer Link State

The Transport Layer sees a Link in one of three states: Active, Low Power, or Inactive. Adapter states are mapped into Transport Layer Link States as follows:


- For a Single-Lane Link, the Link is in Active state when its Adapter is in CL0 state.
- For a Dual-Lane Link, the Link is in Active state after its Lanes are successfully bonded, both Adapters transition from either the Lane Bonding State or Training State to CL0 state, and the Adapters stop transmitting TS2 Ordered Sets.
- A Link is in Low Power state when its Adapters are in CL2, CL1, or CL0s states.
- A Link is otherwise in Inactive state.

4.3 USB4 Link Encoding

The USB4 Link is responsible for Transport Layer Packet flow through the Logical Layer. Figure 4-12 illustrates the logical flow of Transport Layer Packets (received from the Transport Layer) and Ordered Sets (generated by the Logical Layer) through the Logical Layer. Actual implementations may differ as long as the logical order is maintained.

Note: The interface between the Transport Layer and the Logical Layer is a byte streaming interface. The Logical Layer is not aware of Tunneled Packets.

Figure 4-12. Packet Flow in the Logical Layer



The following steps are taken at the transmitter:

1. Lane Distribution – In a Dual-Lane Link, outgoing Transport Layer Packets are distributed among the Lanes. See Section 4.3.1 for more detail.

Note: Following Lane distribution, both Lanes of a Dual-Lane Link go through the same stages. For example, Ordered Sets are inserted to both Lanes at the same time (within the transmit skew inaccuracy).
2. Ordered Sets – Ordered Sets are inserted into the data flow, used for Logical Layer control purposes. See Section 4.3.3 for more detail.
3. Symbol encoding – If RS_FEC encoding is off, bytes received from the Transport Layer shall be encoded with either 64b/66b encoding (Gen 2) or 128b/132b encoding (Gen 3). See Section 4.3.2 for more detail.
4. Bit Swap – Contents of Ordered Sets and Transport Layer Symbols are rearranged to the order of transmission on the wire. See Section 4.3.4.

5. Scrambling – The stream of symbols is scrambled as defined in Section 4.3.5.
6. SKIP insertion – an optional stage to insert SKIP Ordered Sets (see Section 4.4.3).
7. RS-FEC – If on, Reed-Solomon encoding is performed per Section 4.3.6.
8. Pre-Coding – Performed per Section 4.3.6.2.


The following steps are taken at the receiver:

1. Pre-Coding removal is performed per Section 4.3.6.2.
2. RS-FEC – If on, Reed-Solomon decoding is performed per Section 4.3.6.
3. SKIP removal – SKIP Ordered Sets (if present) are removed.
4. De-scrambling – The stream of symbols is de-scrambled as defined in Section 4.3.5.
5. Bit Swap – Contents of Symbols are rearranged to eliminate the bit swap done by the transmitter.
6. Symbol decoding – Symbols are converted back into Transport Layer bytes.
7. Ordered Sets – any Ordered Sets are extracted from the stream of Symbols.
8. Lane Merging and de-skew – If a USB4 Link consists of more than one Lane, then the byte streams from the different Lanes are de-skewed and combined into a single byte stream per Section 4.4.4.

4.3.1 Lane Distribution


If a USB4 Link operates as a Dual-Lane Link, then distribution of Transport Layer bytes among the Lanes shall alternate as depicted in Figure 4-13.

Figure 4-13. Byte Transmission Order on Lanes



The bytes of a Transport Layer Packet are delivered to the Logical Layer as shown in Figure 4-14.

Figure 4-14. Byte Ordering of Transport Layer Packets to the Logical Layer




Bytes are sent from B0 (first) to B_{N-1} (last)

* SupplID

The bytes of an Idle Packet are delivered to the Logical Layer as shown in Figure 4-15.

Figure 4-15. Byte Ordering of Idle Packets to the Logical Layer



4.3.2 Symbol Encoding

The Logical Layer carries two types of payload traffic: Transport Layer Packets (which are received as a byte stream from the Transport Layer), and Ordered Sets (which are added by the Logical Layer and serve Physical Layer control tasks). The format for Symbol encoding depends on whether RS-FEC is on or off.

4.3.2.1 Symbol Encoding of Transport Layer Bytes


Transport Layer bytes are grouped into Symbols. When RS-FEC is off, Transport Layer Packets are encoded using either 64b/66b encoding (for Gen 2 speed traffic) or 128b/132b encoding (for Gen 3 speed traffic).

For 64b/66b encoding, a Data Symbol is made of 64 bits of Transport Layer bytes and 2 bits to identify the Symbol type (Transport Layer bytes vs. Ordered Set), called Sync Bits. Figure 4-16(a) depicts the encoding of a Data Symbol for 64b/66b.

For 128b/132b encoding, a Data Symbol is made of 128 bits of Transport Layer bytes and 4 bits to identify the Symbol type, called Sync Bits. Figure 4-16(b) depicts the encoding of a Data Symbol for 128b/132b.

Symbol encoding is done on each Lane separately. A Symbol may contain either Transport Layer bytes or Ordered Set, but shall not contain both.

Note: The Transport Layer feeds the Logical Layer with a constant stream of bytes (referred to as "Transport Layer Bytes"). A Transport Layer sends Idle Packets when it does not have any other Transport Layer Packets to feed the Logical Layer, ensuring that there are always Transport Layer bytes to pack in to a Data Symbol.

Figure 4-16. Symbol Encoding of Data Symbols**(a) Data Symbol structure in 64b/66b encoding****(b) Data Symbol structure in 128b/132b encoding**

Note: B0 in the Data Symbol Payload is the first from the Transport Layer

An invalid Sync bit value in a received Symbol generates an Alignment Lock Error (ALE) (see Section 4.4.2).

When RS-FEC is on, the basic transmission unit is a block. See Section 4.3.6 for details.

4.3.3 Ordered Sets

This section defines the structure of Ordered Sets with the exception of SLOS1, SLOS2, CL_WAKE1.X, and CL_WAKE2.X Ordered Sets. Section 4.2.1.3.4 defines the structure of SLOS1 and SLOS2. Section 4.2.1.6.1.1 defines the structure of CL_WAKE1.X, and CL_WAKE2.X Ordered Sets.

Ordered Sets are used to perform Physical Layer control tasks. Ordered Sets may be transmitted in the middle of a Transport Layer Packet. An Ordered Set shall have the structure depicted in Table 4-35.

Table 4-35. Ordered Set Structure

| Bits | Value | Description |
|-------|--|---|
| 63:10 | Various | Ordered Set contents, specific to each Ordered Set |
| 9:0 | 00 1111 0010b (Scrambled) or 11 0000 1101b (Not scrambled) | SCR – indicates whether or not the Ordered Set contents are scrambled. |


When RS-FEC is off, Ordered Sets are encoded using 64b/66b encoding in Gen 2 mode and 128b/132b encoding in Gen 3 mode.

For 64b/66b encoding, an Ordered Set Symbol shall contain a single copy of the Ordered Set and 2 Sync Bits. Figure 4-17(a) depicts the Symbol encoding of an Ordered Set Symbol for 64/66b encoding.


For 128b/132b encoding, an Ordered Set Symbol shall contain two copies of the Ordered Set (i.e. 64 bits followed by a second copy of the same 64 bits) and 4 Sync Bits. Figure 4-17(b) depicts the Symbol encoding of an Ordered Set Symbol for 128/132b encoding.

Note: Receiving two copies of an Ordered Set in an Ordered Set Symbol, is the equivalent of receiving two identical Ordered Sets.

Figure 4-17. Symbol Encoding of Ordered Set Symbols



(a) Ordered Set Symbol structure in 64b/66b encoding



(b) Ordered Set Symbol structure in 128b/132b encoding



IMPLEMENTATION NOTE

When operating with 128b/132b encoding, a receiver can identify an Ordered Set by either the first copy of the Ordered Set, the second copy of the Ordered Set, or both.

For RS-FEC encoding, see Section 4.3.6 for the structure of Ordered Sets.

4.3.4 Bit Swap

Bit Swap is the process by which bits from the Transport Layer and the Logical Layer are rearranged into the order in which they are transmitted on the wire.

Bit Swap of Transport Layer bytes and of Ordered Sets shall be performed prior to scrambling so that bytes and bits are delivered to the scrambler in the order that they are transmitted on the wire.

4.3.4.1 Sync Bits


If RS-FEC is off, all Symbols shall be transmitted Sync Bits first. Sync Bits shall be sent in the order of most significant bit (i.e. bit 65 for 64b/66b encoding or bit 131 for 128b/132b encoding) to least significant bit (i.e. bit 64 for 64b/66b encoding or bit 128 for 128b/132b encoding). Transport Layer bytes or Ordered Sets shall be transmitted after the Sync Bits.

If RS-FEC is on, the order of transmission of Sync Bits is defined in Section 4.3.6.

4.3.4.2 Data Symbol Payload


The payload within a Data Symbol shall be transmitted from left to right as depicted in Figure 4-18. Within each byte of payload, individual bits shall be transmitted from bit 0 to bit 7.

Figure 4-18. Bit and Byte Ordering on the Wire – Data Symbol Payload



Bytes are sent from B0 (first) to B7 (last)
Bits within each Byte are sent b0 (first) to b7 (last)
For example, the order of transmission is: bit 56, bit 57, ..., bit 63, bit 48, bit 49, ..., bit 55, ...

(a) Data Symbol Payload transmission order (Gen 2)



Bytes are sent from B0 (first) to B15 (last)
Bits within each Byte are sent b0 (first) to b7 (last)
For example, the order of transmission is: bit 120, bit 121, ..., bit 127, bit 112, bit 113, ..., bit 119, ...

(b) Data Symbol Payload transmission order (Gen 3)


4.3.4.3 Ordered Set Symbol Payload

When an Ordered Set is 64 bits long, it is transmitted in one Ordered Set Symbol payload.

When an Ordered Set is longer than 64 bits (i.e. SLOS, CL_WAKE1.X, CL_WAKE2.X), it cannot fit into the payload of one Symbol, and shall be divided into multiple Symbol payloads. The Ordered Set shall be transmitted in increasing Symbols, starting with Symbol 0.


Within a Symbol payload, the bytes in an Ordered Set shall be transmitted from left to right as depicted in Figure 4-19. Within each byte, individual bits shall be transmitted from bit 0 to bit 7.

Figure 4-19. Bit and Byte Ordering on the Wire – Ordered Set Symbol Payload



Bytes are sent from B0 (first) to B7 (last)
Bits within each Byte are sent b0 (first) to b7 (last)

(a) Ordered Set Symbol payload transmission order (Gen 2)



Bytes are sent from B0 (first) to B15 (last)
Bits within each Byte are sent b0 (first) to b7 (last)

(b) Ordered Set Symbol payload transmission order (Gen 3)

4.3.5 Scrambling

The Logical Layer performs scrambling of transmit traffic and de-scrambling of receive traffic. Scrambling shall be performed according to the rules in Table 4-36.

Table 4-36. Scrambling Rules

| | Scramble? | Advanced LFSR? |
|----------------------------|-------------------------|-----------------------------|
| Data Symbols | | |
| Sync Bits | No | No |
| Transport Layer bytes | Yes | Yes |
| Ordered Set Symbols | | |
| Sync Bits | No | No |
| Contents | Determined by SCR value | When scrambled |
| SCR | No | When contents are scrambled |

The following scrambling and de-scrambling rules shall be followed for all transmitted and received traffic:

- Scrambling and de-scrambling are performed by passing the encoded bits through an Additive LFSR with a polynomial of $G(X) = X^{23} + X^{21} + X^{16} + X^8 + X^5 + X^2 + 1$.
- The most significant output bit of the LFSR is XORed with the data stream on a per-bit basis. The data stream is scrambled in the order that it is sent on wire.

- The scrambler shall load a new seed on the following transitions:
 1. Transition from LOCK2 sub-state to TS1 sub-state in the Training state.
 - Initial value is 1F EEDDh.
 2. On exit from CL0s state, before the Adapter initiating exit transmits the first TS2 Ordered Set in the direction exiting electrical idle.
 - Initial value is 1F EEDDh.
 3. On transition from any state to CL0 when going to a Dual-Lane Link.
 - When exiting CL0s state, a new seed shall be loaded only in the direction exiting electrical idle.
 - Initial value on the Lane 0 is 1D BFBCCh.
 - Initial value on Lane 1 is 06 07BBh.
 - The per-Lane seeds are used, starting with the first byte after the last de-skew Ordered Set.
- Any single-bit errors in the *SCR* field shall be corrected. If the *SCR* field contains an uncorrectable error, the Logical Layer reports an OSE error as defined in Section 4.4.2.

The following pseudocode example produces N bits of LFSR output, starting with a seed value as defined above.

```
poly = 21 0124h
lfsr = seed
output = 0
for (i=1 to N) {
  msb = lfsr[22]
  lfsr = lfsr << 1
  lfsr[23] = 0
  if (msb == 1) then lfsr = (lfsr XOR poly)
  lfsr = lfsr + msb
  output = (output << 1) + msb
}
```



IMPLEMENTATION NOTE

During Training state or during exit from CL0s state, a receiver can identify the exact time to load a new scrambler seed by detecting an Ordered Set with the SCR (bits [9:0]) equal to 00 1111 0010b (scrambled).

See Appendix A.4 for examples of scrambling Data Symbols and Ordered Set Symbols.

4.3.6 RS-FEC

The Logical Layer employs a Reed-Solomon forward error-correction code (RS-FEC). An Adapter shall support RS-FEC at all speeds. An RS(198,194) code over GF(2⁸) is used:

- The primitive polynomial over GF(2⁸) is $p(x) = X^8 + X^4 + X^3 + X^2 + 1$
- The generating polynomial is $g(x) = X^4 + 15X^3 + 54X^2 + 120X + 64$

Each block of 194 bytes shall be generated in the following manner (see Figure 4-20):

- Transport Layer bytes and Ordered Sets are grouped into 16-byte (128 bit) Symbol payloads. Each Symbol payload may contain either one or more Ordered Set or Transport Layer bytes, but shall not contain both.

- When operating at Gen 2 speed, the 16-byte Ordered Set Symbol payload shall contain:
 - For a SLOS Ordered Set or a CL_WAKE2.X Ordered Set, 128 bits of the Ordered Set.
 - For all other Ordered Sets, two 64-bit Ordered Sets. When only one Ordered Set needs to be sent, the second Ordered Set shall be a SKIP Ordered Set. See Section 4.4.3 for the structure of a SKIP Ordered Set.

Note: CL_WAKE1.X Ordered Sets are not RS-FEC encoded. See Section 4.2.1.6.1 for more information.

- When operating at Gen 3 speed, the 16-byte Ordered Set Symbol payload shall contain:
 - For a SLOS Ordered Set or a CL_WAKE2.X Ordered Set, 128 bits of the Ordered Set.
 - For all other Ordered Sets, two copies of the Ordered Set (i.e. 64 bits followed by a second copy of the same 64 bits).

Note: CL_WAKE1.X Ordered Sets are not RS-FEC encoded. See Section 4.2.1.6.1 for more information.

- The RS-FEC encoder is fed with twelve 16-byte Symbol payloads plus 2 bytes of Sync Bits. Each Symbol is allocated a single Sync Bit, indicating whether it contains Transport Layer bytes (Sync Bit = 0b) or Ordered Set (Sync Bit = 1b).
 - The 2 bytes of Sync Bits contain 12 active bits (one per 16-byte Symbol) and 4 reserved bits.
 - Sync Bits shall be delivered to the encoder in order that they are sent to the wire, from bit 15 to bit 0. The active Sync Bits reside in bits[11:0] of the Word. The Sync Bit corresponding to the oldest 16-byte Symbol (Symbol payload 0 in Figure 4-20) resides in bit 0 of the Sync Bits.
 - The 12 active bits are XORed with 333h before being fed to the RS-FEC encoder. The XORed value is the value seen on the wire.
- The RS-FEC encoder generates 4 bytes of redundancy bits (P3 to P0). P3 is the first byte to be sent on the wire and P0 is the last. Within each byte, bits are sent in descending order where bit 7 is sent first and bit 0 is sent last.

The RS-FEC decoder shall correct a received block with up to two 1-byte errors anywhere in the block.

An error in a received block that is detectable and uncorrectable shall cause an RDE error (see Section 4.4.2).

Table 4-37. START_RS_FEC Bit Sequence

| Bits | Value | Description |
|------|------------|--|
| 31:0 | 0F0F 0F0Fh | Unique sequence indicating start of RS-FEC |

If RS-FEC encoding is activated, a USB4 Port shall deactivate RS-FEC encoding on a Lane in the following cases:

- When transitioning to Training.LOCK1 sub-state, after transmitting n SLOS1 Symbols in LOCK1 sub-state with RS-FEC on, where $32 \leq n \leq 64$ in Gen 2 and $16 \leq n \leq 32$ in Gen 3.
- Entry to Training.LOCK2 sub-state.
- Entry to Disabled state.
- Entry to CLd state.
- Entry to CL0s state, in the direction entering Low Power state.
- Entry to CL2 or CL1 states.

Note: RS-FEC decoding of received bits may be turned off on entry to Training state.

4.3.6.2 Pre-Coding

If Pre-Coding is on, then before each bit is sent on the wire, it shall be XOR'ed with the bit sent before it, using the value of the bit after it was coded. The first bit is XOR'ed with 0b.

For example, if the bit sequence to be sent starts with 1,0,1,1,0,... then the sequence after Pre-Coding is 1,1,0,1,1,...

Pre-Coding shall be turned on with the first bit that is RS_FEC encoded.

Pre-Coding shall be turned off with the first bit that is not RS_FEC encoded.

4.4 USB4 Link Operation

4.4.1 Start of Data

When an Adapter transitions to CL0 state, the first transmitted bytes after the last TS2 Ordered Set shall be either a Transport Layer header, an Idle Packet or any Ordered Set other than SLOS, TS1 or TS2. For a Dual-Lane Link, the first transmitted bytes after the last TS2 Ordered Set shall be a de-skew Ordered Set followed by either a Transport Layer header, an Idle Packet or any Ordered Set other than SLOS, TS1 or TS2. See Section 4.4.4.

4.4.2 Error Cases and Recovery

Table 4-38 lists the events that can cause a Logical Layer error along with how that error case shall be handled and reported.

A Router shall support the Ordered Set Errors (OSE) error case. Support for all other error cases is optional. However, when a Router supports an error case, it shall do so as described in this section. A Router shall support the same error cases on all Lane Adapters.

An Adapter reports Logical Layer errors in the *Logical Layer Errors* field in the Lane Adapter Configuration Capability. The *Logical Layer Errors Enable* field in the Lane Adapter Configuration Capability controls which Logical Layer errors are reported by Notification Packets. A Router may hardwire a *Logical Layer Errors Enable* bit (other than the OSE bit) to 0b to permanently disable reporting of an error by Notification Packets. If the response to the error case is a transition to Training state and the Logical Errors Enable field for the error case is set to 1b, a Lane Adapter that is part of the Upstream Facing Port sends a Notification Packet after it transitions back to CL0.

When an Adapter supports an error case, it shall support that error case in all Adapter states unless specified otherwise.

Table 4-38. Error Cases and Impact on Logical Layer

| Error | Event | Response | Reporting |
|----------------------------|---|--|--|
| Alignment Lock Error (ALE) | Adapter received N number of Symbols in a row with illegal Sync Bits values, where N is a number between 1 and 8 (inclusive) that is chosen by the implementation. | Go to Training.LOCK1 sub-state. | An Adapter shall set the <i>ALE</i> bit in the <i>Logical Layer Errors</i> field to 1b. If the <i>ALE</i> bit in the <i>Logical Layer Errors Enable</i> field is 1b, the Router shall send a Notification Packet with Event Code = ERR_LINK to the Connection Manager (see Section 6.5). If the <i>ALE</i> bit in the <i>Logical Layer Errors Enable</i> field is 0b, the Router shall not send a Notification Packet. |
| Order Set Error (OSE) | Adapter received 2 back-to-back Symbols that contain an Ordered Set that is not defined in this specification and/or have an uncorrectable error in the <i>SCR</i> field. | Go to Training.LOCK1 sub-state. | An Adapter shall set the <i>OSE</i> bit in the <i>Logical Layer Errors</i> field to 1b. If the <i>OSE</i> bit in the <i>Logical Layer Errors Enable</i> field is 1b, the Router shall send a Notification Packet with Event Code = ERR_LINK to the Connection Manager (see Section 6.5). If the <i>OSE</i> bit in the <i>Logical Layer Errors Enable</i> field is 0b, the Router shall not send a Notification Packet. |
| Timeout Error (TE) | Adapter entered Training state from either CL0, CL1, or CL2 state and did not transition to CL0 state within <i>tTrainingError</i> after sending the first SLOS1. | Go to Training.LOCK1 sub-state. | An Adapter shall set the <i>TE</i> bit in the <i>Logical Layer Errors</i> field to 1b. If the <i>TE</i> bit in the <i>Logical Layer Errors Enable</i> field is 1b, the Router shall send a Notification Packet with Event Code = ERR_LINK to the Connection Manager (see Section 6.5). If the <i>TE</i> bit in the <i>Logical Layer Errors Enable</i> field is 0b, the Router shall not send a Notification Packet. |
| Elastic Buffer Error (EBE) | The elastic buffer is full. | Go to Training.LOCK1 sub-state. | An Adapter shall set the <i>EBE</i> bit in the <i>Logical Layer Errors</i> field to 1b. If the <i>EBE</i> bit in the <i>Logical Layer Errors Enable</i> field is 1b, the Router shall send a Notification Packet with Event Code = ERR_LINK to the Connection Manager (see Section 6.5). If the <i>EBE</i> bit in the <i>Logical Layer Errors Enable</i> field is 0b, the Router shall not send a Notification Packet. |
| De-skew Buffer Error (DBE) | Skew is too large, resulting in overflow in the de-skew buffer. | Optionally go to Training.LOCK1 sub-state. | An Adapter shall set the <i>DBE</i> bit in the <i>Logical Layer Errors</i> field to 1b. If the <i>DBE</i> bit in the <i>Logical Layer Errors Enable</i> field is 1b, the Router shall send a Notification Packet with Event Code = ERR_LINK to the Connection Manager (see Section 6.5). If the <i>DBE</i> bit in the <i>Logical Layer Errors Enable</i> field is 0b, the Router shall not send a Notification Packet. |

| Error | Event | Response | Reporting |
|----------------------------|---|---|---|
| RS-FEC Decoder Error (RDE) | The RS-FEC decoder identified an uncorrectable error. | Turn off RS-FEC in both directions of the Link. | An Adapter shall set the <i>RDE</i> bit in the <i>Logical Layer Errors</i> field to 1b. If the <i>RDE</i> bit in the <i>Logical Layer Errors Enable</i> field is 1b, the Router shall send a Notification Packet with Event Code = <i>ERR_LINK</i> to the Connection Manager (see Section 6.5). If the <i>RDE</i> bit in the <i>Logical Layer Errors Enable</i> field is 0b, the Router shall not send a Notification Packet. |
| RX Sync Timeout (RST) | While in LOCK1 sub-state of the Training state, the receiver could not lock on Sync Bits for an implementation-specific period of time. | Remain in LOCK1 sub-state. | An Adapter shall set the <i>RST</i> bit in the <i>Logical Layer Errors</i> field to 1b. If the <i>RST</i> bit in the <i>Logical Layer Errors Enable</i> field is 1b, the Router shall send a Notification Packet with Event Code = <i>ERR_LINK</i> to the Connection Manager (see Section 6.5). If the <i>RST</i> bit in the <i>Logical Layer Errors Enable</i> field is 0b, the Router shall not send a Notification Packet. |



CONNECTION MANAGER NOTE

The Connection Manager reads and clears all error indications after a USB4 Link comes up.

4.4.3 Clock Compensation and SKIP

An Adapter may transmit SKIP Ordered Sets at any time. However, it is recommended that an Adapter not send SKIP Ordered Sets unless explicitly required by this specification. A receiver shall drop any received SKIP Ordered Sets and shall be capable of operating in the absence of any SKIP Ordered sets.

The contents of SKIP Ordered Sets are bit-swapped to the order of transmission on the wire (see Section 4.3.4).

A Transmitter shall not send SKIP Ordered Sets while the Adapter is in the Training.LOCK1 or Training.LOCK2 sub-states.

A SKIP Ordered Set shall have the structure defined in Table 4-39.

Table 4-39. SKIP Ordered Set

| Bits | Value | Description |
|-------|---|--|
| 63:10 | 00 1010 1100 1010 1100 1010 1100 1010 1100 1010 1100 1010 1101 1100b | Contents, specific to the Ordered Set. |
| 9:0 | 11 0000 1101b | SCR – Shall be set to this value to indicate that the Ordered Set contents are not scrambled. |

4.4.4 Dual-Lane Skew

A Router shall operate with skew defined in Section 3.5.1 between the receiving Lanes of a Link when measured at the USB Type-C connector (See TP2 in Section 3.3.1). A transmitter shall not introduce skew more than defined in Section 3.4.1. When the received skew exceeds the capacity of the de-skew buffer, the Router may respond as described in Section 4.4.2.

The de-skew Ordered Set defined in Table 4-40 is used to achieve de-skew. The following rules apply for de-skew:

- A de-skew Ordered Set shall be sent on each Lane after both Adapters transition (from any state) to CL0 state, and the USB4 Port in Dual-Lane Link mode.
 - When exiting CL0s state, a de-skew Ordered Set shall only be sent in the direction exiting electrical idle.
 - When a Port operates at Gen 3 speed, the De-Skew Ordered Set shall be transmitted twice. Otherwise it shall be transmitted once.
- A de-skew Ordered Set shall be sent simultaneously on both Lanes within the permitted transmit skew.
- A de-skew Ordered Set shall be sent on both Lanes in the same locations within the Symbol. For example, when a Port operates at Gen 2 speed with RS-FEC enabled, the de-skew Ordered Set is be sent on both Lanes in either the first 64 bits of the Symbol or in the second 64 bits of the Symbol.
- The de-skew Ordered Set shall be the first bytes sent after the TS2 Ordered Sets. TS2 Ordered Sets shall be transmitted on an Adapter in CL0 state until the de-skew Ordered Set is sent.

Table 4-40. De-Skew Ordered Set

| Bits | Value | Description |
|-------|---|--|
| 63:10 | 11 0111 1001 0110 0011 1000 0011 1100 1011 0111 1001 0110 0011 1000b | Contents, specific to the Ordered Set. |
| 9:0 | 00 1111 0010b | SCR – Shall be set to this value to indicate that the Ordered Set contents are scrambled. |



IMPLEMENTATION NOTE

To ensure that the de-skew Ordered Set is sent on both Lanes in the same locations within the Symbol, the transmitter needs to align the beginning of each Symbol on the two Lanes, starting with the first transmitted SLOS Symbol.

4.4.5 Disconnect

This section defines how a disconnect on a USB4 Port is handled. When a disconnect is required, it takes precedence over all other pending requirements.

Section 4.4.5.1 describes a disconnect on the Upstream Facing Port. Section 4.4.5.2 describes a disconnect on a Downstream Facing Port.

4.4.5.1 Upstream Facing Port Disconnect

4.4.5.1.1 SBRX Goes Low

When an Upstream Facing Port detects SBRX at logical low for $t_{\text{DisconnectRx}}$, the Port is disconnected. The disconnected Port may optionally send an LT_LRoff Transaction.

The Router with the disconnected Port shall:

- Drive SBTX to a logical low state on all USB4 Ports for $t_{\text{DisconnectTx}}$.
- Transition to the Uninitialized Unplugged state (see Section 6.3).

Note: An Adapter might detect Link errors (caused by the Link Partner's Adapter transitioning to the CLd state) before the Upstream Facing Port detects SBRX transitioning to logical low. In this case, the Upstream Facing Port may transition its Adapter to the Training state, before the Router transitions to the Uninitialized Unplugged state.

The following events initiate a disconnect on the Upstream Facing Port of a Router as defined in this section:

- The Upstream Facing Port is Hot Unplugged (see Section 6.8.2.1).
- A Downstream Facing Port Reset on the upstream Link Partner (see Section 6.9).
- The Domain enters Sleep state, the *USB4 Port is Configured* bit in the upstream Link Partner is set to 0b, and the *Enable Wake on Connect* bit of the USB4 Port is 0b (see Section 4.5.1).
- The Link Partner failed to train the Link before the defined timeout (see Section 4.2.1.3.3).
- Lane Bonding did not complete before the defined timeout (See Section 4.2.2.2).

4.4.5.1.2 LT_Fall Transaction is Received

Section 4.4.6.2.1.2 describes a disconnect when an Upstream Adapter receives an LT_Fall Transaction.

4.4.5.1.3 LT_LRoff Transaction is Received on a Upstream Facing Port

When an Upstream Facing Port receives an LT_LRoff transaction and the *Enter Sleep* bit in the Router Configuration Space is set to 0b, the Port is disconnected. The disconnected Port shall send an LT_LRoff Transaction.

The Router with the disconnected Port shall:

- Drive SBTX to a logical low state on all USB4 Ports for tDisconnectTx.
- Transition to the Uninitialized Unplugged state (see Section 6.3).

Note: An Adapter might detect Link errors (caused by the Link Partner's Adapter transitioning to the CLd state) before the Upstream Facing Port detects SBRX transitioning to logical low. In this case, the Upstream Facing Port may transition its Adapter to the Training state, before the Router transitions to the Uninitialized Unplugged state.

The following events shall initiate a disconnect as defined in this section:

- The Domain enters Sleep state, in the Downstream Facing Port of the Link Partner the *USB4 Port is Configured* bit is 0b and the *Enable Wake on Connect* bit of the USB4 Port is 1b.

4.4.5.2 Downstream Port Disconnect

4.4.5.2.1 SBRX Goes Low

When a Downstream Facing Port detects SBRX at logical low for tDisconnectRx and its *Lane Disable* bit in the Lane 0 Adapter Configuration Capability is set to 0b, the Port is disconnected and shall:

- Send an LT_LRoff Transaction.
- Discard any Transport Layer Packets received from the Ingress Port(s) on the Router. The Port discards packets for a Path until the *Valid* bit in Path Configuration Space transitions from 0b to 1b.
- Transition its Adapters to the CLd state.

Note: An Adapter might detect Link errors (caused by the Link Partner's Adapter transitioning to the CLd state) before the Downstream Facing Port detects SBRX transitioning to logical

low. In this case, the Downstream Facing Port may transition its Adapter to the Training state, before transitioning to the CLd state.

The Router with the disconnected Port shall continue to send flow control Packets on the Ingress Port(s) for Transport Layer Packets that the disconnected Port discarded. Flow control credit counts shall be updated as if the discarded packets were dequeued and forwarded to the Egress Adapter.

The Router shall do the following for each enabled Lane Adapter in the disconnected Port:

- Send the Connection Manager a Hot Plug Event Packet with the *UPG* bit set to 1b.
- Load the following fields in the Adapter Configuration Space with their default values:
 - Basic Configuration Registers:
 - *Link Credits Allocated*.
 - *HEC Errors* (optional, recommended).
 - *Invalid HopID Errors* (optional, recommended).
 - *ECC Errors* (optional, recommended).
 - TMU Adapter Configuration Capability:
 - *Inter-Domain Time Initiator*.
 - *EnableUniDirectionalMode*.
 - Lane Adapter Configuration Capability:
 - *Target Link Width*.
 - *CL0s Enable*.
 - *CL1 Enable*.
 - *CL2 Enable*.
 - *Lane Bonding*.
 - *PM Secondary* (optional, recommended).
 - *Logical Layer Errors* (optional, recommended).
 - *Logical Layer Errors Enable* (optional, recommended).
- Load the TxFFE register in the SB Register Space with its default values

The following events initiate a disconnect on a Downstream Facing Port of a Router as defined in this section:

- The Downstream Facing Port detects a Hot Unplug (see Section 6.8.2.2).
- The Downstream Facing Port is reset (see Section 6.9).
- The Domain enters Sleep state and the *USB4 Port is Configured* bit in the Downstream Facing Port is set to 0b, and the *Enable Wake on inter-Domain* bit is set to 0b (see Section 4.5.1).
- The Domain enters Sleep state and, in the Downstream Facing Port, the *USB4 Port is inter-Domain* bit is 1b, and the *Enable Wake on inter-Domain* bit is set to 0b (see Section 4.5.1).
- The Link Partner failed to train the Link before the defined timeout (see Section 4.2.1.3.3).
- Lane Bonding did not complete before the defined timeout (See Section 4.2.2.2).

4.4.5.2.2 LT_LRoff Transaction is Received

When a Downstream Facing Port receives an LT_LRoff Transaction and the *Enter Sleep* bit in the Router Configuration Space is set to 0b, the Port is disconnected and shall:

- Send an LT_LRoff Transaction.
- Discard any Transport Layer Packets received from the Ingress Port(s) on the Router. The Port discards packets for a Path until the *Valid* bit in Path Configuration Space transitions from 0b to 1b.
- Transition its Adapters to the CLd state.

The Router with the disconnected Port shall continue to send flow control Packets on the Ingress Port(s) for Transport Layer Packets that the disconnected Port discarded. Flow control credit counts shall be updated as if the discarded packets were dequeued and forwarded to the Egress Adapter.

The Router shall do the following for each enabled Lane Adapter in the disconnected Port:

- Send the Connection Manager a Hot Plug Event Packet with the *UPG* bit set to 1b.
- Load the following fields in Adapter Configuration Space with their default values:
 - Basic Configuration Registers:
 - *Link Credits Allocated*.
 - *HEC Errors* (optional, recommended).
 - *Invalid HopID Errors* (optional, recommended).
 - *ECC Errors* (optional, recommended).
 - TMU Adapter Configuration Capability:
 - *Inter-Domain Time Initiator*.
 - *EnableUniDirectionalMode*.
 - Lane Adapter Configuration Capability:
 - *Target Link Width*.
 - *CL0s Enable*.
 - *CL1 Enable*.
 - *CL2 Enable*.
 - *Lane Bonding*.
 - *PM Secondary* (optional, recommended).
 - *Logical Layer Errors* (optional, recommended).
 - *Logical Layer Errors Enable* (optional, recommended).
- Load the TxFFE register in SB Register Space with its default values.
- Start Lane Initialization from Phase 2.
- The Router may wait in Phase 3 without sending AT Transactions until it receives AT transactions from its Link Partner.

The following events initiate a disconnect on a Downstream Facing Port of a Router as defined in this section:

- The Domain enters Sleep state, the *USB4 Port is inter-Domain* bit is 1b and the *Enable Wake on inter-Domain* bit is set to 1b (see Section 4.5.1).

Note: The Router will not get a response for its AT transactions in Phase 3 while its Link Partner is in Sleep state. The expected behavior is described in Section 4.1.1.2.5.1.

4.4.6 Lane Adapter Disable and Enable

A Connection Manager uses the *Lane Disable* bit in the Lane Adapter Configuration Capability to disable and enable a Lane Adapter. When the Lane Adapter on one side of a Link is disabled or enabled, it affects the state of the Adapter on the other side of the Link.

The flow for disabling and enabling the Upstream Adapter is different than the flow for disabling and enabling a Lane Adapter that is not the Upstream Adapter. When the Upstream Adapter is disabled, it automatically re-enables itself. When an Adapter that is not the Upstream Adapter is disabled, it does not re-enable itself until the Connection Manager sets the *Lane Disable* bit to 0b.

Section 4.4.6.1 describes what happens when the Upstream Adapter is disabled. Section 4.4.6.2 describes what happens when an Adapter that is not the Upstream Adapter is disabled and enabled.



CONNECTION MANAGER NOTE

A Connection Manager shall not change the value of the Lane Disable bit for a Lane Adapter unless all of the fields in its Lane Adapter Configuration Capability except the following are set to their default values: Lane Disable, Target Link Speed, Adapter State.

A Connection Manager shall not change the value of the Lane Disable bit while the USB4 Port is performing Lane Initialization.

See the Connection Manager Guide for more information on how a Connection Manager disables a Lane.

4.4.6.1 Disabled Adapter is the Upstream Adapter

This section defines Lane Disable when the *Lane Disable* bit in the Upstream Adapter is set to 1b.

After the *Lane Disable* bit in its Upstream Adapter is set to 1b, a Router shall:

1. Send an LT_Fall Transaction to the Link Partner of the Upstream Facing Port to signal transition to the Disabled state (see Section 4.1.1.2.1 for the format of the LT_Fall Transaction).
2. Drive SBTX to a logical low state on all USB4 Ports for tDisconnectTx.

Note: The Adapter at the Link Partner responds to its SBRX driven low with the flow defined in Section 4.4.5.2.1.


3. Transition to the Uninitialized Unplugged state (see Section 6.3.1).

Note: The Lane Disable bit in the Upstream Adapter is set to 0b as result of the Router entering the Uninitialized Unplugged state.

4. Start Lane Initialization from Phase 1.
5. After detecting a Router on the Upstream Facing Port (see Section 4.1.2.2), transition to the Uninitialized Plugged state (see Section 6.3.2).

Figure 4-21 shows the Disable flow between a Lane Adapter (Adapter A) and the Upstream Adapter (Adapter B).

Figure 4-21. Lane Disable of the Upstream Adapter



4.4.6.2 Disabled Adapter is not the Upstream Adapter

This section defines Lane Disable when the *Lane Disable* bit in a Lane Adapter that is not the Upstream Adapter is set to 1b.

4.4.6.2.1 Disable Flow

After the *Lane Disable* bit in the Adapter Configuration Space of the Lane Adapter is set to 1b, the Router shall:

1. Send an LT_Fall Transaction on the USB4 Port with the disabled Adapter to signal transition to the Disabled state (see Section 4.1.1.2.1 for the format of the LT_Fall Transaction).
2. Send a Hot Plug Event Packet for the disabled Adapter with the *UPG* bit set to 1b to the Connection Manager (see Section 6.4.2.9).
3. Transition the disabled Adapter to the Disabled state (see Section 4.2.1.1).
4. If the Router detects that SBRX of the disabled Adapter transitions to a low logical state for more than *tDisconnectRx*, the Router shall perform the disconnect flow defined in Section 4.4.5.2.1.

Note: The Adapter does not start Lane Initialization at this point since its *Lane Disable* bit is still set to 1b. The Adapter starts Lane Initialization when the *Lane Disable* bit is set to 0b.

4.4.6.2.1.1 Link Partner is not the Upstream Adapter

If the Link Partner of the disabled Adapter is not the Upstream Adapter, it shall do the following upon receiving the LT_Fall Transaction:

1. Send a Hot Plug Event Packet with the *UPG* bit set to 1b to the Connection Manager.
2. Load the following fields in the Adapter Configuration Space of the Adapter with their default values:

- Basic Configuration Registers:
 - *Link Credits Allocated.*
 - *HEC Errors* (optional, recommended).
 - *Invalid HopID Errors* (optional, recommended).
 - *ECC Errors* (optional, recommended).
- TMU Adapter Configuration Capability:
 - *Inter-Domain Time Initiator.*
- Lane Adapter Configuration Capability:
 - *Target Link Width.*
 - *CL0s Enable.*
 - *CL1 Enable.*
 - *CL2 Enable.*
 - *Lane Bonding.*
 - *PM Secondary* (optional, recommended).
 - *Logical Layer Errors* (optional, recommended).
 - *Logical Layer Errors Enable* (optional, recommended).


3. Transition the Adapter to the CLd state.

Note: An Adapter might detect Link errors (caused by the Link Partner transitioning to the Disabled state) before it detects the LT_Fall Transaction on the Sideband Channel. In this case, the Adapter that detected the errors may transition to Training state, followed by transition to CLd state.

Note: The Adapter does not start Lane Initialization at this point even though its Sideband Channel is up.

Figure 4-22 shows the Disable flow between two Lane Adapters that are not Upstream Adapters where Adapter A is the disabled Adapter and Adapter B is its Link Partner.

Figure 4-22. Lane Disable Flow



4.4.6.2.1.2 Link Partner is the Upstream Adapter

If the Link Partner of the disabled Adapter is the Upstream Adapter, it shall do the following upon receiving the LT_Fall Transaction:

1. Drive SBTX to a logical low state on all USB4 Ports for tDisconnectTx.

Note: The Adapter at the Link Partner responds to its SBRX driven low with the flow defined in Section 4.4.5.2.1.

2. Transition to the Uninitialized Unplugged state (see Section 6.3.1).

Note: A Lane Adapter might detect Link errors (caused by the Link Partner transitioning to the Disabled state) before it detects the LT_Fall Transaction on the Sideband Channel. In this case, the Adapter that detected the errors may transition to Training state, before the Router transitions to the Uninitialized Unplugged state.

3. Transition to the Uninitialized Plugged state (See Section 6.3.2) as a result of detecting SBRX driven high.
4. Start lane Initialization from Phase 1.

4.4.6.2.2 Enable Flow



CONNECTION MANAGER NOTE

After setting the Lane Disable bit to 0b in a Lane Adapter, the Connection Manager shall initiate a Downstream Facing Port Reset in the Port that contains the Lane Adapter. See Section 6.9 for more details.

4.4.7 Time Sync Notification Ordered Set (TSNOS)

When a Router receives a Time Sync Notification Ordered Set (TSNOS) it shall generate a time stamp as described in Section 7.2. The Time Sync Notification Ordered Set shall have the structure in Table 4-41.

Table 4-41. TSN Ordered Set

| Bits | Value | Description |
|-------|---|--|
| 63:10 | 11 1001 0011 1001 0011 1001 0011 1001 0011 1001 0011 1001 0001 1100b | Contents, specific to the Ordered Set. |
| 9:0 | 11 0000 1101b | SCR – Shall be set to this value to indicate that the Ordered Set contents are not scrambled. |

4.5 Sleep and Wake

This section describes how a Router participates in a system transition to sleep state and how wake events may bring the Router back from sleep.

4.5.1 Entry to Sleep

A Router shall enter sleep state when the *Enter Sleep* bit is set to 1b and one of the following sleep events occur:

- Host Router
 - The Router is a PCIe Host Router and it receives a PCIe PERST# signal that transitions from logical high to logical low. If the Router tunnels PCIe traffic, then it shall send at least 3 PERST Active Tunneled Packets on each Downstream Facing Port before entering Sleep state.
 - The Router receives an implementation-specific signal indicating entry to Sleep state.
- Device Router

- The Router tunnels PCIe traffic and receives a PERST Active Tunneled Packet on the Upstream Facing Port.
- The Router receives an LT_LRoff Transaction on the Sideband Channel of an Upstream Facing Port.

A Router shall not enter sleep state unless the *Enter Sleep* bit is set to 1b before a sleep event occurs.



CONNECTION MANAGER NOTE

If a Connection Manager initiates Transactions on the Sideband Channel, then the Connection Manager shall not set the Enter Sleep bit to 1b until all such Transactions are complete. A Connection Manager shall not initiate a Transaction on the Sideband Channel after the Enter Sleep bit is set to 1b.

After setting the Enter Sleep bit to 1b, a Connection Manager shall not write to a USB4 Port is Configured bit or USB4 Port is inter-Domain bit.

After the *Enter Sleep* bit is set to 1b, the Router shall complete any pending transactions on the Sideband Channel. When the Router is ready for the sleep event it shall set the *Sleep Ready* bit to 1b.

After a sleep event occurs, the Router shall do the following for each USB4 Port:

- If the *USB4 Port is inter-Domain* bit is 0b and the *USB4 Port is Configured* bit is 0b, perform a disconnect by driving its SBTX line low for tDisconnectTx. If the *Enable Wake on Connect* bit of the USB4 Port is 1b, the USB4 Port shall drive its SBTX line high after tDisconnectTx.
- If the *USB4 Port is inter-Domain* bit is 1b and the *Enable Wake on inter-Domain* bit is set to 0b, perform a disconnect by driving its SBTX line low for tDisconnectTx.
- Else:
 - If the Router supports PCIe Tunneling, send at least 3 PERST Active Tunneled Packets on all Downstream PCIe Adapters that have the *Path Enable* bit set to 1b.
 - Send an LT_LRoff Transaction on the Sideband Channel within tLRoffResponse from detecting the sleep event.
 - If the *USB4 Port is Inter-Domain* bit is 0b and the *USB4 Port is Configured* bit is 1b, wait for an LT_LRoff Transaction on the Sideband Channel, unless an LT_LRoff Transaction was already received from the time the *Enter Sleep* bit was set to 1b.
 - If the Downstream Facing Port is expecting an LT_LRoff Transaction and doesn't receive it within 30ms, it is recommended that the Port initiate a disconnect by driving its SBTX line low for tDisconnectTx.
 - Transition the Adapters to CLd state.

If a Downstream Facing Port drives its SBTX line low after a sleep event, it is recommended that it keep it low for the duration of the Sleep. This will prevent the attached Router from issuing Sideband transactions and allow it to enter a lower power state.



CONNECTION MANAGER NOTE

The following is an example of how a Connection Manager should transition its Domain into system sleep:

1. *Connection Manager sets the Enter Sleep bit to 1b in each Router.*
2. *Connection Manager reads the Sleep Ready bit in each Router.*

3. When all Routers have Sleep Ready bit set to 1b, host system initiates a Sleep Event, which causes Routers to transition to Sleep state.
4. Host system goes into system sleep.



IMPLEMENTATION NOTE

It is expected that an LT_LRoff Transaction is received within 30ms from sending the LT_LRoff Transaction.

If the LT_LRoff Transaction is sent over an inter-Domain Link, the Link Partner outside the Domain goes through disconnect as defined in Section 4.4.5.2.2.

4.5.2 Behavior in Sleep State

On entry to sleep state, a Router shall restore all Configuration Spaces to their default values.

If the *Enter Sleep* bit is set to 1b, a Router shall retain a copy of the state information listed in Table 4-42 separate from Configuration Space.

Table 4-42. Router State Retained During Sleep

| Field Copied | Configuration Space / Capability | Value |
|------------------------------|----------------------------------|---|
| Upstream Adapter | Router Configuration Space | <i>Upstream Adapter</i> field prior to sleep entry. |
| USB4 Port is Configured | USB4 Port Capability | <i>USB4 Port is Configured</i> bit prior to sleep entry. |
| USB4 Port is Inter-Domain | USB4 Port Capability | <i>USB4 Port is Inter-Domain</i> bit prior to sleep entry. |
| Enable Wake on Connect | USB4 Port Capability | <i>Enable Wake on Connect</i> bit prior to sleep entry. |
| Enable Wake on Disconnect | USB4 Port Capability | <i>Enable Wake on Disconnect</i> bit prior to sleep entry. |
| Enable Wake on Inter-Domain | USB4 Port Capability | <i>Enable Wake on Inter-Domain</i> bit prior to sleep entry. |
| Enable Wake on PCIe | Router Configuration Space | <i>Enable Wake on PCIe</i> bit prior to sleep entry. |
| Enable Wake on USB3 | Router Configuration Space | <i>Enable Wake on USB3</i> bit prior to sleep entry. |
| Enable Wake on DP | Router Configuration Space | <i>Enable Wake on DP</i> bit prior to sleep entry. |
| Far-end receiver termination | -- | Indicator value of USB far-end receiver termination. See Section 9.1.1.1.2. |
| DP Resource Allocation | -- | Which DP Resources were allocated to which DP IN Adapters prior to sleep entry. |

If a USB4 Port has the *USB4 Port is Inter-Domain* state set to 1b, then the USB4 Port shall ignore any Transactions received on the Sideband Channel while in Sleep state.

4.5.3 Wake Events

A Router shall support the wake events defined in Table 4-43.

Table 4-43. Wake Events

| Type | Description |
|----------------------|--|
| Wake on Connect | A Router shall issue a Wake on Connect if the <i>Enable Wake on Connect</i> bit of a USB4 Port is set to 1b, the <i>USB4 Port is Configured</i> bit is 0b, and it detects either of the following after the <i>Enter Sleep</i> bit is set to 1b: <ul style="list-style-type: none"> • A connection on the USB Type-C connector attached to the USB4 Port. • SBRX is at logic high on the USB4 Port for tConnectRx. |
| Wake on Disconnect | A Router shall issue a Wake on Disconnect if the <i>Enable Wake on Disconnect</i> bit of a USB4 Port is set to 1b, the <i>USB4 Port is Inter-Domain</i> bit is set to 0b, the <i>USB4 Port is Configured</i> bit is set to 1b, and the Router detects either of the following after the <i>Enter Sleep</i> bit is set to 1b: <ul style="list-style-type: none"> • A disconnect on the USB Type-C connector attached to the USB4 Port. • SBRX is at logic low on the USB4 Port for tDisconnectRx. |
| Wake on Inter-Domain | A Router shall issue a Wake on Inter-Domain event if the <i>Enable Wake on Inter-Domain</i> bit is set to 1b, the <i>USB4 Port is Inter-Domain</i> bit is set to 1b, and the Router detects either of the following after the <i>Enter Sleep</i> bit is set to 1b: <ul style="list-style-type: none"> • A disconnect on the USB Type-C connector attached to the USB4 Port. • SBRX is at logic low on the USB4 Port for tDisconnectRx. |
| Wake on PCIe | A Router shall issue a Wake on PCIe event if the <i>Enable Wake on PCIe</i> bit is set to 1b, and it detects a PCIe Wake event from any connected PCIe Endpoint or Switch after a Sleep Event occurs. |
| Wake on USB3 | A Router shall issue a Wake on USB3 event if the <i>Enable Wake on USB3</i> bit is set to 1b, and it detects a USB3 Wake event from any connected USB device after a Sleep Event occurs. |
| Wake on USB4 | A Router shall issue a Wake on USB4 event if the <i>USB4 Port is Inter-Domain</i> bit is set to 0b, the <i>USB4 Port is Configured</i> bit is set to 1b, the <i>Enable Wake on USB4 Wake</i> is set to 1b, and the Router detects at least one transition of SBRX to logical low for tWake time after a Sleep Event occurs. |
| Wake on DP | A Router shall issue a Wake on DP event if the <i>Enable Wake on DP</i> bit is set to 1b, and it detects an HPD change or reception of an HPD IRQ after a Sleep Event occurs. |

4.5.4 Exit from Sleep

A Router exits Sleep state when one of the following occurs:

- An Upstream Facing Port Disconnect (see Section 4.5.4.1)
- Detection of a Wake on USB4 event (see Section 4.5.4.2)

4.5.4.1 Upstream Facing Port Disconnect

When a Router detects a disconnect on the UFP, it shall exit sleep state. See Section 4.4.5.1 for how a Router detects and handles a UFP disconnect.

4.5.4.2 Wake on USB4 Event

A Wake on USB4 event is used to propagate a wake event throughout a USB4 Fabric. A Router shall assert SBTX to logical low for tWake time to indicate a Wake on USB4 event.

After detecting a wake event, a Router shall:

1. Issue a Wake on USB4 event on all connected USB4 Ports by asserting SBTX to logical low for tWake time.
 - If the detected wake event is a Wake on USB4 event, the Router may issue a Wake on USB4 event to the USB4 Port where the Wake on USB4 event arrived, but is not required to do so.

2. Begin Lane Initialization on all connected USB4 Ports.
 - A USB4 Port may ignore any Transactions received before it is ready for Lane Initialization. The transmitting USB4 Port shall retry the Transactions as defined in Section 4.1.1.2.5.
3. For every Adapter that reaches CL0 state, the Router shall send a Hot Plug Event Packet to the Connection Manager with the *UPG* bit set to 0b.

4.6 Timing Parameters

Table 4-44 lists the timing parameters for the Logical Layer.



IMPLEMENTATION NOTE

When both a Min and a Max value are listed in the table below, an implementation can choose to implement that parameter with any value in the given range. Furthermore, an implementation cannot assume that its Link Partner will use the Max value and must be prepared for Link Partners that use any value between Min and Max (inclusive).


Table 4-44. Logical Layer Timing Parameters

| Parameter | Description | Min | Max | Units |
|---------------|--|-----|------|----------|
| tDisconnectTx | The time that SBTX is driven to logical low to identify a disconnect. | 50 | -- | ms |
| tDisconnectRx | The time that SBRX is detected in logical low to identify a disconnect. | 14 | 1000 | μs |
| tConnectRx | The time that SBRX is detected in logical high to identify a new connection. | 25 | -- | μs |
| tCmdResponse | The time between receiving an AT Command and sending an AT Response or between receiving an Addresses RT Command and sending an RT Response. | -- | 50 | ms |
| tATTimeout | The amount of time the originator of an AT Command waits before optionally timing out the AT Command. | 100 | -- | ms |
| tRTTimeout | The amount of time the originator of an Addressed RT Command waits before optionally timing out the Addressed RT Command. | 100 | -- | ms |
| tLTPhase4 | The amount of time that Broadcast RT Transactions are sent after completion of Lane Initialization phase 2. | 25 | -- | ms |
| tWake | The time that SBRX is in logic low to initiate exit from Sleep State. | 1 | 9 | Bit time |
| tDisabled | The minimal time an Adapter stays in a Disabled State after entering the state. | 10 | -- | ms |

| Parameter | Description | Min | Max | Units |
|-----------------|---|-----------------------------|---------------------------------------|-------------|
| tTrainingError | The time an Adapter has to transition from Training state to CL0 state before a Timeout Error is generated. The time is measured from when the Adapter sends the first SLOS1 to when it transitions to CL0 state. | -- | 500 | μs |
| tTrainingAbort1 | The amount of time in Training state following Lane Initialization. | 1 | -- | sec |
| tTrainingAbort2 | <ol style="list-style-type: none"> 1. The amount of time in Training state following any transition to Training state other than from CLd state. 2. The time to send LFPS when exiting CL1 or CL2. 3. The time to send SLOS1 in CLOs exit. | 100 | -- | ms |
| tLaneParams | The time interval between transmissions of Broadcast RT Transactions. | 1 | 5 | ms |
| tPollTXFFE | The rate of polling during the TxFFE flows. | 1 | 5 | ms |
| tBonding | The time to transition to CL0 state from Lane Bonding state before failure to bond is determined. | 40 | 100 | μs |
| tCLxIdleRx | The time for a receiver to wait after the last LFPS cycle received before starting calibration. | 130 | 300 | ns |
| tCLxRetry | The time that a requester waits after receiving the last bit of a CL_NACK Ordered Set before sending the first bit of another CL2_REQ Ordered Set or CL1_REQ Ordered Set. | 1.5 | -- | μs |
| tCLxRequest | Time from receiving the last bit of a Request Ordered Set to sending the first bit of a CL2_ACK, CL1_ACK, or CL0s_ACK Ordered Set. | -- | 13 (Host Router) 2 (Device Router) | μs |
| tCLxResponse | Time from detection of a CL2_ACK, CL1_ACK, or CL0s_ACK Ordered Set to sending the first CL_OFF Ordered Set. Also, the time from detection of a Link error to sending the first SLOS. | -- | 120 | Symbol Time |
| tTxOff | Time between sending the last CL_OFF Ordered Set and shutting off the transmitters of a requesting USB4 Port when entering a CLx state. | -- | 100 | ns |
| tEnterLFPS1 | Time between shutting off the transmitters of a requesting USB4 Port and enabling transmission and detection of LFPS in CL1 or CL2 states. | 400 | 500 | ns |
| tEnterLFPS2 | Time between shutting off the receivers of a responding USB4 Port after reception of a CL_OFF Ordered Set and enabling transmission and detection of LFPS in CL1 or CL2 states. | 375 Symbol Times + 100ns | 5000 | ns |

| Parameter | Description | Min | Max | Units |
|---------------------|---|--------------------------|--------------------------------------|-------|
| tEnterLFPS3 | Time between shutting off the receivers of a responding USB4 Port after timeout of CL_OFF Ordered Set reception and enabling transmission and detection of LFPS in CL1 or CL2 states. | 240 Symbol Times + 100ns | -- | ns |
| tEnterLFPS4 | Time between shutting off the transmitters of a requesting USB4 Port and enabling transmission of LFPS in CL0s state. | 240 Symbol Times + 100ns | 2000 | ns |
| tEnterLFPS5 | Time between shutting off the receivers of the responding USB4 Port after reception a CL_OFF Ordered Set and enabling detection of LFPS in CL0s states. | 375 Symbol Times + 100ns | 375 Symbol Times + 600ns | ns |
| tLFPSDuration | Time between transmitting the first LFPS cycle and transmitting the last LFPS cycle. | -- | 2560 | ns |
| tCL0sSwitch | Time to abort normal CL0s exit flow if a CL_WAKE1.(X+1) Ordered Set Symbol does not arrive after CL_WAKE1.X Ordered Set Symbols. | 500 | -- | μs |
| tWarmUpCL0s | When exiting CL0s state, time between receiving the first LFPS cycle and achieving the first Symbol lock. | -- | 50 | μs |
| tWarmUpCL1 | When exiting CL1 state, time between receiving the first LFPS cycle and sending the first LFPS. | -- | 10 | μs |
| tWarmUpCL2 | When exiting CL2 state, time between receiving the first LFPS cycle and sending the first LFPS. | -- | 100 | μs |
| tRxLock | Time to achieve the first Symbol lock during exit from CLx states. | -- | 50 | μs |
| tWakeResponse | Time between receiving the last bit of a CL_WAKE1/2.X Ordered Set Symbol and sending the first bit of the relevant response (CL_WAKE2.X or SLOS). | -- | 7 (Host Router) 1 (Device Router) | μs |
| tTrainingTransition | Time to transition Training Sub-State where no errors on the receive signal | -- | 7 (Host Router) 1 (Device Router) | μs |
| tSymbolLock | Time to achieve Symbol alignment lock in CL0s/1/2 exit flow. | -- | 1 | μs |
| tLRoffResponse | Time between a sleep event and sending an LT_LRoff Transaction. | -- | 20 | ms |
| tTS2Timeout | While exiting CL0s state, the time to transition to Training state after sending TS2 Ordered Sets and not receiving TS2 Ordered Sets. | 9 | -- | μs |

5 Transport Layer




5.1 Transport Layer Packets

5.1.1 Bit/Byte Conventions

The packet formats defined in this chapter are shown with the Doubleword layout in Figure 5-1 where the least-significant bit is on the right. Fields within each Doubleword are also shown using this convention.

Figure 5-1. Convention for Transport Layer Diagrams




* SupplD

5.1.2 Format

All Transport Layer Packets shall start with the 4-byte header described in Section 5.1.2.1. All Transport Layer Packets except Idle Packets shall carry between 1 and 256 bytes (inclusive) of

payload. Section 5.1.3.3.1 described Idle Packets. Figure 5-2 shows the format of a Transport Layer Packet with payload.

Figure 5-2. Transport Layer Packet Format



5.1.2.1 Header

Table 5-1. Transport Layer Packet Header Format

| Bits | Field | Description |
|-------|----------------|--|
| 7:0 | <i>HEC</i> | Header Error Control – See Section 5.1.2.1.1. |
| 15:8 | <i>Length</i> | Length – Shall contain the payload size in bytes excluding the padding size. A value of 00h is used to indicate a payload size of 256 bytes. |
| 26:16 | <i>HopID</i> | HopID – Uniquely identifies a Path in the context of a Link. See Section 5.2.2. |
| 27 | <i>SupplID</i> | Supplemental ID – Used to distinguish certain types of Link Management Packets that need additional distinction. Shall be set to 0b in a Tunneled Packet. |
| 31:28 | <i>PDF</i> | Protocol Defined Field – Contents depend on packet type. Tunneled Packets: defined by the Protocol Adapter Layer. Link Management and Control Packets: together with the <i>SupplID</i> and <i>HopID</i> fields, identifies the Link Management or Control Packet type. |

5.1.2.1.1 Header Error Control (HEC)

The *HEC* field in a Transport Layer Packet header shall cover bits[31:8] of the Transport Layer Packet header. It shall not cover any payload. The *HEC* field consists of 8 redundancy bits, which shall be calculated from bit 31 to bit 8 as follows:

- Width: 8
- Poly: 07h
- Init: 00h
- RefIn: False
- RefOut: False
- XorOut: 55h

See Appendix A for examples of HEC calculations.

When a Router receives a Transport Layer Packet, it shall verify the *HEC* field value in the packet. The Router shall correct any single-bit errors in the Transport Layer Packet header. After correcting an error, a Router shall continue on as if the error had never occurred.

When an Ingress Adapter that is a Lane Adapter detects an uncorrectable HEC error, it shall:

- Drop the packet with the error.
- Set the *HEC Error* bit in the Adapter Configuration Space to 1b.
- Increment the *HEC Errors* field in Adapter Configuration Space.

If the Ingress Adapter that detected the uncorrectable HEC error is a part of a Downstream Facing Port:

1. The Ingress Adapter shall send a Notification Packet upstream if the *HEC Error Enable* bit in the Adapter Configuration Space is set to 1b. The Notification Packet shall contain Event Code = ERR_HEC (see Section 6.5).
2. The Lane Adapters in the USB4™ Port with the Ingress Adapter shall enter the Training state.

If the Ingress Adapter that detected the uncorrectable HEC error is a Lane Adapter that is a part of an Upstream Facing Port:

1. The Lane Adapter(s) in the USB4™ Port with the Ingress Adapter shall enter the Training state.
2. When the Lane Adapter enters CL0 state, it shall send a Notification Packet upstream if the *HEC Error Enable* bit in the Adapter Configuration Space is set to 1b. The Notification Packet shall contain Event Code = ERR_HEC (see Section 6.5).

5.1.2.2 Payload Padding

The size of a Transport Layer Packet is always a multiple of 4 bytes. The Protocol Adapter Layer of a Source Adapter shall add between 0 and 3 bytes of padding to the payload of a Tunneled Packet to ensure that the Tunneled Packet is of a size that is a multiple of 4 bytes. The Protocol Adapter Layer of the Destination Adapter shall remove any bytes of padding.

The content of the padding bytes is implementation-specific and cannot be assumed to have any specific values. However, it is recommended that zero-byte padding be used.

5.1.2.3 Error Correction Code (ECC)

When a Transport Layer Packet contains an *ECC* field, the ECC shall be calculated as described in this section.

The *ECC* field is calculated over a portion of the payload within a Transport Layer Packet. The *ECC* field consists of 8 redundancy bits, which shall be calculated from most significant bit to least significant bit as follows:

- Width: 8
- Poly: 07h
- Init: 00h
- RefIn: False
- RefOut: False
- XorOut: 00h

See Appendix A for examples of ECC calculations.

5.1.3 Transport Layer Packets

Transport Layer Packets are either Tunneled Packets from Protocol Adapter Layers, Control Packets from the Configuration Layer, or generated within Transport Layer (namely, Link Management Packets generated within the Transport Layer).

The format and types of Link Management Packets except for Time Sync Packets are defined in this section. Time Sync Packets are defined in Section 7.3.3 of this specification.

5.1.3.1 Tunneled Packets

Tunneled Packets are generated by the Protocol Adapter Layer of a Source Adapter and handed off to the Transport Layer. A Tunneled Packet shall have the header defined for Transport Layer Packets in Table 5-1. The Protocol Adapter Layer of a Source Adapter shall fragment Protocol traffic larger than 256 bytes into multiple Tunneled Packets. Re-assembly of Protocol traffic from Tunneled Packets shall be performed by the Protocol Adapter Layer of the Destination Adapter.

A Transport Layer may modify the *HopID* and *HEC* fields of a Tunneled Packet. It shall not modify any other fields in a Tunneled Packet header and it shall not modify the payload.

5.1.3.2 Control Packets

Control Packets are defined in Section 6.4.2.

5.1.3.3 Link Management Packets


Link Management Packets are confined to a Link, originating from a Transport Layer at one side of the Link and terminate at the Transport Layer at the other side of the Link.

5.1.3.3.1 Idle Packets

Idle Packets ensure that the Transport Layer feeds the Logical Layer with a continuous byte stream. When an Adapter is in CL0 state, the Transport Layer shall insert Idle Packets at the transmitting end of a USB4 Link if there are no other Transport Layer Packets to be transmitted. A Router shall remove Idle Packets at the receiving end of the USB4 Link.

An Idle Packet consists of a Transport Layer Packet header with no payload. An Idle Packet shall have the format shown in Figure 5-3.

Figure 5-3. Idle Packet Contents



5.1.3.3.2 Credit Grant Packet


This packet is used to transmit credit information from an Ingress Adapter to an Egress Adapter. A Credit Grant Packet shall include the header in Table 5-2 followed by one or more of the Credit Grant Records defined in Table 5-3. A Credit Grant Packet shall not contain more than 64 Credit Grant Records. When more than one Credit Grant Record is sent in the same Credit Grant Packet, they shall be processed in the order received.

Table 5-2. Credit Grant Packet Header

| Bits | Field | Description |
|-------|---------------|--|
| 7:0 | <i>HEC</i> | See Section 5.1.2.1.1 |
| 15:8 | <i>Length</i> | Number of Credit Grant Records * 4 or 00h if packet is carrying 64 Credit Grant Records (256 bytes) |
| 26:16 | <i>HopID</i> | 001h |
| 27 | <i>SupplD</i> | 0b |
| 31:28 | <i>PDF</i> | 1h |

Table 5-3. Credit Grant Record Format

| Bits | Field | Description |
|-------|--------------------|--|
| 7:0 | <i>ECC</i> | Error Correction – Calculated over bits [31:8] of the Credit Grant Record. See Section 5.1.2.3 for calculation. |
| 15:8 | <i>Credits</i> | Flow Control Credits – Indicates the total number (modulo 256) of flow control credits granted to the Egress Adapter for the specified Path or Shared Buffer since initialization. For a Path, it carries the value of the PCA state variable (see Section 5.3.2.1.2). For a Shared Buffer, it carries the value of the SCA state variable (see Section 5.3.2.1.2). |
| 26:16 | <i>CreditHopID</i> | HopID – Indicates the HopID of the Path for which credit grant shall be applied. This field shall only be valid if L Flag= 0b. |
| 30:27 | <i>Rsvd</i> | Reserved |
| 31 | <i>L</i> | L Flag – Identifies whether the Credit Grant Record applies to a Path or a Shared Buffer. Shall be set to 0b for a Path or 1b for a Shared Buffers. |

Figure 5-4. Credit Grant Packet Format**5.1.3.3.3 Path Credit Sync Packet**


This packet is used by an Egress Adapter to maintain credit count synchronization for a given Path with an Ingress Adapter. A Path Credit Sync Packet shall consist of the header in Table 5-4 followed by the payload defined in Table 5-5. See also Figure 5-5.

Table 5-4. Path Credit Sync Packet Header

| Bits | Field | Description |
|-------|---------------|--|
| 7:0 | <i>HEC</i> | See Section 5.1.2.1.1 |
| 15:8 | <i>Length</i> | 04h |
| 26:16 | <i>HopID</i> | HopID corresponding to the target Path |
| 27 | <i>SuppID</i> | 1b |
| 31:28 | <i>PDF</i> | 0h |

Table 5-5. Path Credit Sync Packet Payload

| Bits | Field | Description |
|-------|-------------|--|
| 7:0 | <i>ECC</i> | Error Correction – Calculated over bits [31:8] of the Path Credit Sync Packet payload. See Section 5.1.2.3 for calculation. |
| 15:8 | <i>PCC</i> | Path Credits Consumed – Represents the number (modulo 256) of credits consumed by the Egress Adapter for the Path identified by the <i>HopID</i> field up to (and including) the last packet transmitted on the Path prior to this Path Credit Sync Packet. |
| 31:16 | <i>Rsvd</i> | Reserved. |

Figure 5-5. Path Credit Sync Packet Format**5.1.3.3.4 Shared Buffers Credit Sync Packet**


This packet is used by an Egress Adapter to maintain credit count synchronization with an Ingress Adapter for the Shared Buffer. A Shared Buffers Credit Sync Packet shall consist of the header in Table 5-6 followed by the payload defined in Table 5-7. See also Figure 5-6.

Table 5-6. Shared Buffers Credit Sync Packet Header

| Bits | Field | Description |
|-------|---------------|-----------------------|
| 7:0 | <i>HEC</i> | See Section 5.1.2.1.1 |
| 15:8 | <i>Length</i> | 04h |
| 26:16 | <i>HopID</i> | 001h |
| 27 | <i>SuppID</i> | 0b |
| 31:28 | <i>PDF</i> | 2h |

Table 5-7. Shared Buffers Credit Sync Packet Payload

| Bits | Field | Description |
|-------|-------------|--|
| 7:0 | <i>ECC</i> | Error Correction – Calculated over bits [31:8] of the Shared Buffers Credit Sync Packet payload. See Section 5.1.2.3 for calculation. |
| 15:8 | <i>SCC</i> | Shared Credits Consumed – Represents the number (modulo 256) of Shared Buffer credits consumed by the Egress Adapter up to (and including) the last packet transmitted prior to this Shared Buffers Credit Sync Packet. |
| 31:16 | <i>Rsvd</i> | Reserved. |

Figure 5-6. Shared Buffers Credit Sync Packet Format

5.1.4 Effect of Link State on Transport Layer Packets

Table 5-8 describes how the Logical Layer Link states defined in Section 4.2.3 affect Transport Layer behavior with regards to Transport Layer Packet transport.

Table 5-8. Transport Layer Behavior per Link State

| Link State | Tunneled Packets and Control Packets | Credit Grant Packets | Credit Sync Packets | Time Sync Packets |
|---|---|--|---|--|
| Inactive | Shall not be sent to or received from the Logical Layer. Packets may be dropped by the transmitting Transport Layer. | Shall not be sent to or received from the Logical Layer. | Shall not be sent to or received from the Logical Layer. | Shall not be sent to or received from the Logical Layer. |
| Active | May be sent to Logical Layer and received from Logical Layer. | May be sent to Logical Layer and received from Logical Layer. | May be sent to Logical Layer and received from Logical Layer. | May be sent to Logical Layer and received from Logical Layer. |
| Low Power | Shall trigger transition to Active state. | Credit Grant Packets shall only be sent as a result of a the following: <ul style="list-style-type: none"> Increment on a packet dequeue. Update on a Credit Sync due to packet loss. Initial credits allocation to a Path. Sending a Credit Grant Packet shall trigger transition to Active state. | Shall not be sent. | Shall trigger transition to Active state in time to send the Time Sync Packet. |
| <p><i>Note: CLOs is an asymmetric Low Power state and transition to the Active Link state can only be initiated by the Port that has its Transmitter shut down.</i></p> | | | | |

5.1.5 Minimum Headers Gap


A Router needs to maintain a minimum gap between two non-Idle Transport Layer Packet Headers. The size of the gap depends on the Link type. A Router shall insert Idle Packets to meet the requirements defined in Table 5-9. Figure 5-7 shows an example of two 64-bit Data Symbols sent concurrently with RS-FEC off.

Table 5-9. Minimum Transport Layer Header Gap Requirements

| Link Type | Requirement |
|--------------------|--|
| Gen 2, Single-Lane | A Router shall send no more than one non-Idle Transport Layer Header in a 64-bit Data Symbol. <i>Note: This requirement is always satisfied, as the minimal length of a non-Idle Transport Layer Packet is 64 bits.</i> |
| Gen 2, Dual-Lane | A Router shall send no more than one non-Idle Transport Layer Header in the two 64-bit Data Symbols that are sent concurrently on the two Lanes of a Dual-Lane Link. |
| Gen 3, Single-Lane | A Router shall send no more than one non-Idle Transport Layer Header in a 128-bit Data Symbol. |

| | |
|------------------|---|
| Gen 3, Dual-Lane | A Router shall send no more than one non-Idle Transport Layer Header in the two 128-bit Data Symbols that are sent concurrently on the two Lanes of a Dual-Lane Link. |
|------------------|---|

Figure 5-7. Two Concurrent Data Symbols Example



IMPLEMENTATION NOTE

If a Router receives two or more Transport Layer Packet Headers that violate the rules defined in Table 5-9, it may drop any of the Transport Layer Packets.

5.2 Routing

The Path assigned to a Transport Layer Packet determines how that packet is routed through the USB4 Fabric. This section describes how Transport Layer Packets traverse a Path as forwarded by the Routers along that Path.

5.2.1 Adapter Numbering Rules

A Router can support up to 64 Adapters (including the Control Adapter). The following rules determine how the Adapters within a Router are numbered:

- Each Adapter shall be assigned a different 6-bit Adapter Number.
- The Control Adapter shall be assigned Adapter Number 0.

- For a Device Router, the Lane 0 Adapter in the Upstream Facing Port shall be assigned the lowest Adapter Number among all Lane Adapters.
- A USB4 Port shall have two Lane Adapters. The Lane Adapter Numbers within a USB4 Port shall be consecutive. The Lane 0 Adapter shall have a lower number than the Lane 1 Adapter.
- If an Adapter Number less than the Max Adapter is unused, a Router shall use one of the following methods to indicate that the Adapter is unused:
 - Assign a value of “Unsupported Adapter” to the *Adapter Type* field in the Adapter Configuration Space of the unused Adapter Number. In this case, the Adapter needs to implement the basic configuration registers in its Adapter Configuration Space. It is recommended that a Router use this method.
 - Respond to a Read Request or Write Request that targets the Adapter Configuration Space of the unused Adapter with a Notification Packet with Event Code = ERR_ADDR. In this case, the unused Adapter is not required to implement an Adapter Configuration Space.
- If a Device Router tunnels PCIe traffic, then the Upstream PCIe Adapter shall be assigned the lowest Adapter number among all PCIe Adapters.
- If a Device Router tunnels USB3 traffic, then the Upstream USB3 Adapter shall be assigned the lowest Adapter number among all USB3 Adapters.

5.2.2 HopID Rules

HopIDs are used by Routers to direct packets between Links on a Path. HopIDs are configured by the Connection Manager.



CONNECTION MANAGER NOTE

When setting up a Path, a Connection Manager sets the Output HopID field in the Path Configuration Space of the Ingress Adapter to allocate a HopID for that Path. The Output HopID in an Egress Adapter becomes the Input HopID of the next Ingress Adapter on the Path.

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

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

5.2.3 Routing Tables

An Ingress Adapter uses a Routing Table to determine the Egress Adapter and Egress HopID value for a received Transport Layer Packet. The Routing Table contains an entry for each Ingress HopID that the Ingress Adapter supports. A Routing Table entry contains the Egress Adapter Number and the Egress HopID that correspond to the Ingress HopID. Figure 5-8 depicts a schematic structure of a Routing Table.

Figure 5-8. Routing Table

| | | |
|----------------------------|----------------|--------------|
| HopID = n* | Egress Adapter | Egress HopID |
| | | |
| Ingress HopID → | · | · |
| | · | · |
| HopID = Max Input HopID | Egress Adapter | Egress HopID |

* The first HopID in the table is 1 for Host Interface Adapters and 8 otherwise

A Routing Table is populated according to the Path Configuration Space of the Ingress Adapter. For each Ingress HopID, the Path Configuration Space entry at offset (2 * Ingress HopID) is used as follows:

- The Egress Adapter in the Routing Table entry shall equal the *Output Adapter* field in the Path Configuration Space entry.
- The Egress HopID in the Routing Table entry shall equal the *Output HopID* field in the Path Configuration Space entry.

A Host Interface Adapter shall contain Routing Table entries for Ingress HopIDs 1 through *Max Input HopID*. All other Adapters shall contain Routing Table entries for Ingress HopIDs 8 through *Max Input HopID*.

5.2.4 Routing Rules

Each Ingress Adapter shall have its own Routing Table. For a Single-Lane Link, the Routing Table of the Ingress Adapter that a Transport Layer Packet arrives on shall be used to route the packet. For a Dual-Lane Link, the Routing Table of the Lane 0 Adapter of the Ingress USB4 Port that a Transport Layer Packet arrives on shall be used to route the packet.

5.2.4.1 Control Packets

Control Packets use HopID 0. A Lane Adapter and a Host Interface Adapter shall forward a Transport Layer Packet with a HopID value of 0 to the Control Adapter. The Control Adapter shall forward the packet to an Egress Adapter as defined in Section 6.4.3.2.

A USB3 Adapter, PCIe Adapter, and a DP Adapter do not send or receive Control Packets.

5.2.4.2 Link Management Packets

Link Management Packets use HopIDs 1 through 7 (inclusive). Only Lane Adapters route Link Management Packets – a Protocol Adapter does not send or receive Link Management Packets.

The following rules determine how a Lane Adapter routes a Transport Layer Packet with HopID 1 through 7:

- A Transport Layer Packet with a HopID value of 1 shall be forwarded to the Transport Layer for credit management processing. It shall not be forwarded to an Egress Adapter.
- A Transport Layer Packet with a HopID value of 2 shall be dropped and no further action shall be taken on its behalf.
- A Transport Layer Packet with a HopID value of 3 shall be forwarded to the TMU. It shall not be forwarded to an Egress Adapter.

- A Transport Layer Packet with a HopID value of 4 through 7 shall be dropped and no further action shall be taken on its behalf.

5.2.4.3 Tunneled Packets

Tunneled Packets use HopIDs 8 through Max Input HopID (inclusive). For a Host Interface Adapter only, Tunneled Packets can also use HopIDs 1 through 7 (inclusive).

The following rules determine how an Ingress Adapter routes a Tunneled Packet:

- If the *Valid* bit in the Path Configuration Space for the Path of the packet is 0b, the packet shall be dropped and no further action shall be taken on its behalf.
- If the Transport Layer Packet has an Ingress HopID that is greater than Max Input HopID of the Ingress Adapter, the packet shall be dropped and no further action shall be taken on its behalf.
- If the Routing Table entry corresponding to the Ingress HopID of the Transport Layer Packet contains an Egress HopID greater than the *Max Output HopID* of the Egress Adapter, the packet shall be dropped by the Router and no further action shall be taken on its behalf.
- If the Routing Table entry corresponding to the Ingress HopID of the Transport Layer Packet contains an Egress Adapter that is greater than the *Max Adapter* field in Router Configuration Space, the packet shall be dropped and no further action shall be taken on its behalf.
- Else, the Ingress Adapter shall:
 1. Replace the Ingress HopID value in the Tunneled Packet with the Egress HopID in the Routing Table entry that corresponds to the Ingress HopID.
 2. Update the *HEC* field in the Tunneled Packet.
 3. Forward the Tunneled Packet to the Egress Adapter in the Routing Table entry that corresponds to the Ingress HopID of the Tunneled Packet.



CONNECTION MANAGER NOTE

Routing Tables entries for HopIDs greater than 7 are populated by the Connection Manager during Path setup by programming the Path Configuration Space of the Ingress Adapter.

When configuring a Routing Table to direct outgoing traffic to a Dual-Lane Link, a Connection Manager shall use the Adapter Number of the Lane 0 Adapter associated with the Dual-Lane Link.

5.2.4.4 Routing Example

Figure 5-9 shows an example of how Routing Tables are used to route Transport Layer Packets along a Path. In the example, there are two Paths (one in each direction). The two Paths are independent of each other.


For Path X, a Router (not shown) injects a Transport Layer Packet onto Path X with a HopID of 10. The packet is received on Adapter 1 (Ingress Adapter) of Router A. Adapter 1 remaps the HopID to 11 and forwards the packet to Adapter 2 (Egress Adapter). The packet is then received on Adapter 1 (Ingress Adapter) of Router B. Adapter 1 remaps the HopID to 17 and forwards the packet to Adapter 3 (Egress Adapter).

For Path Y, a Router (not shown) injects a Transport Layer Packet onto Path Y with a HopID of 18. The packet is received on Adapter 3 (Ingress Adapter) of Router B. Adapter 3 remaps the HopID to 9 and forwards the packet to Adapter 1. The packet is then received on Adapter 2 (Ingress Adapter) of Router A. Adapter 2 remaps the HopID to 14 and forwards the packet to Adapter 1.

- USB4 Ports and PCIe Adapters are paired in order of increasing Adapter numbers. A Router shall be able to forward a packet received on the Lane 0 Adapter of a USB4 Port to the paired PCIe Adapter. A Router shall be able to forward a packet received on a PCIe Adapter to the Lane 0 Adapter of the paired USB4 Port.
- A Router shall be able to forward a packet received on an Upstream USB3 Adapter to the Upstream Adapter.
- A Router shall be able to forward a packet received on the Upstream Adapter to the Upstream USB3 Adapter.
- USB4 Ports and USB3 Adapters are paired in order of increasing Adapter numbers. A Router shall be able to forward a packet received on the Lane 0 Adapter of a USB4 Port to the matching USB3 Adapter. A Router shall be able to forward a packet received on a USB3 Adapter to the Lane 0 Adapter of the matching USB4 Port.

Figure 5-10 depicts an example of a USB3 Adapter connectivity scheme in a Router with three DFP. Connectivity for PCIe Adapters follows the same scheme.

Figure 5-10. Example of Connectivity for USB3 Adapters



Note: This section defines the minimum connectivity rules for a Router. A Router implementation may or may not offer additional connectivity. Router behavior is undefined if a Connection Manager configures a Path that assumes connectivity beyond what is defined in this specification.

5.3 Quality of Service (QoS)

5.3.1 Packet Ordering

A Router shall transmit Transport Layer Packets for a Path in the same order that they are received. The ordering of packets on a single Path is guaranteed. However, ordering of packets transmitted on multiple Paths is not guaranteed.

The ordering of Transport Layer Packets on one Path shall not affect the ordering of packets on any other Path.

5.3.2 Flow Control

The Transport Layer employs a per-Link, credit-based flow control mechanism to prevent overflow of receive buffers due to congestion. Each enabled Path has its own flow control scheme. There are four flow control schemes defined in this specification:

- **Flow Control Disabled** – Transport Layer Packets are not flow controlled. An Ingress Adapter drops packets if there is insufficient space in the receive buffer.
- **Dedicated Flow Control** – The Ingress Adapter has buffer space allocated for the Path. The Egress Adapter can only send a Transport Layer Packet when the Ingress Adapter has space in the path buffer.
- **Shared Flow Control** – The Ingress Adapter has a shared buffer for all Paths that use shared flow control. The Egress Adapter can only send a Transport Layer Packet when there is space in the shared buffer.
- **Restricted Shared Flow Control** – The Ingress Adapter uses the shared buffer (same as used for Shared Flow Control), but the Path can only use a limited amount of space in the shared buffer. The Egress Adapter can only send a Transport Layer Packet when there is space in the shared buffer and it has not exceeded its space in the shared buffer.

Flow control is managed individually for each Link. Flow control is managed by the Ingress Adapter and the Egress Adapter at either end of a Link. Section 5.3.2.1 defines flow control management for an Ingress Adapter and Section 5.3.2.2 defines flow control management for an Egress Adapter.

Credits are used to track the number of Transport Layer Packets that an Ingress Adapter can receive. One credit corresponds to one Transport Layer Packet of up to the maximum size. For example, if an Ingress Adapter advertises three credits, it has buffer space to accept three maximum-sized Transport Layer Packets (even though the actual Transport Layer Packets it receives may be less than the maximum size).

Link Management Packets are not subject to flow control and shall not be stored in any of the Flow Control Buffers defined in this section. See the respective sections for handling each type of Link Management Packets.



CONNECTION MANAGER NOTE

For each enabled Path, the Connection Manager shall configure both ends of a Link to have the same flow control scheme.

A Connection Manager shall only use the Flow Control Disabled scheme for the Main-Link Path of DisplayPort Tunneling.

5.3.2.1 Ingress Adapter

The *IFC Flag* and *ISE Flag* in Path Configuration Space determine which flow control scheme is used for a Path. Table 5-10 defines the flow control schemes that are used for each set of *IFC Flag* and *ISE Flag* values.

Table 5-10. Ingress Adapter Flow Control Schemes

| Scheme | IFC Flag | ISE Flag |
|--------------------------------|----------|----------|
| Flow Control Disabled | 0b | 0b |
| Dedicated Flow Control | 1b | 0b |
| Shared Flow Control | 0b | 1b |
| Restricted Shared Flow Control | 1b | 1b |

An Ingress Adapter shall always use the Dedicated Flow Control scheme for a Path that corresponds to HopID 0 (i.e. for Control Packets).

All other Paths shall be configurable during Path Setup. A configurable Path shall use the flow control scheme as determined by its *IFC Flag* and *ISE Flag*.

5.3.2.1.1 Buffer Allocation

An Ingress Lane Adapter shall have a buffer space that is used exclusively for incoming packets. A Connection Manager can divide the buffer space into three types of buffers:

- A Flow Control Disabled Buffer that is used for Paths with Flow Control Disabled.
- A set of Dedicated Buffers for Paths that use the Dedicated Flow Control scheme. There shall be one Dedicated Buffer for each Path that uses the Dedicated Flow Control scheme.
- A Shared Buffer that is used for all Paths that use either the Shared Flow Control scheme or the Restricted Shared Flow Control scheme.

For a Single-Lane Link, the *Total Buffers* field in the Adapter Configuration Space of the Lane 0 Adapter contains the total size of the buffer space available to the Lane 0 Adapter. For a Dual-Lane Link, the *Total Buffers* field in the Adapter Configuration Space of the Lane 0 Adapter contains the total size of the buffer space for both Lane Adapters in the USB4 Port.

**CONNECTION MANAGER NOTE**

A Connection Manager uses the parameters provided by the Buffer Allocation Request Operation (defined in Section 8.3.1.3.4) to allocate Buffers for each Path. A Connection Manager can allocate buffer space up to what it reads from the Total Buffers field.

Table 5-11 defines the buffer allocation parameters a Router reports through the Buffer Allocation Request Operation. See Appendix E for buffer space consideration and example.

Table 5-11. Buffer Allocation Parameters

| Name | Description | Egress Adapter | Requirements |
|-------------|--|----------------|---|
| baMaxUSB3 | The buffer allocation requested for a USB3 Path to achieve maximum target bandwidth. | USB3 | Shall be present if Router has a USB3 Adapter. |
| baMinDPaux | The minimum buffer allocation requested for a DP AUX Path. | Lane or DP | Shall be present if Router has a DP Adapter or multiple USB4 Ports. |
| baMinDPmain | The minimum buffer allocation requested for each DP Main-Link Path. | Lane or DP OUT | Shall be present if Router has a DP OUT Adapter or multiple USB4 Ports. |
| baMaxPCIe | The buffer allocation requested for a PCIe Path to achieve maximum target bandwidth. | PCIe | Shall be present if Router has a PCIe Adapter. |
| baMaxHI | The buffer allocation requested for the sum of all Host Interface Paths to achieve maximum target bandwidth. | HI | Shall be present if Router is a Host Router. |

**CONNECTION MANAGER NOTE**

A Connection Manager shall not setup a DP Path if it cannot allocate at least baMinDPmain buffers for the DP Main-Link Path and at least baMinDPaux buffers for the DP AUX Path. The Connection Manager can setup a USB3, PCIe and HI Paths with less than the maximum requested.

5.3.2.1.1.1 Flow Control Disabled Buffer

An Ingress Adapter shall store Transport Layer Packets arriving on Paths that use the Flow Control Disabled scheme in the Flow Control Disabled Buffer.

The Flow Control Disabled Buffer shall be the size set in the *Non Flow Controlled Buffers* field in Adapter Configuration Space.

5.3.2.1.1.2 Dedicated Flow Control Buffer

If a Path is defined with the Dedicated Flow Control Buffer scheme, an Ingress Adapter shall store any Transport Layer Packets arriving on that Path in the Dedicated Buffer for that Path.

A Dedicated Buffer shall be the size set in the *Path Credits Allocated* field in Path Configuration Space.

5.3.2.1.1.3 Shared Flow Control Buffer

If a Path is defined with the Shared Flow Control Buffer scheme, an Ingress Adapter shall store any Transport Layer Packets arriving on that Path in the Shared Buffer.

The size of the Shared Buffer shall be the size set in the *Link Credits Allocated* field in Adapter Configuration Space.

5.3.2.1.1.4 Restricted Shared Flow Control Buffer

If a Path is defined with the Restricted Shared Flow Control Buffer scheme, an Ingress Adapter shall store any packets arriving on that Path in the Shared Buffer. The Path shall not use more space in the Shared Buffer than is set forth in the *Path Credits Allocated* field in Path Configuration Space.

5.3.2.1.2 Credit Tracking

The following Rules apply for an Ingress Adapter tracking credits:

- If the *IFC Flag* field is set to 0b, credits shall not be tracked for the Path.
- If the *IFC Flag* field is set to 1b, credits shall be tracked for the Path.
- If the *ISE Flag* field is set to 0b, credits shall not be tracked for the Path in the Shared Buffer.
- If the *ISE Flag* field is set to 1b, credits shall be tracked for the Path in the Shared Buffers.
- For each Path with the IFC Flag set to 1b, the Ingress Adapter shall initially allocate the number of credits specified in the *Path Credits Allocated* field in the Path Configuration Space.
 - The Path corresponding to HopID 0 shall be provisioned with at least 2 initial credits.
- For the Shared Buffer, if the *Shared Buffering Capable* bit is set to 1b, the Ingress Adapter shall initially allocate the number of credits in the *Link Credits Allocated* field in the Adapter Configuration Space.

If an Ingress Adapter receives a packet on a flow controlled Path and the appropriate buffer (dedicated or shared) has no space for the packet, then the packet shall be discarded, the *Flow Control Error* bit in the Adapter Configuration Space shall be set to 1b, and the flow control state shall not be affected. If the *Flow Control Error Enable* bit in the Adapter Configuration Space is 1b, then a Notification Packet with Event Code = ERR_FC shall be sent upstream (see Section 6.5).

- Each Ingress Adapter shall track credits individually for its Shared Buffer and all of its Dedicated Buffers.
- When an Ingress Adapter drops a packet (e.g. due to a HEC error), it shall not account for the dropped packet in its credit tracking counters.
- Link Management Packets shall not cause credit counts to increment or decrement when received.

It is recommended that the following state variables be used to track credits for Paths with the *IFC Flag* field set to 1b. Credits are tracked individually for each Path. Other implementations are possible as long as credit tracking and synchronization are maintained:

- Path Credits Received (PCR)
 - Contains the total number (modulo 256) of Transport Layer Packets received on the Path since initialization.
 - Set to 0 on Path setup.

Note: Path 0 is setup upon exit from CLd state.
 - Incremented when a Transport Layer Packet is received on the Path: $PCR = (PCR + 1) \bmod 256$.
 - When a Path Credit Sync Packet is consumed, and after the PCA counter is updated (see below), the PCR counter is updated as follows: $PCR = \text{PCC value from the Path Credit Sync Packet}$.
- Path Credits Allocated (PCA)
 - Contains the total number (modulo 256) of credits allocated to Path since initialization.
 - On Path setup, set to the value in the *Path Credits Allocated* field of the Path Configuration space.

Note: Path 0 is setup upon exit from CLd state.
 - Included in the *Credits* field of a Credit Grant Record for the Path each time a Credit Grant Packet is sent.
 - Incremented when a Transport Layer Packet is dequeued from the Adapter buffers: $PCA = (PCA + 1) \bmod 256$.
 - Each time a Path Credit Sync Packet is consumed, the PCA counter is updated as follows:
 - $PCA = (PCA + (\text{PCC value from the Path Credit Sync Packet} - PCR + 256)) \bmod 256$.

It is recommended that the following state variables be used to track credits for Paths with the *ISE Flag* field set to 1b. Credits are shared between Paths and tracked per Adapter. Other implementations are possible as long as credit tracking and synchronization are maintained:

- Shared Credits Received (SCR)
 - Contains the total number (modulo 256) of Transport Layer Packets received on all Paths with the *ISE Flag* field set to 1b.
 - Set to 0 when the first shared Path is setup.
 - Only valid when the *Shared Buffering Capable* bit is set to 1b.
 - Incremented when a Transport Layer Packet is received on a Path with *ISE Flag* field set to 1b: $SCR = (SCR + 1) \bmod 256$.
 - When a Shared Credits Sync Packet is consumed, and after the SCA counter is updated (see below), the SCR counter is updated as follows: $SCR = \text{SCC value from the Shared Credit Sync Packet}$.
- Shared Credits Allocated (SCA)
 - Contains the total number (modulo 256) of credits allocated to the Shared Buffers since initialization.
 - Only valid when the *Shared Buffering Capable* bit is set to 1b.
 - Set to the value in the *Link Credits Allocated* field when either of the following occur:
 - The value in the *Link Credits Allocated* field changes.
 - The first Path that uses shared flow control is setup.

- Included in the *Credits* field of a Credit Grant Record for the Shared Buffer when a Credit Grant Packet is sent.
- Incremented when a Transport Layer Packet is dequeued from the Shared Buffer and the *ISE Flag* field for the Path is set: $SCA = (SCA + 1) \bmod 256$.
- When a Shared Credits Sync Packet is consumed, the SCA counter is updated as follows:
 - $SCA = (SCA + (SCC \text{ value from the Shared Credit Sync Packet} - SCR + 256)) \bmod 256$.



CONNECTION MANAGER NOTE

A Connection Manager shall not change the Link Credits Allocated field in Adapter Configuration Space if there is an enabled Path with ISE Flag field set to 1b.

5.3.2.1.3 Credit Grant Packets

Credit information is transferred from an Ingress Adapter to an Egress Adapter using Credit Grant Packets. An Ingress Adapter shall send Credit Grant Packets according to the following rules:

- If the *IFC Flag* field in Path Configuration Space is set to 0b, Credit Grant Packets shall not be sent for the Path.
- If the *IFC Flag* field in Path Configuration Space is set to 1b for a Path, an Ingress Adapter shall send Credit Grant Packets for that Path after Transport Layer Packets are dequeued. The policy for how frequently to send Credit Grant Packets after Transport Layer Packets are dequeued is implementation specific.
- If the *IFC Flag* field in Path Configuration Space is set to 1b for the Path, an Ingress Adapter shall send a Credit Grant Packet with a Credit Grant Record for the Path when the Path is first enabled.
- If the *ISE Flag* field in Path Configuration Space is set to 0b, the Path shall not affect Credit Grant Packets sent for the Shared Flow Control Buffer.
- If an Ingress Adapter has a Path with the *ISE Flag* field in Path Configuration Space set to 1b, the Ingress Adapter shall send Credit Grant Packets for its Shared Flow Control Buffer after Transport Layer Packets are dequeued. The policy for how frequently to send Credit Grant Packets after Transport Layer Packets are dequeued is implementation specific.
- If the *ISE Flag* field in Path Configuration Space is set to 1b for the Path, an Ingress Adapter shall send a Credit Grant Packet with a Credit Grant Record for the Shared Buffer when the Path is first enabled.
- When its Link is in the Active state, an Ingress Adapter sends Credit Grant Packets for its Shared Buffer and all of its Dedicated Buffers. Credit Grant Packets shall be sent at least every tCredits.
- When its Link is in a low-power state, an Ingress Adapter sends Credit Grant Packets as defined in Table 5-8.

Note: Sending a Credit Grant Packet causes the Link to transition to the Active State

- When a Link first becomes Active, an Ingress Adapter that is a Lane 0 Adapter shall send a Credit Grant Packet with a Credit Grant Record for HopID 0.

**IMPLEMENTATION NOTE**

If a Router does not send a Credit Grant Packet immediately after a packet is dequeued, it is recommended that the Ingress Adapter buffers are sized to compensate for the delay.

When an Egress Adapter receives a Credit Grant Packet, it shall process each Credit Grant Record in the order received as follows:

- The Egress Adapter shall verify the *ECC* field value in the Credit Grant Record.
 - The Egress Adapter shall correct any single-bit errors. After correcting an error, the Egress Adapter shall continue on as if the error had never occurred.
 - If an uncorrectable error is detected, the Credit Grant Record shall be dropped, and the *ECC Error* field in the Adapter Configuration Registers shall be incremented.
- If the HopID in a Credit Grant Record does not match an enabled Path in the Egress Adapter, the Credit Grant Record shall be dropped and no further actions shall be taken.
- Section 5.3.2.2.1 describes how the information in a Credit Grant Record is used.

5.3.2.2 Egress Adapter

Table 5-12 defines the flow control schemes for an Egress Adapter. The *EFC Flag* and *ESE Flag* in Path Configuration Space determine which flow control scheme is used for a Path.

Table 5-12. Egress Adapter Flow Control Schemes

| Scheme | EFC Flag | ESE Flag |
|--------------------------------|----------|----------|
| Flow Control Disabled | 0b | 0b |
| Dedicated Flow Control | 1b | 0b |
| Shared Flow Control | 0b | 1b |
| Restricted Shared Flow Control | 1b | 1b |

The Path corresponding to HopID 0 shall always use the Dedicated Flow Control scheme. For an Adapter that is not a Host Interface Adapter, the Paths that correspond to Paths 1 through 7 shall always use the Flow Control Disabled scheme. All other Paths are configurable and shall use the flow control scheme that corresponds to the *EFC Flag* and *ESE Flag* in Path Configuration Space.

5.3.2.2.1 Credit Tracking

It is recommended that the following state variables be used to track credits for a Path with the *EFC Flag* field set to 1b. Credits are tracked individually for each Path. Other implementations are possible as long as credit tracking and synchronization are maintained:

- Path Credits Consumed (PCC)
 - Contains a count (modulo 256) of the total number of credits consumed by Transport Layer Packet transmissions on the Path since the Path was initialized.
 - Set to 0 on Path setup.

Note: Path 0 is setup upon exit from CLd state.
 - Incremented each time a Transport Layer Packet is transmitted over the Path:
PCC = (PCC + 1) mod 256.

- Path Credit Limit (PCL)
 - Contains the most recent number of credits advertised by the Ingress Adapter for the Path.
 - Set to 0 on Path setup.

Note: Path 0 is setup upon exit from CLd state.

- Updated each time a Credit Grant Record is received for the Path by overwriting with the value contained in the *Credits* field.

It is recommended that the following state variables be used to track credits for Paths with the *ESE Flag* field set to 1b. Credits are shared between Paths and tracked per Adapter. Other implementations are possible as long as credit tracking and synchronization are maintained:

- Shared Credits Consumed (SCC)
 - Contains a count (modulo 256) of the total number of credits consumed by Transport Layer Packet transmissions on Paths with the *ESE Flag* field enabled.
 - Set to 0 when the first shared Path is setup.
 - Incremented each time a Transport Layer Packet is transmitted over a Path that has the *ESE Flag* field enabled: $SCC = (SCC + 1) \bmod 256$.
- Shared Credit Limit (SCL)
 - Contains the most recent number of credits advertised by the Ingress Adapter for the shared buffers.
 - Set to 0 when the first shared Path is setup.
 - Updated each time a Credit Grant Record is received with a record for the shared buffer by overwriting with the value contained in the *Credits* field.

5.3.2.2.2 Transmission Rules

If a Path uses the Flow Control Disabled scheme ($EFC = 0b$ and $ESE = 0b$), then the Egress Adapter shall not require any credits to transmit a Transport Layer Packet on that Path.

If a Path uses the Dedicated Flow Control scheme ($EFC = 1b$ and $ESE = 0b$), then the Egress Adapter shall require the following condition to be true before transmitting a Transport Layer packet on the Path:

- $[(PCL - PCC > 0) \text{ and } (PCL - PCC < 128)] \text{ or } [(PCL - PCC < 0) \text{ and } (PCC - PCL > 128)]$.

If a Path uses the Shared Flow Control scheme ($EFC = 0b$ and $ESE = 1b$), then the Egress Adapter shall require the following condition to be true before transmitting a Transport Layer packet on the Path:

- $[(SCL - SCC > 0) \text{ and } (SCL - SCC < 128)] \text{ or } [(SCL - SCC < 0) \text{ and } (SCC - SCL > 128)]$.

If a Path uses the Restricted Shared Flow Control scheme ($EFC = 1b$ and $ESE = 1b$), then the Egress Adapter shall require both of the following conditions be true before transmitting a Transport Layer packet on the Path:

- $[(PCL - PCC > 0) \text{ and } (PCL - PCC < 128)] \text{ or } [(PCL - PCC < 0) \text{ and } (PCC - PCL > 128)]$.
- $[(SCL - SCC > 0) \text{ and } (SCL - SCC < 128)] \text{ or } [(SCL - SCC < 0) \text{ and } (SCC - SCL > 128)]$.

5.3.2.3 Credit Counter Synchronization

If a flow controlled Transport Layer Packet is dropped or otherwise lost for any reason, the Egress Adapter credit counters lose synchronization with the Ingress Adapter. In order to re-establish synchronization, an Egress Adapter sends Credit Sync Packets periodically to the Ingress Adapter.

There are two types of Credit Sync Packets: Path Credit Sync Packets, which are defined in Section 5.1.3.3.3, and Shared Buffers Credit Sync Packets, which are defined in Section 5.1.3.3.4.

An Egress Adapter shall send Credit Sync Packets according to the following rules:

- An Egress Adapter shall send a Path Credit Sync Packet every tSync for a Path with the *Egress Flow Control (EFC) Flag* field set to 1b and the *Valid* bit set to 1b.
- If the *Egress Shared Buffering Enable (ESE) Flag* field in Path Configuration Space is set to 1b for at least one enabled Path, an Egress Adapter shall send a Shared Buffers Credit Sync Packet every tSync.
- The credit count in the *PCC* field of a Path Credit Sync Packet shall be based on the number of flow controlled Transport Layer Packets sent on the Path prior to the Path Credit Sync Packet and shall not include flow controlled Transport Layer Packets which have not yet been sent.
- The credit count in the *SCC* field of a Shared Credit Sync Packet shall be based on the number of Transport Layer Packets sent on all Paths that use the Shared Buffer prior to the Shared Credit Sync Packet and shall not include Transport Layer Packets which have not yet been sent.
- An Egress Adapter shall not send Path Credit Sync Packets for a Path that uses the Flow Control Disable scheme.
- An Egress Adapter shall not send a Credit Sync Packet while in a Low Power state.

Note: When an Egress Adapter transitions back to the *CL0* state after recovering from an error, it is recommended that the Egress Adapter send a Credit Sync Packet for each enabled Path that uses flow control.

When an Ingress Adapter receives a Credit Sync Packet, it shall verify the *ECC* field value in the Credit Sync Packet payload as follows:

- The Ingress Adapter shall correct any single-bit errors. After correcting an error, the Ingress Adapter shall continue on as if the error had never occurred.
- If an uncorrectable error is detected, the Credit Sync Packet shall be dropped, and the *ECC Error* field in the Adapter Configuration Registers shall be incremented.

5.3.3 Bandwidth Arbitration and Priority

Bandwidth arbitration and prioritization of traffic in a Domain is done in a distributed manner. Each Egress Adapter contains a configurable traffic manager that does bandwidth arbitration and prioritization for all Paths through that Adapter. Configuration of the traffic manager is done as part of Path setup. A Router shall enable bandwidth arbitration for a given Path when the *Valid* bit in the Path Configuration Space is set to 1b.




IMPLEMENTATION NOTE

The Priority field in the Path Configuration Space is global to a Router. It can be used to determine the priority between queues in an Ingress Adapter, between different Ingress Adapters, between different Egress Adapters, and in each Egress Adapter.

5.3.3.1 Scheduling

The traffic manager for an Egress Adapter shall use the 3-layer scheduling scheme described in this section to schedule outgoing packets. This scheme is summarized in Figure 5-11.

Figure 5-11. Egress Adapter Scheduler

5.3.3.1.1 Path Schedulers

The first layer of the traffic manager consists of a number of schedulers (called Path Schedulers) that prioritize traffic for a given Priority Group. A Priority Group is a set of Paths that share the same scheduling priority value. A Path is assigned a scheduling priority value by writing to the *Priority* field in the Path Configuration Space.

There shall be one Path Scheduler for each Priority Group. A Path Scheduler operates according to the following rules:

- A weighted round-robin (WRR) scheduling scheme shall be implemented among the Paths that share the same Priority Group.
- The Path Scheduler shall support weights in the range of 1-255. The weight assigned to a Path shall be determined by the *Weight* field in the Path Configuration Space.
- When the *Weight* field changes for an enabled Path, the Path Scheduler shall use the new weight.
- When a Path is assigned a weight value of X, the Path Scheduler shall schedule X packets from that Path in one round, where a round refers to one complete iteration over all the Paths in the Priority Group. If less than X packets are available for arbitration, the scheduler shall schedule the number of available packets.

HopID 0 traffic shall be assigned to Priority Group 0. No other traffic shall be assigned to Priority Group 0.

5.3.3.1.2 Priority Group Scheduler

The second layer of the traffic manager consists of a Priority Group Scheduler that prioritizes traffic from each of the Path Schedulers. The Priority Group Scheduler shall employ a strict priority scheme between 8 Priority Groups, where Priority Group 0 has the highest priority and Priority Group 7 has the lowest priority. A request for scheduling by a higher Priority Group always takes precedence over requests of lower Priority Groups.

5.3.3.1.3 Link Scheduler

The Third Layer of the traffic manager consists of a Link Scheduler that multiplexes traffic from the Priority Group Scheduler with other packets generated within the Transport Layer. The Link Scheduler shall schedule traffic according to a strict priority scheme where the following priorities (from highest to lowest) are observed:

- Flow Control Packets.
- Time Sync Packets.
- Packets from the Priority Group Scheduler.

5.3.4 Packet Forwarding Delay Jitter

When a Router routes a Tunneled Packet along a Path, from an Ingress Adapter to an Egress Adapter, it generates Packet Forwarding Delay (PFD) Jitter. The PFD is the delay between the time when the last bit of a Tunneled Packet is received at the Ingress Adapter to the time when the first bit of the Tunneled Packet is sent by the Egress Adapter. The PFD Jitter is the difference between the maximum and the minimum PFD between two Adapters along a Path in a Router. The PFD Jitter in a Router shall be no more than $t_{\text{TunneledPacketJitter}}$.

Note: PFD Jitter does not include the additional delay due to another Transport Layer Packet or Ordered Set having priority in being transmitted. It also does not include the delay in transmission due to insufficient flow control credits.

5.4 Path Tear-down

A Connection Manager tears down a Path starting from the Destination Adapter and ending with the Source Adapter. A Connection Manager initiates Path tear-down by setting the *Valid* bit in the Path Configuration Space to 0b. After the *Valid* bit in a Path Configuration Space changes from 1b to 0b, a Router shall respond to the Write Request that sets the *Valid* bit to 0b and then tear down the Path at its Egress Adapter and Ingress Adapter. Section 5.4.1 describes the flow for tearing down a Path in the Egress Adapter and Section 5.4.2 describes the flow for tearing down a Path in the Ingress Adapter.

Note: The Ingress Adapter and the Egress Adapter mentioned in this section refer to Adapters in the same Router.

Note: The tear down flows described in this section are applicable for all type of Adapters. The rules that involve Credits are applicable for Lane Adapters only.

5.4.1 Egress Adapter

For the Egress Adapter of the Path being torn down, a Router shall perform the following steps in the order listed:

1. If the *ESE Flag* field in the Path Configuration Space is set to 1b for the Path, then the Router shall send a Shared Buffers Credit Sync Packet to the Link Partner.
2. While the *Valid* bit in Path Configuration Space is 0b:
 - The Router shall not send any Path Credit Sync Packets for the Path.
 - The Router shall ignore any Path credit updates for the Path received on the Egress Adapter.
3. The Router may transmit packets on the Path when the *Pending Packets* bit in the Path Configuration Space is set to 1b and for t_{Teardown} after the Router sets the *Pending Packets* bit to 0b.
4. After t_{Teardown} time, the Router shall discard any remaining packets for the Path.
5. The Router shall block the transmission of any packets on the Path after t_{Teardown} has elapsed since it set the *Pending Packets* bit to 0b and until the *Valid* bit in the Path Configuration Space is set again to 1b.

**CONNECTION MANAGER NOTE**

A Connection Manager shall wait at least t_{Teardown} after a Router sets the Pending Packets bit to 0b before setting the Valid bit for the Path to 1b again.

5.4.2 Ingress Adapter

For the Ingress Adapter of the Path being torn down, a Router shall perform the following steps in the order listed:

1. The Router shall drop any packets received on the Path after the *Valid* bit was set to 0b.
2. The Router shall dequeue all packets for the Path that are queued in the flow control buffers. The Router may transmit a dequeued packet as a whole on its designated Egress Adapter or it may discard the packet. A Router shall not transmit a partial packet.
 - If the *ISE Flag* field in Path Configuration Space is set to 1b, the Ingress Adapter shall continue to increment the SCA state variable and send Credit Grant Packets. The SCA variable shall increment each time a packet is dequeued, regardless of whether the packet was discarded or transmitted.
3. While the *Valid* bit in Path Configuration Space is 0b:
 - The Router shall discard any Path Credit Sync Packets received for the Path.
 - The Router shall stop sending Path credits updates for the Path.


5.5 Timing Parameters

Table 5-13 lists the timing parameters for the Transport Layer.

Table 5-13. Transport Layer Timing Parameters

| Parameter | Description | Min | Max | Units |
|-----------------------|--|-----|-----|-------|
| tSync | The time interval at which Credit Sync Packets are periodically sent for a Path. | -- | 64 | ms |
| tCredits | The time interval at which an Ingress Adapter sends Credit Grant Records for a Path. | -- | 200 | μs |
| tTeardown | Duration during which packet transmission is permitted on a removed Path after the <i>Pending Packets</i> bit was set to 0b. | -- | 2 | μs |
| tTunneledPacketJitter | The maximum PFD jitter of Tunneled Packets in a Router. | -- | 100 | ns |

6 Configuration Layer



6.1 Domain Topology


A Spanning Tree contains up to six levels (depths 0 through 5). The Host Router is at the top of the Spanning Tree (depth = 0). Device Routers are connected at depths 1 through 5 (inclusive). A Connection Manager accesses a Domain through the Host Router of that Domain.

6.2 Router Addressing

A Connection Manager assigns each Router in its Domain a unique topological address called a TopologyID. The TopologyID represents the position of the Router within the Domain's Spanning Tree.

The TopologyID is a sequence of seven Adapter numbers representing an Adapter in the Downstream Facing Ports at each level of the Spanning Tree between the Host Router and the Router. The TopologyID of a Host Router is always 0,0,0,0,0,0. The TopologyID for a Device Router at depth X (where X is from 1 to 5) is denoted as 0,...,0,P_{X-1},P_{X-2},...,P₀ where P_n is the Adapter Number of the Adapter in the Downstream Facing Port at level n. Figure 6-1 shows an example of a Spanning Tree along with the assignment of TopologyID values to each Router.

Figure 6-1. Example of TopologyID Assignment



6.3 Router States

Figure 6-2 describes a Host Router state machine.

Figure 6-2. Host Router State Machine





Figure 6-3 describes a Device Router state machine.

Figure 6-3. Device Router State Machine



6.3.1 Uninitialized Unplugged State

A Device Router enters this state upon completion of the following events:

- Power on.
- Upstream Facing Port is disconnected.

Upon entering this state, a Device Router does the following:

- Discards all Transport Layer Packets.
- Removes all Paths.
- Restores all Configuration Spaces to their default values.

Note: All Lane Adapters are in CLd state while a Router is in the Uninitialized Unplugged state.

A Device Router exits this state after detecting a Router on its Upstream Facing Port (see Section 4.1.2.2).

6.3.2 Uninitialized Plugged State

A Host Router enters this state:

- After Power On.
- From the Sleep state on a wake event (see Section 4.5).

A Device Router enters this state:

- From the Uninitialized Unplugged state when it detects a Router on its Upstream Facing Port (see Section 4.1.2.2).
- From the Sleep state on a wake event (see Section 4.5).

While in this state, a Router performs Lane Initialization on all Connected Adapters.

A Host Router exits this state when the *TopologyID Valid* bit is set to 1b.

A Device Router exits this state upon the following events:

- *TopologyID Valid* bit is set to 1b.
- The Router is hot unplugged.

When a Router is in the Uninitialized Plugged state, the values in the *Depth* and *TopologyID* fields of the Router Configuration Space are not valid.

6.3.3 Sleep State

Section 4.5 defines how a Router enters and exits Sleep and how it behaves while in Sleep state.

6.3.4 Enumerated State

A Router enters this state from the Uninitialized Plugged state when its *TopologyID Valid* bit is set to 1b.

A Router in this state supports tunneling of Routed Protocols through the USB4 Fabric.

A Router exits this state when its Upstream Facing Port is disconnected.

See Section 6.7 for more details about Router enumeration.

6.4 Control Packet Protocol**6.4.1 Control Adapter**


A Router shall support an internal Control Adapter that is used solely for transmitting and receiving Control Packets to and from the Transport Layer. Unlike the non-Control Adapters, the Control Adapter does not connect directly to a Link and thus does not have a Physical Layer associated with it.

6.4.2 Control Packets**6.4.2.1 Bit/Byte Conventions**

Control Packets are defined using the bit and byte conventions defined in Section 5.1.1.

6.4.2.2 Format


A Control Packet is a Transport Layer Packet. A Control Packet has the generic structure shown in Figure 6-4. The first DW is the Transport Layer Packet header. The remaining DWs are payload and are defined in Table 6-1.

Figure 6-4 Control Packet Format**Table 6-1. Control Packet Payload**

| DW | Field | Description |
|----|--------------------------|--|
| 1 | <i>Route String High</i> | <p>Route String High – All Control Packets shall include a Route String. The format of the Route String is shown in Figure 6-5.</p> <p>For Control Packets that originate from the Connection Manager and target a Router:</p> <p>TopologyID [55:32] – Shall contain the high 24 bits of the TopologyID of the target Router.</p> <p>Rsvd [62:56] – Shall be set to 0.</p> <p>CM [63] – Shall be set to 0b.</p> <p>For Control Packets that originate from a Router and target the Connection Manager:</p> <p>TopologyID [55:32] – Shall contain the high 24 bits of the TopologyID of the Router that originates the Control Packet.</p> <p>Rsvd [62:56] – Shall be set to 0.</p> <p>CM [63] – Shall be set to 1b.</p> |

| DW | Field | Description |
|------------------|-------------------------|---|
| 2 | <i>Route String Low</i> | <p>Route String Low – All Control Packets shall include a Route String. The format of the Route String is shown in Figure 6-5.</p> <p>For Control Packets that originate from the Connection Manager and target a Router:</p> <p>TopologyID [31:0] – Shall contain the low 32 bits of the TopologyID of the target Router.</p> <p>For Control Packets that originate from a Router and target the Connection Manager:</p> <p>TopologyID [31:0] – Shall contain the low 32 bits of the TopologyID of the Router that originates the Control Packet.</p> |
| (n+2):3 (if n>0) | <i>Control Data</i> | <p>Control Data – n DWs (where n >=0) carrying additional information specific to each type of Control Packet.</p> |
| n+3 | <i>CRC</i> | <p>CRC – A 32-bit Cyclic Redundancy Code that protects the Route String and Control Data. The CRC does not cover the packet header.</p> <p>The CRC shall be calculated in increasing DW order, starting with the <i>Route String High</i> DW. Within each DW, CRC shall be calculated from bit[31] to bit[0]. The following CRC shall be used:</p> <ul style="list-style-type: none"> • Width: 32 • Poly: 1EDC6F41h • Init: FFFFFFFFh • RefIn: True • RefOut: True • XorOut: FFFFFFFFh <p>See Appendix A for examples of CRC calculations for several Control Packets.</p> |

Figure 6-5. Route String Format



6.4.2.3 Read Request

A Connection Manager uses Read Requests to read from a Configuration Space. A Read Request shall have the format shown in Figure 6-6. The fields in a Read Request shall be as defined in Table 6-2.


Table 6-2. Content of a Read Request

| DW | Bits | Field | Description |
|----|------|--------------------------|--|
| 0 | 31:0 | <i>Header</i> | <p>Header – Transport Layer Packet header</p> <p>HEC – See Section 5.1.2.1.1</p> <p>Length – 10h</p> <p>HopID – 000h</p> <p>SuppID – 0b</p> <p>PDF – 1h</p> |
| 1 | 31:0 | <i>Route String High</i> | Route String High – See Section 6.4.2.2. |
| 2 | 31:0 | <i>Route String Low</i> | Route String Low – See Section 6.4.2.2. |

Table 6-3. Content of a Read Response

| DW | Bits | Field | Description |
|-----------------|-------|---------------------------------|---|
| 0 | 31:0 | <i>Header</i> | Header – Transport Layer Packet header. HEC – See Section 5.1.2.1.1 Length – See Section 5.1.2 HopID – 000h SuppID – 0b PDF – 1h |
| 1 | 31:0 | <i>Route String High</i> | Route String High – See Section 6.4.2.2. |
| 2 | 31:0 | <i>Route String Low</i> | Route String Low – See Section 6.4.2.2. |
| 3 | 12:0 | <i>Address</i> | Address – DW address of the first configuration register to be read. This field shall contain the same value as the associated Read Request. |
| | 18:13 | <i>Read Size</i> | Read Size – Number of Doublewords that were requested to be read. This field shall contain the same value as the associated Read Request. |
| | 24:19 | <i>Adapter Num</i> | Adapter Num – For a Response to a Request that targets Router Configuration Space, this field shall contain the Adapter Number on which the associated Read Request arrived. For a Response to a Request that targets other Configuration Spaces, this field shall contain the <i>Adapter Num</i> value in the associated Read Request. |
| | 26:25 | <i>Configuration Space (CS)</i> | Configuration Space – Identifies the target Configuration Space. This field shall contain the same value as the associated Read Request. |
| | 28:27 | <i>Sequence Number (SN)</i> | Sequence Number – Used to associate a Response Packet with a Request Packet. This field shall contain the same value as the associated Read Request. |
| | 31:29 | <i>Rsvd</i> | Reserved. |
| (3+Read Size):4 | 31:0 | <i>Read Data</i> | Read Data – Data read from the target Configuration Space. The size of this field shall match the number of DWs in the <i>Read Size</i> field. Data shall be structured in increasing address order with bit 0 of each DW containing bit 0 of the corresponding configuration register. |
| 4+Read Size | 31:0 | <i>CRC</i> | CRC – See Table 6-1. |

Figure 6-7. Read Response



6.4.2.5 Write Request


A Connection Manager uses Write Requests to write to a Configuration Space. A Write Request shall have the format shown in Figure 6-8. The fields in a Write Request shall be as defined in Table 6-4.

Table 6-4. Content of a Write Request

| DW | Bits | Field | Description |
|----|-------|---------------------------------|--|
| 0 | 31:0 | <i>Header</i> | Header – Transport Layer Packet header. HEC – See Section 5.1.2.1.1 Length – See Section 5.1.2 HopID – 000h SuppID – 0b PDF – 2h |
| 1 | 31:0 | <i>Route String High</i> | Route String High – See Section 6.4.2.2. |
| 2 | 31:0 | <i>Route String Low</i> | Route String Low – See Section 6.4.2.2. |
| 3 | 12:0 | <i>Address</i> | Address – DW address of the first configuration register to be written. |
| | 18:13 | <i>Write Size</i> | Write Size – Number of Doublewords that were requested to be written. The <i>Write Size</i> field shall be greater than 0 and less than or equal to 60. |
| | 24:19 | <i>Adapter Num</i> | Adapter Num – Adapter number of the Adapter whose Configuration Space is being written or 0 if Router Configuration Space is being written. |
| | 26:25 | <i>Configuration Space (CS)</i> | Configuration Space – Identifies the target Configuration Space. 00b – Path Configuration Space 01b – Adapter Configuration Space 10b – Router Configuration Space 11b – Counters Configuration Space |

| DW | Bits | Field | Description |
|------------------|-------|-----------------------------|--|
| 3 | 28:27 | <i>Sequence Number (SN)</i> | Sequence Number – Used to associate a Response Packet with a Request Packet. The sequence number is assigned by the Connection Manager and is returned in the Response Packet by the destination Router. |
| | 31:29 | <i>Rsvd</i> | Reserved. |
| (3+Write Size):4 | 31:0 | <i>Write Data</i> | Write Data – Data to be written to Configuration Space. The size of this field shall match the number of DWs in the <i>Write Size</i> field. Data shall be structured in increasing address order with bit 0 of each DW containing bit 0 of the corresponding configuration register. |
| 4+Write Size | 31:0 | <i>CRC</i> | CRC – See Table 6-1. |

Figure 6-8. Write Request



CONNECTION MANAGER NOTE

A Connection Manager shall set the Adapter Num field in a Write Request to 0b when the Configuration Space field is 10b.


6.4.2.6 Write Response

A Router uses a Write Response to respond to a Write Request. A Write Response shall have the format shown in Figure 6-9. The fields in a Write Response shall be as defined in Table 6-5.

Table 6-5. Content of a Write Response

| DW | Bits | Field | Description |
|----|-------|---------------------------------|--|
| 0 | 31:0 | <i>Header</i> | Header – Transport Layer Packet header. HEC – See Section 5.1.2.1.1 Length – 10h HopID – 000h SupplD – 0b PDF – 2h |
| 1 | 31:0 | <i>Route String High</i> | Route String High – See Section 6.4.2.2. |
| 2 | 31:0 | <i>Route String Low</i> | Route String Low – See Section 6.4.2.2. |
| 3 | 12:0 | <i>Address</i> | Address – DW address of the first configuration register the associated Request is targeting. This field shall contain the same value as the associated Write Request. |
| | 18:13 | <i>Write Size</i> | Write Size – Number of Doublewords that were requested starting from the Address field value. This field shall contain the same value as the associated Write Request. |
| | 24:19 | <i>Adapter Num</i> | Adapter Num – This field shall contain the Adapter Num value in the associated Write Request. |
| | 26:25 | <i>Configuration Space (CS)</i> | Configuration Space – Identifies the target Configuration Space. This field shall contain the same value as the associated Write Request. |
| | 28:27 | <i>Sequence Number (SN)</i> | Sequence Number – Used to associate a Response Packet with a Request Packet. This field shall contain the same value as the associated Write Request. |
| | 31:29 | <i>Rsvd</i> | Reserved. |
| 4 | 31:0 | <i>CRC</i> | CRC – See Table 6-1. |

Figure 6-9. Write Response



6.4.2.7 Notification Packet


A Notification Packet shall have the format shown in Figure 6-10 and the fields defined in Table 6-6.

A Notification Packet carrying a Hot Plug Acknowledgment (see Table 6-11) is called a Hot Plug Acknowledgment Packet.

Table 6-6. Content of a Notification Packet

| DW | Bits | Field | Description |
|----|-------|--------------------------|--|
| 0 | 31:0 | <i>Header</i> | Header – Transport Layer Packet header. HEC – See Section 5.1.2.1.1 Length – 10h HopID – 000h SuppID – 0b PDF – 3h |
| 1 | 31:0 | <i>Route String High</i> | Route String High – See Section 6.4.2.2. |
| 2 | 31:0 | <i>Route String Low</i> | Route String Low – See Section 6.4.2.2. |
| 3 | 7:0 | <i>Event Code</i> | Event Code – Identifies the type of event. See Table 6-11. |
| | 13:8 | <i>Adapter Num</i> | Adapter Num – Adapter number where the event was detected. |
| | 29:14 | <i>Rsvd</i> | Reserved. |
| | 31:30 | <i>PG</i> | PG – Differentiates Hot Plug and Hot Unplug Events in a Hot Plug Acknowledgment Packet. 00b – Notification Packet is not a Hot Plug Acknowledgment 01b – Rsvd 10b – Hot Plug Event 11b – Hot Unplug Event |
| 4 | 31:0 | <i>CRC</i> | CRC – See Table 6-1. |

Figure 6-10. Notification Packet



6.4.2.8 Notification Acknowledgement Packet


A Notification Acknowledgment Packet shall have the format shown in Figure 6-11 and the fields defined in Table 6-7.

Table 6-7. Content of a Notification Acknowledgement Packet

| DW | Bits | Field | Description |
|----|------|---------------|--|
| 0 | 31:0 | <i>Header</i> | Header – Transport Layer Packet header. HEC – See Section 5.1.2.1.1 Length – 0Ch HopID – 000h SuppID – 0b PDF – 4h |

| DW | Bits | Field | Description |
|----|------|--------------------------|--|
| 1 | 31:0 | <i>Route String High</i> | Route String High – See Section 6.4.2.2 |
| 2 | 31:0 | <i>Route String Low</i> | Route String Low – See Section 6.4.2.2 |
| 3 | 31:0 | <i>CRC</i> | CRC – See Table 6-1. |

Figure 6-11. Notification Acknowledgement Packet




6.4.2.9 Hot Plug Event Packet

This packet is used by a Router to notify a Connection Manager that a Hot Plug or Hot Unplug Event has occurred. A Hot Plug Event Packet shall have the structure defined in Table 6-8 and Figure 6-12.

Table 6-8. Content of a Hot Plug Event Packet


| DW | Bits | Field | Description |
|----|------|--------------------------|--|
| 0 | 31:0 | <i>Header</i> | Header – Transport Layer Packet header. HEC – See Section 5.1.2.1.1 Length – 10h HopID – 000h SupplD – 0b PDF – 5h |
| 1 | 31:0 | <i>Route String High</i> | Route String High – See Section 6.4.2.2. |
| 2 | 31:0 | <i>Route String Low</i> | Route String Low – See Section 6.4.2.2. |
| 3 | 5:0 | <i>Adapter Num</i> | Adapter Num – Adapter number that experienced the Hot Plug or Hot Unplug Event. |
| | 30:6 | <i>Rsvd</i> | Reserved. |
| | 31 | <i>UPG</i> | UPG – Shall be set to 0b for a Hot Plug Event or 1b for a Hot Unplug Event. |
| 4 | 31:0 | <i>CRC</i> | CRC – See Table 6-1. |

Figure 6-12. Hot Plug Event Packet**6.4.2.10 Inter-Domain Request**

An Inter-Domain Request is used for Inter-Domain communication between two Connection Managers. An Inter-Domain Request shall have the format shown in Figure 6-13 and fields as defined in Table 6-9.

Table 6-9. Content of an Inter-Domain Request


| DW | Bits | Field | Description |
|---------|------|--------------------------|--|
| 0 | 31:0 | <i>Header</i> | Header – Transport Layer Packet header. HEC – See Section 5.1.2.1.1 Length – See Section 5.1.2 HopID – 000h SuppID – 0b PDF – 6h |
| 1 | 31:0 | <i>Route String High</i> | Route String High – See Section 6.4.2.2 and Figure 6-15. The Route String shall include the TopologyID of the Inter-Domain Router (i.e. the Router in the Domain that interfaces to the Inter-Domain Link). The <i>CM</i> bit shall be set to 0b. |
| 2 | 31:0 | <i>Route String Low</i> | Route String Low – See Section 6.4.2.2 and Figure 6-15. The Route String shall include the TopologyID of the Inter-Domain Router (i.e. the Router in the Domain that interfaces to the Inter-Domain Link). |
| (N+3):3 | 31:0 | <i>InterDomain Data</i> | InterDomain Data – N number of DWs to be transferred to the other Connection Manager. The contents of this field are outside the scope of this specification. See the USB4 Inter-Domain Specification for definition. |
| N+4 | 31:0 | <i>CRC</i> | CRC – See Table 6-1. |

Figure 6-13. Inter-Domain Request**6.4.2.11 Inter-Domain Response**

An Inter-Domain Response shall have the format shown in Figure 6-14 and fields as defined in Table 6-10.

Table 6-10. Content of an Inter-Domain Response

| DW | Bits | Field | Description |
|---------|------|--------------------------|--|
| 0 | 31:0 | <i>Header</i> | Header – Transport Layer Packet header. HEC – See Section 5.1.2.1.1 Length – See Section 5.1.2 HopID – 000h SuppID – 0b PDF – 7h |
| 1 | 31:0 | <i>Route String High</i> | Route String High – See Section 6.4.2.2 and Figure 6-15. The Route String shall include the TopologyID of the Inter-Domain Router (i.e. the Router in the Domain that interfaces to the Inter-Domain Link). The <i>CM</i> bit shall be set to 0b. |
| 2 | 31:0 | <i>Route String Low</i> | Route String Low – See Section 6.4.2.2 and Figure 6-15. The Route String shall include the TopologyID of the Inter-Domain Router (i.e. the Router in the Domain that interfaces to the Inter-Domain Link). |
| (N+3):3 | 31:0 | <i>InterDomain Data</i> | InterDomain Data – N number of DWs to be transferred to the other Connection Manager. The contents of this field are outside the scope of this specification. See the USB4 Inter-Domain Specification for definition. |
| N+4 | 31:0 | <i>CRC</i> | CRC – See Table 6-1. |

Figure 6-14. Inter-Domain Response

6.4.3 Control Packet Routing

6.4.3.1 Upstream-Bound Packets

An Uninitialized Router shall discard a Control Packet with the *CM* bit set to 1b and shall not send any packets in response. An Enumerated Router shall forward a Control Packet with the *CM* bit set to 1b to its Upstream Adapter.

6.4.3.2 Downstream-Bound Packets

When a Router receives a Control Packet with the *CM* bit set to 0b on its Upstream Adapter, it shall route the packet as follows:

- If the *TopologyID Valid* bit in Router Configuration Space is set to 0b, then the Router shall process the packet as follows:
 - If the packet is a Read Request or a Write Request that targets Router Configuration Space, the Router shall process the packet as described in Section 6.4.3.3.
 - Else, Router shall drop the packet and shall not send any packets in response.
- Else the Router shall extract the Egress Adapter number from the Route String that corresponds to the Router's depth in the Spanning Tree (as present in the *Depth* field in the Router Configuration Space). For example, given a Route String with a TopologyID of 0,0,0,0,3,1,8 and a Router with a depth of 2, the Egress Adapter number is 3.
 - If the extracted Adapter number is 0, the Control Adapter of the Router shall consume the packet. The Router shall process the packet using the Enumerated Router Flow in Section 6.4.3.2.1.
 - If the extracted Adapter number refers to a Protocol Adapter, the packet shall be dropped and the Router shall send the Connection Manager a Notification Packet with Event Code = ERR_ADP as defined in Table 6-11.
 - If the extracted Adapter number refers to a disconnected or disabled Adapter, the Router shall drop the packet and shall send the Connection Manager a Notification Packet with Event Code = ERR_CONN as defined in Table 6-11.
 - If the extracted Adapter number refers to a connected Adapter and the *Lock* bit in the Adapter Configuration Space is set to 1b, the Router shall drop the packet and shall send the Connection Manager a Notification Packet with Event Code = ERR_LOCK as defined in Table 6-11.


- Else, the Router shall forward the packet over the Egress Adapter that matches the extracted Adapter number.

When a Router receives a Control Packet with the *CM* bit set to 0b on a Lane Adapter that is not the Upstream Adapter, it shall route the packet as follows:

- If the *TopologyID Valid* bit in Router Configuration Space is set to 0b, then the Router shall process the packet as follows:
 - If the packet is a Read Request or a Write Request, then the Router shall drop the packet and shall send the Adapter that originated the Request a Notification Packet with Event Code = ERR_NUA as defined in Table 6-11.
 - Else, Router shall drop the packet and shall not send any packets in response.
- Else:
 - If the packet is an Inter-Domain Request or an Inter-Domain Response, then the Router shall modify the packet as follows, and then send the packet over the Upstream Adapter.
 - Replace the Route String in the packet with the Route String of the receiving Router within the receiving Domain, then add the Ingress Adapter number of the Adapter connected to the Inter-Domain Link.

For example, Figure 6-15 shows two Domains that are connected by an Inter-Domain Link. When the Host Router in Domain 1 sends an Inter-Domain Packet to the Host Router in Domain 2, the Route String in the Packet has a TopologyID of [0,...,0,3,5]. The Host Router in Domain 2 then updates the Route String in the packet to [0,...,0,0,5].
 - Set the *CM* bit to 1b.
 - Update the *CRC* field.
 - If the packet is a Read Request or a Write Request, then the Router shall drop the packet and shall send the Adapter that originated the Request a Notification Packet with Event Code = ERR_ENUM as defined in Table 6-11.
 - Else, Router shall drop the packet and shall not send any packets in response.

Figure 6-15. Example of Control Packet Routing Between Domains





CONNECTION MANAGER NOTE

The Connection Manager shall set the Lock bit in a Downstream Facing Port to 0b before it can enumerate a Router connected to the Downstream Facing Port.

6.4.3.2.1 Enumerated Router Flow

The following describes how an enumerated Router handles a Control Packet:

- If the Control Packet is either a Read Request or a Write Request, the Router shall process the packet as described in Section 6.4.3.3.
- If the Control Packet is a Hot Plug Acknowledgment Packet, the Router shall process the packet as described in Section 6.8.
- If the Control Packet is a Notification Acknowledgment Packet, the Router shall process the packet as described in Section 6.6.
- Else, the Router shall drop the Control Packet and shall not send any packets in response.

6.4.3.3 Processing of Read and Write Requests

A Router that is the target of a Read Request or a Write Request shall process the packet according to the following rules:

- If the packet addresses any Configuration Space other than the Router Configuration Space, and if the *Adapter Num* field in the packet exceeds the value of the *Max Adapter* field in the Router Configuration Space:
 - The read or write operation shall not be performed and a Response Packet shall not be sent.
 - The Router shall send the Connection Manager a Notification Packet with Event Code = ERR_ADP as defined in Table 6-11.
- Else, if the packet is a Write Request and the *Write Size* field in the packet is zero, then the write operation shall not be performed. The Router shall send a Write Response.
- Else, if the packet is a Read Request or a Write Request, and it targets a Lane Adapter that is unused (see Section 5.2.1), the Router may send the Connection Manager a Notification Packet with Event Code = ERR_ADDR as defined in Table 6-11.
- Else, if the packet is a Write Request for which the *Length* field in the packet header does not equal the expected length $[(Write\ Size + 4) * 4]$:
 - The Router shall not perform a write operation and shall not send a Write Response.
 - The Router shall send the Connection Manager a Notification Packet with Event Code = ERR_LEN as defined in Table 6-11.
- Else, if the packet is a Write Request and if the *Address* and *Write Size* fields in the packet extend beyond the address range supported:
 - The part of the write data that fits within the supported address range shall be written. The part of the write data that fits outside the supported address range shall be dropped.
 - A Write Response shall not be sent.
 - The Router shall send the Connection Manager a Notification Packet with Event Code = ERR_ADDR as defined in Table 6-11.
- Else, if the packet is a Read Request and if the *Address* and *Read Size* fields in the packet extend beyond the address range supported:
 - A Read Response shall not be sent.

- The Router shall send a Notification Packet on its UFP with Event Code = ERR_ADDR as defined in Table 6-11.
- Else, if the packet is a Read Request and the *Read Size* field in the packet is zero then the Router shall send a Read Response without a *Read Data* field.
- Else, if the packet is a Read Request and *Read Size* field in the packet contains a value larger than 60:
 - A Read Response shall not be sent.
 - The Router shall send the Connection Manager a Notification Packet with Event Code = ERR_LEN as defined in Table 6-11.
- Else, process the packet and send a Response Packet.
 - A Router shall send a Write Response for a Write Request to a Path Configuration Space only after it has executed the Write Request, including setting the entry in the Routing Table and in the Egress Arbiter.

Note: Unless stated otherwise, sending a Write Response does not indicate that the Router has taken any action yet related to the contents of the Write Request.

6.4.4 Control Packet Reliability

The following rules provide reliable transport for Control Packets:

- Each Router along the Path of a Control Packet shall check the validity of the *CRC* field. If a packet fails the CRC check, the Router shall discard the packet.
- Unless otherwise specified, a Router that is the target of a Read Request shall send a Read Response within *tCPResponse* of receiving the Request.
- Unless otherwise specified, a Router that is the target of a Write Request shall send a Write Response within *tCPResponse* of receiving the Request.
- A Router forwarding a Control Packet shall send the packet on an Egress Adapter not later than *tCPForward* from the time the packet was received on an Ingress Adapter.
- A Router that receives a Read Request or a Write Request may drop the packet if it is received before the Router has responded to the previous Read Request or Write Request.



CONNECTION MANAGER NOTE

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



CONNECTION MANAGER NOTE

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

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

6.5 Notification Events

Table 6-11 lists the events that generate a Notification Packet. All Event Codes not listed in Table 6-11 are reserved.

Table 6-11. Notification Events

| Event Code | Reference | Initiator | Event Code | Adapter Num |
|---------------------------------|--|--------------------|------------|--|
| Error Events | | | | |
| ERR_CONN | Section 6.4.3 | Router | 0 | The unconnected Adapter. |
| ERR_LINK | Section 4.4.2 | Router | 1 | The Adapter that experienced the error. |
| ERR_ADDR | Section 6.4.3.3 | Router | 2 | The Adapter that is the target of the Read/Write request. |
| ERR_ADP | 1. Section 6.4.3.3 2. Section 6.4.3 | Router | 4 | 1. The Adapter Num value in the Read or Write Request; or 2. The non-Lane Adapter that triggered the event. |
| ERR_ENUM | Section 6.4.3 | Router | 8 | 00h. |
| ERR_NUA | Section 6.4.3 | Router | 9 | 00h. |
| ERR_LEN | Section 6.4.3.3 | Router | 11 | The Adapter Num value in the Request Packet. |
| ERR_HEC | Section 5.1.2.1.1 | Router | 12 | The Ingress Adapter that experienced the error. |
| ERR_FC | Section 5.3.2.1.2 | Router | 13 | The Ingress Adapter that experienced the error. |
| ERR_PLUG | Section 6.8.1.1 | Router | 14 | The Lane 0 Adapter of the hot plugged USB4 Port. |
| ERR_LOCK | Section 6.4.3.2 | Router | 15 | 00h. |
| Hot Plug Acknowledgement | | | | |
| HP_ACK | Section 6.8 | Connection Manager | 7 | The <i>Adapter Num</i> field from the Hot Plug Event Packet. |
| Protocol Adapter Events | | | | |
| DP_BW | Section 10.7 | Router | 32 | The DP IN Adapter. |

6.6 Notification Acknowledgement

A Router expects a Notification Acknowledgment Packet from the Connection Manager in response to the following packets:

- A Notification Packet with Event Code = ERR_LINK.
- A Notification Packet with Event Code = ERR_HEC
- A Notification Packet with Event Code = ERR_FC
- A Notification Packet with Event Code = ERR_PLUG
- A Notification Packet with Event Code = DP_BW

A Router shall retransmit a Notification Packet that requires a Notification Acknowledgment Packet if a Notification Acknowledgment Packet is not received within the time specified by the *Notification Timeout* field in the Router Configuration Space.

A Router shall not send another packet that requires a Notification Acknowledgment while a previous packet that requires a Notification Acknowledgment is pending (i.e. before a Notification Acknowledgment Packet is received or a timeout occurs).

Note: If a Notification Acknowledgment Packet is received after the Notification Timeout expired, it is possible that the Router associates the Notification Acknowledgment Packet with a later Notification Packet. That can result in the Router not retransmitting a Notification Packet even though it did not reach the Connection Manager.



CONNECTION MANAGER NOTE

A Connection Manager shall 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*

6.7 Router Enumeration and Initialization

Router enumeration is the process by which a Connection Manager detects a connected Router and assigns it a TopologyID. After a Connection Manager enumerates a Router, the Router is part of the Connection Manager's Domain.



CONNECTION MANAGER NOTE

A USB4 Connection Manager sets the Connection Manager USB4 Version field to indicate rev. 1.0 support. The below assumes that the Connection Manager USB4 Version field is set to indicate rev. 1.0 support.

On transition to the Uninitialized state, a TBT3-Compatible Router shall:

- Expose USB4 Ports and PCIe Adapters as defined in Section 13.3.1.
- Set its sleep and wake behavior as defined in Section 13.2.4.
- Expose the additional registers defined in Section 13.6.

When the *TopologyID Valid* bit is set to 1b, the Router is enumerated, and a TBT3-Compatible Router shall:

- Expose all its USB4 Ports and PCIe Adapters (if any) to the Connection Manager.
- Set the *Lock* bit to 1b on all Downstream Facing Ports.
- Set its sleep and wake behavior as defined in Section 4.5.
- Disable access to the additional registers defined in Section 13.6.

The TBT3-Compatible Router shall then set the *Router Ready* bit to 1b.

On transition to the Uninitialized state, a Router that is not TBT3-Compatible shall:

- Expose all its USB4 Ports and PCIe Adapters (if any) to the Connection Manager.
- Set the *Lock* bit to 1b on all Downstream Facing Ports.
- Set sleep and wake behavior to default as defined in Section 4.5.

The Router that is not TBT3-Compatible shall then set the *Router Ready* bit to 1b.



CONNECTION MANAGER NOTE

A Connection Manager detects a connected Router by either:

- *Receiving a Hot Plug Event Packet from the Router with the UPG bit set to 0b (see Section 6.8); or*
- *Issuing a Read Request to the Router and receiving a Read Response.*

A Connection Manager shall abide by the following rules when enumerating a Router:

- *The Connection Manager shall use a single Write Request to enumerate a Router. The Write Request sets the Depth field, the TopologyID field, the Connection Manager USB4 Version field, and the TopologyID Valid bit (the latter to 1b).*
- *The Connection Manager may enumerate an Uninitialized Router at any time but is not required to do so.*
- *The Connection Manager shall not enumerate a Router that is connected to a Downstream Facing Port via an Adapter with an Adapter Type that equals Unsupported Adapter.*
- *The Connection Manager may issue Read Requests to an Uninitialized Router.*
- *The Connection Manager shall not issue a Write Request to an Uninitialized Router other than the Write Request that enumerates the Router.*
- *The Connection Manager may identify Inter-Domain Links by sending a Read Request or a Write Request and receiving a Notification Packet with Event Code = ERR_ENUM or Event Code = ERR_NUA.*

After enumerating a Router, the Connection Manager may read the contents of the Router's DROM. If, after reading the contents of DROM, the Connection Manager decides that it does not want the Router in its Domain, it may issue a Downstream Port Reset to the Router as described in Section 6.9. It is recommended that the Connection Manager keep a record of the Router's UUID so that on a subsequent Hot Plug event, the Connection Manager can compare the stored UUID value to the UUID value of the plugged Router to determine whether to enumerate the hot-plugged Router.

After setting the *Router Ready* bit to 1b, a Device Router waits for the Connection Manager to set the *Configuration Valid* bit to 1b. When the *Configuration Valid* bit in Router Configuration Space is set to 1b, a Device Router shall:

- If the *USB3 Tunneling On* bit is set to 1b, establish USB3 tunneling functionality
- If the *PCIe Tunneling On* bit is set to 1b, establish PCIe tunneling functionality

The Device Router shall then set the *Configuration Ready* bit in Router Configuration Space to 1b.



IMPLEMENTATION NOTE

A Device Router may decide to establish USB3 tunneling functionality when its Upstream Facing Port operates in USB4 Mode.

6.8 Hot Plug and Hot Unplug Events

There are two types of Hot Plug Events:

- **Router Hot Plug** – Occurs when a Router detects that another Router is hot-plugged into a Downstream Facing Port. Router Hot Plug is further defined in Section 6.8.1.
- **Adapter Hot Plug** – Occurs when a Protocol Adapter detects the Hot Plug of a device behind the Protocol Adapter. The methods for communicating a DP Adapter Hot Plug to a Connection Manager are defined in 10.3.3. No other Protocol Adapters report Adapter Hot Plug Events.

There are two types of Hot Unplug Events:

- **Router Hot Unplug** – Occurs when a Router detects that another Router is hot-removed from a Downstream Facing Port. Router Hot unplug is further defined in Section 6.8.2.
- **Adapter Hot Unplug** – Occurs when a Protocol Adapter detects the Hot Unplug of a device behind the Protocol Adapter. The methods for communicating a DP Adapter Hot Unplug to a Connection Manager are defined in 10.3.3. No other Protocol Adapters report Adapter Hot Unplug Events.

A Router expects a Hot Plug Acknowledgment Packet from the Connection Manager in response to a Hot Plug Event Packet. A Router shall retransmit a Hot Plug Event Packet if a Hot Plug Acknowledgment Packet acknowledging the Hot Plug or the Hot Unplug Event is not received within the time specified by the *Notification Timeout* field in the Router Configuration Space.

A Router only reports one Hot Plug or Unplug event at a time. A Router shall not send a Hot Plug Event Packet for a new Hot Plug Event from any Adapter until it receives a Hot Plug Acknowledgment Packet for the previous Hot Plug/Unplug Event. A Router shall not send a Hot Unplug Event Packet for a new Hot Plug Event from any Adapter until it receives a Hot Plug Acknowledgment Packet for the previous Hot Plug/Unplug Event.

After receiving a Hot Plug Acknowledgment Packet, a Router shall not send any additional Hot Plug Event Packets for that Hot Plug/Unplug Event. A Router shall ignore a Hot Plug Acknowledgment Packet for a Hot Plug/Unplug Event that was already acknowledged.

A Router shall not generate two consecutive Hot Plug Events or two consecutive Hot Unplug Events for a given Adapter. The next event after a Hot Plug Event for a given Adapter shall always be a Hot Unplug Event. Similarly, the next event after a Hot Unplug Event for a given Adapter shall always be a Hot Plug Event.

A Router shall always report a Hot Plug Event or a Hot Unplug Event. When a Hot Plug Event Packet cannot be sent, the Router shall store the event and shall send Hot Plug Event Packet when conditions allow.



IMPLEMENTATION NOTE

Some Hot Plug Events can cancel each other out and thus do not require storage. For example, if an Adapter detects a Hot Plug then a Hot Unplug before a Hot Plug Event Packet was sent for the first Hot Plug, there is no need to send Hot Plug Event Packets for either the hot plug or the hot unplug.



CONNECTION MANAGER NOTE

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

A Connection Manager can send one Hot Plug Acknowledgement Packet per Hot Plug or Hot Unplug Event (regardless of the number of Hot Plug or Hot Unplug Event Packets received for that event) or it can send a Hot Plug Acknowledgement Packet for each Hot Plug/Unplug Packet. A Connection Manager shall respond to Hot Plug Event Packets from a Router in the order received.

6.8.1 Router Hot Plug

Router Hot Plug handling for Routers that detect a Router Hot Plug on a Downstream Facing Port is described in Sections 6.8.1.1 and 6.8.1.2 respectively for Routers in the Enumerated and Uninitialized states. Section 6.8.1.3 describes how a newly connected Router handles a Router Hot Plug.

6.8.1.1 Enumerated Routers

When a Router in the Enumerated state detects a Router Hot Plug on one of its Downstream Facing Ports (See Section 4.1.2.2), it shall perform the following steps:

1. Perform Lane Initialization on the Lanes of the Downstream Facing Port with the Hot Plugged Router. See Section 4.1.2 for Lane Initialization and Section 4.2 for Adapter states.
2. For each Adapter in the USB4 Port that reaches CL0 state, send a Hot Plug Event Packet to the Connection Manager with the *UPG* bit set to 0b.
3. If none of the Adapters in the USB4 Port reaches CL0 state within *tTrainingAbort1* after entering the Training state, and if the “Hot Plug Failure Indication” capability is enabled in the Router (see Section 8.3.1.3.3.1), then the Router shall send the Connection Manager a Notification Packet with Event Code = *ERR_PLUG* as defined in Table 6-11.

6.8.1.2 Uninitialized Routers

When a Router in the Uninitialized state detects a Router Hot Plug, it shall not send a Hot Plug Event Packet until it transitions to the Enumerated state. After transitioning to the Enumerated state, the Router shall follow the procedure in Section 6.8.1.1 for each connected Adapter in the Downstream Facing Port.

6.8.1.3 Hot Plugged Router

A hot plugged Router shall enable the following for HopID 0:

- Forwarding of Control Packets to and from the Control Adapter.
- Egress scheduling.

6.8.2 Router Hot Unplug

Section 6.8.2.1 describes how a Router handles a Hot Unplug on the Upstream Facing Port. Section 6.8.2.2 describes how a Router handles a Hot Unplug on a Downstream Facing Port.

6.8.2.1 Hot Unplug on the Upstream Facing Port

If a Router is still powered on after being unplugged, it shall initiate a disconnect on the Upstream Facing Port by driving its SBTX line low (see Section 4.4.5.1.1).

6.8.2.2 Hot Unplug on a Downstream Facing Port

When a Router detects a Router Hot Unplug on a Downstream Facing Port, it shall initiate a disconnect on the Downstream Facing Port by driving its SBTX line low (see Section 4.4.5.2.1).

If the Router is not enumerated, it shall not send a Hot Plug Event Packet.

6.9 Downstream Facing Port Reset

When the *Downstream Port Reset* bit in a Downstream Facing Port is set to 1b, a Router shall discard any pending Sideband transactions and initiate a disconnect on the Downstream Facing Port by driving its SBTX line low. See Section 4.4.5.1.1 for disconnect when Link Partner is an Upstream Facing Port and Section 4.4.5.2.1 for disconnect when Link Partner is a Downstream Facing Port.

The Router shall drive the SBTX line high when the *Downstream Port Reset* bit of the Downstream Facing Port is set to 0b. Lane Initialization then takes place.



CONNECTION MANAGER NOTE

A Connection Manager should use a Downstream Port Reset to change Link Parameters. To change Link Parameters, the Connection manager sets one or more of the following fields in the Adapter Configuration Space of the Downstream Facing Port, then initiates a Downstream Facing Port Reset:

- *Target Link Speed field initiates a change in speed.*
- *Request RS-FEC Gen 2 bit initiates a change in RS-FEC at Gen 2 speeds.*
- *Request RS-FEC Gen 3 bit initiates a change in RS-FEC at Gen 3 speeds.*



CONNECTION MANAGER NOTE

It is recommended that a Connection Manager execute the following steps when performing a Downstream Facing Port Reset:

1. *The Connection Manager sets the Downstream Port Reset bit of a Downstream Facing Port to 1b.*
2. *The Connection Manager waits for 10ms.*
3. *The Connection Manager sets the Downstream Port Reset bit of the Downstream Facing Port to 0b to allow Lane Initialization to take place.*

6.10 Timing Parameters

Table 6-12 lists the timing parameters for the Configuration Layer.

Table 6-12. Configuration Layer Timing Parameters

| Parameter | Description | Min | Max | Units |
|-------------|---|-----|-----|-------|
| tCPResponse | The time between receiving a Control Packet and sending a response. | -- | 2 | ms |
| tCPForward | The time between receiving a Control Packet and forwarding it on an Egress Adapter. | -- | 500 | μs |

7 Time Synchronization

A Router shall support the USB4™ Time Synchronization Protocol described in this chapter. Support for Inter-Domain Time Synchronization is optional for a Host Router. The Time Synchronization Protocol provides a mechanism for synchronizing the real-time clocks and absolute time of connected Routers to a high degree of accuracy and precision.

7.1 Time Synchronization Architecture

The Time Synchronization Protocol is a distributed protocol that defines how the real-time clocks in a USB4 Fabric synchronize with each other. The real-time clocks are organized into a hierarchy with the Host Router at the top of the hierarchy determining the reference time for the entire Domain. Clock synchronization is achieved by exchanging Ordered Sets and Time Sync Packets where downstream devices use both local timestamps and the timing information in the Time Sync Packets to adjust their clocks to the time of the Host Router.


When Inter-Domain time synchronization is disabled, the Time Synchronization Protocol executes within the scope of a single Domain. All Time Sync Packets, state machines and other entities are associated with a single Domain. The time established within one Domain by the protocol is independent of the time in other Domains.

When Inter-Domain time synchronization is enabled, the Time Synchronization Protocol executes within the scope of an interconnected set of Domains. The Connection Managers of the Domain establish an Inter-Domain clock synchronization hierarchy by selecting one of the Host Routers to be the time source for all Domains. The Time Synchronization Protocol then synchronizes the clocks of the other Domains to the Inter-Domain Host Router clock.

7.1.1 Synchronization Hierarchy

7.1.1.1 Intra-Domain Hierarchy

Figure 7-1 shows an example of the time synchronization hierarchy within a Domain. The time synchronization hierarchy follows the Spanning Tree of the Domain. The Host Router of the Domain provides the clock for the Domain.

Figure 7-1. Time Synchronization Hierarchy within a Domain (Informative)

7.1.1.2 Inter-Domain Hierarchy

When multiple Domains are connected together, each Domain is synchronized to its own Host Router Clock. The Connection Managers of the Domains then enable time synchronization across the Domains. If there is more than one Inter-Domain Link between two Domains, then only one Inter-Domain Link is enabled to perform the Time Sync handshake.

When time synchronization is enabled across an Inter-Domain Link, one Domain is the Time Source and the other Domain follows. The Time Source Domain has the *IDTR* (Inter-Domain Time Responder) bit in the TMU Adapter Configuration Capability set to 1b and *IDTI* (Inter-Domain Time Initiator) bit set to 0b. The Domain that follows the Inter-Domain Host Router Time, has the *IDTI* bit set to 1b and *IDTR* bit set to 0b.

In the case where there are multiple daisy-chained Domains connected to one another, the overall Inter-Domain topology follows a Spanning Tree similar to the Intra-Domain hierarchy. This hierarchy between Domains is determined by the *IDTR* and *IDTI* bits.

7.1.2 Time Sync Parameters

7.1.2.1 Local Time


A Router shall provide a free-running clock for use in capturing the time stamps needed for the Time Synchronization Protocol. This free-running clock is called the Local Clock. The Local Clock shall run at a frequency of 125 MHz or greater with an accuracy of ± 100 ppm relative to the nominal Local Clock frequency. The Local Clock shall not be spread-spectrum.

A Router shall use an 80-bit Local Time counter with the format shown in Figure 7-2 to track Local Time. The *Nanoseconds* field represents the number of whole nanoseconds and the

Fractional Nanoseconds field represents a fraction of a nanosecond that is less than 1. For example, a Local Time value of 524.25 is stored as 0...020C4000h.

The Local Time counter shall be incremented up with the Local Clock. The Local Time counter holds an unsigned fractional value.

Figure 7-2. Local Time Counter Format



Note: If the local clock frequency is such that each tick cannot be represented fully with the Local Time Counter, it affects the overall accuracy of the Local Time Counter. For example, if the local clock frequency is 300 MHz, the period is 3.333... ns. The accuracy of each tick with ideal clock would be ~1.53 ppm.

7.1.2.2 Time Offset

The time offset between Local Time and Host Router Time is stored in the *TimeOffsetFromHR* register in the TMU Adapter Configuration Capability in Router Configuration Space. The *TimeOffsetFromHR* register shall have the format shown in Figure 7-3.

Figure 7-3. TimeOffsetFromHR Register Format



The *TimeOffsetFromHR* register holds a signed fractional value and uses 2's complement representation. For example, 2.5 ns is represented as 0000 0000 0002 8000h and -2.5 ns is represented as FFFF FFFF FFFD 8000h.

Note: A value of 7FFF FFFF FFFF FFFFh is used to indicate that the time offset is too big to be represented.

The *TimeOffsetFromHR* register shall be updated at the conclusion of every Time Sync Handshake.


7.1.2.3 Frequency Offset

The frequency offset between the Local Clock and the Host Router Clock is stored in the *FreqOffsetFromHR* register in the TMU Adapter Configuration Capability in Router Configuration Space. The *FreqOffsetFromHR* register holds the frequency offset as a signed fractional value represented using 2's complement notation. The frequency offset shall be computed according to Equation 7-2.

For example, a frequency offset of 1ppm is represented as F = 218DEFh in the *FreqOffsetFromHR* register.

The *FreqOffsetFromHR* register shall be updated at the conclusion of every Time Sync Handshake.

Figure 7-4. FreqOffsetFromHR Register Format




7.2 Time Stamp Measurement

A USB4 Port shall generate a time stamp whenever it either sends or receives a Time Sync Notification Ordered Set (TSNOS). A USB4 Port shall use the value in the Local Time counter to capture time stamps.

The start of the first bit of the TSNOS is referred to as the Time Stamp Point. This applies to both Single-Lane Links and Dual-Lane Links. A time stamp shall be taken at the Time Stamp Point of the First TSNOS. If one or more back-to-back TSNOS are received immediately after the first TSNOS, they shall be ignored. The same transmit reference plane shall be used for all transmitted TSNOS and the same receive reference plane shall be used for all received TSNOS. However, the reference plane may be different for transmit and receive Paths through the Physical Layer. The time stamp measurement shall have a resolution of at least 8 ns (i.e. the period of the Local Clock).

Figure 7-5. Time Measurement Model for 64/66b Encoding



The time stamp in a Time Sync Packet shall have the format shown in Figure 7-2.

7.2.1 Asymmetry Corrections

Time stamps shall be corrected for asymmetry between transmit and receive paths. An Upstream Facing Port (UFP) shall correct for asymmetry by performing the following computations:

$$t1 = \text{Delay Request Sent Time Stamp} + \text{TxTimeToWire}$$

$$t4 = \text{Delay Response Received Time Stamp} - \text{RxTimeToWire}$$

where, TxTimeToWire is the value in the *TxTimeToWire* field of the TMU_ADP_CS_1 register of the UFP and RxTimeToWire is the value in the *RxTimeToWire* field of the TMU_ADP_CS_2 register of the UFP.

Note: The reference plane is a virtual representation of the RX and TX pipeline depth. The reference plane is constant regardless of whether or not RS-FEC is activated.

A Downstream Facing Port (DFP) shall correct for asymmetry by performing the following computations:

$$t2 = \text{Delay Request Received Time Stamp} - \text{RxTimeToWire}$$

$$t3 = \text{Delay Response Sent Time Stamp} + \text{TxTimeToWire}$$

where, TxTimeToWire is the value in the *TxTimeToWire* field of the TMU_ADP_CS_1 register of the DFP and RxTimeToWire is the value in the *RxTimeToWire* field of the TMU_ADP_CS_2 register of the DFP.

The time duration between when a USB4 Port generates a time stamp and when it transmits first bit of a TSNOS on the wire shall be equal to the value in the *TxTimeToWire* field of the TMU_ADP_CS_1 register. The time duration between when a USB4 Port receives the first bit of a TSNOS on the wire and when it generates a time stamp shall be equal to the value in the *RxTimeToWire* field of the TMU_ADP_CS_2 register. The values in the *TxTimeToWire* and *RxTimeToWire* fields are vendor defined.

7.3 Time Sync Protocol

7.3.1 Time Sync Handshake

The Time Sync Handshake is used to measure the time offset between two USB4 Ports.

There are two types of Time Sync Handshakes: Bi-Directional and Uni-Directional. Bi-Directional Time Sync Handshakes are described in Section 7.3.1.1. Uni-Directional Time Sync Handshakes are described in Section 7.3.1.2.

Time Sync Handshakes can use HiFi Mode or LowRes Mode. Mode is configured using the *TSPacketInterval* field. A Router shall support Bi-Directional Time Sync Handshakes in HiFi Mode. A Router shall support Uni-Directional Time Sync Handshakes in HiFi Mode and LowRes Mode.



CONNECTION MANAGER NOTE

When using Bi-Directional Handshakes, a Router only support HiFi Mode. A Connection Manager shall not configure a Router in LowRes Mode when using Bi-Directional Handshakes.



IMPLEMENTATION NOTE

A Link cannot enter CLx states when Bi-Directional HiFi Mode Time Sync Handshakes are enabled (the handshake frequency is too high to allow time for low power entry and exit). Because of this, a Device Router should only request Bi-Directional HiFi Mode if it is absolutely required for device functionality (e.g. a USB4 device that uses PCIe PTM for audio and requires higher accuracy than Uni-Directional HiFi Mode can support). A Device Router uses DROM to indicate what type of Time Sync Handshakes it needs to operate. See the USB4 DROM Specification for more details.

The Time Sync Handshake uses three types of Time Sync Messages: Delay Request Message, Delay Response Message and Follow-Up Packet. Delay Request and Delay Response Messages are implemented using the TSNOS described in Section 7.3.3.1. A Follow-Up Packet is a Transport Layer Packet and has the format described in Section 7.3.3.2.

A Time Sync Handshake takes place between two USB4 Ports as follows:

- When the USB4 Ports are connected by a Single-Lane Link, a Time Sync Handshake shall occur over that Link.
- When the USB4 Ports are connected by a Dual-Lane Link, the Delay Request and Delay Response Ordered Sets in the Time Sync Handshake shall be sent on both the Lane 0 and Lane 1. The Follow-Up Packet shall be sent using both Lanes (i.e. alternating bytes across Lanes). The timestamp shall be taken when the first Ordered Set from the Link Partner is received.
- A Receiver shall ignore a TSNOS that arrives back-to-back after another TSNOS.




CONNECTION MANAGER NOTE

When configuring a USB4 Port for a Single-Lane Link, a Connection Manager shall disable the Lane 1 Adapter, before enabling Time Sync Handshakes.

7.3.1.1 Bi-Directional Time Sync Handshake

Figure 7-6 shows the Bi-Directional Time Sync Handshake.

Figure 7-6. Bi-Directional Time Sync Handshake



A DFP shall not initiate a Bi-Directional Time Sync Handshake.

A UFP shall send a Delay Request to the DFP at the interval specified in the *TSPacketInterval* field in the *TMU_RTR_CS_3* register in Router Configuration space. A DFP shall transmit a Delay Response Message within *RespTimeout* of receiving a Delay Request Message. A DFP shall transmit a Follow-Up Packet within *SendTimeout* of transmitting the associated Delay Response Message.

Note: A Connection Manager can control the Time Sync Handshake rate by writing to the TSPacketInterval field in the TMU_RTR_CS_3 register in Router Configuration space.

A UFP shall generate time stamp t_1 upon transmission of a Delay Request and shall generate time stamp t_4 upon reception of a Delay Response. A DFP shall generate time stamp t_2 upon receipt of a Delay Request and shall generate time stamp t_3 upon transmission of a Delay Response. Section 7.2 defines how time stamps are measured. Section 7.2.1 defines how timestamps are corrected for asymmetry. Transmission and reception of a Follow-Up Packet does not result in the generation of time stamps.

If an error occurs during the transmission or reception of a Time Sync Packet, the entire Time Sync Handshake shall be voided (i.e. neither time stamps nor values from the Follow-Up Packet shall be used).

Figure 7-7 shows the state machine that is recommended for a UFP using Bi-Directional Time Sync Handshake. The state machine uses the following timeout values:

- ReqTimeout is the maximum time between a Time Sync event and transmission of Delay Request. See Table 7-1 for recommended value.

Note: A Router assumes that the transmission of Delay Requests is constant and consistent. However, any variance in the time it takes to transmit Delay Request adds to overall noise. The purpose of ReqTimeout is to add robustness and make sure the state machine doesn't get stuck as a result of such noise.

- RespTimeout is the maximum time between sending a Delay Request and receiving a Delay Response. See Table 7-1 for recommended value.
- FPTimeout is the maximum time between receiving a Delay Response and receiving a Follow-Up Packet. See Table 7-1 for recommended value.

Table 7-1. Bidirectional UFP Timeout Values

| Timeout Name | Timeout Recommended Value |
|--------------|---------------------------|
| ReqTimeout | 1 μ s |
| RespTimeout | 2 μ s |
| FPTimeout | 10 μ s |

Note: In the state machine shown in Figure 7-7 and Figure 7-8, a Link Event occurs when the USB4 Port receives an Ordered Set related to Link negotiation or training. A Time Sync Event occurs once every TSPacketInterval and indicates the start of a new Time Sync Handshake.

Figure 7-7. UFP State Machine for Bi-Directional Time Sync Handshake (Recommended)





Figure 7-8 shows the state machine that is recommended for a DFP using Bi-Directional Time Sync Handshake. The state machine uses the following timeout values:

- **RespTimeout** is the maximum time between receiving a Delay Request and sending a Delay Response. See Table 7-2 for recommended value.
- **SendTimeout** is the maximum time between sending a Delay Response and sending a Follow-Up Packet. See Table 7-2 for recommended value.

Table 7-2. Bidirectional DFP Timeout Values

| Timeout Name | Timeout Recommended Value |
|--------------|---------------------------|
| RespTimeout | 1 μ s |
| SendTimeout | 8 μ s |

Figure 7-8. DFP State Machine for Bi-Directional Time Sync Handshake (Recommended)




When using Bi-Directional Time Sync Handshake, a DFP shall do the following upon receiving a Delay Request:

1. Take the t_2 timestamp.
2. Send a Delay Response.
 - a. Delay Response shall be sent within RespTimeout time or receiving the Delay Request.
 - b. The DFP shall take the t_3 timestamp upon Delay Response transmission.
3. Compute the updated TimeOffsetFromHR parameter according to Equation 7-6 using the t_3 timestamp from Step 2b.
4. Send Follow-Up Packet with the TimeOffsetFromHR calculated in Step 3.

7.3.1.2 Uni-Directional Time Sync Handshake

The Uni-Directional Time Sync Handshake is shown in Figure 7-9.

Figure 7-9. Uni-Directional Time Sync Handshake



When using Uni-Directional Time Sync Handshake, only the DFP shall initiate a Time Sync Handshake.

When the Link is not in a CLx state, a DFP shall send a Delay Response to the UFP at the interval specified in the *TSPacketInterval* field in the TMU_RTR_CS_3 register in Router Configuration space. A DFP shall transmit a Follow-Up Packet within SendTimeout after transmitting the associated Delay Response Packet.

When the Link is in a CLx state, a DFP may “pause” Time Sync Handshakes by not sending a Delay Response to the UFP. The DFP may pause Time Sync Handshakes for any period of time while the Link is in CLx state. The DFP shall resume Time Sync Handshakes upon exiting the CLx state.



IMPLEMENTATION NOTE

Pausing Time Sync Handshakes will affect time synchronization accuracy in a UFP. When resuming Time Sync Handshakes, the value of the timestamp in the first Follow Up Packet may not be as expected (for example, the Local Counter may restart from 0 upon exiting CLx state). Accuracy is not guaranteed until tConvergeTime after Time Sync Handshakes are resumed. It is important that a Router implementation take this into consideration when pausing and resuming Time Sync Handshakes in order to not disrupt tunneled traffic.

When using Uni-Directional Time Sync Handshakes, the propagation delay (T_{pd}) between the DFP and UFP cannot be computed. Therefore, Uni-Directional Time Sync Handshake is only intended for use when a constant time offset between Routers is not needed (e.g. video tunneling). The format of the Follow-Up Packet is the same as for Bi-Directional Time Sync Handshake, keeping the *RequestReceiptTS* (t_2) and *ResponseOriginTS* (t_3) equal.

Uni-Directional Time Sync Handshake shall be used when the following conditions are true:

- Both Link Partners support Uni-Directional Time Sync Handshake.
- Both Link Partners have the *EnableUniDirectionalMode* field in the TMU_AD_P_CS_3 register in Adapter Configuration Space set to 1b.



CONNECTION MANAGER NOTE

When configuring a Link to use Uni-Directional Time Sync Handshakes, it is recommended that the EnableUniDirectionalMode field in the UFP be set to 1 before setting the EnableUniDirectionalMode field to 1 in the DFP.

Figure 7-10 shows the state machine that is recommended for a DFP using Uni-Directional Time Sync Handshake.

Note: In the state machine shown in Figure 7-10, a Link Event occurs when the USB4 Port transmits an Ordered Set related to Link negotiation or training.

Figure 7-10. DFP State Machine for Uni-Directional Time Sync Handshake (Recommended)





Figure 7-11 shows the state machine that is recommended for a UFP using Uni-Directional Time Sync Handshake.

Note: In the state machine shown in Figure 7-11, a Link Event occurs when the USB4 Port receives an Ordered Set related to Link negotiation or training.

Figure 7-11. UFP State Machine for Uni-Directional Time Sync Handshake (Recommended)



7.3.2 Inter-Domain Time Sync

Time synchronization between multiple Domains is possible when the Connection Managers of the Domains establish an Inter-Domain clock synchronization hierarchy. The Inter-Domain clock synchronization hierarchy is defined by the *IDTR* and *IDTI* bits as follows:

- A Domain is the Inter-Domain Time Source when the *IDTR* bit in the *TMU_ADAP_CS_3* register of the TMU Adapter Configuration Capability is set to 1b in the Port that is part of the Inter-Domain Link. The Port with the *IDTR* bit set to 1b is referred to as the “IDTR Port” in this section.
- A Domain follows the Inter-Domain Time Source when the *IDTI* bit in the *TMU_ADAP_CS_3* register of the TMU Adapter Configuration Capability is set to 1b in the Port that is part of the Inter-Domain Link. The Port with the *IDTI* bit set to 1b is referred to as the “IDTI Port” in this section.

A USB4 Port shall perform Time Sync Handshakes as described in Section 7.3.1 across the Inter-Domain Link when either the *IDTR* or *IDTI* bit is set to 1b.

After completing a Time Sync Handshake across an Inter-Domain Link, the IDTI Port shall calculate the following:

- The Inter-Domain time stamp (see Section 7.4.2.1).
- The Inter-Domain frequency offset (see Section 7.4.2.2).
- The Inter-Domain time offset (see Section 7.4.2.3).

After calculating the Inter-Domain time offset, Inter-Domain frequency offset and Inter-Domain time stamp, the Router with the IDTI Port shall update the *InterDomainTimeStamp*, *FreqOffsetFromInterDomainHR* and *TimeOffsetFromInterDomainHR* fields in Router Configuration Space. If the IDTI Port is part of a Device Router, it shall prepare an Inter-Domain Time Stamp Packet as described in Section 7.3.3.3 where:

- The *IDTimeStamp* field contains the calculated Inter-Domain time stamp.

- The *FreqOffsetFromInterDomainHR* field contains calculated the Inter-Domain frequency offset.
- The *TimeOffsetFromInterDomainHR* field contains the calculated Inter-Domain time offset.

Note: The Inter-Domain Time Stamp Packet is sent up the Spanning Tree via the UFP and then forwarded by each Router on the way until it reaches Host Router.

If the *TSInterDomainInterval* field in Router Configuration Space is 0, the IDTI Port shall send the Inter-Domain Time Stamp Packet after each Inter-Domain Time Sync Handshake. If the *TSInterDomainInterval* field in Router Configuration Space is not 0, the IDTI Port shall send the Inter-Domain Time Stamp Packet at time intervals equal to $(TSInterDomainInterval + 1) * TSPacketInterval$ number of microseconds.

When a Host Router receives an Inter-Domain Time Stamp Packet:

- If the *IDE* bit in the *TMU_RTR_CS_0* register of the Host Router's TMU Router Configuration Capability is set to 1b, the Host Router shall update its *TimeOffsetFromInterDomainHR* and *FreqOffsetFromInterDomainHR* registers using the time offset and frequency offset respectively contained in the Inter-Domain Time Stamp Packet.

Note: The updated values in the Host Router's TimeOffsetFromInterDomainHR and FreqOffsetFromInterDomainHR registers are propagated in the Domain via Follow-Up Packets using the corresponding fields of the packet: IDTimeStamp, TimeOffsetFromInterDomain and FreqOffsetFromInterDomain. This results in all the Routers in the Domain synchronizing their time to the Inter-Domain Host Router Clock.

- If the *IDE* bit in the Host Router is set to 0b, the Host Router shall drop the Inter-Domain Time Stamp Packet and shall not update its *TimeOffsetFromInterDomainHR* or *FreqOffsetFromInterDomainHR* registers.

If the IDTI Port is part of a Host Router, there are no Inter-Domain Time Stamp Packets and Inter-Domain parameters are computed directly when Inter-Domain Time Sync is enabled.



CONNECTION MANAGER NOTE

It is recommended that a Connection Manager not set the IDE bit to 1b until after computations over the Inter-Domain Link are stable.

At any instant of time, a Router can compute the Inter-Domain Host Router Time as described in Section 7.4.2.4.


Figure 7-12 shows an example of Inter-Domain time synchronization between two Domains where the following takes place:

1. The Connection Manager detects an Inter-Domain Link between Domain 1 and Domain 2.
2. The Connection Manager sets *IDTI* bit to 1b in Domain 1 and *IDTR* bit to 1b in Domain 2, making Domain 2 the Inter-Domain Time Source with Domain 1 following the Domain 2 time.
3. Either Bi-Directional or Uni-Directional Time Sync handshakes take place across the Inter-Domain Link.
4. Domain 1 calculates the time and frequency offsets between the IDTI Port and the Inter-Domain Time Source.

5. The Router with the IDTI Port sends an Inter-Domain Time Stamp Packet to the Host Router in Domain 1.
6. Host Router in Domain 1 receives the Inter-Domain Time Stamp Packet and updates its time and frequency offsets.
7. The updated time is propagated through Domain 1 via Time Sync Handshakes within the Domain.

The red line in Figure 7-12 designates the route on which an Inter-Domain Time Stamp Packet is sent to the top of the tree.

Figure 7-12. Inter-Domain Time Sync Protocol (Informative)



7.3.3 Packet Formats

7.3.3.1 Time Sync Notification Ordered Set Format

Both a Delay Request and a Delay Response shall consist of the Time Sync Notification Ordered Set (TSNOS) defined in Section 4.4.7.

7.3.3.2 Follow-Up Packet Format

A Follow-Up Packet shall have the format shown in Figure 7-13.

Figure 7-13. Follow-Up Packet Format

| | | | | | | |
|---------------------------------|---------|----------------|------------------------|-------------|-----|---|
| 31 | PDF = 1 | Suppl D = 0 | HopID = 3 | Length = 60 | HEC | 0 |
| ResponseOriginTS 79:64 | | | RequestReceiptTS 79:64 | | | |
| RequestReceiptTS 63:32 | | | | | | |
| RequestReceiptTS 31:0 | | | | | | |
| ResponseOriginTS 63:32 | | | | | | |
| ResponseOriginTS 31:0 | | | | | | |
| TimeOffsetFromHR 63:32 | | | | | | |
| TimeOffsetFromHR 31:0 | | | | | | |
| FreqOffsetFromHR 31:0 | | | | | | |
| Reserved | | | IDTimeStamp 79:64 | | | |
| IDTimeStamp 63:32 | | | | | | |
| IDTimeStamp 31:0 | | | | | | |
| TimeOffsetFromInterDomain 63:32 | | | | | | |
| TimeOffsetFromInterDomain 31:0 | | | | | | |
| FreqOffsetFromInterDomain 31:0 | | | | | | |
| CRC32 | | | | | | |

A Follow-Up Packet shall have the *PDF* field set to 1, the *HopID* set to 3, and the *Length* set to 60. The payload shall contain the fields in Table 7-3.

Table 7-3. Follow-Up Packet Payload

| DW | Bits | Field | Description |
|----|-------|-------------------------|---|
| 0 | 15:0 | <i>RequestReceiptTS</i> | For Bi-Directional Time Sync Handshakes: This field contains the time stamp from when the corresponding Delay Request Packet was received (i.e. time stamp t_2). The time stamp shall include the asymmetry corrections performed at the DFP as specified in Section 7.2.1. This field shall have the format shown in Figure 7-2. For Uni-Directional Time Sync Handshakes: This field shall contain the same value as the <i>ResponseOriginTS</i> field. |
| 1 | 31:0 | | |
| 2 | 31:0 | | |
| 0 | 31:16 | <i>ResponseOriginTS</i> | This field contains the time stamp from when the corresponding Delay Response Packet was sent (i.e. time stamp t_3). The time stamp shall include the asymmetry corrections performed at the DFP as specified in Section 7.2.1. This field shall have the format shown in Figure 7-2. |
| 3 | 31:0 | | |
| 4 | 31:0 | | |
| 5 | 31:0 | <i>TimeOffsetFromHR</i> | This field contains the time offset between the DFP and the Host Router. This field shall have the format shown in Figure 7-3. |
| 6 | 31:0 | | |

| DW | Bits | Field | Description |
|----|-------|------------------------------------|---|
| 7 | 31:0 | <i>FreqOffsetFromHR</i> | This field contains the frequency offset between the DFP and the Host Router. This field shall have the format shown in Figure 7-4. |
| 8 | 31:16 | <i>Reserved</i> | This field shall be set to zero. |
| | 15:0 | <i>IDTimeStamp</i> | For a Host Router: If <i>IDE</i> bit is set to 1b, this field contains the most recent value of the Inter-Domain time stamp received from the IDTI Port. If <i>IDE</i> bit is set to 0b, then this field shall be set to 0. For a Device Router: This field shall contain the <i>IDTimeStamp</i> value from the last Follow-Up Packet that the IDTI Port received. |
| 9 | 31:0 | | |
| 10 | 31:0 | | |
| 11 | 31:0 | <i>TimeOffsetFromInterDomainHR</i> | For a Host Router: If <i>IDE</i> bit is set to 1b, this field contains value of the <i>TimeOffsetFromInterDomainHR</i> field in the last received Inter-Domain Time Stamp Packet. If <i>IDE</i> bit is set to 0b, then this field shall be set to 0. For a Device Router: This field shall contain the <i>TimeOffsetFromInterDomainHR</i> value in the last Follow-Up Packet that the IDTI Port received. |
| 12 | 31:0 | | |
| 13 | 31:0 | <i>FreqOffsetFromInterDomainHR</i> | For a Host Router: If <i>IDE</i> bit is set to 1b, this field contains value of the <i>FreqOffsetFromInterDomainHR</i> field in the last received Inter-Domain Time Stamp Packet. If <i>IDE</i> bit is set to 0b, then this field shall be set to 0. For a Device Router: This field shall contain the <i>FreqOffsetFromInterDomainHR</i> value in the last Follow-Up Packet that the IDTI Port received. |
| 14 | 31:0 | <i>CRC32</i> | CRC32 remainder value computed over the entire packet payload. The CRC32 computation shall be based on the following specification: <ul style="list-style-type: none"> • Width: 32 • Poly: 1EDC 6F41h • Init: FFFF FFFFh • RefIn: True • RefOut: True • XorOut: FFFF FFFFh |

7.3.3.3 Inter-Domain Time Stamp Packet

An Inter-Domain Time Stamp Packet shall have the format shown in Figure 7-14.

Figure 7-14. Inter-Domain Time Stamp Packet Format

| | | | | |
|-----------------------------------|-------------|-------------------|-------------|-----|
| PDF = 2 | SupplID = 0 | HopID = 3 | Length = 28 | HEC |
| Reserved | | IDTimestamp 79:64 | | |
| IDTimestamp 63:32 | | | | |
| IDTimestamp 31:0 | | | | |
| TimeOffsetFromInterDomainHR 63:32 | | | | |
| TimeOffsetFromInterDomainHR 31:0 | | | | |
| FreqOffsetFromInterDomainHR | | | | |
| CRC32 | | | | |

An Inter-Domain Time Stamp Packet shall have the PDF set to 2, the HopID set to 3, and the Length set to 28. The payload shall contain the fields in Table 7-4.

Table 7-4. Inter-Domain Time Stamp Packet Payload

| DW | Bits | Field | Description |
|-------|-------|------------------------------------|--|
| 0 | 31:16 | <i>Reserved</i> | This field shall be set to zero. |
| 0,1,2 | 79:0 | <i>IDTimestamp</i> | This field contains the value of the Inter-Domain timestamp. It is computed using Equation 7-11. This field shall have the format shown in Figure 7-2. |
| 3,4 | 63:0 | <i>TimeOffsetFromInterDomainHR</i> | This field contains the time offset between the Host Router in the Domain with the IDTI Port and the Inter-Domain Time Source. It is computed using Equation 7-16. This field shall have the format shown in Figure 7-3. |
| 5 | 31:0 | <i>FreqOffsetFromInterDomainHR</i> | This field contains the frequency offset between the Host Router in the Domain with the IDTI Port and the Inter-Domain Time Source. It is computed using Equation 7-13. This field shall have the format shown in Figure 7-4. |
| 6 | 31:0 | <i>CRC32</i> | CRC32 remainder value computed over the entire packet payload. The CRC32 computation shall be based on the following specification: <ul style="list-style-type: none"> • Width: 32 • Poly: 1EDC 6F41h • Init: FFFF FFFFh • RefIn: True • RefOut: True • XorOut: FFFF FFFFh |

7.4 Time Computations

A Router shall be able to compute the Host Router Time at any instant in time. A Router can do this by maintaining a continuous notion of Host Router Time or it can use some other means that are implementation specific.

Section 7.4.1 describes how to compute Host Router Time within a single Domain. Section 7.4.2 describes how to compute Host Router Time within multiple interconnected Domains.

Table 7-5 lists the variables used for the time computations in this section. Table 7-6 lists the index notations and their meanings.


Table 7-5. Definition of Variables

| Variable | Description |
|-----------------|--|
| $f(U,D)$ | Frequency ratio between the UFP (U) and DFP (D). |
| $f(U,H)$ | Frequency ratio between the UFP and Host Router (H). |
| $F(U,D)$ | Frequency offset between the UFP and DFP. |
| $F(U,H)$ | Frequency offset between the UFP and Host Router. |
| $F(D,H)$ | Frequency offset between the DFP and Host Router. |
| T_{pd} | Propagation delay between the UFP and DFP. |
| $O_{last}(D,H)$ | Most recent Time offset between the DFP and Host Router. |
| $t_{last}(D)$ | Time when $O_{last}(D, H)$ was computed at the DFP. |

Table 7-6. Index Notation

| Index | Description |
|-------|--|
| n | Current computation cycle. |
| N | Current computation cycle of the frequency offset. This computation occurs every w number of Time Sync Handshakes. For example, if $w = 800$, then the first frequency offset computation occurs when $n = 800$, $N = 1$, and the next occurs when $n = 1600$, $N = 2$, and so on. |
| w | Number of Time Sync Handshakes over which the frequency offset is computed. This value is obtained from the <i>Freq Measurement Window</i> field in the <code>TMU_RTR_CS_0</code> register in the TMU Router Configuration Capability in Router Configuration Space. |
| U_i | A UFP in Domain i . |
| D_i | A DFP in Domain i . |
| H_i | The Host Router in Domain i . |
| t_h | Host Router Time. |

Figure 7-15 shows an example Inter-Domain topology for time synchronization. The Host Router of a Domain i is indicated by using the notation H_i . The DFP and UFP in Domain i are indicated as D_i and U_i respectively. In the following sections the subscript i refers to the ITDI Domain, the subscript j refers to the ITDR Domain and the subscript g refers to the Time Source Domain. Note that a Router can have both its UFP and DFP involved in the Time Sync Handshake and computations. In addition, a Router may also have the *ITDI* and *ITDR* bits set to 1, which means that it may perform Inter-Domain calculations both as a Time Source and a follower.

Figure 7-15. Inter-Domain Topology (Informative)

7.4.1 Intra-Domain Equations

A Router shall use the following series of computations to deduce the current Host Router Time within a single Domain:

- The UFP uses Equation 7-1 to compute the frequency ratio between itself and its Link Partner.
- The UFP uses Equation 7-2 to compute the frequency offset. The UFP uses the frequency ratio obtained in Step 1 as input to the computation.
- The UFP uses Equation 7-3 to compute the frequency ratio between itself and the Host Router. The UFP uses the *FreqOffsetFromHR* value in the last Follow-Up Packet received from its Link Partner as input to the computation.
- The UFP uses Equation 7-4 to compute the frequency offset between itself and the Host Router. The UFP use the frequency ratio computed in Step 3 as input to the computation.
- The UFP uses Equation 7-5 to compute the propagation delay of the cable between the UFP and its Link Partner.
- The Link Partner (i.e. the DFP upstream of the UFP) uses Equation 7-6 to compute the most updated time offset from the Host Router. The Link Partner uses the most recent result from Equation 7-9, as input to the computation.

For example, in a Spanning Tree with 3 routers (A->B->C), Router A is the Host Router and is therefore the Domain Time Source. Because Router A is the Domain Time Source, the value of Equation 7-6 is 0 at Router A and when Router A sends a Follow-Up Packet to Router B, the *TimeOffsetFromHR* field will be 0. When Router B receives a Follow-Up Packet from Router A, it computes the time offset from the Host Router using Equation 7-9 and stores it for further usage. When Router B sends a Follow-Up Packet to Router C, the time offset from the Host Router previously computed by Router B is out-of-date. Because of this, Router B uses Equation 7-6 to compute the updated time offset from Host Router between itself and Router A, before sending a Follow-Up Packet with the updated time offset to Router C.

- The UFP uses Equation 7-7 (Bi-Directional Time Sync Handshake) or Equation 7-8 (Uni-Directional Time Sync Handshake) to compute the time offset from its Link Partner.
- The UFP uses Equation 7-9 to compute the time offset from the Host Router. The UFP uses the time offset from Equation 7-7 or Equation 7-8, and the value in the last received Follow-Up Packet as input to the computation.
- The UFP uses Equation 7-10 to compute the current time offset from the Host Router. The UFP uses the time offset computed in Equation 7-9 and the frequency offset computed in Equation 7-3 as input to the computation.

Equation 7-1. Frequency Ratio between UFP and DFP

$$f(U_i, D_i)[N] = \frac{t_4[n] - t_4[n-w]}{t_3[n] - t_3[n-w]}$$

Equation 7-2. Frequency Offset between UFP and DFP

$$F(U_i, D_i)[N] = \{f(U_i, D_i)[N] - 1\} \cdot 2^{41}$$

Equation 7-3. Frequency Ratio between UFP and Host Router

$$f(U_i, H_i)[N] = f(U_i, D_i) \cdot \left\{ 1 + \frac{F(D_i, H_i)[N]}{2^{41}} \right\}$$

where $F(D_i, H_i)[N]$ is the value in the *FreqOffsetFromHR* field of the most recently received Follow-Up Packet.

Equation 7-4. Frequency Offset between UFP and Host Router

$$F(U_i, H_i)[N] = \{f(U_i, H_i)[N] - 1\} \cdot 2^{41}$$

Equation 7-5. Propagation Delay (Bi-Directional Time Sync Handshake only)

$$T_{pd}[n] = \frac{(t_4[n] - t_1[n]) - (t_3[n] - t_2[n])}{2}$$

Equation 7-6. Time Offset that DFP Sends in Follow-Up Packet

$$O(D_i, H_i)[n] = O_{last}(D_i, H_i) - (t_3[n] - t_{last}(D_i)[n]) \cdot \frac{F(D_i, H_i)[N]}{2^{41} + F(D_i, H_i)[N]}$$

where:

- $O_{last}(D_i, H_i)$ is the most recent computed time offset from Host Router by the DFP (see Equation 7-7).
- $t_{last}(D_i)[n]$ is the most recent t_4 time stamp in the DFP when the last Time Sync Handshake was performed at the UFP of the Router.
- $F(D_i, H_i)[N]$ is the most recent frequency offset from the Host Router computed by the DFP (see Equation 7-4).

Note: The value of $F(D_i, H_i)[N]$ is the same as $F(U_i, H_i)[N]$ in the same Router.

Equation 7-7. Time Offset from DFP for Bi-Directional Time Sync Handshake

$$O(U_i, D_i)[n] = t_3[n] - t_4[n] + T_{pd}[n]$$

Equation 7-8. Time Offset from DFP for Uni-Directional Time Sync Handshake

$$O(U_i, D_i)[n] = t_3[n] - t_4[n] + Const$$

where *Const* is implementation specific (helps to minimize time offset caused by Uni-Directional Time Sync Handshake where propagation delay is not computed).

Equation 7-9. Time Offset from Host Router

$$O(U_i, H_i)[n] = O(U_i, D_i)[n] + O(D_i, H_i)[n]$$

Equation 7-10. Host Router Time at Any Given Time (t_{now})

$$t_h = t_{last}(U_i)[n] + O(U_i, H_i)[n] + (t_{now} - t_{last}(U_i)[n]) \cdot \frac{2^{41}}{2^{41} + F(U_i, H_i)[n]}$$

where $F(U_i, H_i)[N]$ is the frequency offset of the UFP from the Host Router.

Note: Equation 7-10 is a subset of Equation 7-17 (Inter-Domain Host Router Time).

7.4.2 Inter-Domain Equations

The purpose of an Inter-Domain Time Sync Handshake is to determine the time and frequency offsets from the Inter-Domain Host Router. Using this information, a Router in any Domain can compute the Inter-Domain Host Router time at any moment.

A Domain that contains an IDTI Port performs the following steps to compute the Inter-Domain Host Router time:

1. Compute time synchronization parameters the same as a UFP in a single Domain using Equation 7-9.

Note: Inter-Domain computations are done in parallel to the Intra-Domain computations and the two processes are independent of each other, although some parameters from the Intra-Domain computations are used in Inter-Domain time synchronization.

2. Compute the following three parameters:
 - Inter-Domain time stamp (Section 7.4.2.1).
 - Inter-Domain frequency offset (Section 7.4.2.2).
 - Inter-Domain time offset (Section 7.4.2.3).
3. Use the Parameters in Step 2 to compute the Inter-Domain Host Router time (Section 7.4.2.4).

Note: In this section, t_1 , t_2 , t_3 and t_4 refer to the timestamps from the handshake between an IDTI Port and an IDTR Port.

7.4.2.1 Inter-Domain Time Stamp Computation

The Inter-Domain Time Stamp is a translated $t_4[n]$ time stamp from the local clock of a Router (S_i) to the time units of Domain i . The computations in this section are used to convert time stamps taken at a Router's local clock into the equivalent Host Router time.

The following formula is used to compute an Inter-Domain time stamp.

Equation 7-11. Inter-Domain Time Stamp

$$t_{last-id}[n] = t_{last}(U_i)[n] + O(U_i, H_i)[n] + (t_4[n] - t_{last}(U_i)[n]) \cdot \frac{2^{41}}{2^{41} + F(U_i, H_i)[N]}$$

where:

- i represents the current domain.
- $t_{last}(U_i)[n]$ is the time stamp at the instant the last computation of the IDTI Port was executed, which is the t_4 time stamp in the UFP time sync handshake.
- $t_4[n]$ is the Time stamp from the Inter-Domain Time Sync Handshake.
- $O(U_i, H_i)[n]$ is the most updated time offset computed by this Router in its own Domain i .

7.4.2.2 Inter-Domain Frequency Offset Computation.

The frequency ratio between the Host Router in the Domain with the IDTI Port and the Inter-Domain Time Source is computed using the following formula.

Equation 7-12. Frequency Ratio between the Host Router of the Domain with the IDTI Port and the Inter-Domain Time Source

$$f(H_i, H_g)[N] = \frac{f(U_i, D_j)[N] \cdot f(D_j, H_j)[N] \cdot f(H_j, H_g)[N]}{f(U_i, H_i)[N]}$$

where:

- $f(U_i, H_i)[N]$ is the most recent frequency ratio computed as part of the Intra-Domain flow (Equation 7-3).
- $f(U_i, D_j)[N]$ is the frequency ratio between the Local Clock of the IDTR Port and the Local Clock of the IDTI Port, computed as part of the Inter-Domain flow using Equation 7-3.
- $f(D_j, H_j)[N]$ is the frequency ratio derived from the *FreqOffsetFromHR* field in the last Follow-Up Packet sent by the IDTR Port via the Inter-Domain Link.
- $f(H_j, H_g)[N]$ is the frequency ratio derived from the *FreqOffsetFromInterDomainHR* field in the last Follow-Up Packet sent by the IDTR Port via the Inter-Domain Link.

Equation 7-13. Frequency Offset of Inter-Domain from Inter-Domain Host Router

$$F(H_i, H_g)[N] = \{f(H_i, H_g)[N] - 1\} \cdot 2^{41}$$

7.4.2.3 Inter-Domain Time Offset Computation

The time offset of the IDTR Port that is sent in the Follow-Up Packet at the end of an Inter-Domain Time Sync Handshake is computed using the following equation.

Equation 7-14. Time Offset from the IDTR Port

$$\phi_j[n] = O(D_j, H_j)[n] - (t_3[n] - t_{last}(U_j)[n]) \cdot \frac{F(D_j, U_j)[n]}{2^{41} + F(D_j, H_j)[n]}$$

where:

- $F(D_j, U_j)[N]$ is the frequency offset of the IDTR Port in its own Domain.
- $O(D_j, H_j)[n]$ is the most recent computed time offset from Host Router by the DFP (see Equation 7-7) of the IDTR Port in its own Domain.
- $t_3[n]$ is the time stamp of the IDTR Port during the current Inter-Domain Time Sync Handshake.
- $t_{last}(U_i)[n]$ is the last computation time stamp of the IDTR Port in its own Domain as part of the UFP calculations.

The time offset between the IDTI Port and the Inter-Domain Time Source is computed using the following equation.

Equation 7-15. Time Offset between the IDTI Port and Inter-Domain Time Source

$$O(D_i, H_j)[n] = O(U_i, D_j)[n] + \phi_j[n]$$

where:

- $O(U_i, D_j)[n]$ is the time offset from the IDTR Port as computed by the IDTR Port at the end of an Inter-Domain Time Sync Handshake (Equation 7-7).
- $\phi_j[n]$ is the Inter-Domain time offset sent by the IDTR Port in the Follow-Up Packet (TimeOffsetFromHR).

The time offset between the Host Router of the Domain with the IDTI Port and the Inter-Domain Time Source is computed using the following equation.

Equation 7-16. Time Offset between the Host Router of the Domain with the IDTI Port and the Inter-Domain Time Source

$$O(D_i, H_j)[n] + O(H_j, H_g)[n] - O(U_i, H_i)[n] + \left((t_4[n] - t_{last}(U_i)[n]) \cdot \frac{F(U_i, H_i)[N]}{2^{41} + F(U_i, H_i)[N]} \right) - \left((t_4[n] - t_{last-id}[n] + O(U_i, H_j)[n]) \cdot \frac{F(H_j, H_g)[N]}{2^{41} + F(H_j, H_g)[N]} \right)$$

where:

- $F(U_i, H_i)[N]$ is the frequency offset from the Host Router in the Domain of the IDTI Port.
- $F(H_j, H_g)[N]$ is the frequency offset from the Inter-Domain Time Source as sent by the IDTR Port in the Follow-Up Packet (*FreqOffsetFromInterDomainHR*).
- $t_{last}(U_i)[n]$ is the time stamp of last computation of the IDTI Port in its own Domain.
- $t_{last-id}[n]$ is the Inter-Domain time stamp supplied in the Follow-Up Packet (*IDTimeStamp*).
- $t_4[n]$ is the time stamp of the current Inter-Domain Time Sync Handshake between the IDTI Port and the IDTR Port.
- $O(U_i, H_i)[n]$ is the time offset from the Host Router in the Domain of the IDTI Port.
- $O(H_j, H_g)[n]$ is the time offset from the Inter-Domain Time Source as sent by the IDTR Port in a Follow-Up Packet (*TimeOffsetFromInterDomainHR*).
- $O(D_i, H_j)[n]$ is the time offset from the Inter-Domain Time Source as computed in Equation 7-15.

7.4.2.4 Inter-Domain Host Router Time Computation

A Router can compute the Inter-Domain Host Router Time (t_{idg}) using the following equation.

Equation 7-17. Inter-Domain Host Router Time

$$t_{idg} = \frac{t_{last-id}[n] + O(H_i, H_g)[n] + \frac{t_{last}(U_i)[n] + O(U_i, H_i)[n] + \left\{ \frac{t_{now} - t_{last}(U_i)[n]}{1 + \frac{F(U_i, H_i)[N]}{2^{41}}} \right\} - t_{last-id}[n]}{1 + \frac{F(H_i, H_g)[N]}{2^{41}}}}$$

where:

- t_{now} is the time at which the computation is being performed.
- $t_{last-id}[n]$ is the Inter-Domain time stamp in the *IDTimeStamp* field of the last received Follow-Up Packet.
- $O(H_i, H_g)[n]$ is that last computed Inter-Domain time offset from Inter-Domain Host Router.
- $t_{last}(U_i)[n]$ is the time stamp of the last UFP computation in the current Router (t_4).
- $O(U_i, H_i)[n]$ is the most updated time offset from the Host Router of the current Router's Domain.
- $F(U_i, H_i)[N]$ is the frequency offset from the Host Router computed in the current Router.
- $F(H_i, H_g)[N]$ is the frequency offset from the Inter-Domain Host Router in the *FreqOffsetFromInterDomainHR* field of the last received Follow-Up Packet.

If the IDTI Port is part of a Host Router, all the components for Equation 7-17 are computed in Host Router (i.e. Equation 7-15 and Equation 7-16 are not forwarded in the Follow-Up Packet).

7.4.3 Filtering

In order to achieve the required accuracy, low pass filters are applied on some of the computations. The parameters that need to be filtered are Propagation Delay (Tpd), Time Offsets, and Frequency Offset. The recommended filter is IIR filter as follows:

Equation 7-18. IIR Low Pass Filter


$$Y[n] = \left(1 - \frac{1}{2^p}\right) \cdot Y[n-1] + \frac{1}{2^p} \cdot X[n]$$

where:

- X[n] is the new input to filter.
- Y[n-1] is the previous value of the filter.
- Y[n] is the new value of the filter.
- p is the filter strength.

As shown in Figure 7-16, applying an IIR low pass filter on running counter-like time stamps results in attenuation of the value.

Figure 7-16. Filter Attenuation



In order to compensate for filter attenuation, the error value is computed as follows and passed through IIR low pass filter:

$$E[n] = X[n] - Y[n]$$

$$E_{filtered}[n] = IIR(E[n])$$

The filtered error is then added back to the filtered time stamp value.

The values in the following registers in Router Configuration Space can be used for filtering time stamps, error, frequency offset, and propagation delay:

- FreqAvgConst.
- DelayAvgConst.
- OffsetAvgConst.
- ErrorAvgConst.

It is recommended that a USB4 Port apply IIR filters with strength (p) of at least 4.

7.5 Time Synchronization Accuracy Requirements

Time synchronization accuracy is measured in two ways: Paired Measurement and Standalone Measurement. Section 7.5.1 describes Paired Measurement. Section 7.5.2 describes Standalone Measurement.

When a Router is configured for HiFi Mode Time Sync Handshakes, time synchronization accuracy is $t_{HiAccuracy}$. When a Router is configured for LowRes Mode Time Sync Handshakes, time synchronization accuracy is $t_{LowAccuracy}$.


A Router shall reach the required time synchronization accuracy within $t_{ConvergeTime}$ after Time Sync Handshakes are enabled.

If Time Sync Handshakes are paused, a UFP shall reach the required time synchronization accuracy within $t_{ConvergeTime}$ after receiving a Delay Response Packet.

7.5.1 Paired Measurement

Paired Measurement is performed between two Routers (either in the same Domain or different Domains) that participate in a Time Sync Handshake. There are two parameters that affect time synchronization accuracy: Static Offset and Dynamic Noise. Figure 7-17 illustrates these parameters.

Figure 7-17. Dynamic Noise Types



For Bi-Directional Time Sync Handshake, the Static Offset between two Routers connected to one another shall not add more than 8 ns per Link.


For Uni-Directional Time Sync Handshake, the Static Offset depends on cable length and the amount of Re-timers on the Link.

In HiFi Mode, the Dynamic Noise between two Routers connected to one another shall not add more than $t_{HiAccuracy}$. In LowRes Mode, the Dynamic Noise between two Routers connected to one another shall not add more than $t_{LowAccuracy}$.

7.5.2 Standalone Measurement

In this method the time noise is measured compared to an ideal line, extrapolated from the measurement points shown in Figure 7-18.

Figure 7-18. Standalone Measurement Points



In HiFi Mode, the noise between measurement points and the ideal line shall not be more than $t_{HiAccuracy}$. In LowRes Mode, the noise between measurement points and the ideal line shall not be more than $t_{LowAccuracy}$.

Note: This measurement is also used to analyze the frequency response of the noise during compliance testing.

In order to achieve the goal of 1 ns static noise, the budget for TxTimeToWire uncertainty shall not be more than 6.4 ns dynamic noise with no more than 6.4 ns static offset that changes after each power up. RxTimeToWire shall not be more than 12.8 ns dynamic noise with no more than 6.4 ns static offset that changes after each power up.

Dynamic uncertainty in TxTimeToWire and RxTimeToWire shall be filtered out during calculations by IRR or other filtering.

7.5.3 Measuring Method


When talking about time measurement, it is important to define the points when the time measurement is done. The computed Host Router Time is the most accurate on the computation edge of the clock.

As shown in Figure 7-19, placing a Time Measurement Event at a predictable and consistent distance from a Time Computation Event helps reduce noise to the time measurement.

Sending a Serial Time Link Protocol (STLP) Packet is an example of a Time Measurement Event. An STLP Packet contains the current Host Router Time at the point of transmission of the first bit of the STLP Packet. See Appendix D for more information on the Serial Time Link Protocol.


External events that cannot be predicted and need to use the Host Router time for timestamps (i.e. ITP, PTM), might have lower accuracy than internally generated events, depending on the method of maintaining the Host Router time and the method of timestamp management.

Figure 7-19. Time Events



The recommended setup for performing time measurements is shown in Figure 7-20.

Figure 7-20. Measuring Method



7.5.4 Accuracy Parameters

Table 7-7 lists the accuracy parameters for time synchronization.

Table 7-7. Time Synchronization Accuracy Parameters

| Parameter | Description | Min | Max | Units |
|---------------|--|-----|-----|-------|
| tHiAccuracy | Time synchronization accuracy in HiFi Mode. | -- | 1 | ns |
| tLowAccuracy | Time synchronization accuracy in LowRes Mode. | -- | 4 | ns |
| tConvergeTime | Time, after Time Sync Handshakes are enabled, to reach the required time synchronization accuracy. | -- | 50 | ms |

7.6 Software Configuration

7.6.1 Intra-Domain Time Synchronization Setup

Within a single Domain, the Time Synchronization Protocol is disabled by default.

A Router that is hot-plugged to a Domain shall only initiate the Time Sync Handshake on its UFP when the following conditions are true:

- The physical layer has established the Link between the hot-plugged Router and the Domain.
- The *TSPacketInterval* field in the TMU_RTR_CS_3 register in Router Configuration space is greater than 0.

7.6.2 Inter-Domain Time Synchronization Setup

When multiple Domains are connected together, the Time Synchronization Protocol is explicitly enabled by the Connection Manager.



CONNECTION MANAGER NOTE

A Connection Manager enables Inter-Domain time synchronization by:

1. *Selecting one Domain to act as the Inter-Domain Time Source. The Local Clock of the Host Router of the selected Domain becomes the Inter-Domain Time Source Clock.*
2. *Creating an Inter-Domain time synchronization hierarchy by creating a Spanning Tree of inter-connected Domains.*
3. *Enabling the Time Synchronization Protocol across each Inter-Domain Link in the Spanning Tree of Inter-connected Domains by setting the IDTR bit to 1b at one end of a Link and the IDTR bit to 1b at the other end of the Link.*
4. *Updating the LocalTime register of each Router in each Domain with an IDTI Port to the new Inter-Domain Time Source using the time posting registers.*
5. *Enabling the processing of Inter-Domain Time Stamp Packets by setting the Inter-Domain Enable register to 1 in the TMU Router Configuration Capability of the Host Router in each Domain with a ITDI Port.*

7.6.3 Post Time Mechanism

The Post Time Mechanism is used to update the local time of a Router. A Router shall update its local time when the Post Time is greater than zero and is less than or equal to the *Nanoseconds* field in the Host Router Time Register. Time Posting shall only be activated when the Post Time High register is written. The Router updates its local time by writing the Post Local Time to the *Nanoseconds* field of its LocalTime register.

The Post Time is set by the Connection Manager in the *Post Time High* and *Post Time Low* fields in the TMU Router Configuration Capability in Router Configuration Space. The Post Local Time

is set by the Connection Manager in the *Post Local Time Low* and *Post Local Time High* fields in the TMU Router Configuration Capability in Router Configuration Space.



CONNECTION MANAGER NOTE

The following is an example of how a Connection Manager uses the Post Time Mechanism to update a Router's local time:

1. *The Connection Manager reads the LocalTime register in the Host Router of the Domain (if doing Intra-Domain Time Sync) or the Inter-Domain Host Router (if doing Inter-Domain Time Sync).*
2. *The Connection Manager writes the Nanoseconds field of the Host Router Time to the Post Local Time registers in the TMU Router Configuration Capability of the Router.*
3. *The Connection Manager writes a value of 0x1 to the Post Time field in the TMU Router Configuration Capability of the Router, which causes the Router to update its local time to the value in the Post Local Time registers.*
4. *The Connection Manager periodically reads 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.*



CONNECTION MANAGER NOTE

When a Connection Manager receives a Hot Plug Event Packet for a Device Router, it uses the Post Time Mechanism to update the local time of the Device Router. It is recommended that a Connection Manager update the local time in the Device Router before enabling Time Synchronization in order to minimize time disruptions.

7.6.4 Time Disruption Bit

There are several system events that can cause a TMU to experience disruption in time. A Connection Manager uses the *Time Disruption* bit in Router Configuration Space to indicate to the Router when a time disruption event is occurring.



CONNECTION MANAGER NOTE

A Connection Manager shall set the Time Disruption bit in Router Configuration Space to 1b before any of the following time disruption events occur:

- *TMU mode changes (TSPacketInterval, Direction, Filters, Frequency Measurement Window).*
- *Inter-Domain time sync is enabled.*
- *Time Posting is applied.*

After the disruptive behavior has passed, the Connection Manager shall clear the Time Disruption bit to 0b to indicate that time is reliable now.

8 Configuration Spaces

This chapter provides a detailed description of the configuration registers within the Configuration Spaces defined by the USB4™ architecture. The Configuration Spaces are accessed via Read and Write Request packets as defined in Section 6.4.3.3.

8.1 Configuration Fields Access Types

Table 8-1 defines the access types that are allowed for a particular configuration register field.

Table 8-1. Configuration Register Fields Access Types

| Access Type | Description |
|-------------|--|
| R/W | Read/Write. A field with this access type shall be capable of both read and write operations. The value read from this field shall reflect the last value written to it unless the field was reset in the interim. |
| R/W S | Read/Write Status. A field with this access type shall be capable of both read and write operations. The value read from this field may or may not reflect the last value written. |
| RO | Read Only. A write to a field with this access type shall have no effect. A read shall return a meaningful value. |
| R/Clr | Read Clear. A field with this access type shall be cleared to 0 after it is read. A write to a field with this attribute shall have no effect on its value. |
| W/Clr | Write Clear. A field with this access type shall be cleared to 0 after it is written to. A read shall return a meaningful value. |
| R/W SC | Read/Write Self Clearing. When set to 1b a field with this access type causes an action to be initiated. A field with this attribute shall read as 0b after the action is complete. The value returned prior to completion of the action is vendor defined. |
| Rsvd | Reserved. Reserved for future implementation. A write to this field shall have no effect. Unless defined otherwise, a read shall return 0. |
| RsvdZ | Reserved and Zero. Reserved for future implementation. A read shall return 0. |
| VD | Vendor Defined. A Vendor may put any value in this field. |



CONNECTION MANAGER NOTE

A Connection Manager shall not change the value in a Vendor Defined field. A Connection Manager shall read a Vendor Defined field before writing to it, so it can write the correct value.

A Connection Manager shall only write 0 to a field marked as RsvdZ.

8.2 Configuration Spaces

A Router shall implement the following Configuration Spaces:

- **Router Configuration Space:** This Configuration Space contains configuration information at a Router level. See Section 8.2.1 for details of the registers that make up the Router Configuration Space.
- **Adapter Configuration Space:** This Configuration Space contains configuration information at the Adapter level. Each Adapter other than the Control Adapter has its own Adapter Configuration Space. See Section 8.2.2 for details of the registers that make up the Adapter Configuration Space.
- **Path Configuration Space:** This Configuration Space contains configuration information at the Path level. Each Adapter other than the Control Adapter has its own Path Configuration Space with an entry in it for each of its Paths. See Section 8.2.3 for details of the registers that make the Path Configuration Space.

A Router may also optionally implement the following Configuration Space:

- **Counters Configuration Space:** This Configuration Space contains performance statistics information at the Adapter level. See Section 8.2.4 for details of the registers that make the Counters Configuration Space. The Counters Configuration Space (CCS) Flag in the Adapter Configuration Space indicates whether an Adapter contains a Counters Configuration Space.

Note: The DW columns in the tables in this chapter refer to the Doubleword index of the register. In the case of a Capability, the index is relative to the beginning of the Capability.

All fields in a Configuration Space that are not Read Only (RO) shall contain their Default Values until a different value is written by a Connection Manager. Default Values are defined by the tables in this Chapter. A value of "Vendor Defined" in the Default Value column means that the default value is vendor defined and may vary.

8.2.1 Router Configuration Space

A Router Configuration Space shall have the format and contain the register fields depicted in Figure 8-1. The first set of Doublewords in Figure 8-1 describe the basic attributes of a Router. The rest of the Configuration Space is populated with a linked list of Capabilities.

Figure 8-1. Structure of the Router Configuration Space




Table 8-2 lists the Capabilities supported by the Router Configuration Space. A Capability listed as “Required” shall be present in Router Configuration Space. A Capability listed as “Optional” may be present in Router Configuration Space.

Table 8-2. List of Router Configuration Capabilities

| Capability | Required / Optional | Capability ID |
|--------------------------|---------------------|---------------|
| TMU Router Configuration | Required | 03h |
| Vendor Specific | Optional | 05h |

There are two types of Vendor Specific Configuration Capabilities: the Vendor Specific Capability and the Vendor Specific Extended Capability. The number of Vendor Specific Configuration Capabilities and the length of each Vendor Specific Configuration Capability is implementation specific. A Router that implements Vendor Specific Configuration Capabilities shall not depend on a Connection Manager's support for the Vendor Specific Configuration Capabilities.

Section 8.2.1.3 defines the required fields for a Vendor Specific Capability. Section 8.2.1.4 defines the required fields for a Vendor Specific Extended Capability.

Capabilities shall be linked in the following order:

- Required Capabilities.
- Optional Capabilities.
- Vendor Specific Capabilities.
- Vendor Specific Extended Capabilities.

8.2.1.1 Basic Configuration Registers

The registers in Router Configuration Space shall have the structure and fields described in Table 8-3.

Table 8-3. Router Configuration Space Basic Attributes

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|--|------|----------------|
| 0 | ROUTER_CS_0 | 15:0 | Vendor ID This field shall identify the manufacturer of the Router silicon. It is assigned by the USB-IF. | RO | Vendor ID |
| | | 31:16 | Product ID This field shall contain a value that is assigned by the manufacturer of the Router silicon to identify the type of the Router. | RO | Vendor Defined |
| 1 | ROUTER_CS_1 | 7:0 | Next Capability Pointer This field shall contain the Doubleword index of the first Capability in the Router Configuration Space. | RO | Vendor Defined |
| | | 13:8 | Upstream Adapter This field contains the Adapter number of the Adapter that routes Control Packets to and from the Connection Manager. For a Host Router, the Host Interface Adapter is the Upstream Adapter. For a Device Router, the Upstream Adapter is the Lane 0 Adapter in the Upstream Facing Port. | RO | Vendor Defined |
| | | 19:14 | Max Adapter This shall contain the Adapter number of the highest numbered Adapter in the Router. | RO | Vendor Defined |
| | | 22:20 | Depth This field contains the number of hops from the Router to the Host Router in the Spanning Tree topology. A Router shall support Depths up to and including 5. | R/W | 0 |
| 1 | ROUTER_CS_1 | 23 | Reserved | Rsvd | 0 |
| | | 31:24 | Revision Number This field shall contain the value assigned by the manufacturer to identify the revision number of the Router. | RO | Vendor Defined |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value | | | | | | | | | |
|-------------------|---|-------------------------------|--|-------------------------------|-------------------------------|------------------------------|-------------------|--------------|------|--------------|----|-----|-----|---|
| 2 | ROUTER_CS_2 | 31:0 | TopologyID Low This field contains the least significant 32 bits of the TopologyID assigned to the Router by the Connection Manager. | R/W | 0 | | | | | | | | | |
| 3 | ROUTER_CS_3 | 23:0 | TopologyID High This field contains the most significant 24 bits of the TopologyID assigned to the Router by the Connection Manager. | R/W | 0 | | | | | | | | | |
| | | 30:24 | Reserved | Rsvd | 0 | | | | | | | | | |
| | | 31 | TopologyID Valid (V) This bit indicates whether or not the Router is enumerated. 0b – Router is not enumerated 1b – Router is enumerated | R/W | 0 | | | | | | | | | |
| 4 | ROUTER_CS_4 | 7:0 | Notification Timeout This field contains the timeout value in milliseconds used by the Router to retry Hot Plug Event Packets and packets that require a Notification Acknowledgment. | R/W | 0Ah | | | | | | | | | |
| | | 15:8 | Connection Manager USB4 Version (CMUV) This field identifies which version of the USB4 specification is supported by the Connection Manager where: <table border="1" data-bbox="568 965 1129 1151"> <thead> <tr> <th></th> <th>Major Version (Bits 15:12)</th> <th>Minor Version (Bits 11:8)</th> </tr> </thead> <tbody> <tr> <td>Reserved for TBT3</td> <td>0</td> <td>x</td> </tr> <tr> <td>USB4 Rev 1.0</td> <td>1</td> <td>0</td> </tr> </tbody> </table> All other values are reserved. | | Major Version (Bits 15:12) | Minor Version (Bits 11:8) | Reserved for TBT3 | 0 | x | USB4 Rev 1.0 | 1 | 0 | R/W | 0 |
| | | | Major Version (Bits 15:12) | Minor Version (Bits 11:8) | | | | | | | | | | |
| | | Reserved for TBT3 | 0 | x | | | | | | | | | | |
| | | USB4 Rev 1.0 | 1 | 0 | | | | | | | | | | |
| 23:16 | Reserved | Rsvd | 0 | | | | | | | | | | | |
| 31:24 | USB4 Version This field identifies which version of the USB4 specification is supported by the Router where: <table border="1" data-bbox="568 1368 1129 1554"> <thead> <tr> <th></th> <th>Major Version (Bits 31:29)</th> <th>Minor Version (Bits 28:24)</th> </tr> </thead> <tbody> <tr> <td>Reserved for TBT3</td> <td>000b</td> <td>Not applicable</td> </tr> <tr> <td>USB4 Rev 1.0</td> <td>001b</td> <td>00000b</td> </tr> </tbody> </table> All other values are reserved. A Router shall set this field to 20h. | | Major Version (Bits 31:29) | Minor Version (Bits 28:24) | Reserved for TBT3 | 000b | Not applicable | USB4 Rev 1.0 | 001b | 00000b | RO | 20h | | |
| | Major Version (Bits 31:29) | Minor Version (Bits 28:24) | | | | | | | | | | | | |
| Reserved for TBT3 | 000b | Not applicable | | | | | | | | | | | | |
| USB4 Rev 1.0 | 001b | 00000b | | | | | | | | | | | | |
| 0 | Enter Sleep (SLP) The Connection Managers sets this bit to 1b to transition the Router into Sleep state. | R/W | 0 | | | | | | | | | | | |
| 1 | Enable Wake on PCIe (WoP) When set to 1b, a PCIe Wake indication from a PCIe device connected to a PCIe downstream port causes the Router to exit from sleep. When set to 0b, a PCIe Wake indication from a PCIe device connected to a PCIe downstream port does not cause the Router to exit from sleep. A Host Router shall hardwire this bit to 0b. | R/W | 0 | | | | | | | | | | | |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|------|---------------|
| 5 | ROUTER_CS_5 | 2 | Enable Wake on USB3 (WoU) When set to 1b, a USB3 Wake indication from a USB device causes the Router to exit from sleep. When set to 0b, a USB3 Wake indication from a USB device does not cause the Router to exit from sleep. A Host Router shall hardwire this bit to 0b. | R/W | 0 |
| | | 3 | Enable Wake on DP (WoD) When set to 1b, an HPD change or reception of an HPD IRQ causes the Router to exit from sleep. When set to 0b, an HPD change or reception of an HPD IRQ does not cause the Router to exit from sleep. A Host Router shall hardwire this bit to 0b. | R/W | 0 |
| | | 22:4 | Reserved | Rsvd | 0 |
| | | 23 | CM TBT3 Not Supported (CNS) This bit identifies if a USB4 Connection Manager is TBT3-Compatible: 0b – TBT3-Compatible 1b – Not TBT3-Compatible A TBT3 Connection Manager does not write to this bit. A Router shall ignore this bit if the <i>Configuration Valid</i> bit is set to 0b. This field does not apply to Host Routers. | R/W | 0b |
| | | 24 | PCIe Tunneling On (PTO) A Connection Manager uses this bit to let a Device Router know that it intends to enable PCIe Tunneling: 0b – PCIe tunneling will not be enabled 1b – PCIe tunneling will be enabled A Device Router shall ignore this bit if the <i>Configuration Valid</i> bit is set to 0b. This field does not apply to Host Routers. | R/W | 0b |
| | | 25 | USB3 Tunneling On (UTO) A Connection Manager uses this bit to let a Device Router know that it intends to enable USB3 Tunneling: 0b – USB3 tunneling will not be enabled 1b – USB3 tunneling will be enabled A Device Router shall ignore this bit if the <i>Configuration Valid</i> bit is set to 0b. This field does not apply to Host Routers. | R/W | 0b |
| | | 26 | Internal Host Controller On (HCO) A Connection Manager uses this bit to let a Device Router know that it intends to enable the internal host controller: 0b – Internal host controller will not be enabled 1b – Internal host controller will be enabled A Device Router shall ignore this bit if the <i>Configuration Valid</i> bit is set to 0b. This field does not apply to Host Routers. A Connection Manager shall set this bit to 0b if the USB4 Upstream operates at USB4 mode and may set it to 1b if the USB4 Upstream Port operates at TBT3 Mode. | R/W | 0b |
| | | 30:27 | Reserved | Rsvd | 0 |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|------|----------------|
| | | 31 | <p>Configuration Valid (CV) A Connection Manager uses this bit to let a Device Router know whether or not the <i>PTO</i>, <i>UTO</i>, and <i>HCO</i> fields are valid: 0b – <i>PTO</i>, <i>UTO</i>, and <i>HCO</i> fields are not valid 1b – <i>PTO</i>, <i>UTO</i>, and <i>HCO</i> fields are valid This field does not apply to Host Routers.</p> | R/W | 0 |
| 6 | ROUTER_CS_6 | 0 | <p>Sleep Ready (SLPR) A Router sets this bit to 1b after the Enter Sleep bit is set to 1b and the Router is Ready for a sleep event. A Router sets this bit to 0b after entering Sleep state.</p> | RO | 0 |
| | | 1 | <p>TBT3 Not Supported (TNS) A Router shall set this field to 1b if it does not support the TBT3-Compatible behavior defined in Chapter 13. A Router shall set this field to 0b if it supports the TBT3 Compatible behavior defined in Chapter 13.</p> | RO | Vendor Defined |
| | | 2 | <p>Wake on PCIe Status A Router shall set this bit to 1b when a PCIe Wake indication from a PCIe device connected to a PCIe downstream port causes the Router to exit from sleep. A Router shall set this bit to 0b upon entry to sleep.</p> | RO | 0b |
| | | 3 | <p>Wake on USB3 Status A Router shall set this bit to 1b when a USB3 Wake indication causes the Router to exit from sleep. A Router shall set this bit to 0b upon entry to sleep.</p> | RO | 0b |
| | | 4 | <p>Wake on DP Status A Router shall set this bit to 1b when a DP Wake indication causes the Router to exit from sleep. A Router shall set this bit to 0b upon entry to sleep.</p> | RO | 0b |
| | | 17:5 | Reserved | Rsvd | 0 |
| | | 18 | <p>Internal Host Controller Implemented (HCI) A Router shall set this bit to 0b if it does not implement an internal host controller. A Router shall set this bit to 1b if it implements an internal host controller. This field does not apply to Host Routers.</p> | RO | Vendor Defined |
| | | 23:19 | Reserved | Rsvd | 0 |
| | | 24 | <p>Router Ready (RR) A Router sets this bit to 1b after configuring any hardware based on the value of the <i>Connection Manager USB4 Version</i> field. See Section 6.7 and Section 13.4.1.</p> | RO | 0 |
| | | 25 | <p>Configuration Ready (CR) A Device Router shall set this bit to 1b when it is ready for the Protocol Tunneling enabled by the Connection Manager. A Host Router shall not set this bit to 1b.</p> | RO | 0 |
| | | 31:26 | Reserved | Rsvd | 0 |
| 7 | ROUTER_CS_7 | 31:0 | <p>UUID (High) This field contains bits 63:32 of the UUID value described in Section 8.2.1.1.1.</p> | RO | Vendor Defined |
| 8 | ROUTER_CS_8 | 31:0 | <p>UUID (Low) This field contains bits 31:0 of the UUID value described in Section 8.2.1.1.1.</p> | RO | Vendor Defined |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|--|------|---------------|
| 9 | ROUTER_CS_9 | 31:0 | Data[0] This field contains a DW optionally written or read by the Connection Manager as part of a Router Operation. | R/W | 0 |
| 10 | ROUTER_CS_10 | 31:0 | Data[1] This field contains a DW optionally written or read by the Connection Manager as part of a Router Operation. | R/W | 0 |
| 11 | ROUTER_CS_11 | 31:0 | Data[2] This field contains a DW optionally written or read by the Connection Manager as part of a Router Operation. | R/W | 0 |
| 12 | ROUTER_CS_12 | 31:0 | Data[3] This field contains a DW optionally written or read by the Connection Manager as part of a Router Operation. | R/W | 0 |
| 13 | ROUTER_CS_13 | 31:0 | Data[4] This field contains a DW optionally written or read by the Connection Manager as part of a Router Operation. | R/W | 0 |
| 14 | ROUTER_CS_14 | 31:0 | Data[5] This field contains a DW optionally written or read by the Connection Manager as part of a Router Operation. | R/W | 0 |
| 15 | ROUTER_CS_15 | 31:0 | Data[6] This field contains a DW optionally written or read by the Connection Manager as part of a Router Operation. | R/W | 0 |
| 16 | ROUTER_CS_16 | 31:0 | Data[7] This field contains a DW optionally written or read by the Connection Manager as part of a Router Operation. | R/W | 0 |
| 17 | ROUTER_CS_17 | 31:0 | Data[8] This field contains a DW optionally written or read by the Connection Manager as part of a Router Operation. | R/W | 0 |
| 18 | ROUTER_CS_18 | 31:0 | Data[9] This field contains a DW optionally written or read by the Connection Manager as part of a Router Operation. | R/W | 0 |
| 19 | ROUTER_CS_19 | 31:0 | Data[10] This field contains a DW optionally written or read by the Connection Manager as part of a Router Operation. | R/W | 0 |
| 20 | ROUTER_CS_20 | 31:0 | Data[11] This field contains a DW optionally written or read by the Connection Manager as part of a Router Operation. | R/W | 0 |
| 21 | ROUTER_CS_21 | 31:0 | Data[12] This field contains a DW optionally written or read by the Connection Manager as part of a Router Operation. | R/W | 0 |
| 22 | ROUTER_CS_22 | 31:0 | Data[13] This field contains a DW optionally written or read by the Connection Manager as part of a Router Operation. | R/W | 0 |
| 23 | ROUTER_CS_23 | 31:0 | Data[14] This field contains a DW optionally written or read by the Connection Manager as part of a Router Operation. | R/W | 0 |
| 24 | ROUTER_CS_24 | 31:0 | Data[15] This field contains a DW optionally written or read by the Connection Manager as part of a Router Operation. | R/W | 0 |
| 25 | ROUTER_CS_25 | 31:0 | Metadata A Connection Manager uses this field to pass Router Operation metadata to the Router. The contents of this field vary with the Router Operation type and are defined in the Section 8.3.1. | R/W | 0 |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|--|------|---------------|
| 26 | ROUTER_CS_26 | 15:0 | Opcode A Connection Manager uses this field to indicate the Router Operation to be executed. See Table 8-22 for the list of supported Router Operations. | R/W | 0 |
| | | 23:16 | Reserved | Rsvd | 0 |
| | | 29:24 | Status A Router uses this field to indicate the status of a Router Operation after completion. The values for this field vary with Router Operation type and are defined in the Section 8.3.1. | R/W | 0 |
| | | 30 | Operation Not Supported (ONS) A Router uses this field to indicate whether or not the initiated Router Operation is supported: 0 - Supported 1 - Not Supported | R/W | 0 |
| | | 31 | Operation Valid (OV) This field is set to 1b by the Connection Manager to indicate that a Router Operation is posted for processing. A Router shall process the Router Operation in the <i>Opcode</i> field when the value in this field changes from 0b to 1b. A Router shall set this field to 0b after it finishes processing the Router Operation. | R/W | 0b |



CONNECTION MANAGER NOTE

When writing to Router Configuration Space, a Connection Manager shall abide by the following rules:

- *When writing to the Router Configuration Space of a Host Router, a Connection Manager shall not change the value in the following fields:*
 - *PCIe Tunneling On.*
 - *USB3 Tunneling On.*
 - *Internal Host Controller On.*
 - *Configuration Valid.*
- *The Connection Manager shall read the Upstream Adapter field before writing to DW1 of Router Configuration Space (ROUTER_CS_1). When the Connection Manager writes to the Upstream Adapter field, it needs to write the same value that was read so that the value in the Upstream Adapter field does not change.*
- *The Connection Manager shall write 0 in the Depth field of a Host Router.*
- *The Connection Manager shall not set the Notification Timeout field to 0.*
- *The Connection Manager shall only write 20h to the USB4 Version field.*
- *The Connection Manager shall not write to DW26 of Router Configuration Space (ROUTER_CS_26) when the Operation Valid bit is 1b.*
- *For a Device Router:*
 - *The Connection Manager shall set the Configuration Valid field to 1b after setting the PTO, UTO, and HCO fields to the desired values.*
 - *The Connection Manager shall not set the UTO and HCO fields to 1b at the same time. At least one of the fields must be set to 0b.*
 - *The Connection Manager shall set the Configuration Valid field to 1b before setting up any Paths through the Router.*
 - *A Connection Manager shall not setup USB3 and/or PCIe Tunneling until the Configuration Ready bit is set to 1b.*

8.2.1.1.1 UUID

Each Router contains a Universally Unique ID (UUID) assigned by the Router vendor. The UUID shall have the format shown in Figure 8-2 where:

- Vendor ID is a 16-bit ID assigned by the USB-IF, which identifies the silicon vendor. It shall contain the same value as the *Vendor ID* field in Router Configuration Space.
- Component ID is a 44-bit ID that is unique to a single piece of silicon that contains one or more Routers. Routers that reside in the same silicon shall have the same Component ID. Routers with the same Vendor ID that reside in separate silicon shall have different Component IDs.
- Router ID is a 4-bit ID that indicates the Router instance within the product.

- A product containing multiple Router instances shall increment the Router ID for each Router instance, starting at 0. A product containing a single Router instance shall set this field to 0.

Figure 8-2. UUID Format



8.2.1.2 TMU Router Configuration Capability

A TMU Router Configuration Capability shall have the structure depicted in Figure 8-3 and the fields defined in Table 8-4.

Any field that spans across multiple Doublewords (e.g. LocalTime Low, LocalTime Middle, and LocalTime High) shall use the Register Locking Mechanism defined in Section 8.2.1.2.1 and the Register Group Locking Mechanism defined in Section 8.2.1.2.2.

Figure 8-3. Structure of the TMU Router Configuration Capability

| | | | | | | | | | | |
|-----------------------|------------------|------------------|------------------|----------------|---|-----------------|-------------------------|-----------------------------|--|----|
| 31 | | | | | 0 | | | | | |
| I D E | U C A P | R S V d | R S V d | T D | Freq Measurement Window | Capability ID | Next Capability Pointer | | | 0 |
| | | | | | LocalTime Low 31:0 | | | | | 1 |
| | | | | | LocalTime High 63:32 | | | | | 2 |
| | | | | | TSPacketInterval | LocalTime 79:64 | | | | 3 |
| | | | | | TimeOffsetFromHR 31:0 | | | | | 4 |
| | | | | | TimeOffsetFromHR 63:32 | | | | | 5 |
| | | | | | TimeOffsetFromDFP 31:0 | | | | | 6 |
| | | | | | TimeOffsetFromDFP 63:32 | | | | | 7 |
| | | | | | FreqOffsetFromHR | | | | | 8 |
| | | | | | FreqOffsetFromDFP | | | | | 9 |
| | | | | | Propagation Delay 31:0 | | | | | 10 |
| | | | | | Propagation Delay 63:32 | | | | | 11 |
| | | | | | Computation Timestamp 31:0 | | | | | 12 |
| | | | | | Computation Timestamp 63:32 | | | | | 13 |
| | | | | | Reserved | | | Computation Timestamp 79:64 | | 14 |
| TSInterDomainInterval | | ErrorAvgConst | | OffsetAvgConst | | DelayAvgConst | | FreqAvgConst | | 15 |
| | | | | | InterDomain Computation Timestamp 31:0 | | | | | 16 |
| | | | | | InterDomain Computation Timestamp 63:32 | | | | | 17 |
| I D S | Reserved | | | | InterDomain Computation Timestamp 79:64 | | | | | 18 |
| | | | | | TimeOffsetFromInterDomainHR 31:0 | | | | | 19 |
| | | | | | TimeOffsetFromInterDomainHR 63:32 | | | | | 20 |
| | | | | | FreqOffsetFromInterDomainHR | | | | | 21 |
| | | | | | Post Local Time 31:0 | | | | | 22 |
| | | | | | Post Local Time 63:32 | | | | | 23 |
| | | | | | Post Time 31:0 | | | | | 24 |
| | | | | | Post Time 63:32 | | | | | 25 |

Table 8-4. TMU Router Configuration Capability Fields

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|------|----------------|
| 0 | TMU_RTR_CS_0 | 7:0 | Next Capability Pointer This field shall contain the Doubleword index of the next Capability in the Router Configuration Space. A Router shall set this field to 00h if the TMU Router Configuration Capability is the final Capability in the linked list of Capabilities in the Router Configuration Space. | RO | Vendor Defined |
| | | 15:8 | Capability ID A Router shall set this field to 03h indicating this is the start of a TMU Router Configuration Capability. | RO | 03h |
| | | 26:16 | Freq Measurement Window This field contains the number of Time Sync Handshakes that occur before the frequency ratio and frequency offset are computed. | R/W | 800 |
| | | 27 | Time Disruption (TD) The Connection Manager sets this bit to 1b before any of the following time disruptions: <ul style="list-style-type: none"> • TMU mode changes (TSPacketInterval, Direction, Filters, Frequency Measurement Window) • Inter-Domain time sync is enabled • Time Posting is applied After the time disruption has passed, the Connection Manager sets this bit to 0b. | R/W | 0 |
| | | 29:28 | Reserved | Rsvd | 0 |
| | | 30 | Uni-Directional Capability (UCAP) This field shall be 0b if Uni-Directional Time Sync Handshakes are not supported. This field shall be 1b if Uni-Directional Time Sync Handshakes are supported. | RO | Vendor Defined |
| | | 31 | Inter-Domain Enable (IDE) For a Host Router: This field is set to 1 to enable the Host Router of a Domain to start receiving Inter-Domain Time Stamp packets. Otherwise, it shall be set to 0. For a Device Router: This field shall be set to 0b and shall have no effect. | R/W | 0 |
| 1 | TMU_RTR_CS_1 | 31:0 | LocalTime Low This field contains the least significant 32 bits of the 80-bit LocalTime counter. | RO | 0 |
| 2 | TMU_RTR_CS_2 | 31:0 | LocalTime Middle This field contains the middle 32 bits of the 80-bit LocalTime counter. | RO | 0 |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|--|------|---------------|
| 3 | TMU_RTR_CS_3 | 15:0 | LocalTime High This field contains the most significant 16 bits of the 80-bit LocalTime counter. | RO | 0 |
| | | 31:16 | TSPacketInterval This field contains the time interval between successive transmissions of the Delay Request TSNOS on a Link. The time interval is specified in units of 1 μ s with 100ppm accuracy. There are three allowed values for this register: 0: Shall disable Time Sync Handshake initiation by the Router 16: "HiFi" Mode (16 μ s) 1000 "LowRes" Mode (1ms) Other values shall not be used. <i>Note: The Default value of this field is 0, which means that Time Sync Handshake initiation is disabled after wake up.</i> | R/W | 0 |
| 4 | TMU_RTR_CS_4 | 31:0 | TimeOffsetFromHR Low This field contains the least significant 32 bits of the computed time offset between the Local Clock and the Host Router's Local Clock. A Router shall calculate the time offset and described in Equation 7-9. | RO | 0 |
| 5 | TMU_RTR_CS_5 | 31:0 | TimeOffsetFromHR High This field contains the most significant 32 bits of the computed time offset between the Local Clock and the Host Router's Local Clock. A Router shall calculate the time offset as described in Equation 7-9. | RO | 0 |
| 6 | TMU_RTR_CS_6 | 31:0 | TimeOffsetFromDFP Low This field contains the least significant 32 bits of the computed time offset between the DFP and UFP clocks. A Router shall calculate the time offset as described in Equation 7-7 for Bi-Directional Time Sync Handshakes or Equation 7-8 for Uni-Directional Time Sync Handshakes. | RO | 0 |
| 7 | TMU_RTR_CS_7 | 31:0 | TimeOffsetFromDFP High This field contains the most significant 32 bits of the computed time offset between the DFP and UFP clocks. A Router shall calculate the time offset as described in Equation 7-7 for Bi-Directional Time Sync Handshakes or Equation 7-8 for Uni-Directional Time Sync Handshakes. | RO | 0 |
| 8 | TMU_RTR_CS_8 | 31:0 | FreqOffsetFromHR This field contains the computed frequency offset between the Local Clock and the Host Router's Local Clock, represented using 2's complement format. A Router shall calculate the frequency offset as described in Equation 7-4. | RO | 0 |
| 9 | TMU_RTR_CS_9 | 31:0 | FreqOffsetFromDFP This field contains the computed frequency offset between the DFP and the UFP clocks, represented using 2's complement format. A Router shall calculate the frequency offset as described in Equation 7-2. | RO | 0 |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|------|---------------|
| 10 | TMU_RTR_CS_10 | 31:0 | Propagation Delay Low This field contains the least significant 32 bits of the computed time delay between the Router and its Upstream Link Partner. This field shall have the same format as the TimeOffsetFromHR register. A Router shall calculate the time delay as described in Equation 7-5. | RO | 0 |
| 11 | TMU_RTR_CS_11 | 31:0 | Propagation Delay High This field contains the most significant 32 bits of the computed time delay between the Router and its Upstream Link Partner. This field shall have the same format as the TimeOffsetFromHR register. A Router shall calculate the time delay as described in Equation 7-5 | RO | 0 |
| 12 | TMU_RTR_CS_12 | 31:0 | Computation Time Stamp Low This field shall contain the least significant 32 bits of the most recent value of the t_4 time stamp ($t_4[n]$). | RO | 0 |
| 13 | TMU_RTR_CS_13 | 31:0 | Computation Time Stamp Middle This field shall contain the middle 32 bits of the most recent value of the t_4 time stamp ($t_4[n]$). | RO | 0 |
| 14 | TMU_RTR_CS_14 | 15:0 | Computation Time Stamp High This field shall contain the most significant 16 bits of the most recent value of the t_4 time stamp ($t_4[n]$). | RO | 0 |
| | | 31:16 | Reserved | Rsvd | 0 |
| 15 | TMU_RTR_CS_15 | 5:0 | FreqAvgConst This field contains the IIR filter co-efficient that shall be used to average the frequency ratio. The IIR filter co-efficient is given by $(1/2^{\text{FreqAvgConst}})$. | R/W | 8 |
| | | 11:6 | DelayAvgConst This field contains the IIR filter co-efficient that shall be used to average the propagation delay. The IIR filter co-efficient is given by $(1/2^{\text{DelayAvgConst}})$. | R/W | 8 |
| | | 17:12 | OffsetAvgConst This field contains the IIR filter co-efficient that shall be used to average the time offset. The IIR filter co-efficient is given by $(1/2^{\text{OffsetAvgConst}})$. | R/W | 8 |
| | | 23:18 | ErrorAvgConst This field contains the IIR filter co-efficient that shall be used to average the time offset averaging error. The IIR filter co-efficient is given by $(1/2^{\text{ErrorAvgConst}})$. | R/W | 8 |
| | | 31:24 | TSInterDomainInterval This field controls the time interval between transmitting Inter-domain Time Stamp Packets. The time interval is given in microseconds and is equal to $(\text{TSInterDomainInterval} + 1) * \text{TSPacketInterval}$. | R/W | 0 |
| 16 | TMU_RTR_CS_16 | 31:0 | Inter-Domain Time Stamp Low For an Inter-Domain Time Initiator: This field shall contain the least significant 32 bits of the computed value of the Inter-Domain time stamp. The time stamp shall be computed as described in Section 7.4.2.1. For all other Routers: This field shall contain the least significant 32 bits of the IDTimeStamp value contained in the last received Follow-Up Packet. | RO | 0 |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|------|--|
| 17 | TMU_RTR_CS_17 | 31:0 | <p>Inter-Domain Time Stamp Middle</p> <p>For an Inter-Domain Time Initiator:</p> <p>This field shall contain the middle 32 bits of the computed value of the Inter-Domain time stamp. The time stamp shall be computed as described in Section 7.4.2.1.</p> <p>For all other Routers:</p> <p>This field shall contain the middle 32 bits of the IDTimeStamp value contained in the last received Follow-Up Packet.</p> | RO | 0 |
| 18 | TMU_RTR_CS_18 | 15:0 | <p>Inter-Domain Time Stamp High</p> <p>For an Inter-Domain Time Initiator:</p> <p>This field shall contain the most significant 32 bits of the computed value of the Inter-Domain time stamp. The time stamp shall be computed as described in Section 7.4.2.1.</p> <p>For all other Routers:</p> <p>This field shall contain the most significant 32 bits of the IDTimeStamp value contained in the last received Follow-Up Packet.</p> | RO | 0 |
| | | 30:16 | Reserved | Rsvd | 0 |
| | | 31 | <p>Inter-Domain Not Supported (IDNS)</p> <p>This field shall indicate whether or not a Router supports Inter-Domain Time Synchronization:</p> <p>0 – Supported</p> <p>1 – Not Supported</p> <p>A Host Router may set this field to either 1 or 0.</p> <p>A Device Router shall set this field to 0.</p> | RO | Host Router: Vendor Defined Device Router: 0 |
| 19 | TMU_RTR_CS_19 | 31:0 | <p>TimeOffsetFromInterDomainHR Low</p> <p>For an Inter-Domain Time Initiator:</p> <p>This field shall contain the least significant 32 bits of the computed time offset between the local Host Router Clock and the Inter-Domain Time Source. This field shall have the same format as shown in Figure 7-3. The time offset shall be computed as described in Section 7.4.2.3.</p> <p>For all other Routers:</p> <p>This field shall contain the most recent value of the TimeOffsetFromInterDomainHR field contained in the last received Follow-Up Packet.</p> | RO | 0 |
| 20 | TMU_RTR_CS_20 | 31:0 | <p>TimeOffsetFromInterDomainHR High</p> <p>For an Inter-Domain Time Initiator:</p> <p>This field contains the most significant 32 bits of the computed time offset between the local Host Router Clock and the Inter-Domain Time Source. The format of this register is shown in Figure 7-3. The time offset shall be computed as described in Section 7.4.2.3.</p> <p>For all other Routers:</p> <p>This field contains the most recent value of the TimeOffsetFromInterDomainHR field contained in the last received Follow-Up Packet.</p> | RO | 0 |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|--|------|---------------|
| 21 | TMU_RTR_CS_21 | 31:0 | <p>FreqOffsetFromInterDomainHR</p> <p>For an Inter-Domain Time Initiator:</p> <p>This field shall contain the computed frequency offset between the local Host Router Clock and the Inter-Domain Time Source, represented using 2's complement format.</p> <p>The frequency offset shall be computed as described in Section 7.4.2.2.</p> <p>For all other Routers:</p> <p>This field shall contain the most recent value of the <i>FrequencyOffsetFromInterDomainHR</i> field contained in the last received Follow-Up Packet.</p> | RO | 0 |
| 22 | TMU_RTR_CS_22 | 31:0 | <p>Post Local Time Low</p> <p>This field contains the least significant 32 bits of the Post Local Time. The Post Local Time is used to update a Router's local time as part of the Post Time Mechanism. See Section 7.6.3.</p> | R/W | 0 |
| 23 | TMU_RTR_CS_23 | 31:0 | <p>Post Local Time High</p> <p>This field contains the most significant 32 bits of the Post Local Time. The Post Local Time is used to update a Router's local time as part of the Post Time Mechanism. See Section 7.6.3.</p> | R/W | 0 |
| 24 | TMU_RTR_CS_24 | 31:0 | <p>Post Time Low</p> <p>This field contains the least significant 32 bits of the Post Time. The Post Time is when a Router updates its local time as part of the Post Time Mechanism. See Section 7.6.3. A Router shall set this field to 0 after updating its local time.</p> | R/W | 0 |
| 25 | TMU_RTR_CS_25 | 31:0 | <p>Post Time High</p> <p>This field contains the most significant 32 bits of the Post Time. The Post Time is when a Router updates its local time as part of the Post Time Mechanism. See Section 7.6.3. A Router shall set this field to 0 after updating its local time.</p> <p>The Post Time Mechanism is only activated if this register is written.</p> | R/W | 0 |



CONNECTION MANAGER NOTE

A Connection Manager uses the following values to configure the *TSPacketInterval* field:

- *TMU_OFF* = 0.
- *HiFi* = 16.
- *LowRes* = 1000.

8.2.1.2.1 Register Locking Mechanism

Some of the fields in a TMU Router Configuration Capability span more than 1 DW. In order to keep consistent values in these fields, a Router shall update the value in the entire field (i.e. Low, Middle, and High DWs) when the Connection Manager reads the Low DW of the field. A Router shall not change the value in the Middle and High DWs until the next time the Low DW is read. For example, when a Connection Manager reads the LocalTime {Low, Medium, High} registers, it reads the LocalTime Low register first to lock in the LocalTime value. It then reads the LocalTime Medium register and the LocalTime High register, which each return the value that was locked when the LocalTime Low register was first read.

The Register Locking Mechanism shall be implemented for following registers:

- LocalTime {Low, Middle, High}.
- TimeOffsetFromHR {Low, Middle, High}.
- Inter-Domain Time Stamp {Low, Middle, High}.



CONNECTION MANAGER NOTE

In order to guarantee a consistent value across all registers, a Connection Manager shall read the Low DW of the following fields before reading the Medium DW and High DW:

- *LocalTime {Low, Medium, High}.*
- *TimeOffsetFromHR {Low, Medium, High}.*
- *Inter-Domain Time Stamp {Low, Medium, High}.*

8.2.1.2.2 Register Group Locking Mechanism

There are some register groups that are locked together when a particular DW is read. This DW is called the Triggering DW. The value of a locked register group shall change only when the Triggering DW is accessed.

Table 8-5 lists the register groups that shall be locked.

Table 8-5. Locked Registers Groups

| Group Name | Triggering DW | Locked Registers |
|--------------------------------|-----------------------------|--|
| Time Computation Group | TimeOffsetFromHR Low | <ul style="list-style-type: none"> • TimeOffsetFromHR Low • TimeOffsetFromHR High • TimeOffsetFromDFP Low • TimeOffsetFromDFP High • FrequencyOffsetFromHR • FrequencyOffsetFromDFP • Propagation Delay Low • Propagation Delay High • Computation Time Stamp Low • Computation Time Stamp Middle • Computation Time Stamp High |
| Inter-Domain Computation Group | Inter-Domain Time Stamp Low | <ul style="list-style-type: none"> • Inter-Domain Time Stamp Low • Inter-Domain Time Stamp Middle • Inter-Domain Time Stamp High • TimeOffsetFromInterDomainHR Low • TimeOffsetFromInterDomainHR High • FrequencyOffsetFromInterDomainHR |

8.2.1.3 Vendor Specific Capability (VSC)

A Vendor Specific Capability shall have the format and fields shown in Figure 8-4. The remaining fields within the Vendor Specific Registers are vendor defined.

Figure 8-4. Structure of a Vendor Specific Capability




Table 8-6 describes the fields that a Vendor Specific Capability shall contain.

Table 8-6. Vendor Specific Capability Fields

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|--|------|----------------|
| 0 | VSC_CS_0 | 7:0 | Next Capability Pointer This field shall contain the Doubleword index of the next Capability in the Router Configuration Space. A Router shall set this field to 00h if the Vendor-Specific Capability is the final Capability in the linked list of Capabilities in the Router Configuration Space. | RO | Vendor Defined |
| | | 15:8 | Capability ID A Router shall set this field to 05h indicating this is the start of a Vendor-Specific Capability. | RO | 05h |
| | | 23:16 | VSC ID This field shall contain the vendor-defined ID number that identifies the nature and format of the VSC structure. | RO | Vendor Defined |
| | | 31:24 | VSC Length This field shall contain the total number of Doublewords in the VSC structure including Doubleword 0 and the Vendor-Specific Doublewords that follow it. | RO | Vendor Defined |

8.2.1.4 Vendor Specific Extended Capability (VSEC)

The Vendor Specific Extended Capability allows larger Capabilities to be defined, beyond the limitations of the Vendor Specific Capability. A Vendor Specific Extended Capability shall have the format and fields shown in Figure 8-5.

Figure 8-5. Structure of a Vendor Specific Extended Capability




Table 8-7 describes the fields that a Vendor Specific Extended Capability shall contain.

Table 8-7. Vendor Specific Extended Capability Fields

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|------|----------------|
| 0 | VSEC_CS_0 | 7:0 | Reserved | Rsvd | 0 |
| | | 15:8 | Capability ID A Router shall set this field to 05h indicating this is the start of a Vendor-Specific Capability. | RO | 05h |
| | | 23:16 | VSEC ID This field shall contain the vendor-defined ID number that identifies the nature and format of the VSEC structure. | RO | Vendor Defined |
| | | 31:24 | VSEC Header A Router shall set this field to 00h to indicate that the Capability is an Extended Capability. | RO | 00h |
| 1 | VSEC_CS_1 | 15:0 | Next Capability Pointer This field shall contain the Doubleword index of the next Capability in the Router Configuration Space. A Router shall set this field to 00h if the Vendor-Specific Capability is the final Capability in the linked list of Capabilities in the Router Configuration Space. | RO | Vendor Defined |
| | | 31:16 | VSEC Length This field shall contain the total number of Doublewords in the VSEC structure including Doubleword 0, Doubleword 1, and the Vendor-Specific Doublewords that follow. | RO | Vendor Defined |

A Vendor Specific Extended Capability may reside anywhere in the 8192 DWs that a Read or Write Request packet can address. It is recommended that the first Vendor Specific Extended Capability starts at DW address 255.

8.2.2 Adapter Configuration Space

Every Adapter (except for a Control Adapter) shall have its own Adapter Configuration Space. An Adapter Configuration Space shall have the structure depicted in Figure 8-6. The Adapter Configuration Space structure depicted in Figure 8-6 begins with a set of Doublewords describing the basic attributes of an Adapter. The rest of the space is populated with a linked list of Capabilities.

A Connection Manager reads from or writes to Adapter Configuration Space using the Read Requests and Write Requests defined in Section 6.4.2. A Router shall allow a Connection Manager to access Adapter Configuration Space regardless of whether or not the Adapter is connected.

Figure 8-6. Structure of the Adapter Configuration Space




Table 8-8 lists the Configuration Capabilities supported by Adapter Configuration Space. A Capability listed as “Required” shall be present in Adapter Configuration Space. A Capability listed as “Optional” may be present in Adapter Configuration Space.

Table 8-8. List of Adapter Configuration Capabilities

| Capability | Required / Optional | Capability ID |
|------------------------------|--|---------------|
| TMU Adapter Configuration | A TMU Adapter Configuration Capability is required for Lane Adapters. Shall not be present for any other Adapter. | 03h |
| Lane Adapter Configuration | Required for Lane Adapters. Shall not be present for any other Adapter. | 01h |
| USB4 Port Capability | Required for the Lane 0 Adapter in a USB4 Port. Shall not be present for any other Adapter. | 06h |
| PCIe Adapter Configuration | Required for PCIe Adapters. Shall not be present for any other Adapter. | 04h |
| DP IN Adapter Configuration | Required for DP IN Adapters. Shall not be present for any other Adapter. | 04h |
| DP OUT Adapter Configuration | Required for DP OUT Adapters. Shall not be present for any other Adapter. | 04h |
| USB3 Adapter Configuration | Required for USB3 Adapters. Shall not be present for any other Adapter. | 04h |

| | | |
|-----------------|-----------|-----|
| Vendor Specific | Optional. | 05h |
|-----------------|-----------|-----|

The number of Vendor Specific Capabilities and Vendor Specific Extended Capabilities is implementation specific. Section 8.2.1.3 defines the structure of the Vendor Specific Capability. Section 8.2.1.4 defines the structure of the Vendor Specific Extended Capability.

A Router’s operation must not depend on a Connection Manager’s support for the Vendor Specific Capabilities and Vendor Specific Extended Capabilities.

Capabilities shall be linked in the following order:

- Required Capabilities.
- Optional Capabilities.
- Vendor Specific Capabilities.
- Vendor Specific Extended Capabilities.

8.2.2.1 Basic Configuration Registers

The first 9 Doublewords in an Adapter Configuration Space shall have the format and fields described in Figure 8-7 and Table 8-9.

Figure 8-7. Basic Configuration Registers of the Adapter Configuration Space




Table 8-9. Adapter Configuration Space Basic Attributes

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|------|----------------|
| 0 | ADP_CS_0 | 31:0 | Vendor Defined | VD | Vendor Defined |
| 1 | ADP_CS_1 | 7:0 | Next Capability Pointer This field shall contain the Doubleword index of the first Capability in the Adapter Configuration Space. | RO | Vendor Defined |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|--|-------|----------------|
| 1 | ADP_CS_1 | 18:8 | Max Counter Sets This field shall contain the number of counter sets provided by the Adapter in Counters Configuration Space. The value in this field shall be at least 1 if the <i>CCS Flag</i> is set to 1b. | RO | Vendor Defined |
| | | 19 | Counters Configuration Space (CCS) Flag An Adapter shall set this flag to 1b if it supports Counters Configuration Space. Otherwise this flag shall be set to 0b. | RO | Vendor Defined |
| | | 31:20 | Reserved | Rsvd | 0 |
| 2 | ADP_CS_2 | 7:0 | Adapter Type Sub-type This field shall identify the Adapter sub-type using the Sub-Type encodings in Table 8-10. | RO | See Table 8-10 |
| | | 15:8 | Adapter Type Version This field shall identify the Adapter version using the version encodings in Table 8-10. | RO | See Table 8-10 |
| | | 23:16 | Adapter Type Protocol This field shall identify the Adapter protocol type using the Protocol encodings in Table 8-10. | RO | See Table 8-10 |
| | | 31:24 | Reserved This field shall be set to 01h. | Rsvd | 01h |
| 3 | ADP_CS_3 | 19:0 | Reserved | Rsvd | 0 |
| | | 25:20 | Adapter Number This field shall contain the Adapter number for the Adapter. | RO | Vendor Defined |
| | | 28:26 | Reserved | Rsvd | 0 |
| | | 29 | HEC Error (HE) A Lane 0 Adapter sets this bit to 1b to indicate the reception of a Transport Layer Packet with an uncorrectable HEC error in the header. The Lane 1 Adapter shall set this bit to 0b. This field is reserved in a Protocol Adapter and shall be set to 0b. | R/Clr | 0 |
| | | 30 | Flow Control Error (FCE) A Lane 0 Adapter sets this bit to 1b to indicate the reception of a Packet on a flow controlled Path when the appropriate buffer (dedicated or shared) has no space for the packet. The Lane 1 Adapter shall set this bit to 0b. This field is reserved in a Protocol Adapter and shall be set to 0b. | R/Clr | 0 |
| | | 31 | Shared Buffering Capable (SBC) This bit indicates whether or not a USB4 Port is capable of sharing flow control buffers among flow controlled Paths. A Lane 0 Adapter shall set this bit to 1b if shared buffering is supported. A Lane 0 Adapter shall set this bit to 0b if shared buffering is not supported. The Lane 1 Adapter shall set this bit to 0b. This field is ignored in a Protocol Adapter. | RO | Vendor Defined |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|------|----------------|
| 4 | ADP_CS_4 | 9:0 | <p>Non-Flow Controlled Buffers</p> <p>A Connection Manager uses this field to configure the number of ingress buffers allocated for non-flow-controlled Paths in a Lane Adapter.</p> <p>A Connection Manager may change this field dynamically when setting up or tearing down a Path.</p> <p>This field is not used for a Protocol Adapter.</p> | R/W | 0 |
| | | 19:10 | Reserved | Rsvd | 0 |
| | | 29:20 | <p>Total Buffers</p> <p>This field shall contain the total number of ingress buffers available to a Lane Adapter as defined in Section 5.3.2.1.1.</p> <p><i>Note: The value in the Total Buffers field includes any buffers for Path 0.</i></p> <p>This field is ignored in a Protocol Adapter.</p> | RO | Vendor Defined |
| | | 30 | <p>Plugged</p> <p>This field indicates whether or not the Adapter is plugged.</p> <p>0b – Adapter is not plugged 1b – Adapter is plugged</p> <p>This field is reserved in a USB3 Adapter, a DP IN Adapter, a PCIe Adapter, and a Host Interface Adapter, and shall be set to 0.</p> | RO | 0 |
| | | 31 | <p>Lock (LCK)</p> <p>This bit controls whether or not a Connection Manager can access a Router that is downstream of the Adapter. This bit is only used for the Adapters in a Downstream Facing Port.</p> <p>When the bit is 1b, the Adapter is “locked”, which means that Control Packets are not forwarded to the downstream Router. When the bit is 0b, the Adapter is “unlocked” and Control Packets can be forwarded to the downstream Router.</p> <p>An Adapter shall set this bit to 1b after the Adapter goes through a disconnect.</p> <p>An Adapter may ignore a write to this bit if the Adapter does not have a Router connected downstream.</p> <p>See Section 6.7 for more information on the functionality of this bit.</p> | R/W | 0b |
| 5 | ADP_CS_5 | 10:0 | <p>Max Input HopID</p> <p>This field shall contain the highest HopID value the Adapter supports for incoming packets.</p> | RO | Vendor Defined |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|-------|----------------|
| 5 | ADP_CS_5 | 21:11 | Max Output HopID This field shall contain the highest HopID value the Adapter supports for outgoing packets. | RO | Vendor Defined |
| | | 28:22 | Link Credits Allocated A Connection Manager uses this field to configure the initial number of credits to be allocated to the Shared Flow Control Buffer used by the Ingress Lane Adapter. This field is not used for a Protocol Adapter. | R/W | 0 |
| | | 29 | HEC Error Enable (HEE) A Connection Manager uses this field to enable HEC error reporting in a Lane Adapter as follows: 1b – HEC error reporting enabled 0b – HEC error reporting disabled This bit is Reserved for a Lane 1 Adapter. This field is only used for Lane Adapters. It is ignored for Protocol Adapters. | R/W | 0b |
| | | 30 | Flow Control Error Enable (FCEE) A Connection Manager uses this field to enable Flow Control error reporting in a Lane Adapter as follows: 1b – FCEE error reporting enabled 0b – FCEE error reporting disabled This bit is Reserved for a Lane 1 Adapter. This field is only used for Lane Adapters. It is ignored for Protocol Adapters. | R/W | 0b |
| | | 31 | Disable Hot Plug Events (DHP) A Connection Manager uses this bit to configure whether or not the Router sends a Hot Plug Event Packet as a result of a Hot Plug Event or a Hot Unplug Event on this Adapter. 1b – Hot Plug Event Packet not sent 0b – Hot Plug Event Packet is sent Only a Lane Adapter, a DP IN Adapter, or a DP OUT Adapter may optionally implement this bit. All other Adapters shall hardwire this bit to 0b. | R/W | 0b |
| 6 | ADP_CS_6 | 31:0 | HEC Errors This field shall contain the number of ingress Transport Layer packets dropped due to HEC errors. A Lane Adapter shall increment the counter in this field from 0 and shall stop counting at FFFF FFFFh. A Protocol Adapter shall not increment the counter and shall set this field to 0. | W/Clr | 0 |
| 7 | ADP_CS_7 | 31:0 | Invalid HopID Errors This field shall contain the number of ingress Transport Layer packets with a HopID outside the supported range or a HopID that does not belong to an enabled Path. An Adapter shall increment the counter in this field from 0 and shall stop counting at FFFF FFFFh. | W/Clr | 0 |
| 8 | ADP_CS_8 | 31:0 | ECC Errors This field shall contain the number of Credit Sync Packets and Credit Grant Records dropped due to ECC errors. A Lane Adapter shall increment the counter in this field from 0 and shall stop counting at FFFF FFFFh. A Protocol Adapter shall not increment the counter and shall set this field to 0. | W/Clr | 0 |

Table 8-10. Adapter Types

| Type | Protocol | Version | Sub-Type |
|-------------------------|----------|---------|----------|
| Unsupported Adapter | 00h | 00h | 00h |
| Lane Adapter | 00h | 00h | 01h |
| Host Interface Adapter | 00h | 00h | 02h |
| Downstream PCIe Adapter | 10h | 01h | 01h |
| Upstream PCIe Adapter | 10h | 01h | 02h |
| DP OUT Adapter | 0Eh | 01h | 02h |
| DP IN Adapter | 0Eh | 01h | 01h |
| Downstream USB3 Adapter | 20h | 01h | 01h |
| Upstream USB3 Adapter | 20h | 01h | 02h |



CONNECTION MANAGER NOTE

When writing to Adapter Configuration Space of a Protocol Adapter, a Connection Manager shall not change the value in the following fields:

- *Non-Flow Controlled Buffers.*
- *Link Credits Allocated.*

8.2.2.2 TMU Adapter Configuration Capability

A TMU Adapter Configuration Capability shall have the structure depicted in Figure 8-8 and shall contain the fields defined in Table 8-11.

For a USB4 Port with two enabled Lane Adapters, the values in the TMU Adapter Configuration Capability of both Adapters shall be identical. When a value in the TMU Adapter Configuration Capability of one Lane Adapter is written to, the other Lane Adapter in the USB4 Port shall update its value to match.

Figure 8-8. Structure of the TMU Adapter Configuration Capability

| | | | | | |
|-------------------|--------------------|---------------|---------------------|-------------------------|-----------------------|
| 31 | | | | 0 | |
| Reserved | | Capability ID | | Next Capability Pointer | |
| TxTimeToWire | | | | | |
| RxTimeToWire | | | | | |
| ID TI | ID TR | U D M | Reserved | | |
| TX TSNOS counter | | | RX TSNOS counter | | |
| TX Packet counter | | | RX Packet counter | | |
| Vendor Defined | | | Reserved | | D T S V D |
| Rsvd | Bad Packet Counter | | Lost Packet Counter | | Lost TSNOS Counter |

Table 8-11. TMU Adapter Configuration Capability Fields

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|-------|----------------|
| 0 | TMU_ADP_CS_0 | 7:0 | Next Capability Pointer This field shall contain the Doubleword index of the next Capability in the Adapter Configuration Space. A Router shall set this field to 00h if the TMU Adapter Configuration Capability is the final Capability in the linked list of Capabilities in the Adapter Configuration Space. | RO | Vendor Defined |
| | | 15:8 | Capability ID An Adapter shall set this field to 03h indicating this is the start of a TMU Adapter Configuration Capability. | RO | 03h |
| | | 31:16 | Reserved | Rsvd | 0 |
| 1 | TMU_ADP_CS_1 | 31:0 | TxTimeToWire This field shall contain the time duration from the instant the time stamp is taken at the Physical Layer to the instant when the first bit of the TSNOS is transmitted on the wire. The time shall be specified in nanoseconds multiplied by 2 ¹⁶ . For example, 2.5ns is represented as 0000 0002 8000h. | RO | Vendor Defined |
| 2 | TMU_ADP_CS_2 | 31:0 | RxTimeToWire This field shall contain the time duration from the instant the first bit of the TSNOS is received at the wire to the instant when the time stamp is taken at the Physical Layer. The time shall be specified in nanoseconds multiplied by 2 ¹⁶ . For example, 2.5 ns is represented as 0000 0002 8000h. | RO | Vendor Defined |
| 3 | TMU_ADP_CS_3 | 28:0 | Reserved | Rsvd | 0 |
| | | 29 | EnableUniDirectionalMode (UDM) A Connection Manager uses this bit to enable Uni-Directional Time Sync Handshakes. The TMU Adapter shall set this bit to 0b when its USB4 Port is disconnected. | R/W | 0 |
| | | 30 | Inter-Domain Time Responder (IDTR) A Connection Manager uses this bit to configure the USB4 Port as an Inter-Domain Time Responder (i.e. the responder to Time Sync Handshakes across an Inter-Domain Link). If set to 1b, the USB4 Port shall respond to Time Sync Handshakes over the Interdomain Link as initiated by the ITDI Port. Otherwise this bit shall be set to 0b. | R/W | 0 |
| | | 31 | Inter-Domain Time Initiator (IDTI) A Connection Manager uses this bit to configure the USB4 Port as an Inter-Domain Time Initiator (i.e. the initiator of Time Sync Handshakes across an Inter-Domain Link). If set to 1b, the USB4 Port shall Initiate Time Sync Handshakes over the Inter-Domain Link. Otherwise this bit shall be set to 0b. | R/W | 0 |
| 4 | TMU_ADP_CS_4 | 15:0 | RX TSNOS Counter This field shall contain the number of TSNOS received by TMU. The counter shall not increment past FFFFh. | R/Clr | 0 |
| | | 31:16 | TX TSNOS Counter This field shall contain the number of TSNOS sent by TMU. The counter shall not increment past FFFFh. | R/Clr | 0 |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|-------|----------------|
| 5 | TMU_ADP_CS_5 | 15:0 | RX Packet Counter This field shall contain the number of Time Sync Packets received by TMU. The counter shall not increment past FFFFh. | R/Clr | 0 |
| | | 31:16 | TX Packet Counter This field shall contain the number of Time Sync Packets sent by TMU. The counter shall not increment past on FFFFh. | R/Clr | 0 |
| 6 | TMU_ADP_CS_6 | 0 | Vendor Defined | VD | Vendor Defined |
| | | 1 | Disable Time Sync (DTS) A Connection Manager uses this bit to disable Time Synchronization Handshake initiation. If set to 1b, the Adapter shall not send any Delay Requests or Delay Responses. If set to 0b, the Adapter may send Delay Requests or Delay Responses. | R/W | 0 |
| | | 3:2 | Vendor Defined | VD | Vendor Defined |
| | | 15:4 | Reserved | RsvdZ | 0 |
| | | 31:16 | Vendor Defined | VD | Vendor Defined |
| 7 | TMU_ADP_CS_7 | 9:0 | Lost TSNOS Counter This field shall contain the number of times that a Delay Response was expected during a Time Sync Handshake but not received. The counter shall not increment past 3FFh. <i>Note: Only a Upstream Facing Port increments this counter (the Downstream Facing Port does not receive Delay Responses).</i> | R/Clr | 0 |
| | | 19:10 | Lost Packet Counter This field shall contain the number of times that a Follow-Up Packet was expected during a Time Sync Handshake but not received. The counter shall not increment past 3FFh. <i>Note: Only a Upstream Facing Port increments this counter (the Downstream Facing Port does not receive Follow-Up Packets).</i> | R/Clr | 0 |
| | | 29:20 | Bad Packet Counter This field shall contain the number of Follow-Up Packets and Inter-Domain Packets received with bad CRC. The counter shall not increment past 3FFh. | R/Clr | 0 |
| | | 31:30 | Reserved | Rsvd | 0 |

8.2.2.3 Lane Adapter Configuration Capability

A Lane Adapter Configuration Capability shall have the structure depicted in Figure 8-9 and shall have the fields defined in Table 8-12.

Figure 8-9. Structure of the Lane Adapter Configuration Capability

| | | | | | | | | | | | |
|----|----------|-----|---------------|-----------------------------|-----------------------|---------------|----|------|-------------------------|-------------------|-------------------|
| 31 | | | | | | | | | | | 0 |
| 0 | Rsvd | | CLx Support | Supported Link Widths | Supported Link Speeds | Capability ID | | | Next Capability Pointer | | |
| 1 | Rsvd | PMS | Adapter State | Negotiated Link Width | Current Link Speed | LB | LD | Rsvd | CLx Enable | Target Link Width | Target Link Speed |
| 2 | Reserved | | | Logical Layer Errors Enable | | Reserved | | | Logical Layer Errors | | |

Table 8-12. Contents of the Lane Adapter Configuration Capability

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|------|----------------|
| 0 | LANE_ADP_CS_0 | 7:0 | Next Capability Pointer This field shall contain the Doubleword index of the next Capability in the Adapter Configuration Space. An Adapter shall set this field to 00h if this Capability is the final Capability in the linked list of Capabilities in the Adapter Configuration Space. | RO | Vendor Defined |
| | | 15:8 | Capability ID An Adapter shall set this field to 01h indicating this is a Lane Adapter Configuration Capability. | RO | 01h |
| | | 19:16 | Supported Link Speeds This field indicates which Link speed(s) are supported by the Adapter. Bit definitions within this field are: Bit [17:16] Rsvd Bit 18 Gen 3 Bit 19 Gen 2 For a USB4 host or peripheral device: An Adapter shall set bit 18 to 1b if it supports Gen 3 speed. Otherwise, bit 18 shall be 0b. An Adapter shall set bit 19 to 1b to indicate support for Gen 2 speed. For a USB4 hub: An Adapter shall set bit 18 to 1b to indicate support for Gen 3 speed. An Adapter shall set bit 19 to 1b to indicate support for Gen 2 speed. The Lane 1 Adapter in a USB4 Port shall declare the same value as the Lane 0 Adapter. | RO | Vendor Defined |
| | | 25:20 | Supported Link Widths This field indicates which Link widths are supported by the Adapter (xN – corresponding to N Lanes). Bit definitions within this field are: Bit 20 x1 Bit 21 x2 Bit [25:22] Rsvd An Adapter shall set bit 20 to 1b to indicate support for x1 operation. An adapter shall set bit 21 to 1b to indicate support for x2 operation. Lane 1 Adapter in a USB4 Port shall declare the same value as the Lane 0 Adapter. | RO | 000011b |
| | | 26 | CL0s Support An Adapter shall set this field to 1b if it supports CL0s Low Power. Otherwise this bit shall be set to 0b. | RO | Vendor Defined |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|------|---|
| | | 27 | CL1 Support An Adapter shall set this field to 1b if it supports CL1 Low Power state. Otherwise this bit shall be set to 0b. | RO | Vendor Defined |
| | | 28 | CL2 Support An Adapter shall set this bit to 1b if it supports CL2 Low Power state. Otherwise this bit shall be set to 0b. | RO | Vendor Defined |
| | | 31:29 | Reserved | Rsvd | 000b |
| 1 | LANE_ADP_CS_1 | 3:0 | Target Link Speed A Connection Manager uses this field to indicate which Link speed the Router attempts to establish for the USB4 Port. Defined encodings are: 1000b – Router shall attempt Gen 2 speed 1100b – Router shall attempt Gen 3 speed All other encodings are reserved. | R/W | Host Router: Vendor Defined Device Router: Shall match the highest speed supported by the Lane |
| | | 9:4 | Target Link Width A Connection Manager uses this field to indicate the Link width that the Router attempts to establish for the USB4 Port. Defined encodings are: 0000 01b – Establish two Single-Lane Links 0000 11b – Establish a Dual-Lane Link All other encodings are reserved. | R/W | 000 01b |
| | | 10 | CL0s Enable A Connection Manager uses this bit to control whether or not the Adapter can enter CL0s state. 0b – entry to CL0s state is disabled 1b – entry to CL0s state is enabled | R/W | 0b |
| | | 11 | CL1 Enable A Connection Manager uses this bit to control whether or not the Adapter can enter CL1 state. 0b – entry to CL1 state is disabled 1b – entry to CL1 state is enabled | R/W | 0b |
| | | 12 | CL2 Enable A Connection Manager uses this bit to control whether or not the Adapter can enter CL2 state. 0b – entry to CL2 state is disabled 1b – entry to CL2 state is enabled | R/W | 0b |
| | | 13 | Reserved | Rsvd | 0b |
| | | 14 | Lane Disable (LD) A Connection Manager sets this bit to 1b to transition the Adapter to the Disabled state. A Connection Manager sets this bit to 0b to transition the Adapter out of the Disabled state. <i>Note: Writing to this bit is immediately reflected in the value read from the bit, regardless of actual Adapter state.</i> | R/W | 0 |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|-----------|----------------|
| | | 15 | <p>Lane Bonding (LB)</p> <p>A Connection Manager sets this bit to 1b in either Adapter of a USB4 Port to transition the Adapters to the Lane Bonding state.</p> <p>Writing 0b to this bit shall have no effect.</p> <p>If the Adapter is already in the Lane Bonding state when this bit is set to 1b, the Adapter may reenter the Lane Bonding state after Lane Bonding is complete but is not required to do so.</p> | R/W SC | 0 |
| | | 19:16 | <p>Current Link Speed</p> <p>This field shall indicate the negotiated Link speed.</p> <p>Defined encodings are:</p> <p>1000b Gen 2</p> <p>0100b Gen 3</p> <p>All other encodings are reserved.</p> <p>The Lane 1 Adapter in a USB4 Port shall contain the same value as the Lane 0 Adapter.</p> <p>This field is only valid when the Link is in the Active state.</p> | RO | Vendor Defined |
| | | 25:20 | <p>Negotiated Link Width</p> <p>This field shall indicate the negotiated Link width (xN – corresponding to N Lanes).</p> <p>Defined encodings are:</p> <p>000001b x1</p> <p>000010b x2</p> <p>All other encodings are reserved.</p> <p>The Lane 1 Adapter in a USB4 Port shall contain the same value as the Lane 0 Adapter.</p> <p>This field is only valid when the Link is in the Active state.</p> | RO | Vendor Defined |
| | | 29:26 | <p>Adapter State</p> <p>This field shall indicate the current Adapter state.</p> <p>Defined encodings are:</p> <p>0000b: Disabled state</p> <p>0001b: Training or Lane Bonding state</p> <p>0010b: CL0 state</p> <p>0011b: Transmitter in CL0s state</p> <p>0100b: Receiver in CL0s state</p> <p>0101b: CL1 state</p> <p>0110b: CL2 state</p> <p>0111b: CLd state</p> <p>All other encodings are reserved.</p> | RO | 7h |
| | | 30 | <p>PM Secondary (PMS)</p> <p>A Connection Manager uses this field to indicate whether or not the Adapter is a PM Secondary Adapter.</p> <p>0b – Adapter is not a PM Secondary Adapter</p> <p>1b – Adapter is a PM Secondary Adapter</p> | R/W | 1b |
| | | 31 | Reserved | RsvdZ | 0b |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|-------|---------------|
| 2 | LANE_ADP_CS_2 | 6:0 | Logical Layer Errors An Adapter uses this field to indicate the occurrence of a Logical Layer Error. See Section 4.4.2. Bit 0 – Aligner Lock Error (ALE) Bit 1 – Order Set Error (OSE) Bit 2 – Timing Error (TE) Bit 3 – Elastic Buffer Error (EBE) Bit 4 – De-Skew Buffer Error (DBE) Bit 5 – RS-FEC Decoder Error (RDE) Bit 6 – RX Sync Timeout (RST) | R/Clr | 00h |
| | | 15:7 | Reserved | Rsvd | 0 |
| | | 22:16 | Logical Layer Errors Enable A Connection Manager uses this field to enable Logical Layer error reporting. See Section 4.4.2. When a bit is set to 1b, the corresponding error is reported by a Notification Packet: Bit 0 – Alignment Lock Error (ALE) Bit 1 – Order Set Error (OSE) Bit 2 – Timing Error (TE) Bit 3 – Elastic Buffer Error (EBE) Bit 4 – De-Skew Buffer Error (DBE) Bit 5 – RS-FEC Decoder Error (RDE) Bit 6 – RX Sync Timeout (RST) | R/W | 00h |
| | | 31:23 | Reserved | Rsvd | 0 |

8.2.2.4 USB4 Port Capability

A USB4 Port Capability shall have the structure depicted in Figure 8-10 and shall have the fields defined in Table 8-13.

Figure 8-10. Structure of USB4 Port Capability




Table 8-13. USB4 Port Capability Fields

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|------|----------------|
| 0 | PORT_CS_0 | 7:0 | <p>Next Capability Pointer</p> <p>This field shall contain the Doubleword index of the next Capability in the Adapter Configuration Space.</p> <p>An Adapter shall set this field to 00h if the Capability is the final Capability in the linked list of Capabilities in the Adapter Configuration Space.</p> | RO | Vendor Defined |
| | | 15:8 | <p>Capability ID</p> <p>An Adapter shall set this field to 06h indicating this is USB4 Port Capability.</p> | RO | 06h |
| | | 31:16 | <p>Reserved</p> | Rsvd | 0 |
| 1 | PORT_CS_1 | 7:0 | <p>Address</p> <p>A Connection Manager uses this field when it initiates a read or write to SB Register Space.</p> <p>This field indicates the 8-bit address of the register in the SB Register Space that is the target of the read/write.</p> | R/W | 0 |
| | | 15:8 | <p>Length</p> <p>A Connection Manager uses this field when it initiates a read or write to SB Register Space.</p> <p>This field indicates the number of bytes to read/write.</p> <p>After executing a read or write to the SB Register Space, a Router shall set this field to the value of the <i>LEN</i> field in the AT Response, the RT Response, or the local access.</p> | R/W | 0 |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|--|--------|--|------|---------------|
| 1 | PORT_CS_1 | 18:16 | <p>Target</p> <p>A Connection Manager uses this field when it initiates a read or write to SB Register Space.</p> <p>This field defines which SB Register Space to access:</p> <ul style="list-style-type: none"> 000b – Router (via local access) 001b – Link Partner (via AT Transaction) 010b – Re-timer (via RT Transaction) <p>All other encodings are reserved.</p> | R/W | 0 |
| | | 19 | Reserved | Rsvd | 0 |
| | | 23:20 | <p>Re-timer Index</p> <p>A Connection Manager uses this field when it initiates a read or write to SB Register Space.</p> <p>When the <i>Target</i> field is 010b, this field contains the Re-timer Index of the target Re-timer.</p> <p>Otherwise, this field is reserved and set to 0.</p> | R/W | 0 |
| | | 24 | <p>WnR</p> <p>A Connection Manager uses this field to indicate whether it is initiating a read or a write:</p> <ul style="list-style-type: none"> 0b – Read 1b – Write | R/W | 0b |
| | | 25 | <p>No Response (NR)</p> <p>A Router uses this field when a Connection Manager initiates a read or write to SB Register Space.</p> <p>A Router shall set this bit to 1b if it did not receive a response for the read/write (including after any retransmissions).</p> <p>This field is only valid when the <i>Pending</i> bit is set to 0b.</p> | R/W | 0b |
| | | 26 | <p>Result Code (RC)</p> <p>A Router uses this field when a Connection Manager initiates a read or write to SB Register Space.</p> <p>For a Read operation:</p> <ul style="list-style-type: none"> A Router shall set this bit to 0b if the <i>LEN</i> field in the AT Response or the RT Response is greater than zero, or if a local access completes successfully. A Router shall set this bit to 1b if the <i>LEN</i> field in the AT Response or the RT Response is 0, or if a local access completes unsuccessfully. <p>For a Write operation:</p> <ul style="list-style-type: none"> A Router shall set this bit to the value of the Result Code in the AT Response or the RT Response. For a local access, a Router shall set this bit to 0b if the access completes successfully. A Router shall set this bit to 1b if the access completes unsuccessfully. <p>This field is only valid when the <i>Pending</i> bit is set to 0b.</p> | R/W | 0b |
| | | 30:27 | Reserved | Rsvd | 0 |
| 31 | <p>Pending (PND)</p> <p>A Connection Manager sets this bit to 1b to initiate a read or write to SB Register Space.</p> <p>A Router shall set this bit to 0b after it finishes the SB Register Space read/write.</p> | R/W | 0b | | |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|------|---------------------------|--------|--|------|---------------|
| 17:2 | PORT_CS_17 - PORT_CS_2 | 31:0 | <p>Data [15:0]</p> <p>A Router uses these fields when a Connection Manager initiates a read or write to SB Register Space.</p> <p>For a write:</p> <p style="padding-left: 20px;">The Connection Manager sets these fields to contain the Doublewords to be written to the SB Register Space.</p> <p>For a read:</p> <p style="padding-left: 20px;">The Router shall set these fields to contain the Doublewords read from the SB Register Space.</p> <p>Doublewords shall be arranged in increasing address order, starting at DW2 of the USB4 Port Capability and ending with the last Doubleword written/read.</p> | R/W | 0 |
| 18 | PORT_CS_18 | 7:0 | <p>Cable USB4 Version</p> <p>This field shall identify which version of the USB4 specification is supported by the USB Type-C® Cable where:</p> <p style="padding-left: 20px;">Bits 7:4 identify the major version</p> <p style="padding-left: 20px;">Bits 3:0 identify the minor version</p> <p>Allocated values:</p> <p style="padding-left: 20px;">0.x - Reserved for TBT3</p> <p style="padding-left: 20px;">1.0 - Version 1.0</p> <p>All other encodings are reserved.</p> | RO | 10h |
| | | 8 | <p>Bonding Enabled (BE)</p> <p>An Adapter shall set this bit to 1b when the conditions for Lane bonding are met (See Section 4.1.2.3). Otherwise, this bit shall be set to 0b.</p> | RO | 0 |
| | | 9 | <p>TBT3-Compatible Mode (TCM)</p> <p>An Adapter shall set this bit to 1b when the Link is operating in TBT3-Compatible Mode.</p> <p>This bit is set to 0b otherwise.</p> | RO | 0 |
| | | 10 | <p>CLx Protocol Support (CPS)</p> <p>This bit indicates that the Router supports the Low Power protocol. When a Router supports the Low Power protocol, it either accepts or rejects entry to a CLx state as defined in Section 4.2.1.6.</p> <p>A Router shall set this bit to 1b.</p> | RO | 1 |
| | | 11 | <p>RS-FEC Enabled (Gen 2) (RE2)</p> <p>An Adapter shall set this bit to 1b when the USB4 Port operates in Gen 2 and RS-FEC is enabled.</p> <p>This bit shall be set to 0b otherwise.</p> | RO | 0 |
| | | 12 | <p>RS-FEC Enabled (Gen 3) (RE3)</p> <p>An Adapter shall set this bit to 1b when the USB4 Port operates in Gen 3 and RS-FEC is enabled.</p> <p>This bit shall be set to 0b otherwise.</p> | RO | 0 |
| | | 13 | <p>Router Detected (RD)</p> <p>An Adapter shall set this bit to 1b when the USB4 Port detects a connected Router (see Section 4.1.2.2)</p> <p>An Adapter shall set this bit to 0b upon a disconnect.</p> | RO | 0 |
| | | 15:14 | Reserved | Rsvd | 0 |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|------|---------------|
| 18 | PORT_CS_18 | 16 | <p>Wake on Connect Status</p> <p>An Adapter shall set this bit to 1b after a wake event is generated by the USB4 Port as a result of a connect to the USB4 Port.</p> <p>This bit shall not be set to 1b unless the <i>Enable Wake on Connect</i> bit is 1b.</p> <p>This bit shall be set to 0b on entry to sleep.</p> | RO | 0 |
| | | 17 | <p>Wake on Disconnect Status</p> <p>An Adapter shall set this bit to 1b after a wake event is generated by the USB4 Port as a result of a disconnect from the USB4 Port.</p> <p>This bit shall not be set to 1b unless the <i>Enable Wake on Disconnect</i> bit is 1b.</p> <p>This bit shall be set to 0b on entry to sleep.</p> | RO | 0 |
| | | 18 | <p>Wake on USB4 Wake Status</p> <p>An Adapter shall set this bit to 1b after a wake event is generated by the USB4 Port as a result of a USB4 Wake.</p> <p>This bit shall not be set to 1b unless the <i>Enable Wake on USB4 Wake</i> bit is 1b.</p> <p>This bit shall be set to 0b on entry to sleep.</p> | RO | 0 |
| | | 19 | <p>Wake on Inter-Domain Status</p> <p>An Adapter shall set this bit to 1b after a wake event is generated by the USB4 Port as a result of an Inter-Domain Wake.</p> <p>This bit shall not be set to 1b unless the <i>Enable Wake on Inter-Domain</i> bit is 1b.</p> <p>This bit shall be set to 0b on entry to sleep.</p> | RO | 0 |
| | | 31:20 | Reserved | Rsvd | 0 |
| 19 | PORT_CS_19 | 0 | <p>Downstream Port Reset (DPR)</p> <p>A Connection Manager uses this bit to reset a Downstream Facing Port.</p> <p>A Downstream Facing Port shall initiate a Downstream Port Reset when this bit is set to 1b.</p> <p>Setting this bit to 0b shall transition the Adapters in the Downstream Facing Port out of the CLd state.</p> <p>An Adapter in an Upstream Facing Port shall ignore this bit and may hardwire the bit to 0b.</p> | R/W | 0b |
| | | 1 | <p>Request RS-FEC Gen 2 (RS2)</p> <p>A Connection Manager uses this bit to request enabling of RS-FEC encoding at Gen 2 speeds. If a Link is Active, the Link must be re-initialized before RS-FEC can be enabled.</p> <p>If this bit is set to 0b, the USB4 Port shall disable RS-FEC at Gen 2 speeds during the next Lane Initialization.</p> <p>If this bit is set to 1b, then the Link Partner response determines whether or not RS-FEC is enabled at Gen 2 speeds.</p> | R/W | 1b |
| | | 2 | <p>Request RS-FEC Gen 3 (RS3)</p> <p>A Connection Manager uses this bit to request enabling of RS-FEC encoding at Gen 3 speeds. If a Link is Active, the Link must be re-initialized before RS-FEC can be enabled.</p> <p>If this bit is set to 0b, the USB4 Port shall disable RS-FEC at Gen 3 speeds during the next Lane Initialization.</p> <p>If this bit is set to 1b, then the Link Partner response determines whether or not RS-FEC is enabled at Gen 3 speeds.</p> | R/W | 1b |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|--|------|---------------|
| 19 | PORT_CS_19 | 3 | USB4 Port is Configured (PC) A Connection Manager sets this bit as follows: 1b – Router connected to the USB4 Port is configured and therefore supports sleep state 0b – Router connected to the USB4 Port is not configured and therefore does not support sleep | R/W | 0b |
| | | 4 | USB4 Port is Inter-Domain (PID) A Connection Manager sets this bit as follows: 1b – USB4 Port is part of an Inter-Domain Link 0b – USB4 Port is not part of an Inter-Domain Link | R/W | 0b |
| | | 15:5 | Reserved | Rsvd | 0 |
| | | 16 | Enable Wake on Connect A Connection Manager uses this bit to configure whether or not a connection on the USB Type-C connector attached to this USB4 Port causes the Router to exit the sleep state. 1b – Connection causes exit from sleep 0b – Connection does not cause exit from sleep | R/W | 0 |
| | | 17 | Enable Wake on Disconnect A Connection Manager uses this bit to configure whether or not a disconnect on the USB Type-C connector attached to this USB4 Port causes the Router to exit sleep state. 1b – Disconnect causes exit from sleep 0b – Disconnect does not cause exit from sleep | R/W | 0 |
| | | 18 | Enable Wake on USB4 Wake A Connection Manager uses this bit to configure whether or not a USB4 Wake causes the Router to exit the sleep state. 1b – USB4 Wake causes exit from sleep 0b – USB4 Wake does not cause exit from sleep | R/W | 1 |
| | | 19 | Enable Wake on Inter-Domain A Connection Manager uses this bit to configure whether or not an Inter-Domain Wake causes the Router to exit the sleep state. 1b – Inter-Domain Wake causes exit from sleep 0b – Inter-Domain Wake does not cause exit from sleep | R/W | 0 |
| | | 31:20 | Reserved | Rsvd | 0 |



CONNECTION MANAGER NOTE

When writing to a USB4 Port Capability, a Connection Manager shall abide by the following rules:

- The Target Link Speed field in the Lane Adapters of a USB4 Port shall be set to the same value.
- The Target Link Width field in the Lane Adapters of a USB4 Port shall be set to the same value.
- A Connection Manager shall not write to any of DW1 through DW17 (inclusive) of the USB4 Port Capability (PORT_CS_1) when the Pending bit is 1b.
- A Connection Manager shall not set the Length field to a value greater than 64.
- When the Target field is 010b, the Connection Manager shall set the Re-timer Index field to the index of the target Re-timer. When the Target field is not 010b, the Connection Manager shall set the Re-timer Index field to 0.

8.2.2.5 USB3 Adapter Configuration Capability

A USB3 Adapter Configuration Capability shall have the structure depicted in Figure 8-11 and shall have the fields defined in Table 8-14.

Figure 8-11. Structure of USB3 Adapter Configuration Capability




Table 8-14. USB3 Adapter Configuration Capability Fields

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|------|----------------|
| 0 | ADP_USB3_CS_0 | 7:0 | Next Capability Pointer This field shall contain the Doubleword index of the next Capability in the Adapter Configuration Space. An Adapter shall set this field to 00h if the Capability is the final Capability in the linked list of Capabilities in the Adapter Configuration Space. | RO | Vendor Defined |
| | | 15:8 | Capability ID An Adapter shall set this field to 04h indicating this is a Protocol Adapter Configuration Capability. | RO | 04h |
| | | 29:16 | Reserved | Rsvd | 0 |
| | | 30 | Valid (V) A Connection Manager sets this bit to 1b when the contents of the Capability are valid. When this bit is set to 0b, the contents of the Capability are not valid. | R/W | 0b |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|--|------|---------------|
| 0 | ADP_USB3_CS_0 | 31 | <p>Path Enable (PE)</p> <p>A Connection Manager uses this bit to enable transmission and reception of Tunneled Packets.</p> <p>When set to 1b, the Port may send and receive Tunneled Packets.</p> <p>When set to 0b, the Port shall not send or receive Tunneled Packets.</p> | R/W | 0b |
| 1 | ADP_USB3_CS_1 | 11:0 | <p>Consumed Upstream Bandwidth</p> <p>This field shall contain the amount of upstream bandwidth consumed for isochronous USB3 traffic. The upstream bandwidth is given as the maximal number of bytes consumed within each USB3 frame of 1 ms in multiple of $(512 * 2^{\text{Scale}})$ bytes.</p> <p>A Router shall not update this field when the <i>Host Controller Ack</i> bit is set to 1b.</p> <p>This field shall be hardwired to 0 for a Device Router.</p> | RO | 0 |
| | | 23:12 | <p>Consumed Downstream Bandwidth</p> <p>This field shall contain the amount of downstream bandwidth consumed for isochronous USB3 traffic. The downstream bandwidth is given as the maximal the number of bytes consumed within each USB3 frame of 1 ms in multiple of $(512 * 2^{\text{Scale}})$ bytes.</p> <p>A Router shall not update this field when the <i>Host Controller Ack</i> bit is set to 1b.</p> <p>This field shall be hardwired to 0 for a Device Router.</p> | RO | 0 |
| | | 30:24 | Reserved | Rsvd | 0 |
| | | 31 | <p>Host Controller Ack (HCA)</p> <p>A Router shall set this bit to 1b when a Connection Manager is allowed to read the <i>Consumed Upstream Bandwidth</i> and <i>Consumed Downstream Bandwidth</i> fields or update the <i>Allocated Upstream Bandwidth</i> or <i>Allocated Downstream Bandwidth</i> fields. Otherwise, this bit shall be set to 0b.</p> <p>This bit shall be hardwired to 0b for a Device Router.</p> | RO | 0b |
| 2 | ADP_USB3_CS_2 | 11:0 | <p>Allocated Upstream Bandwidth</p> <p>This field shall contain the amount of upstream bandwidth allocated for isochronous USB3 traffic. The upstream bandwidth is given as the maximal number of available bytes within each USB3 frame of 1 ms in multiple of $512 * 2^{\text{Scale}}$ bytes.</p> <p>This field shall be hardwired to 0 for a Device Router.</p> <p>A Connection Manager shall not set this field to a higher bandwidth than the Maximum Supported Link Rate</p> | R/W | 0 |
| | | 23:12 | <p>Allocated Downstream Bandwidth</p> <p>This field shall contain the amount of downstream bandwidth allocated for isochronous USB3 traffic. The downstream bandwidth is given as the maximal number of available bytes within each USB3 frame of 1 ms in multiple of $512 * 2^{\text{Scale}}$ bytes.</p> <p>This field shall be hardwired to 0 for a Device Router.</p> <p>A Connection Manager shall not set this field to a higher bandwidth than the Maximum Supported Link Rate</p> | R/W | 0 |
| | | 30:24 | Reserved | Rsvd | 0 |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|--|------|----------------|
| | | 31 | <p>Connection Manager Request (CMR)</p> <p>The Connection Manager sets this bit to 1b to request permission to read the <i>Consumed Upstream Bandwidth</i> and <i>Consumed Downstream Bandwidth</i> fields or update the <i>Allocated Upstream Bandwidth</i> or <i>Allocated Downstream Bandwidth</i> fields.</p> <p>The Connection Manager sets this bit to 0b after it has completed the read/update.</p> <p>The Connection Manager set this bit to 0 when the internal host controller is not enabled in the Router.</p> | R/W | 0b |
| 3 | ADP_USB3_CS_3 | 5:0 | <p>Scale</p> <p>A Connection Manager uses this field to set the granularity of bandwidth negotiated in the <i>Consumed Upstream Bandwidth</i> field, the <i>Consumed Downstream Bandwidth</i> field, the <i>Available Upstream Bandwidth</i> field, and the <i>Available Downstream Bandwidth</i> field.</p> <p>A Device Router shall hardwire this field to 0.</p> <p>A Connection Manager shall set this field to 0 if the Maximum Supported Link Rate equals to 0.</p> | R/W | 0 |
| | | 31:6 | Reserved | Rsvd | 0 |
| 4 | ADP_USB3_CS_4 | 6:0 | <p>Actual Link Rate</p> <p>0h: 10 Gbps (Gen 2 - Single-Lane) 1h: 20 Gbps (Gen 2 - Dual-Lane) 7Fh - 2h: Reserved</p> | RO | 0 |
| | | 7 | <p>USB3 Link Valid (ULV)</p> <p>0: <i>Actual Link Rate</i> field is not valid 1: <i>Actual Link Rate</i> field is valid</p> | RO | 0 |
| | | 11:8 | <p>Port Link State (PLS)</p> <p>This field shall indicate the port link state of the internal USB3 device port connected to the USB3 Adapter Layer:</p> <p>0h: U0 state 1h: Reserved 2h: U2 state 3h: U3 state 4h: Disabled state 5h: RxDetect state 6h: Inactive state 7h: Polling state 8h: Recovery state 9h: Hot Reset state Ah - Eh: Reserved Fh: Resume state</p> | RO | Vendor Defined |
| | | 18:12 | <p>Maximum Supported Link Rate</p> <p>0h: 10 Gbps (Gen 2 - Single-Lane) 1h: 20 Gbps (Gen 2 - Dual-Lane) 7Fh - 2h: Reserved</p> | RO | Vendor Defined |
| | | 31:19 | Reserved | Rsvd | 0 |

8.2.2.6 DP Adapter Configuration Capability

A DP IN Adapter Configuration Capability shall have the structure depicted in Figure 8-12 and shall have the fields defined in Table 8-15. A DP OUT Adapter Configuration Capability shall have the structure depicted in Figure 8-13 and shall have the fields defined in Table 8-16.

| | | | | | |
|-------|--|-------|--|-------|---|
| | | 31 | Video Enable (VE) A Connection Manager uses this bit to control when the DP IN Adapter can send and receive packets on the Video Path. When set to 1, the Adapter may send packets on the Video Path. When set to 0, the Adapter shall not send packets on the Video Path. | R/W | 0 |
| 1 | ADP_DP_CS_1 | 10:0 | AUX Tx HopID An Adapter shall set this field to 8. | RO | 8 |
| | | 21:11 | AUX Rx HopID An Adapter shall set this field to 8. | RO | 8 |
| | | 31:22 | Reserved | Rsvd | 0 |
| 2 | ADP_DP_CS_2 | 2:0 | NRD Max Lane Count (NRD MLC) The Non-Reduced Max Lane Count is set by a Connection Manager to reflect the highest common Max Lane Count between the two DP Adapters, regardless of bandwidth availability. | R/W | 0 |
| | | 3 | SW Link Init (SWLI) A Connection Manager uses this bit to initiate Link-Initialization. When this bit transitions from 0 to 1, the Adapter shall initiate Link-Init as described in Section 10.4.13. | R/W | 0 |
| | | 5:4 | Reserved | RsvdZ | 0 |
| | | 6 | HPD Status This field shall contain the HPD value received from the DP OUT Adapter. | RO | 0 |
| | | 9:7 | NRD Max Link Rate (NRD MLR) The Non-Reduced Max Link Rate is set by a Connection Manager to reflect the highest common Max Link Rate between the two DP Adapters, regardless of bandwidth availability. | R/W | 0 |
| | | 10 | CM Ack (CA) A Connection Manger sets this field to 1b to indicate that it completed the bandwidth allocation process. A Connection Manager sets this bit to 0b after the DP IN Adapter sets <i>DPTX Req</i> field to 0b. | R/W | 0 |
| | | 12:11 | Granularity (GR) A Connection Manager sets this field to indicate the bandwidth granularity for the <i>Requested BW</i> , <i>Allocated BW</i> and <i>Estimated BW</i> : 0: 0.25 1: 0.5 2: 1.0 3: Reserved The units are in Gbps. | R/W | 0 |
| | | 15:13 | Group_ID A Connection Manager uses this field to indicate the group this Path is associated with: 0h – The Path is not associated to any group. 1-7h – The group number the Path is associated with. | R/W | 0 |
| 19:16 | CM_ID A Connection Manager uses this field to indicate its own index, in case there are multiple CMs in the Host system. | R/W | 0 | | |

| | | | | | |
|---|--------------|-------|--|-------|----------------|
| | | 20 | CM BW Allocation Mode Support (CMMS) 0: Not supported 1: Supported | R/W | Vendor Defined |
| | | 23:21 | Reserved | RsvdZ | 0 |
| | | 31:24 | Estimated BW A Connection Manager uses this field to indicate the estimated available bandwidth for the DP IN Adapter. | R/W | 0 |
| 3 | ADP_DP_CS_3 | 8:0 | Vendor Defined | VD | Vendor Defined |
| | | 9 | HPD Output Clear (HPDC) A Connection Manager uses this field to clear the HPD output. When this bit is 1b, an Adapter shall drive HPD low to cause a single event of HPD output clear. | R/W | 0 |
| | | 10 | HPD Output Set (HPDS) A Connection Manager uses this field to set the HPD output. When this bit is 1b, a Port shall drive HPD high to cause a single event of HPD output set. | R/W | 0 |
| | | 31:11 | Vendor Defined | VD | Vendor Defined |
| DP_LOCAL_CAP Register reflects the static local capabilities of the DP IN Adapter per connection. For example, Connecting a DP IN Adapter as DP ALT-Mode in Multi-Function operation limits the <i>Maximal Lane Count</i> field in this register. | | | | | |
| 4 | DP_LOCAL_CAP | 3:0 | Protocol Adapter Version This field shall identify which version of the USB4 Specification the Adapter supports. 0h – 2h: Reserved 3h: Reserved for TBT3 4h: Version 1.0 5h – Fh: Reserved | RO | 4 |
| | | 7:4 | Maximal DPCD Rev 0: DPCD r1.1 1: DPCD r1.2 2: DPCD r1.3 3: DPCD r1.4a 4 – Fh: Reserved | RO | Vendor Defined |
| | | 11:8 | Maximal Link Rate 0: 1.62GHz 1: 2.7GHz 2: 5.4GHz 3: 8.1GHz 4 – Fh: Reserved | RO | Vendor Defined |
| | | 14:12 | Maximal Lane Count 0: 1 lane 1: 2 lanes 2: 4 lanes 3 – 7: Reserved If the DP IN Adapter was connected as part of MFDP, this field shall not indicate 4 lanes. | RO | Vendor Defined |
| | | 15 | MST Capability 0: Not supported 1: Supported | RO | Vendor Defined |
| | | 21:16 | Reserved | Rsvd | 0 |

| | | | | | |
|--|---------------|-------|--|------|----------------|
| | | 22 | TPS3 Capability 0: Not supported 1: Supported | RO | Vendor Defined |
| | | 23 | Reserved | Rsvd | 1 |
| | | 24 | TPS4 Capability 0: Not supported 1: Supported | RO | Vendor Defined |
| | | 25 | FEC Not Supported 0: Supported 1: Not supported | RO | Vendor Defined |
| | | 26 | Secondary Split Capability 0: Not Supported 1: Supported | RO | Vendor Defined |
| | | 27 | LTTPR Not Supported 0: Supported 1: Not supported | RO | Vendor Defined |
| | | 28 | DP IN BW Allocation Mode Support 0: Not supported 1: Supported | Rsvd | Vendor Defined |
| | | 29 | DSC Not Supported 0: Supported 1: Not supported | RO | Vendor Defined |
| | | 31:30 | Reserved | Rsvd | 0 |
| The DP_REMOTE_CAP register reflects the local capabilities of the paired DP OUT Adapter. The Values in this register are set by the Connection Manager during Path Configuration. A DP IN Adapter shall reset the fields in this register to their default values when the DP OUT Adapter is unpaired. | | | | | |
| 5 | DP_REMOTE_CAP | 3:0 | Protocol Adapter Version Identifies what version of the USB4 Specification is supported by the paired DP OUT Adapter 0h – 2h: Reserved 3h: Reserved for TBT3 4h: Version 1.0 5h – Fh: Reserved | R/W | 0 |
| | | 7:4 | Maximal DPCD Rev 0h: DPCD r1.1 1h: DPCD r1.2 2h: DPCD r1.3 3h: DPCD r1.4a 4h – Fh: Reserved | R/W | 0 |
| | | 11:8 | Maximal Link Rate 0h: 1.62GHz 1h: 2.7GHz 2h: 5.4GHz 3h: 8.1GHz 4h – Fh: Reserved | R/W | 0 |
| | | 14:12 | Maximal Lane Count 0: 1 lane 1: 2 lanes 2: 4 lanes 3 – 7: Reserved | R/W | 0 |

| | | | | | |
|---|-----------|-------|--|-------|---|
| | | 15 | MST Capability 0: Not supported 1: Supported | R/W | 0 |
| | | 21:16 | Reserved | RsvdZ | 0 |
| | | 22 | TPS3 Capability 0: Not supported 1: Supported | R/W | 0 |
| | | 23 | Reserved | RsvdZ | 0 |
| | | 24 | TPS4 Capability 0: Not supported 1: Supported | R/W | 0 |
| | | 25 | FEC Not Supported 0: Supported 1: Not supported | R/W | 0 |
| | | 26 | Secondary Split Capability 0: Not supported 1: Supported | R/W | 0 |
| | | 27 | LTPR Not Supported 0: Supported 1: Not Supported | R/W | 0 |
| | | 28 | Reserved | RsvdZ | 0 |
| | | 29 | DSC Not Supported 0: Supported 1: Not supported | R/W | 0 |
| | | 31:30 | Reserved | RsvdZ | 0 |
| The <i>Lane Count</i> and <i>Link Rate</i> fields are only valid after the DP Link is established. | | | | | |
| 6 | DP_STATUS | 2:0 | Lane Count 1: 1 lane 2: 2 lanes 4: 4 lanes 0; 3; 5 - 7: Reserved | RO | 0 |
| | | 7:3 | Reserved | Rsvd | 0 |
| | | 11:8 | Link Rate 0h: 1.62GHz 1h: 2.7GHz 2h: 5.4GHz 3h: 8.1GHz 4h - Fh: Reserved | RO | 0 |
| | | 16:12 | Reserved | Rsvd | 0 |
| | | 23:17 | Reserved | RsvdZ | 0 |
| | | 31:24 | Allocated BW A Connection Manager uses this field to indicate the allocated bandwidth for the DP IN Adapter. | R/W | 0 |
| DP_COMMON_CAP fields reflects the lowest common capability between DP_LOCAL_CAP fields and DP_REMOTE_CAP fields. The DP_COMMON_CAP fields shall be updated any time the DP_REMOTE_CAP fields are updated. | | | | | |

| | | | | | |
|----|-----------------|-------|---|------|---|
| 7 | DP_COMMON_CAP | 3:0 | Protocol Adapter Version This field shall identify the highest common version of the USB4 Specification that is supported by both the DP IN Adapter and the DP OUT Adapter. 0h – 2h: Reserved 3h: Reserved for TBT3 4h: Version 1.0 5h – Fh: Reserved | RO | 0 |
| | | 7:4 | Maximal DPCD Rev 0h: DPCD r1.1 1h: DPCD r1.2 2h: DPCD r1.3 3h: DPCD r1.4a 4h – Fh: Reserved | RO | 0 |
| | | 11:8 | Maximal Link Rate 0h: 1.62GHz 1h: 2.7GHz 2h: 5.4GHz 3h: 8.1GHz 4h – Fh: Reserved | RO | 0 |
| | | 14:12 | Maximal Lane Count 0: 1 lane 1: 2 lanes 2: 4 lanes 3 – 7: Reserved | RO | 0 |
| | | 15 | MST Capability 0: Not supported 1: Supported | RO | 0 |
| | | 21:16 | Reserved | Rsvd | 0 |
| | | 22 | TPS3 Capability 0: Not supported 1: Supported | RO | 0 |
| | | 23 | Reserved | Rsvd | 0 |
| | | 24 | TPS4 Capability 0: Not supported 1: Supported | RO | 0 |
| | | 25 | FEC Not Supported 0: Supported 1: Not supported | RO | 0 |
| | | 26 | Secondary Split Capability 0: Not supported 1: Supported | RO | 0 |
| | | 27 | LTPR Not Supported 0: Supported 1: Not Supported | RO | 0 |
| | | 28 | Reserved | Rsvd | 0 |
| | | 29 | DSC Not Supported 0: Supported 1: Not supported | RO | 0 |
| 30 | Reserved | RO | 0 | | |

| | | | | | |
|---|-------------|------|---|----|---|
| | | 31 | DPRX Capabilities Read Done 0: Not Done 1: Done A DP Adapter shall set the value of this field after DPCD addresses 00001h and 00002h are read. See Section 10.4.6.2. | RO | 0 |
| A Connection Manager does not address this register if the DP IN Adapter does not support the DP BW Allocation Mode (i.e. DP_LOCAL_CAP.DPIN BW Allocation Mode Support = 1b). | | | | | |
| 8 | ADP_DP_CS_8 | 7:0 | Requested BW A DP IN Adapter uses this field to reflect the requested bandwidth by the DPTX to the Connection Manager. | RO | 0 |
| | | 29:8 | Reserved | RO | 0 |
| | | 30 | DPTX BW Allocation Mode Enable (DPME) 0: Disabled (default) 1: Enabled | RO | 0 |
| | | 31 | DPTX Req(DR) A DP IN Adapter sets this field to 1b when DPTX writes to the REQUESTED_BW register. A DP IN Adapter sets this field to 0b after the Connection Manager sets the <i>CM Ack</i> bit to 1b. | RO | 0 |

Figure 8-13. Structure of DP OUT Adapter Configuration Capability

| | |
|----|---|
| 31 | 0 |
| 0 | VE AE RSV Video HopID Capability ID Next Capability pointer |
| 1 | Reserved AUX Rx HopID AUX Tx HopID |
| 2 | Reserved Maximum Accumulation Cycles RSV RSV RSV RSV |
| 3 | Vendor Defined |
| 4 | LOCAL Capabilities |
| 5 | REMOTE Capabilities |
| 6 | Reserved UFS CMHS Reserved LR Reserved LC |
| 7 | COMMON Capabilities |
| 8 | Reserved |

Table 8-16. DP OUT Adapter Configuration Capability Fields

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|------|---------------|
| 0 | ADP_DP_CS_0 | 7:0 | Next Capability Pointer This field shall contain the Doubleword index of the next Capability in the Adapter Configuration Space. An Adapter shall set this field to 00h if the Capability is the final Capability in the linked list of Capabilities in the Adapter Configuration Space. | RO | (VSEC) |

| | | | | | |
|--|-------------|-------|--|------|----------------|
| | | 15:8 | Capability ID An Adapter shall set this field to 04h indicating this is a Protocol Adapter Configuration Capability. | RO | 04h |
| | | 26:16 | Video HopID An Adapter shall set this bit to 9. | RO | 9 |
| | | 29:27 | Reserved | Rsvd | 0 |
| 0 | ADP_DP_CS_0 | 30 | AUX Enable (AE) A Connection Manager uses this bit to control when the DP OUT Adapter can send and receive packets on the AUX Path. When set to 1, the Adapter may send and receive packets on the AUX Path. When set to 0, the Adapter shall not send or receive packets on the AUX Path. | R/W | 0 |
| | | 31 | Video Enable (VE) A Connection Manager uses this bit to control when the DP OUT Adapter can receive packets on the Video Path. When set to 1, the Adapter may receive packets on the Video Path. When set to 0, the Adapter shall not receive packets on the Video Path. | R/W | 0 |
| 1 | ADP_DP_CS_1 | 10:0 | AUX Tx HopID An Adapter shall set this field to 8. | RO | 8 |
| | | 21:11 | AUX Rx HopID An Adapter shall set this field to 8. | RO | 8 |
| | | 31:22 | Reserved | Rsvd | 0 |
| 2 | ADP_DP_CS_2 | 2:0 | Reserved | Rsvd | 0 |
| | | 3 | SW Link Init (SWLI) A Connection Manager uses this bit to initiate Link-Initialization. When this bit transitions from 0 to 1, the Adapter shall initiate Link-Init as described in Section 10.4.13. | R/W | 0 |
| | | 5:4 | Reserved | Rsvd | 0 |
| | | 6 | HPD Status This field shall contain the HPD value sent to DP IN Adapter. | RO | 0 |
| | | 7 | Reserved | Rsvd | 0 |
| | | 23:8 | Maximum Accumulation Cycles The number of DP Link clock cycles that a DP OUT Adapter accumulates Main-Link Data when operating at the highest supported DP Link Rate. | RO | Vendor Defined |
| | | 31:24 | Reserved | Rsvd | 0 |
| 3 | ADP_DP_CS_3 | 31:0 | Vendor Defined | VD | Vendor Defined |
| DP_LOCAL_CAP Register reflects the static local capabilities of the DP OUT Adapter per connection. For example, Connecting a DP OUT Adapter as DP ALT-Mode in Multi-Function operation limits the Maximal Lane Count in this register. | | | | | |

| | | | | | |
|-------|-----------------|-------|--|------|----------------|
| 4 | DP_LOCAL_CAP | 3:0 | Protocol Adapter Version This field shall identify which version of the USB4 Specification the DP OUT Adapter supports 0h – 2h: Reserved 3h: Reserved for TBT3 4h: Version 1.0 5h – Fh: Reserved | RO | 4 |
| | | 7:4 | Maximal DPCD Rev 0h: DPCD r1.1 1h: DPCD r1.2 2h: DPCD r1.3 3h: DPCD r1.4a 4h – Fh: Reserved | RO | Vendor Defined |
| 4 | DP_LOCAL_CAP | 11:8 | Maximal Link Rate 0h: 1.62GHz 1h: 2.7GHz 2h: 5.4GHz 3h: 8.1GHz 4h – Fh: Reserved | RO | Vendor Defined |
| | | 14:12 | Maximal Lane Count 0: 1 lane 1: 2 lanes 2: 4 lanes 3 – 7: Reserved If the DP OUT Adapter was connected as part of MFD, this field shall not indicate 4 lanes. | RO | Vendor Defined |
| | | 15 | MST Capability 0: Not supported 1: Supported | RO | Vendor Defined |
| | | 21:16 | Reserved | Rsvd | 0 |
| | | 22 | TPS3 Capability 0: Not supported 1: Supported | RO | Vendor Defined |
| | | 23 | Reserved | Rsvd | 1 |
| | | 24 | TPS4 Capability 0: Not supported 1: Supported | RO | Vendor Defined |
| | | 25 | FEC Not Supported 0: Supported 1: Not supported | RO | Vendor Defined |
| | | 26 | Secondary Split Capability 0: Not supported 1: Supported | RO | Vendor Defined |
| | | 27 | LTPR Not Supported 0: Supported 1: Not supported | RO | Vendor Defined |
| | | 28 | Reserved | Rsvd | 0 |
| | | 29 | DSC Not Supported 0: Supported 1: Not supported | RO | Vendor Defined |
| 31:30 | Reserved | Rsvd | 0 | | |

| | | | | | |
|---|-----------------|-------|---|-------|---|
| DP_REMOTE_CAP register reflects the local capabilities of the paired DP IN Adapter. The Values in this register are set by the Connection Manager during Path Configuration. A DP OUT Adapter shall reset the fields in this register to their default values when the DP IN Adapter is unpaired. | | | | | |
| 5 | DP_REMOTE_CAP | 3:0 | Protocol Adapter Version Identifies what version of the USB4 Specification is supported by the paired DP IN Adapter 0h – 2h: Reserved 3h: Reserved for TBT3 4h: Version 1.0 5h – Fh: Reserved | R/W | 0 |
| 5 | DP_REMOTE_CAP | 7:4 | Maximal DPCD Rev 0h: DPCD r1.1 1h: DPCD r1.2 2h: DPCD r1.3 3h: DPCD r1.4a 4h – Fh: Reserved | R/W | 0 |
| | | 11:8 | Maximal Link Rate 0h: 1.62GHz 1h: 2.7GHz 2h: 5.4GHz 3h: 8.1GHz 4h – Fh: Reserved | R/W | 0 |
| | | 14:12 | Maximal Lane Count 0: 1 lane 1: 2 lanes 2: 4 lanes 3 – 7: Reserved | R/W | 0 |
| | | 15 | MST Capability 0: Not supported 1: Supported | R/W | 0 |
| | | 21:16 | Reserved | RsvdZ | 0 |
| | | 22 | TPS3 Capability 0: Not supported 1: Supported | R/W | 0 |
| | | 23 | Reserved | RsvdZ | 0 |
| | | 24 | TPS4 Capability 0: Not supported 1: Supported | R/W | 0 |
| | | 25 | FEC Not Supported 0: Supported 1: Not supported | R/W | 0 |
| | | 26 | Secondary Split Capability 0: Not supported 1: Supported | R/W | 0 |
| | | 27 | LTTTPR Not Supported 0: Supported 1: Not supported | R/W | 0 |
| | | 28 | Reserved | RsvdZ | 0 |
| | | 29 | DSC Not Supported 0: Supported 1: Not supported | R/W | 0 |
| 31:30 | Reserved | RsvdZ | 0 | | |

| | | | | | |
|--|-----------------|-------|--|-------|---|
| 6 | DP_STATUS_CTRL | 2:0 | Lane Count 0: DP Main Link is not Active 1: 1 Lane 2: 2 Lanes 4: 4 Lanes 3; 5 - 7: Reserved | RO | 0 |
| | | 7:3 | Reserved | Rsvd | 0 |
| 6 | DP_STATUS_CTRL | 11:8 | Link Rate 0h: 1.62GHz 1h: 2.7GHz 2h: 5.4GHz 3h: 8.1GHz 4h - Fh: Reserved | RO | 0 |
| | | 16:12 | Reserved | Rsvd | 0 |
| | | 24:17 | Reserved | RsvdZ | 0 |
| | | 25 | CM Handshake (CMHS) A Connection Manager uses this bit for the handshake defined in Section 10.4.2.1. | R/W | 0 |
| | | 26 | DP IN Adapter USB4 Flag (UF) A Connection Manager uses this bit to indicate whether the DP IN Adapter is a USB4 DP IN Adapter or a TBT3 DP IN Adapter. 0: TBT3 1: USB4 | R/W | 0 |
| 31:27 | Reserved | RsvdZ | 0 | | |
| This register reflects the lowest common capability between DP_LOCAL_CAP fields and DP_REMOTE_CAP fields. The DP_COMMON_CAP fields shall be updated any time the DP_REMOTE_CAP fields are updated. | | | | | |
| 7 | DP_COMMON_CAP | 3:0 | Protocol Adapter Version This field shall identify the highest common version of the USB4 Specification that is supported by both the DP OUT Adapter and the DP IN Adapter 0h - 2h: Reserved 3h: Reserved for TBT3 4h: Version 1.0 5h - Fh: Reserved | RO | 0 |
| | | 7:4 | Maximal DPCD Rev 0h: DPCD r1.1 1h: DPCD r1.2 2h: DPCD r1.3 3h: DPCD r1.4a 4h - Fh: Reserved | RO | 0 |
| | | 11:8 | Maximal Link Rate 0h: 1.62GHz 1h: 2.7GHz 2h: 5.4GHz 3h: 8.1GHz 4h - Fh: Reserved | RO | 0 |
| | | 14:12 | Maximal Lane Count 0: 1 lane 1: 2 lanes 2: 4 lanes 3 - 7: Reserved | RO | 0 |

| | | | | | |
|---|---------------|-------|---|------|---|
| | | 15 | MST Capability 0: Not supported 1: Supported | RO | 0 |
| | | 21:16 | Reserved | Rsvd | 0 |
| | | 22 | TPS3 Capability 0: Not supported 1: Supported | RO | 0 |
| | | 23 | Reserved | Rsvd | 0 |
| 7 | DP_COMMON_CAP | 24 | TPS4 Capability 0: Not supported 1: Supported | RO | 0 |
| | | 25 | FEC Not Supported 0: Supported 1: Not supported | RO | 0 |
| | | 26 | Secondary Split Capability 0: Not supported 1: Supported | RO | 0 |
| | | 27 | LTTPR Not Supported 0: Supported 1: Not supported | RO | 0 |
| | | 28 | Reserved | Rsvd | 0 |
| | | 29 | DSC Not Supported 0: Supported 1: Not supported | RO | 0 |
| | | 31:30 | Reserved | Rsvd | 0 |
| 8 | ADP_DP_CS_8 | 31:0 | Reserved | Rsvd | 0 |

8.2.2.7 PCIe Adapter Configuration Capability

A PCIe Adapter Configuration Capability shall have the structure depicted in Figure 8-14 and shall have the fields defined in Table 8-17.

Figure 8-14. Structure of PCIe Adapter Configuration Capability



Table 8-17. PCIe Adapter Configuration Capability Fields

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|--|------|----------------|
| 0 | ADP_PCIE_CS_0 | 7:0 | Next Capability Pointer This field shall contain the Doubleword index of the next Capability in the Adapter Configuration Space. An Adapter shall set this field to 00h if the Capability is the final Capability in the linked list of Capabilities in the Adapter Configuration Space. | RO | Vendor Defined |
| | | 15:8 | Capability ID An Adapter shall set this field to 04h indicating this is a Protocol Adapter Configuration Capability. | RO | 04h |

| | | | | | |
|---|---------------|-------|---|----|----------------|
| | | 16 | <p>Link</p> <p>An Adapter shall set this bit to indicate the LinkUp state of the PCIe Physical Layer Logical Sub-block above the PCIe Adapter:</p> <p>0: Link is down 1: Link is up</p> | RO | 0b |
| | | 17 | <p>TX EI</p> <p>An Adapter shall set this bit to indicate whether or not the PCIe Physical Layer Logical Sub-block above the PCIe Adapter is in Electrical Idle state for its transmitter:</p> <p>0b: Transmitter is not in Electrical Idle state 1b: Transmitter is in Electrical Idle state</p> | RO | 1b |
| 0 | ADP_PCIE_CS_0 | 18 | <p>RX EI</p> <p>An Adapter shall set this bit to indicate whether or not the PCIe Physical Layer Logical Sub-block above the PCIe Adapter is in Electrical Idle state for its receiver:</p> <p>0b: Receiver is not in Electrical Idle state 1b: Receiver is in Electrical Idle state</p> | RO | 1b |
| | | 19 | <p>RST</p> <p>An Adapter shall set this bit to indicate whether or not the attached PCIe Switch Port is in PCIe Warm Reset/PCIe domain is active:</p> <p>0b: PCIe Switch Port is not in reset 1b: PCIe Switch Port is in reset</p> | RO | 1b |
| | | 24:20 | Vendor Defined | VD | Vendor Defined |
| | | 28:25 | <p>LTSSM</p> <p>An Adapter shall set this bit to indicate the LTSSM state in the PCIe Physical Layer Logical Sub-block above the PCIe Adapter:</p> <p>0h: Detect state 1h: Polling state 2h: Configuration state 3h: Configuration.Idle state 4h: Recovery state 5h: Recovery.idle state 6h: L0 state 7h: L1 state 8h: L2 state 9h: Disabled state Ah: Hot Reset state Bh - Fh: Reserved</p> | RO | 0h |
| | | 30:29 | Vendor Defined | VD | Vendor Defined |
| | | 31 | <p>Path Enable (PE)</p> <p>A Connection Manager uses this bit to control when a PCIe Adapter can send PCIe Tunneled Packets across the USB4 Fabric. It also controls the in-band presence indication to the internal PCIe Switch:</p> <p>0b: PCIe Packets shall not be sent. In-band presence is set to 0b 1b: PCIe Tunneled Packets may be sent. In-band presence is set to 1b</p> | RW | 0b |

8.2.3 Path Configuration Space

The Path Configuration Space in an Adapter contains a 2-Doubleword entry for each Path supported by the Adapter. Section 8.2.3.1 defines entry for Path 0. Section 8.2.3.2 defines the Path entries for a Lane Adapter. Section 8.2.3.3 defines the Path entries for a Protocol Adapter.

8.2.3.1 Path 0 Entry

A Lane Adapter and a Host Interface Adapter shall support a Path for HopID 0 (referred to as "Path 0"). Figure 8-15 shows the Path Configuration Space entry for Path 0. Table 8-18 defines the fields for Path 0.

Figure 8-15. Structure of Path 0 Entry Configuration Space




Table 8-18. Contents of Path 0 Entry

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|--|------|--------------------------------|
| 0 | PATH_CS_0 | 5:0 | Vendor Defined | VD | Vendor Defined |
| | | 16:6 | Reserved | Rsvd | 0h |
| | | 23:17 | Path Credits Allocated This field shall contain the initial value of the <i>Path Credits Allocated</i> state variable for the Ingress Adapter of the Path. | RO | Vendor Defined (Minimum 2h) |
| | | 30:24 | Reserved | Rsvd | 0 |
| | | 31 | Vendor Defined | VD | Vendor Defined |
| 1 | PATH_CS_1 | 11:0 | Vendor Defined | VD | Vendor Defined |
| | | 22:12 | Counter ID A Connection Manager uses this field to set the ID number of the counter set that is used to collect statistics for the Path. The Counter ID shall be less than the <i>Max Counter Sets</i> field for the Adapter. This field is valid only when the <i>Counter Enable (CE)</i> bit is set to 1b. | R/W | 0 |
| | | 23 | Counter Enable (CE) A Connection Manager uses this bit to enable the counters in Counter Configuration Space for the Path. 1b: Counter set for Path is enabled 0b: Counter set for Path is disabled When this bit is 1b, the Adapter shall increment the counter set for the Path as defined in Table 8-21. | R/W | 0 |
| | | 31:24 | Vendor Defined | VD | Vendor Defined |

8.2.3.2 Lane Adapters

A Lane Adapter shall support Paths from HopID 8 to Max Input HopID (inclusive). Each entry shall be formatted as shown in Figure 8-16. Table 8-19 defines the fields for a Path entry. The entries for the Paths with HopIDs 1 through 7 are undefined and the space is reserved.

When a Path is configured to route Tunneled Packets from a USB4 Port to a Protocol Adapter the following Path entry fields shall be ignored by the Router:

- *Weight.*
- *Egress Flow Control Flag.*
- *Egress Shared Buffering Enable Flag.*

Figure 8-16. Structure of Path Entry ‘n’ in Path Configuration Space at Lane Adapter




Table 8-19. Contents of Path Entry in Path Configuration Space at Lane Adapter

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|------|---------------|
| 0 | PATH_CS_0 | 10:0 | Output HopID A Connection Manager uses this field to set the Egress HopID for the Path. | R/W | 000h |
| | | 16:11 | Output Adapter A Connection Manager uses this field to set the Adapter number of the Egress Adapter for the Path. | R/W | 00h |
| | | 23:17 | Path Credits Allocated A Connection Manager uses this field to set the initial value of the Path Credits Allocated state variable for the Ingress Adapter of the Path. | R/W | 0 |
| | | 30:24 | Reserved | Rsvd | 0 |
| | | 31 | Valid A Connection Manager uses this field to enable or disable the Path. 1b: Path is enabled and can send/receive packets 0b: Path is disabled and cannot send/receive packets | R/W | 0b |
| 1 | PATH_CS_1 | 7:0 | Weight A Connection Manager uses this field to set the WRR scheduler weight for the Path. | R/W | 00h |
| | | 10:8 | Priority A Connection Manager uses this field to set the Priority Group for the Path. | R/W | 0h |
| | | 11 | Reserved | Rsvd | 0b |

| | | | | | |
|---|-----------|-------|--|-------|-----|
| | | 22:12 | Counter ID A Connection Manager uses this field to set the ID number of the counter set that is used to collect statistics for the Path. The Counter ID shall be less than the <i>Max Counter Sets</i> field for the Adapter. This field is valid only when the <i>Counter Enable (CE)</i> bit is set to 1b. | R/W | 0 |
| | | 23 | Counter Enable (CE) A Connection Manager uses this bit to enable the counters in Counter Configuration Space for the Path. 1b: Counter set for Path is enabled 0b: Counter set for Path is disabled When this bit is 1b, the Adapter shall increment the counter set for the Path as defined in Table 8-21. | R/W | 0 |
| | | 24 | Ingress Flow Control (IFC) Flag A Connection Manager uses this bit in combination with the ISE Flag to configure the Ingress Flow Control scheme. See Table 5-10. | R/W | 0b |
| 1 | PATH_CS_1 | 25 | Egress Flow Control Flag (EFC) Flag A Connection Manager uses this bit in combination with the ESE Flag to configure the Egress Flow Control scheme. See Table 5-10. | R/W | 0b |
| | | 26 | Ingress Shared Buffering Enable (ISE) Flag A Connection Manager uses this bit in combination with the IFC Flag to configure the Ingress Flow Control scheme. See Table 5-10. | R/W | 0b |
| | | 27 | Egress Shared Buffering Enable (ESE) Flag A Connection Manager uses this bit in combination with the EFC Flag to configure the Egress Flow Control scheme. See Table 5-10. | R/W | 0b |
| | | 28 | Pending Packets (PP) An Adapter shall set this field to 1b when one or more packets that belong to the Path are waiting to be dequeued. Otherwise it shall be set to zero. | RO | 0b |
| | | 29 | Reserved | RsvdZ | 0b |
| | | 31:30 | Reserved | Rsvd | 00b |

8.2.3.3 Protocol Adapters

Host Interface Adapter shall support Paths from HopID 1 to Max Input HopID (Inclusive). USB3/PCIe/DP Adapter shall support Paths from HopID 8 to Max Input HopID (Inclusive). Each entry shall be formatted as shown in Figure 8-17. Table 8-20 defines the fields for a Path entry.

Figure 8-17. Structure of Path Entry ‘n’ in Path Configuration Space of a Protocol Adapter




Table 8-20. Contents of Path Entry in Path Configuration Space of a Protocol Adapter

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|----------------------------|------|---------------|
|----|---------------|--------|----------------------------|------|---------------|

| | | | | | |
|---|-----------|-------|--|------|----------------|
| 0 | PATH_CS_0 | 10:0 | Output HopID A Connection Manager uses this field to set the Egress HopID for the Path. | R/W | 000h |
| | | 16:11 | Output Adapter A Connection Manager uses this field to set the Adapter number of the Egress Adapter for the Path. | R/W | 00h |
| | | 23:17 | Path Credits Allocated It is recommended that a Router set this field to 0b | VD | Vendor Defined |
| | | 30:24 | Reserved | Rsvd | 0 |
| 0 | PATH_CS_0 | 31 | Valid A Connection Manager uses this field to enable or disable the Path. 1b: Path is enabled and can send/receive packets 0b: Path is disabled and cannot send/receive packets | R/W | 0b |
| 1 | PATH_CS_1 | 7:0 | Weight A Connection Manager uses this field to set the WRR scheduler weight for the Path. | R/W | 00h |
| | | 10:8 | Priority A Connection Manager uses this field to set the Priority Group for the Path. | R/W | 0h |
| | | 11 | Reserved | Rsvd | 0b |
| | | 22:12 | Counter ID A Connection Manager uses this field to set the ID number of the counter set that is used to collect statistics for the Path. The Counter ID shall be less than the <i>Max Counter Sets</i> field for the Adapter. This field is valid only when the <i>Counter Enable (CE)</i> bit is set to 1b. | R/W | 0 |
| | | 23 | Counter Enable (CE) A Connection Manager uses this bit to enable the counters in Counter Configuration Space for the Path. 1b: Counter set for Path is enabled 0b: Counter set for Path is disabled When this bit is 1b, the Adapter shall increment the counter set for the Path as defined in Table 8-21. | R/W | 0 |
| | | 24 | Ingress Flow Control (IFC) Flag It is recommended that an Adapter set this field to 0b. | VD | Vendor Defined |
| | | 25 | Egress Flow Control Flag (EFC) Flag A Connection Manager uses this bit in combination with the ESE Flag to configure the Egress Flow Control scheme. See Table 5-10. | R/W | 0b |
| | | 26 | Ingress Shared Buffering Enable (ISE) Flag It is recommended that an Adapter set this field to 0b. | VD | Vendor Defined |
| | | 27 | Egress Shared Buffering Enable (ESE) Flag A Connection Manager uses this bit in combination with the EFC Flag to configure the Egress Flow Control scheme. See Table 5-10. | R/W | 0b |

| | | | | | |
|--|--|-------|---|-------|-----|
| | | 28 | Pending Packets (PP) An Adapter shall set this field to 1b when one or more packets that belong to the Path are waiting to be dequeued. Otherwise it shall be set to zero. | RO | 0b |
| | | 29 | Reserved | RsvdZ | 0b |
| | | 31:30 | Reserved | Rsvd | 00b |



CONNECTION MANAGER NOTE

When writing to the Path Configuration Space of a Lane Adapter or Protocol Adapter, a Connection Manager shall abide by the following rules:

- The Connection Manager needs to read a Path Configuration Space register before writing to it. The Connection Manager shall not change the value of any fields that are defined as RsvdZ or VD. See Section 8-1 for more information.
- The Weight field for a Path cannot be 0 if the Valid bit is 1b.
- The Priority field for a Path needs to be set to a value between 1 and 7 (inclusive).
- The Output Adapter field for a Path needs to be set to an Adapter that has a Path Configuration Space.

8.2.3.4 Path Configuration Space Access

When a Connection Manager configures a Path, within a Router, from Ingress Adapter N, HopID 'n' to Egress Adapter M, it issues the following single write access:

- The *Adapter Num* field shall be equal to N.
- The *Address* field shall be equal to 2 x 'n'.
- The *Length* field shall be equal to 2.

Note: To configure the above Path, no access is issued to Adapter M.



CONNECTION MANAGER NOTE

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

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



CONNECTION MANAGER NOTE

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



CONNECTION MANAGER NOTE

After setting the Valid bit to 1b in a path Configuration Space, a Connection Manager shall not modify any fields except the Weight field and the Valid bit.


8.2.3.4.1 Path Configuration Example

The entry for a Path contains fields that apply to the Ingress Adapter of the Path and fields that apply to the Egress Adapter of the Path, however, all parameters are written to the Ingress Adapter. In the example shown in Figure 8-18, configuring the Routers along Path X requires one configuration per Router:

- Router A: Single Path access to Path 10 at Adapter Number 1.
- Router B: Single Path access to Path 11 at Adapter Number 3.

Note: The Path Configuration Spaces of Adapter 2 in Router A and Adapter 1 in Router B are independent of Path X and therefore do not reflect any of the above configuration.

Figure 8-18. Configuration of a Path



8.2.4 Counters Configuration Space

An Adapter with the CCS Flag in the Adapter Configuration Space set to 1b shall implement the Counters Configuration Space depicted in Figure 8-19.

A Counter Configuration Space shall contain the number of counter sets specified in the *Max Counter Sets* field from the Adapter Configuration Space. Each counter set shall consist of the three counters described in Table 8-21.

A Path is associated with a counter set through the *Counter ID* field in the Path Configuration Space. A Path can only be associated with one counter set. However a single counter set may count packets for more than one Path.

Figure 8-19. Structure of the Counters Configuration Space




Table 8-21. Counter Set Fields

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|-------|---------------|
| 0 | CNT_CS_0 | 31:0 | Received Packets Low This field shall contain the lower 32 bits of a 64-bit received packets counter. An Ingress Adapter shall increment the received packets counter by 1 each time it receives a packet on a Path that uses this counter set. The counter shall increment from 0 and shall stop counting at FFFF FFFF FFFF FFFFh. | W/Clr | 0 |
| 1 | CNT_CS_1 | 31:0 | Received Packets High This field contains the upper 32 bits of a 64-bit received packets counter. An Ingress Adapter shall increment the received packets counter by 1 each time it receives a packet on a Path that uses this counter set. The counter shall increment from 0 and shall stop counting at FFFF FFFF FFFF FFFFh. | W/Clr | 0 |
| 2 | CNT_CS_2 | 31:0 | Dropped Packets This field contains a 32-bit dropped packets counter. An Ingress Adapter shall increment the dropped packets counter by 1 for every packet that is dropped due to insufficient buffer space for a Path that uses this counter set. The counter shall increment from 0 and shall stop counting at FFFF FFFFh. | W/Clr | 0 |

**CONNECTION MANAGER NOTE**

It is recommended that a Connection Manager read the Received Packets Low field and the Received Packets High field in a single Read Request in order to guarantee that the received packets counter values are from the same counter snapshot.

8.3 Operations

A Router supports a Router-level interface and a Port-level interface that allow a Connection Manager to initiate various Operations.

The Router-level interface initiates Router Operations, which perform Router-wide tasks such as NVM read/write and DisplayPort™ resource management. Section 8.3.1 defines Router Operations.

The Port-level interface initiates Port Operations, which can be used to initiate Port-level tasks such as compliance tests, NVM access, and receiver Lane margining tests. Section 8.3.2 defines Port Operations.

A Router shall handle Router Operations and Port Operations concurrently.

8.3.1 Router Operations

A Connection Manager writes to the ROUTER_CS_9 through ROUTER_CS_26 registers in Router Configuration Space to initiate a Router Operation.

A Router shall process a Router Operation when the *Operation Valid* bit changes from 0b to 1b. The Router shall execute the specific Router Operation indicated by the *Opcode* field as defined in the sections below.

When a Router Operation is defined to include metadata information, the Router shall fetch the information from the *Metadata* field in Router Configuration Space.

When a Router Operation is defined to include additional information, the Router shall fetch the information from the *Data DWs* in Router Configuration Space.

Once the Router Operation is complete, the Router shall:

1. If the Router Operation returns completion metadata information, write the metadata information to the *Metadata* field in Router Configuration Space.
2. If the Router Operation returns additional completion information, write the additional information to the *Data DWs* in Router Configuration Space.
3. Set the *Operation Not Supported* bit to 0b if the Router supports the Operation. Set the *Operation Not Supported* bit to 1b if the Router does not support the Operation.
4. If the *Operation Not Supported* bit is 0b, update the *Status* field with the results of the Router Operation.

Note: The Status field is only applicable when the Operation Not Supported bit is set to 0b.

5. Set the *Operation Valid* bit to 0b.



CONNECTION MANAGER NOTE

A Connection Manager shall not issue a new Router Operation until the previous Router Operation has completed.

Table 8-22 lists the Router Operations defined for a Router.

Table 8-22. List of Router Operations

| Router Operation | Opcode | Operation | | Completion | | Reference |
|--------------------------------|-----------|-------------|----------|-------------|----------|-------------------|
| | | Metadata DW | Data DWs | Metadata DW | Data DWs | |
| Reserved | 00h – 0Fh | -- | -- | -- | -- | none |
| Query DP Resource Availability | 10h | 1 | 0 | 1 | 0 | Section 8.3.1.1.1 |
| Allocate DP Resource | 11h | 1 | 0 | 1 | 0 | Section 8.3.1.1.2 |
| De-allocate DP Resource | 12h | 1 | 0 | 1 | 0 | Section 8.3.1.1.3 |
| Reserved | 13h – 1Fh | -- | -- | -- | -- | none |
| NVM Write | 20h | 0 | 1 to 16 | 0 | 0 | Section 8.3.1.2.2 |
| NVM Authenticate Write | 21h | 0 | 0 | 0 | 0 | Section 8.3.1.2.3 |
| NVM Read | 22h | 1 | 0 | 1 | 1 to 16 | Section 8.3.1.2.4 |
| NVM Set Offset | 23h | 1 | 0 | 1 | 0 | Section 8.3.1.2.1 |
| DROM Read | 24h | 1 | 0 | 1 | 1 to 16 | Section 8.3.1.2.5 |
| Get NVM Sector Size | 25h | 0 | 0 | 1 | 0 | Section 8.3.1.2.6 |
| Reserved | 26h – 2Fh | -- | -- | -- | -- | none |

| Router Operation | Opcode | Operation | | Completion | | Reference |
|-----------------------------------|---------------|-------------|----------|-------------|----------|-------------------|
| | | Metadata DW | Data DWs | Metadata DW | Data DWs | |
| Get PCIe Downstream Entry Mapping | 30h | 0 | 0 | 1 | 2 | Section 8.3.1.3.1 |
| Get Capabilities | 31h | 1 | 0 | 1 | 1 to 16 | Section 8.3.1.3.2 |
| Set Capabilities | 32h | 1 | 0 to 16 | 0 | 0 | Section 8.3.1.3.3 |
| Buffer Allocation Request | 33h | 0 | 0 | 0 | 1 to 5 | Section 8.3.1.3.4 |
| Block Sideband Port Operations | 34h | 0 | 0 | 0 | 0 | Section 8.3.1.4.1 |
| Unblock Sideband Port Operations | 35h | 0 | 0 | 0 | 0 | Section 8.3.1.4.2 |
| Get Container-ID | 36h | 0 | 0 | 0 | 4 | Section 8.3.1.3.5 |
| Reserved | 37h - 7FFFh | -- | -- | -- | -- | none |
| Vendor Specific Router Operations | 8000h - FFFFh | -- | -- | -- | -- | none |

8.3.1.1 DP Tunneling Operations

8.3.1.1.1 Query DP Resource Availability (Conditional)

A Router shall support the Query DP Resource Availability Router Operation if it has one or more DP IN Adapters. This Router Operation is not applicable for Routers that do not have a DP IN Adapter.

Operation Initiation

Table 8-23 defines the contents of the *Metadata* field for a Query DP Resource Availability Router Operation.

Table 8-23. Query DP Resource Availability Operation Metadata

| Field | Bit(s) | Description |
|-----------------|--------|--|
| <i>Metadata</i> | 31:0 | DisplayPort Number This field contains the DP IN Adapter number queried by the Router Operation. |

This Router Operation does not write to the *Data DWs*.

Operation Completion

A Router shall return the *Metadata* and *Status* fields defined in Table 8-24.

Table 8-24. Query DP Resource Availability Completion Metadata and Status

| Field | Bit(s) | Description |
|-----------------|--------|--|
| <i>Metadata</i> | 31:0 | DisplayPort Number This field contains the <i>DisplayPort Number</i> value from the Operation Metadata. |
| <i>Status</i> | 29:24 | Status 0h: A DP Source is detected and a resource is available to be allocated to the DP IN Adapter 1h: No DP Source is detected and/or a resource is not available to be allocated to the DP IN Adapter 2h – Fh: Reserved |

This Operation does not return any *Data DWs*.

8.3.1.1.2 Allocate DP Resource (Conditional)

A Router shall support the Allocate DP Resource Availability Router Operation if it has one or more DP IN Adapters. This Router Operation is not applicable for Routers that do not have a DP IN Adapter.

Operation Initiation

Table 8-25 defines the contents of the *Metadata* field for an Allocate DP Resource Router Operation.

Table 8-25. Allocate DP Resource Operation Metadata

| Field | Bit(s) | Description |
|-----------------|--------|---|
| <i>Metadata</i> | 31:0 | DisplayPort Number This field contains the Adapter number of the DP IN Adapter that is the recipient of the allocated DP stream resource. |

This Router Operation does not write to the *Data DWs*.

Operation Completion

A Router shall return the *Metadata* and *Status* fields defined in Table 8-26.

Table 8-26. Allocate DP Resource Completion Metadata and Status

| Field | Bit(s) | Description |
|-----------------|--------|--|
| <i>Metadata</i> | 31:0 | DisplayPort Number This field contains the <i>DisplayPort Number</i> value from the Operation Metadata. |
| <i>Status</i> | 29:24 | Status 0h: Resource was allocated 1h: Resource was not allocated 2h – Fh: Reserved If a resource was already allocated to this DP IN Adapter by a previous Allocate DP Resource Operation, then a Router shall respond with <i>Status</i> = 0h. |

This Operation does not return any *Data DWs*.

8.3.1.1.3 De-allocate DP Resource (Conditional)

A Router shall support the De-Allocate DP Resource Availability Router Operation if it has one or more DP IN Adapters. This Router Operation is not applicable for Routers that do not have a DP IN Adapter.

Operation Initiation

Table 8-27 defines the contents of the *Metadata* field for a De-allocate DP Resource Router Operation.

Table 8-27. De-Allocate DP Resource Operation Metadata

| Field | Bit(s) | Description |
|-----------------|--------|--|
| <i>Metadata</i> | 31:0 | DisplayPort Number This field contains the Adapter number of the DP IN Adapter that is the recipient of the DP stream resource that is being de-allocated. |

This Router Operation does not write to the *Data* DWs.

Operation Completion

A Router shall return the *Metadata* and *Status* fields defined in Table 8-28.

Table 8-28. De-Allocate DP Resource Completion Metadata and Status

| Field | Bit(s) | Description |
|-----------------|--------|--|
| <i>Metadata</i> | 31:0 | DisplayPort Number This field contains the <i>DisplayPort Number</i> value from the Operation Metadata. |
| <i>Status</i> | 29:24 | Status 0h: Resource was de-allocated 1h: Resource was not de-allocated 2h – Fh: Reserved If no resource is allocated to this DP IN Adapter, then a Router shall respond with Status = 0h. |

This Router Operation does not return any *Data* DWs.

8.3.1.2 NVM Operations

NVM Operations are used to read from and write to the Non Volatile Memory (NVM) in a USB4 product.

The following is an example of a Router Operation sequence used by software to update the NVM of a Router:

1. Software issues a Get NVM Sector Size Operation to read the sector size of the NVM.
2. Software issues an NVM Set Offset Operation, which sets the first location in NVM to be written by the following NVM Write Operation.
3. Software issues a sequence of NVM Write Operations, each writing a 64B block of data to NVM.
 - Following an NVM Write Operation, if the value of the *Status* field in the completion status is 1h, software repeats all previous NVM Write Operations. The first Router Operation to be issued is an NVM Set Offset Operation that sets the location in NVM of the first block to be rewritten.
4. Software issues an NVM Authenticate Write Operation to indicate to the target that all data was sent to the target. After receiving the NVM Authenticate Write Operation, the target performs an authentication check over the data written.

The following is an example of a Router Operation sequence used by software to validate the results of the NVM update:

1. Software then waits for at least 5 seconds from the completion of the NVM Authenticate Write Operation.
2. Software reads the results of the NVM Authenticate Write Operation. If the *Status* field in the Completion Metadata is 0h, the update ended successfully. If the *Status* field is not 0h, then software repeats all previous NVM Write Operations. The first Router Operation issued is an NVM Set Offset Operation that sets the location in NVM of the first block to be rewritten, followed by re-issuing one or more NVM Write Operations.

8.3.1.2.1 NVM Set Offset (Conditional)

A Device Router shall support the NVM Set Offset Router Operation. A Host Router may optionally support this Router Operation.

The NVM Set Offset Router Operation sets the first location in NVM to be written by the following NVM Write Router Operation.

Operation Initiation

Table 8-29 defines the contents of the *Metadata* field for an NVM Set Offset Router Operation.

Table 8-29. NVM Set Offset Operation Metadata

| Field | Bit(s) | Description |
|-----------------|--------|--|
| <i>Metadata</i> | 1:0 | Reserved |
| | 23:2 | NVM Offset This field contains the first address to be written relative to the base address of the region being written to. NVM Offset is specified in DWs. |
| | 31:24 | Reserved |

This Router Operation does not write to the *Data* DWs.

Operation Completion

A Router shall return the *Metadata* and *Status* fields defined in Table 8-30.

Table 8-30. NVM Set Offset Completion Metadata and Status

| Field | Bit(s) | Description |
|-----------------|--------|--|
| <i>Metadata</i> | 1:0 | Reserved |
| | 23:2 | NVM Offset This field contains the NVM Offset value from the Operation. |
| | 31:24 | Reserved |
| <i>Status</i> | 29:24 | Status 0h: Router Operation completed successfully 1h: Router Operation failed to execute 2h: NVM is not implemented 3h – 3Fh: Reserved |

This Router Operation does not return any *Data* DWs.

8.3.1.2.2 NVM Write (Conditional)

A Device Router shall support the NVM Write Router Operation. A Host Router may optionally support this Router Operation.

The NVM Write Router Operation writes 64 bytes to NVM, starting at the address equal to the NVM Offset. A Router shall increment its NVM Offset value by 16 after executing an NVM Write Router Operation.

Operation Initiation

This Operation does not write to the *Metadata* field.

Table 8-31 defines the contents of the *Data* DWs.

Table 8-31. NVM Write Operation Data

| DW | Bit(s) | Description |
|-----|--------|---|
| 0 | 31:0 | DW0 The first Doubleword, to be written into NVM at <i>NVM Offset</i> address. |
| 1 | 31:0 | DW1 The second Doubleword, to be written into NVM at the next NVM Offset. |
| ... | 31:0 | ... |
| 15 | 31:0 | DW15 The last Doubleword of this Router Operation, to be written into NVM at the next NVM Offset. |

Operation Completion

A Router shall return the *Status* field defined in Table 8-32.

Table 8-32. NVM Write Completion Status

| Field | Bit(s) | Description |
|---------------|--------|--|
| <i>Status</i> | 29:24 | Status 0h: Router Operation completed successfully 1h: Router Operation failed to execute 2h – 3Fh: Reserved |

This Router Operation does not return any *Metadata* or *Data* DWs.

8.3.1.2.3 NVM Authenticate Write (Conditional)

A Device Router shall support the NVM Authenticate Write Router Operation. A Host Router may optionally support this Router Operation.

The NVM Authenticate Write Router Operation indicates to the target that all data was sent to the target. After receiving an NVM Authenticate Write Router Operation, the target performs an authentication check over the data written. The authentication check is implementation specific and outside the scope of this specification.

Operation Initiation

This Router Operation does not write to the *Metadata* field. It does not write to the *Data* DWs.

Operation Completion

A Router shall return the *Status* field defined in Table 8-33.

Table 8-33. NVM Authenticate Write Completion Status

| Field | Bit(s) | Description |
|---------------|--------|--|
| <i>Status</i> | 29:24 | Status 0h: Authentication completed successfully 1h: Authentication Failed 2h: Retry NVM write 3h: NVM is not implemented 4h – 3Fh: Reserved |

If the *Status* field is 2h, then software writes the last blocks of data again, starting with the block whose offset is the beginning of a sector-aligned page that contains the last successful offset written minus the sector size. The first Router Operation to be issued is an NVM Set Offset Router Operation that sets the location in NVM of the first block to be rewritten, followed by issuing an NVM Authenticate Write Router Operation again.

This Router Operation does not return any *Metadata* or *Data* DWs.

8.3.1.2.4 NVM Read (Conditional)

A Device Router shall support the NVM Read Router Operation. A Host Router may optionally support this Router Operation.

The NVM Read Router Operation reads up to 64 bytes from NVM.

Operation Initiation

Table 8-34 defines the contents of the *Metadata* field for an NVM Read Router Operation.

Table 8-34. NVM Read Operation Metadata

| Field | Bit(s) | Description |
|-----------------|--------|---|
| <i>Metadata</i> | 1:0 | Reserved |
| | 23:2 | NVM Offset This field contains the first address to be read relative to the base address of the region being read. NVM Offset is specified in DWs. |
| | 27:24 | Length Number of Doublewords that shall be read starting from the <i>NVM Offset</i> field value. If this field is zero, then 16 DWs are read. |
| | 31:28 | Reserved |

This Router Operation does not write to the *Data* DWs.

Operation Completion

A Router shall return the *Metadata* and *Status* fields defined in Table 8-35.

Table 8-35. NVM Read Router Completion Metadata

| Field | Bit(s) | Description |
|-----------------|--------|--|
| <i>Metadata</i> | 1:0 | Reserved |
| | 23:2 | NVM Offset This field contains the <i>NVM Offset</i> value from the Operation Metadata. |
| | 27:24 | Length This field contains the <i>Length</i> value from the Operation Metadata. |
| | 31:28 | Reserved |
| <i>Status</i> | 29:24 | Status 0h: Router Operation completed successfully 1h: Router Operation failed to execute 2h – 3Fh: Reserved |

If this Operation is completed successfully, the Router returns Length number of Doublewords of Data as shown in Table 8-36.

Table 8-36. NVM Read Router Completion Data

| DW | Bit(s) | Description |
|-----|--------|--|
| 0 | 31:0 | DW0 The first Doubleword read from NVM at <i>NVM Offset</i> address. |
| 1 | 31:0 | DW1 The second Doubleword read from NVM at the next NVM Offset. |
| ... | 31:0 | ... |
| 15 | 31:0 | DW15 The last Doubleword read from NVM at the next NVM Offset. |

8.3.1.2.5 DROM Read (Conditional)

A Device Router shall support the DROM Read Router Operation. A Standalone AIC Host Router shall support this Router Operation. A Platform Integrated Host Router may optionally support this Router Operation.

The DROM Read Router Operation reads up to 64 bytes from DROM.

Operation Initiation

Table 8-37 defines the contents of the *Metadata* field for a DROM Read Router Operation.

Table 8-37. DROM Read Router Operation Metadata

| Field | Bit(s) | Description |
|-----------------|--------|---|
| <i>Metadata</i> | 1:0 | Reserved |
| | 14:2 | Address DW address in DROM relative to DROM first address. Address = 0 targets the first DW in DROM. |
| | 19:15 | Read Size Number of Doublewords that shall be read starting from the <i>Address</i> field value. The <i>Read Size</i> field shall be greater than 0 and less than or equal to 16. |
| | 31:20 | Reserved |

This Router Operation does not write to the *Data* DWs.

Operation Completion

A Router shall return the *Metadata* and *Status* fields defined in Table 8-38.

Table 8-38. DROM Read Router Completion Metadata and Status

| Field | Bit(s) | Description |
|-----------------|--------|--|
| <i>Metadata</i> | 1:0 | Reserved |
| | 14:2 | Address This field contains the <i>Address</i> value from the Operation Metadata. |
| | 19:15 | Read Size This field contains the <i>Read Size</i> value from the Operation. |
| | 31:20 | Reserved |
| <i>Status</i> | 29:24 | Status 0h: Router Operation completed successfully 1h: Router Operation failed to execute 2h – 3Fh: Reserved |

If this Operation is completed successfully, the Router returns *Read Size* number of Doublewords of Data as shown in Table 8-39.

Table 8-39. DROM Read Router Completion Data

| DW | Bit(s) | Description |
|-----|--------|---|
| 0 | 31:0 | DW0 The first Doubleword read from DROM at <i>Address</i> . |
| 1 | 31:0 | DW1 The second Doubleword read from DROM at the next address. |
| ... | 31:0 | ... |
| 15 | 31:0 | DW15 The last Doubleword read from DROM at the next address. |

8.3.1.2.6 Get NVM Sector Size (Conditional)

Software uses the Get NVM Sector Size Operation to determine the sector size of the Router NVM.

A Device Router shall support the NVM Sector Size Router Operation. A Host Router may optionally support this Router Operation.

Operation Initiation

This Router Operation does not write to the *Metadata* field or *Data* DWs.

Operation Completion

A Router shall return the *Metadata* and *Status* fields defined in Table 8-40.

Table 8-40. Get NVM Sector Size Completion Metadata and Status

| Field | Bit(s) | Description |
|-----------------|--------|--|
| <i>Metadata</i> | 23:0 | Sector Size Equals the sector size of the NVM in bytes. For example, a value of 1000h indicates a sector size of 4KB. |
| | 31:24 | Reserved |
| <i>Status</i> | 29:24 | Status 0h: Router Operation completed successfully 1h: Router Operation failed to execute 2h: NVM is not implemented 3h – 3Fh: Reserved |

This Router Operation does not return any *Data* DWs.

8.3.1.3 Router Discovery Operations

8.3.1.3.1 Get PCIe Downstream Entry Mapping (Conditional)

A Connection Manager uses this Operation to retrieve information about the mapping of a PCIe downstream facing port. The mapping information is returned in an entry, where there is one entry per PCIe downstream facing port. Each execution of this Operation retrieves the entry for one PCIe downstream facing port. Entries are returned in increasing order of their Entry Index.

A Router shall support the Get PCIe Downstream Entry Mapping Router Operation if it supports PCIe Tunneling. This Router Operation is not applicable for a Router that does not support PCIe tunneling.

Operation Initiation

This Router Operation does not write to the *Metadata* field or *Data* DWs.

Operation Completion

If a Router supports the Get PCIe Downstream Entry Mapping Router Operation, it shall return the *Metadata* and *Status* fields defined in Table 8-41.

Table 8-41. Get PCIe Downstream Entry Mapping Completion Metadata and Status

| Field | Bit(s) | Field Name and Description |
|-----------------|--------|--|
| <i>Metadata</i> | 7:0 | Total Number of Entries This field indicates the total number of entries. A Host Router that supports PCIe tunneling shall have one entry per Downstream PCIe Adapter. A Device Router that supports PCIe tunneling shall have one entry per PCIe Downstream Bridge. |
| | 15:8 | Entry Index This field indicates the current entry index. The values of this field shall be zero to <i>Total Number of Entries</i> - 1. |
| | 31:16 | Reserved |
| <i>Status</i> | 29:24 | Status 0h: Operation completed successfully 1h: Operation failed to execute 2h – Fh: Reserved |

The first time this Operation is executed, a Router shall respond with the entry for Entry Index = 0h. On each subsequent execution of the Operation, the Router shall respond with the next entry (*Entry Index = 01h, Entry Index = 02h, etc.*). After the last entry is retrieved, the Router shall restart at the first entry (*Entry Index = 0h*) the next time the Operation is executed.

A Router shall return the entry for a PCIe Downstream mapping in *Data DW0* and *DW1* as defined in Table 8-42.

Table 8-42. Get PCIe Downstream Entry Mapping Completion Data

| DW | Bit(s) | Field Name and Description |
|------|--------|---|
| 0 | 0 | Native PCIe Link 0 - Entry is for a Downstream Bridge connected to a Downstream PCIe Adapter 1 - Entry is for a Downstream Bridge connected through a native PCIe link |
| | 6:1 | PCIe Adapter Number If <i>Native PCIe Link</i> is set to 0, this field shall indicate the Adapter Number of the Downstream PCIe Adapter. Otherwise it shall be set to 0. |
| | 31:7 | Reserved |
| 1 | 15:0 | Non-FPB Routing ID The Routing ID value assigned to this Downstream Bridge according to the non-FPB addressing scheme. A value of 0 indicates that the non-FPB addressing scheme is not used through this Downstream Bridge. <i>Note: When a Host Router sets both the Non-FPB Routing ID and FPB Routing ID fields to zero, it indicates that the non-FPB addressing scheme is used.</i> |
| | 31:16 | FPB Routing ID The Routing ID value assigned to this Downstream Bridge according to the FPB addressing scheme. A value of 0 indicates that the FPB addressing scheme is not used through this Downstream Bridge. |
| 2-15 | 31:0 | Reserved |

8.3.1.3.2 Get Capabilities (Conditional)

A Connection Manager uses this Operation to retrieve information about certain Router capabilities. Every execution of this Operation retrieves information about a single capability.

A Router shall support the Get Capabilities Router Operation if it supports the Set Capabilities Router Operation. Otherwise, a Router shall not support this Router Operation.

Operation Initiation

Table 8-43 defines the contents of the *Metadata* field for a Get Capabilities Router Operation.

Table 8-43. Get Capabilities Operation Metadata

| Field | Bit(s) | Description |
|-----------------|--------|--|
| <i>Metadata</i> | 7:0 | Capability Index This field indicates the capability to be read. The value in this field shall not exceed the <i>Max Capability Index</i> . |
| | 31:8 | Reserved |

This Router Operation does not write to the *Data DWs*.

Operation Completion

A Router that supports the Get Capabilities Operation shall return the *Metadata* field and the *Status* field defined in Table 8-44.


Table 8-44. Get Capabilities Operation Completion Metadata and Status

| Field | Bit(s) | Description |
|-----------------|--------|---|
| <i>Metadata</i> | 7:0 | Max Capability Index This field indicates the highest Capability Index supported by the Router. |
| | 15:8 | Capability Index This field contains the <i>Capability Index</i> value from the Get Capabilities Operation. |
| | 29:16 | Reserved |
| | 30 | Capability Supported 0b - Capability is not supported 1b - Capability is supported This bit shall be set to 0b for <i>Capability Index</i> = 0. |
| | 31 | Capability Enabled 0b - Capability is disabled 1b - Capability is enabled This bit shall be set to 0b for <i>Capability Index</i> = 0. |
| <i>Status</i> | 29:24 | Status 0h: Router Operation completed successfully 1h: Router Operation failed to execute 2h – 3Fh: Reserved |

When a Router receives a Get Capabilities Operation with *Capability Index* = 0, it shall return a list of the capabilities that the Router supports and indicate which capabilities are enabled. Capability information is returned in a series of 2-bit entries, where each entry represents a different capability. Each 2-bit entry consists of the following:

- *Capability Supported (SUP)* bit – The entry for capability n ($n = 0$ to *Max Capability Index*) to is located in DW $\lfloor n / 16 \rfloor$, bit $2*[n - 16 * \lfloor n / 16 \rfloor]$. The entry for *Capability Index* = 0 is reserved.
 - The *Capability Supported* bit shall be set to 0b when the capability is not supported.
 - The *Capability Supported* bit shall be set to 1b when the capability is supported.
- *Capability Enabled (ENA)* bit – The entry for capability n ($n = 0$ to *Max Capability Index*) to is located in DW $\lfloor n / 16 \rfloor$, bit $\{2*[n - 16 * \lfloor n / 16 \rfloor] + 1\}$. The entry for *Capability Index* = 0 is reserved.
 - The *Capability Enabled* bit shall be set to 0b when the capability is disabled.
 - The *Capability Enabled* bit shall be set to 1b when the capability is enabled.

The 2-bit entries are organized as a packed array of Doublewords. The last Doubleword may not necessarily be fully populated. The list of capabilities is returned in the *Data* field and shall be formatted as shown in Figure 8-20.

Figure 8-20. Get Capabilities Operation Data Response for Capability Index 0

The contents of the *Data DWs* returned by the Get Capability Operation for *Capability Index* > 0 are specific to each capability. Table 8-45 lists the capabilities and the number of *Data DWs* returned by each capability.

Table 8-45. List of Capabilities

| Capability Index | Capability Name | # of Data DWs returned | Reference |
|------------------|-----------------------------|------------------------|---------------------|
| 01h | Hot Plug Failure Indication | 0 | Section 8.3.1.3.2.1 |
| 02h - FFh | Reserved | -- | N/A |

8.3.1.3.2.1 Hot Plug Failure Indication

This capability allows a Router to inform the Connection Manager when a Hot Plug Event is detected on a Downstream Facing Port, but none of the Adapters in that USB4 Port succeeds to establish a Link (see Section 6.8.1.1).

This capability does not return any Data DWs.

8.3.1.3.3 Set Capabilities (Conditional)

A Connection Manager uses this Operation to configure a Router's capabilities. Every execution of this Operation configures a single capability.

A Router shall support the Set Capabilities Router Operation if it supports the Get Capabilities Router Operation. Otherwise, a Router shall not support this Router Operation.

Operation Initiation

Table 8-46 defines the contents of the *Metadata* field for a Set Capabilities Router Operation.

Table 8-46. Set Capabilities Operation Metadata

| Field | Bit(s) | Description |
|-----------------|--------|---|
| <i>Metadata</i> | 7:0 | Capability Index This field indicates the capability to be written. The values of this field shall not exceed the <i>Max Capability Index</i> . |
| | 8 | Enable Capability This bit is used to request that a Capability be enabled or disabled 0b – Router shall disable the Capability 1b – Router shall enable the Capability |
| | 31:9 | Reserved |

The *Data DWs* for the Set Capability Operation are specific to each capability. Table 8-47 lists the capabilities and the number of *Data DWs* written by each capability.

Table 8-47. List of Capabilities

| Capability Index | Capability Name | # of Data DWs written | Reference |
|------------------|-----------------------------|-----------------------|---------------------|
| 01h | Hot Plug Failure Indication | 0 | Section 8.3.1.3.3.1 |
| 02h – FFh | Reserved | -- | N/A |

Operation Completion

A Router that supports this Operation shall return the *Status* field defined in Table 8-48.

Table 8-48. Set Capabilities Operation Completion Status

| Field | Bit(s) | Field Name and Description |
|---------------|--------|--|
| <i>Status</i> | 29:24 | Status 0h: Router Operation completed successfully 1h: Router Operation failed to execute 2h – 3Fh: Reserved |

This Operation does not return any Metadata.

This Operation does not return any *Data DWs*.

8.3.1.3.3.1 Hot Plug Failure Indication

A Router shall enable this capability when all of the following conditions are true:

- The Router supports the Get Capabilities Operation and the Set Capabilities Operation.
- The Router supports the “Hot Plug Failure Indication” capability.
- The Router receives a Set Capabilities Operation with *Capability Index* = 1h and *Enable Capability* = 1b.

8.3.1.3.4 Buffer Allocation Request (Required)

A Router shall support this Router Operation.

The Buffer Allocation Request Router Operation reads the Router’s preferences for buffer allocation per Path type.

Operation Initiation

This Router Operation does not write to the *Metadata* field. It does not write to the *Data DWs*.

Operation Completion

A Router shall return *Metadata* and *Status* fields defined in Table 8-49.

Table 8-49. Buffer Allocation Request Router Completion Status and Metadata

| Field | Bit(s) | Description |
|-----------------|--------|--|
| <i>Metadata</i> | 7:0 | Length This field shall be equal to the number of buffer allocation parameters the Router reports. Buffer allocation parameters are defined in Table 8-50. |
| | 31:8 | Reserved |
| <i>Status</i> | 29:24 | Status 0h: Router Operation completed successfully 1h: Router Operation failed to execute 2h – 3Fh: Reserved |

If this Operation is completed successfully, the Router returns *Length* Doublewords of Data where each DW has the structure shown in Table 8-50.

Table 8-50. Buffer Allocation Request Router Completion Data DW Structure

| Bit(s) | Description |
|--------|--|
| 15:0 | Parameter Index 0h: Reserved 1h: baMaxUSB3 2h: baMinDPaux 3h: baMinDPmain 4h: baMaxPCIE 5h: baMaxHI 6h-FFFFh: Reserved |
| 31:16 | Requested Buffer Allocation This field contains the number of buffers requested for the corresponding buffer allocation parameter indicated by the <i>Parameter Index</i> field. |

8.3.1.3.5 Get Container-ID (Conditional)

A USB4 hub shall support this Router Operation. A USB4 peripheral device with an internal USB3 hub shall support this Router Operation. A USB4 peripheral device that does not contain an internal USB3 hub may optionally support this Router Operation. This Router Operation is not applicable for USB4 hosts. The Get Container-ID Router Operation reads the Container-ID value which allows software to associate a Router with its internal USB SuperSpeed Plus hub. The return value for the Container-ID shall be identical to the Container-ID read from the internal USB SuperSpeed Plus hub.

Operation Initiation

This Router Operation does not write to the *Metadata* field. It does not write to the *Data* DWs.

Operation Completion

This Router Operation does not return Metadata.

A Router that supports this operation shall return *Status* field defined in Table 8-51.

Table 8-51. Get Container-ID Router Completion Status

| Field | Bit(s) | Description |
|---------------|--------|--|
| <i>Status</i> | 29:24 | Status 0h: Router Operation completed successfully 1h: Router Operation failed to execute 2h – 3Fh: Reserved |

If this Operation is completed successfully, the Router returns four Doublewords of Data as shown in Table 8-52.

Table 8-52. Get Container-ID Router Completion Data DW Structure

| DW | Bit(s) | Description |
|------|--------|-----------------------------|
| 0 | 31:0 | Container-ID[31:0] |
| 1 | 31:0 | Container-ID[63:32] |
| 2 | 31:0 | Container-ID[95:64] |
| 3 | 31:0 | Container-ID[127:96] |
| 4-15 | 31:0 | Reserved |

8.3.1.4 Port Control Operations

8.3.1.4.1 Block Sideband Port Operations (Optional)

A Router may optionally support the Block Sideband Port Operations Router Operation.

The Block Sideband Port Operations Router Operation disables Port Operations when initiated by Transactions over the Sideband Channel. After receiving a Block Sideband Port Operations Router Operation, a Router shall change the access type for SB Registers 8, 9 and 18 in all its Ports from RW to RO when accessed by Sideband Transactions. Note that these registers are still RW for local accesses done by Connection Manager as described in Section 4.1.1.3.2.

Operation Initiation

This Router Operation does not write to the *Metadata* field. It does not write to the Data DWs.

Operation Completion

A Router that supports this Operation shall return the *Status* field defined in Table 8-53.

Table 8-53. Block Sideband Port Operation Completion Status

| Field | Bit(s) | Description |
|---------------|--------|---|
| <i>Status</i> | 29:24 | Status 0h: Router Operation completed successfully 1h: Router Operation failed to execute 2h – Fh: Reserved |

8.3.1.4.2 Unblock Sideband Port Operation (Conditional)

A Router may optionally support the Unblock Sideband Port Operation Router Operation. If a Router supports the Block Sideband Port Operations Router Operation, it shall support this Router Operation.

The Unblock Sideband Port Operation Router Operation enables Port Operations when initiated by Transactions over the Sideband Channel. After receiving an Unblock Sideband Port

Operations Router Operation, the Router shall change the access type for SB Registers 8, 9 and 18 from RO to RW when accessed by Sideband Transactions.

Operation Initiation

This Router Operation does not write to the *Metadata* field. It does not write to the Data DWs.

Operation Completion

A Router that supports this Operation shall return the *Status* field defined in Table 8-54.

Table 8-54. Unblock Sideband Port Operation Completion Status

| Field | Bit(s) | Description |
|---------------|--------|---|
| <i>Status</i> | 29:24 | Status 0h: Router Operation completed successfully 1h: Router Operation failed to execute 2h – Fh: Reserved |

8.3.2 Port Operations

A Connection Manager uses registers 8 (Opcode), 9 (Metadata), and 18 (Data) in the SB Register Space of a USB4 Port to initiate a Port Operation. Section 4.1.1.3.2 describes how a Connection Manager writes to SB Register Space. A Connection Manager uses the following flow to issue a Port Operation:

1. The Connection Manager optionally writes to the Metadata register in the SB Register Space of the target Port.
2. The Connection Manager optionally writes to the Data register in the SB Register Space of the target Port.
3. The Connection Manager writes to the Opcode register in the SB Register Space of the target Port.



CONNECTION MANAGER NOTE

A Connection Manager shall verify the successful completion of a write to the SB Register Space of the target Port before proceeding with the next write to SB Register Space. A Connection Manager verifies the results of a Port Operation by issuing a read to the Opcode register in the SB Register Space of the target Port. The Connection Manager also reads any Completion Metadata and Completion Data. If the Opcode Register = “!CMD” or “ERR” the contents of the Metadata and Data registers are invalid and shall be ignored.

When the Opcode register in SB Register Space is written, a USB4 Port shall execute the Port Operation associated with the Opcode register using the information in the Metadata and Data registers.

After executing the Port Operation, the USB4 Port updates the Opcode, Metadata, and Data register as follows:

- If the USB4 Port successfully completed the Port Operation, it shall set the Opcode register to 0. The USB4 Port shall update the Metadata register with completion metadata (if the Port Operation is defined to return metadata), and the Data register with completion data (if the Port Operation is defined to return data).
- Else, if the Port Operation is not supported, the USB4 Port shall set the Opcode register to a FourCC value of “!CMD” (444D4321h). The USB4 Port may update the Metadata and Data registers. However, any updates will be ignored by the Connection Manager.

- Else, the USB4 Port shall set the Opcode register to a FourCC value of “ERR ” (20525245h) to indicate that the Port Operation is supported, but could not be completed. The USB4 Port may update the Metadata and Data registers. However, any updates will be ignored by the Connection Manager



CONNECTION MANAGER NOTE

A Connection Manager verifies the results of a Port Operation by issuing a read to the Opcode register in the SB Register Space of the target Port. The Connection Manager also reads any Completion Metadata and Completion Data.

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

Table 8-55 lists the Port Operations defined for a Port. A Port may also support vendor specific Opcodes. The second byte (Opcode 1) of a vendor specific Opcode shall have a value between 61h and 7Ah (inclusive) to distinguish from Opcodes defined in this specification. All unused Opcodes (except for vendor specific Opcodes) are reserved and shall not be used.

Table 8-55. List of Port Operations

| Port Operation | Opcode ¹ | Operation | | Completion | | Reference |
|--|---------------------|-------------|----------|-------------|----------|-------------------|
| | | Metadata DW | Data DWs | Metadata DW | Data DWs | |
| SET_TX_COMPLIANCE | TXCM (4D435854h) | 1 | 0 | 0 | 0 | Section 8.3.2.1.1 |
| SET_RX_COMPLIANCE | RXCM (4D435852h) | 1 | 0 | 0 | 0 | Section 8.3.2.1.2 |
| START_BER_TEST | SBER (52454253h) | 1 | 0 | 0 | 0 | Section 8.3.2.1.3 |
| END_BER_TEST | EBER (52454245h) | 1 | 0 | 0 | 2 | Section 8.3.2.1.4 |
| END_BURST_TEST | BBER (52454242h) | 1 | 0 | 0 | 3 | Section 8.3.2.1.5 |
| READ_BURST_TEST | RBUR (52454252h) | 1 | 0 | 0 | 3 | Section 8.3.2.1.6 |
| ENTER_EI_TEST | EEIT (54494545h) | 1 | 0 | 0 | 0 | Section 8.3.2.1.7 |
| ROUTER_OFFLINE_MODE | LSEN (4E45534Ch) | 1 | 0 | 0 | 0 | Section 8.3.2.2.1 |
| ENUMERATE_RE-TIMERS | ENUM (4D554E45h) | 0 | 0 | 0 | 0 | Section 8.3.2.2.2 |
| READ_LANE_MARGIN_CAP | RDCP (50434452h) | 0 | 0 | 0 | 2 | Section 8.3.2.3.1 |
| RUN_HW_LANE_MARGINING | RHMG (474D4852h) | 1 | 0 | 0 | 2 | Section 8.3.2.3.2 |
| RUN_SW_LANE_MARGINING | RSMG (474D5352h) | 1 | 0 | 0 | 1 | Section 8.3.2.3.3 |
| READ_SW_MARGIN_ERR | RDSW (57534452h) | 0 | 0 | 1 | 0 | Section 8.3.2.3.4 |
| Notes: | | | | | | |
| 1. Byte 0 of the Opcode is the rightmost byte of the hexadecimal representation. | | | | | | |

8.3.2.1 Compliance Port Operations

The Port Operations defined in the subsections below are used to bring transmitters and receivers into compliance mode and to execute bit error tests, burst error rate tests, clock switch tests, TxFFE equalization tests, Electrical Idle tests, and Return Loss tests.

The following is an example Port Operation sequence used for compliance testing. It assumes that compliance software interfaces with test equipment that includes a PD controller (or its equivalent) and a Sideband Channel controller (or its equivalent).

For transmitter compliance testing:

1. Compliance software disconnects the Link to be tested (via the PD Controller embedded in the test equipment).
2. Compliance software sets a USB4 Link or TBT3-Compatible Link (to the PD controller).
3. Compliance software disables AT Transactions in the Port at the test equipment side in order to delay Lane Initialization (see step 8 below).
4. Compliance software connects the Link to be tested (via the PD Controller).
5. The PD Controller sets a USB4 Link or TBT3-Compatible Link.
6. Compliance software issues a SET_TX_COMPLIANCE Port Operation to the Router and to each On-Board Re-timer in the Router Assembly under test to set the transmit parameters.
7. Compliance software sets the Port at the test equipment side to either support Gen 3 or to not support Gen 3.
8. Compliance software enables AT Transactions in the Port at the test equipment side in order to proceed with Lane Initialization.
 - Compliance software disabled RS-FEC during compliance testing.
9. Compliance software performs TxFFE equalization tests with the component next to the USB Type-C connector.
 - Compliance software selects the preset values for the receiver at the test equipment side.
 - Compliance software determines when to complete the test by setting the Rx Lock bit to 1b at the test equipment side.
10. Compliance software performs clock switch test with the On-Board Re-timer next to the USB Type-C connector.
 - Compliance software sets the *Clock Switch Done* bit to 1b at the test equipment side to initiate a clock switch by the Re-timer.
11. Once the Lane under test is up, compliance software performs checks on the pattern received from the Router Assembly under test.
12. Once the transmitter compliance test is done, compliance software disconnects the Link under test.

For receiver compliance testing:

1. Compliance software disconnects the Link to be tested (via the PD Controller embedded in the test equipment).

2. Compliance software sets a USB4 Link or TBT3-Compatible Link (to the PD controller).
3. Compliance software disables AT Transactions in the Port at the test equipment side in order to delay Lane Initialization (see step 8 below).
4. Compliance software connects the Link to be tested (via the PD Controller).
5. The PD Controller sets a USB4 Link or TBT3-Compatible Link.
6. Compliance software issues a SET_RX_COMPLIANCE Port Operation to the Router under test.
7. Compliance software sets the Port at the test equipment side to either support Gen 3 or to not support Gen 3.
8. Compliance software enables AT Transactions in the Port at the test equipment side in order to proceed with Lane Initialization.
 - Compliance software disabled RS-FEC during compliance testing.
9. Compliance software performs TxFFE equalization test with the component next to the USB Type-C connector.
10. Compliance software waits for completion of Lane Initialization.
 - Compliance software performs a bit error rate test, a burst error rate test, or a clock switch test as described below.

For bit error rate test:

1. Compliance software issues a START_BER_TEST Port Operation to start a receiver compliance test using the chosen BER test pattern. The Operation may be issued to a Router, an On-Board Re-timer, or to multiple components.
2. Compliance software issues an END_BER_TEST Port Operation to stop a running receiver BER test pattern and retrieve the results of the test. The Operation may be issued to a Router, an On-Board Re-timer, or to multiple components.
3. Once the test is completed, compliance software disconnects the Link under test.

For burst error rate test:

1. Compliance software issues a START_BER_TEST Port Operation to start a receiver compliance test using the chosen BER test pattern. The Operation may be issued to a Router, an On-Board Re-timer, or to multiple components.
2. Compliance software may issue a READ_BURST_TEST Port Operation to read the counters involved in the test. The Operation may be issued multiple times. The Operation may be issued to a Router, an On-Board Re-timer, or to multiple components.
3. Compliance software issues an END_BURST_TEST Port Operation to stop a running receiver burst error rate test pattern and retrieve the results of the test. The Operation may be issued to a Router, an On-Board Re-timer, or to multiple components.
4. Once the test is completed, compliance software disconnects the Link under test.

For clock switch test:

1. Compliance software issues a START_BER_TEST Port Operation to the component next to the USB Type-C connector.
2. Compliance software sets the appropriate test conditions such as emulating a clock switch.
3. Compliance software issues an END_BER_TEST Port Operation to the component next to the USB Type-C connector.
4. Once the test is completed, compliance software disconnects the Link under test.
5. Software transitions the Router to the Uninitialized state.

For transmitter electrical idle test:

1. Compliance software disconnects the Link to be tested (via the PD Controller embedded in the test equipment).
2. Compliance software connects the Link to be tested (via the PD Controller).
3. Compliance software waits for completion of Lane Initialization.
4. Compliance software issues an ENTER_EI_TEST Port Operation to the component next to the USB Type-C connector.
5. Once the test is completed, compliance software disconnects the Link under test.

For receiver return loss test:

1. Compliance software disconnects the Link to be tested (via the PD Controller embedded in the test equipment).
2. Compliance software sets a USB4 Link or TBT3 compatible Link (to the PD controller).
3. Compliance software connects the Link to be tested (via the PD Controller).
4. The PD Controller sets a USB4 Link or TBT3-compatible Link.
5. Compliance software sets the Port at the test equipment side to either support Gen 3 or to not support Gen 3.
6. Compliance software enables AT Transactions in the Port at the test equipment side in order to proceed with Lane Initialization.
7. Compliance software polls the Rx Enable bit in the Lane under test. When the Rx Enable bit is 1b, compliance software performs the return loss test.
8. Once the test is completed, compliance software disconnects the Link under test.

Note: Compliance Operations do not change the Lane Initialization flow unless specifically written otherwise.

8.3.2.1.1 SET_TX_COMPLIANCE (Required)

The SET_TX_COMPLIANCE Port Operation sets the transmit parameters for a Port.

A USB4 Port shall support the SET_TX_COMPLIANCE Port Operation.

Operation Initiation

This Port Operation does not have Operation Data. The Operation Metadata for the SET_TX_COMPLIANCE Port Operation is defined in Table 8-56.

Table 8-56. SET_TX_COMPLIANCE Operation Metadata

| DW | Bit(s) | Field Name and Description |
|----|--------|---|
| 0 | 5:0 | Port – Identifies the target USB4 Port. For a Router: This field contains the Adapter Number of the Lane 0 Adapter of the target USB4 Port. For a Re-timer: 0h: Target is the USB4 Port whose SB Register Space is written to. 1h: Target is the USB4 Port whose SB Register Space is not written to. |
| | 8:6 | Adapter – Identifies the affected Adapter (s) within the USB4 Port. 000b: Lane 0 Adapter 001b: Lane 1 Adapter 111b: All Adapters All other values are reserved. When this field is set to 111b, the pattern on the Lanes shall have skew between 16 UI and 128 UI. |
| | 12:9 | Pattern – Sets the transmitting pattern. 0000b: PRBS31 - a polynomial $G(x) = x^{31} + x^{28} + 1$ shall be used. 0001b: PRBS15 - a polynomial $G(x) = x^{15} + x^{14} + 1$ shall be used. 0010b: PRBS9 - a polynomial $G(x) = x^9 + x^5 + 1$ shall be used. 0011b: PRBS7 - a polynomial $G(x) = x^7 + x^6 + 1$ shall be used. 0100b: SQ2 - a repeating pattern of bits “101010...”. 0101b: SQ4 - a repeating pattern of bits where the repeating pattern is 2 copies of 1b followed by 2 copies of 0b (“1100...”). 0110b: SQ32 - a repeating pattern of bits where the repeating pattern is 16 copies of 1b followed by 16 copies of 0b. 0111b: SQ128 - a repeating pattern of bits where the repeating pattern is 64 copies of 1b followed by 64 copies of 0b. 1111b: SLOS1. For Example: PRBS7 equals 1000 0011 0000 1010 0011 1100 1000 1011 0011 1010 1001 1111 0100 0011 1000 1001 0011 0110 1011 0111 1011 0001 1010 0101 1101 1100 1100 1010 1011 1111 1000 000b All other values are reserved. |
| | 16:13 | Preset – Sets the preset number for TxFFE parameters at the transmitter (see Table 3-5). |
| | 17 | Set Modifications – Enables non-default values for Preset and for Modifications. 0b: Load default TxFFE parameters and ignore the <i>Modifications</i> field. 1b: Load TxFFE parameters from the <i>Preset</i> field and load any signal shaping parameters based on the <i>Modifications</i> field. |
| | 25:18 | Modifications – Sets other signal shaping parameters. Bit 0: Disable de-emphasis (0b), or enable de-emphasis (1b). Bit 1: Disable pre-shoot (0b), or enable pre-shoot (1b). All other bits are reserved. |
| | 31:26 | Reserved |

Operation Completion

After receiving a SET_TX_COMPLIANCE Port Operation, a USB4 Port shall disable the transition from Training state to CLd state due to timeouts defined in Section 4.2.1.3.3. When the USB4 Port detects SBRX at logical low for tDisconnectRx, it shall re-enable the transition from Training state to CLd state due to timeouts.

The Completion for this Operation does not have Completion Metadata.

The Completion for this Operation does not have Completion Data.

8.3.2.1.2 SET_RX_COMPLIANCE (Required)

The SET_RX_COMPLIANCE Port Operation sets a receiver to compliance test. The Operation Metadata for the SET_RX_COMPLIANCE Port Operation is defined in Table 8-57.

A USB4 Port shall support the SET_RX_COMPLIANCE Port Operation.

After receiving a SET_RX_COMPLIANCE Port Operation, a USB4 Port shall disable the transition from Training state to CLd state due to timeouts defined in Section 4.2.1.3.3. When the USB4 Port detects SBRX at logical low for tDisconnectRx, it shall re-enable the transition from Training state to CLd state due to timeouts.

Operation Initiation

This Port Operation does not have Operation Data.

Table 8-57. SET_RX_COMPLIANCE Operation Metadata

| DW | Bit(s) | Field Name and Description |
|----|--------|---|
| 0 | 5:0 | Port – Identifies the target USB4 Port. For a Router: This field contains the Adapter Number of the Lane 0 Adapter of the target USB4 Port. For a Re-timer: 0h: Target is the USB4 Port whose SB Register Space is written to. 1h: Target is the USB4 Port whose SB Register Space is not written to. |
| | 8:6 | Adapter – Identifies the affected Adapter within the USB4 Port. 000b: Lane 0 Adapter 001b: Lane 1 Adapter All other values are reserved. |
| | 31:9 | Reserved |

Operation Completion

The Completion for this Operation does not have Completion Metadata.

The Completion for this Operation does not have Completion Data.

8.3.2.1.3 START_BER_TEST (Required)

The START_BER_TEST Port Operation starts a receiver compliance test using a chosen BER test pattern. The Operation Metadata for the START_BER_TEST Port Operation is defined in Table 8-58.

A USB4 Port shall support the START_BER_TEST Port Operation.

After receiving this Port Operation, a USB4 Port shall do the following for the Adapter targeted by the Operation:

1. Lock the receiver associated with the Adapter on the BER test pattern defined in the Operation Metadata.
2. Set the *DW Count*, *Bit Error Count*, *Error Capture Count*, and *Burst Restart Count* counters to 0.
3. Continue running the BER test pattern until an END_BER_TEST or an END_BURST_TEST Port Operation is received.

**CONNECTION MANAGER NOTE**

The next Port Operation after a START_BER_TEST Port Operation shall be an END_BER_TEST Port Operation or END_BURST_TEST Port Operation. A Connection Manager shall not initiate another Port Operation while a BER test pattern is running.

Operation Initiation

This Port Operation does not have Operation Data.

Table 8-58. START_BER_TEST Operation Metadata

| DW | Bit(s) | Field Name and Description |
|----|--------|---|
| 0 | 5:0 | Port – Identifies the target USB4 Port. For a Router: This field contains the Adapter Number of the Lane 0 Adapter of the target USB4 Port. For a Re-timer: 0h: Target is the USB4 Port whose SB Register Space is written to. 1h: Target is the USB4 Port whose SB Register Space is not written to. |
| | 8:6 | Adapter – Identifies the affected Adapter within the USB4 Port. 000b: Lane 0 Adapter 001b: Lane 1 Adapter All other values are reserved. |
| | 12:9 | Pattern – Sets the transmitting pattern. 0000b: PRBS31 0001b: PRBS15 0010b: PRBS9 0011b: PRBS7 0100b: SQ2 (optional) 0101b: SQ4 (optional) 0110b: SQ32 (optional) All other values are reserved. |
| | 31:13 | Reserved |

Operation Completion

The Completion for this Operation does not have Completion Metadata.

The Completion for this Operation does not have Completion Data.

8.3.2.1.4 END_BER_TEST (Required)

The END_BER_TEST Port Operation stops a running receiver BER test pattern. The Operation Metadata for the END_BER_TEST Port Operation is defined in Table 8-59.

A USB4 Port shall support the END_BER_TEST Port Operation.

Operation Initiation

This Port Operation does not have Operation Data.

Table 8-59. END_BER_TEST Operation Metadata

| DW | Bit(s) | Field Name and Description |
|----|--------|---|
| 0 | 5:0 | Port – Identifies the target USB4 Port. For a Router: This field contains the Adapter Number of the Lane 0 Adapter of the target USB4 Port. For a Re-timer: 0h: Target is the USB4 Port whose SB Register Space is written to. 1h: Target is the USB4 Port whose SB Register Space is not written to. |
| | 8:6 | Adapter – Identifies the affected Adapter within the USB4 Port. 000b: Lane 0 Adapter 001b: Lane 1 Adapter All other values are reserved. |
| | 31:9 | Reserved |

Operation Completion

After receiving this Port Operation, a USB4 Port shall stop the *DW Count* and *Bit Error Count* counters associated with the Adapter in the Operation Metadata and shall update the Completion Data as defined in Table 8-60.

The Completion for this Operation does not have Completion Metadata.

Table 8-60. END_BER_TEST Completion Data

| DW | Bit(s) | Field Name and Description |
|----|--------|--|
| 0 | 31:0 | DW Count Low - This field contains the lowest 32 bits of the <i>DW Count</i> . The <i>DW Count</i> is a 48-bit field that contains the number of Doublewords received during the test. The counter increments from 0 and shall stop counting at FF...Fh. |
| 1 | 15:0 | DW Count High - This field contains the highest 16 bits of a 48-bit number of bits received during the test. This field contains the highest 16 bits of the <i>DW Count</i> . The <i>DW Count</i> is a 48-bit field that contains the number of Doublewords received during the test. The counter increments from 0 and shall stop counting at FF...Fh. |
| | 31:16 | Bit Error Count – Number of bit errors encountered during the test. The counter increments from 0 and shall stop counting at FFFFh. |

**CONNECTION MANAGER NOTE**

After a USB4 Port executes the END_BER_TEST Operation, the Connection Manager can bring the USB4 Port out of compliance mode by transitioning the Router with the USB4 Port to the Uninitialized state.

8.3.2.1.5 END_BURST_TEST (Conditional)

The END_BURST_TEST Port Operation stops a running receiver burst BER test pattern.

A USB4 Port shall support the END_BURST_TEST Port Operation if it employs DFE with more than one tap (see Section 3.5.3). Otherwise this Port Operation is not applicable.

Operation Initiation

The Operation Metadata for the END_BURST_TEST Port Operation is defined in Table 8-61.

This Port Operation does not have Operation Data.

Table 8-61. END_BURST_TEST Operation Metadata

| DW | Bit(s) | Field Name and Description |
|----|--------|---|
| 0 | 5:0 | Port – Identifies the target USB4 Port. For a Router: This field contains the Adapter Number of the Lane 0 Adapter of the target USB4 Port. For a Re-timer: 0h: Target is the USB4 Port whose SB Register Space is written to. 1h: Target is the USB4 Port whose SB Register Space is not written to. |
| | 8:6 | Adapter – Identifies the affected Adapter within the USB4 Port. 000b: Lane 0 Adapter 001b: Lane 1 Adapter All other values are reserved. |
| | 31:9 | Reserved |

Operation Completion

After receiving this Port Operation, a USB4 Port shall stop the *DW Count*, *Burst Restart Count*, and *Error Capture Count* counters associated with the Adapter in the Operation Metadata and shall update the Completion Data as defined in Table 8-62.

The Completion for this Operation does not have Completion Metadata.

Table 8-62. END_BURST_TEST Completion Data

| DW | Bit(s) | Field Name and Description |
|----|--------|---|
| 0 | 31:0 | DW Count Low - This field contains the lowest 32 bits of the <i>DW Count</i> . The <i>DW Count</i> is a 48-bit field that contains the number of Doublewords received during the test. The counter increments from 0 and shall stop counting at FF...Fh. |
| 1 | 15:0 | DW Count High - This field contains the highest 16 bits of a 48-bit number of bits received during the test. This field contains the highest 16 bits of the <i>DW Count</i> . The <i>DW Count</i> is a 48-bit field that contains the number of Doublewords received during the test. The counter increments from 0 and shall stop counting at FF...Fh. |
| | 31:16 | Burst Restart Count – Number of burst errors encountered during the test. The counter increments from 0 and shall stop counting at FFFFh. See Section 3.5.3 for the definition of the <i>Burst Restart Count</i> counter. |
| 2 | 31:0 | Error Capture Count – Number of error events encountered during the test. The counter increments from 0 and shall stop counting at FF...Fh. See Section 3.5.3 for the definition of the <i>Error Capture Count</i> counter. |

**CONNECTION MANAGER NOTE**

After a USB4 Port executes the *END_BURST_TEST* Operation, the Connection Manager can bring the USB4 Port out of compliance mode by transitioning the Router with the USB4 Port to the Uninitialized state.

8.3.2.1.6 READ_BURST_TEST (Conditional)

The *READ_BURST_TEST* Port Operation reads the *DW Count*, *Burst Restart Count*, and *Error Capture Count* counters. The Operation does not stop execution of the test.

A USB4 Port shall support the *READ_BURST_TEST* Port Operation if it employs DFE with more than one tap (see Section 3.5.3). Otherwise this Port Operation is not applicable.

Operation Initiation

The Operation Metadata for the READ_BURST_TEST Port Operation is defined in Table 8-63.

This Port Operation does not have Operation Data.

Table 8-63. READ_BURST_TEST Operation Metadata

| DW | Bit(s) | Field Name and Description |
|----|--------|---|
| 0 | 5:0 | Port – Identifies the target USB4 Port. For a Router: This field contains the Adapter Number of the Lane 0 Adapter of the target USB4 Port. For a Re-timer: 0h: Target is the USB4 Port whose SB Register Space is written to. 1h: Target is the USB4 Port whose SB Register Space is not written to. |
| | 8:6 | Adapter – Identifies the Adapter within the USB4 Port. 000b: Lane 0 Adapter 001b: Lane 1 Adapter All other values are reserved. |
| | 31:9 | Reserved |

Operation Completion

After receiving this Port Operation, a USB4 Port shall update the Completion Data as defined in Table 8-64.

The Completion for this Operation does not have Completion Metadata.

Table 8-64. READ_BURST_TEST Completion Data

| DW | Bit(s) | Field Name and Description |
|----|--------|---|
| 0 | 31:0 | DW Count Low - This field contains the lowest 32 bits of the DW Count. The DW Count is a 48-bit field that contains the number of Doublewords received during the test. The counter increments from 0 and shall stop counting at FF...Fh. |
| 1 | 15:0 | DW Count High - This field contains the highest 16 bits of a 48-bit number of bits received during the test. This field contains the highest 16 bits of the DW Count. The DW Count is a 48-bit field that contains the number of Doublewords received during the test. The counter increments from 0 and shall stop counting at FF...Fh. |
| | 31:16 | Burst Restart Count – Number of burst errors encountered during the test. The counter increments from 0 and shall stop counting at FFFFh. See Section 3.5.3 for the definition of the Burst Restart Count counter. |
| 2 | 31:0 | Error Capture Count – Number of error events encountered during the test. The counter increments from 0 and shall stop counting at FF...Fh. See Section 3.5.3 for the definition of the Error Capture Count counter. |

8.3.2.1.7 ENTER_EI_TEST (Required)

The ENTER_EI_TEST Port Operation brings a transmitter into electrical idle state.

A USB4 Port shall support the ENTER_EI_TEST Port Operation.

Operation Initiation

The Operation Metadata for the ENTER_EI_TEST Port Operation is defined in Table 8-65.

Table 8-65. ENTER_EI_TEST Operation Metadata

| DW | Bit(s) | Field Name and Description |
|----|--------|---|
| 0 | 5:0 | Port – Identifies the target USB4 Port. For a Router: This field contains the Adapter Number of the Lane 0 Adapter of the target USB4 Port. For a Re-timer: 0h: Target is the USB4 Port whose SB Register Space is written to. 1h: Target is the USB4 Port whose SB Register Space is not written to. |
| | 8:6 | Adapter – Identifies the Adapter within the USB4 Port. 000b: Lane 0 Adapter 001b: Lane 1 Adapter All other values are reserved. |
| | 31:9 | Reserved |

This Port Operation does not have Operation Data.

Operation Completion

After receiving an ENTER_EI_TEST Port Operation, a USB4 Port shall disable the transition from Training state to CLd state due to timeouts (defined in Section 4.2.1.3.3). When the USB4 Port detects SBRX at logical low for tDisconnectRx, it shall re-enable the transition from Training state to CLd state due to timeouts.

A Router that receives an ENTER_EI_TEST Port Operation shall transition the Lane transmitter defined in the Operation into electrical idle state.

The Completion for this Operation does not have Completion Metadata.

The Completion for this Operation does not have Completion Data.

8.3.2.2 Service Port Operations

8.3.2.2.1 ROUTER_OFFLINE_MODE (Required)

The ROUTER_OFFLINE_MODE Port Operation transitions a USB4 Port that is not connected into an offline mode. When in this mode, the USB4 Port shall not perform Lane Initialization. The USB4 Port still processes Operations delivered from software.

A USB4 Port shall support the ROUTER_OFFLINE_MODE Port Operation. A USB4 Port shall execute this Operation when delivered locally. A USB4 Port shall reject this Operation when delivered from the Sideband Channel.

Operation Initiation

The *Target* field in the USB4 Port Capability is set to 000b.

This Port Operation does not have Operation Data.

Table 8-66. ROUTER_OFFLINE_MODE Operation Metadata

| DW | Bit(s) | Field Name and Description |
|----|--------|--|
| 0 | 0 | Enter Offline Mode When set to 0b, the USB4 Port shall enter offline mode. When set to 1b, the USB4 Port shall exit offline mode. |
| | 31:1 | Reserved |

Operation Completion

The Completion for this Operation does not have Completion Metadata.

The Completion for this Operation does not have Completion Data.

8.3.2.2.2 ENUMERATE_RE-TIMERS (Required)

The ENUMERATE_RE-TIMERS Port Operation causes a USB4 Port to send a Broadcast RT Transaction.

A USB4 Port shall support the ENUMERATE_RE-TIMERS Port Operation. A USB4 Port shall execute this Operation when delivered locally. A USB4 Port shall reject this Operation when delivered from the Sideband Channel.

Operation Initiation

This Port Operation does not have Operation Metadata.

This Port Operation does not have Operation Data.

Operation Completion

The Completion for this Operation does not have Completion Metadata.

The Completion for this Operation does not have Completion Data.

8.3.2.3 Receiver Lane Margining Port Operations

The Port Operations defined in the subsections below are used for Receiver Lane Margining tests. See Section 3.8 for more details on Receiver Lane Margining.

8.3.2.3.1 READ_LANE_MARGIN_CAP (Required)

The READ_LANE_MARGIN_CAP Port Operation reads the Receiver Lane Margining capabilities of a USB4 Port.

A USB4 Port shall support the READ_LANE_MARGIN_CAP Port Operation.

Operation Initiation

This Port Operation does not have Operation Metadata.

This Port Operation does not have Operation Data.

Operation Completion

If completed successfully, the Port Operation returns two Doublewords of Completion Data as shown in Table 8-67.

The Completion for this Operation does not have Completion Metadata.

Table 8-67. READ_LANE_MARGIN_CAP Completion Data

| DW | Bit(s) | Field Name and Description |
|----|--------|--|
| 0 | 1:0 | Receiver Lane Margining Modes This field indicates the Receiver Lane Margining modes supported by the USB4 Port. 00b: Reserved 01b: Only HW mode supported 10b: Only SW mode supported 11b: Both HW mode and SW mode supported |

| DW | Bit(s) | Field Name and Description |
|----|--------|---|
| | 2 | <p>Two-Lane Margining</p> <p>This field indicates whether the USB4 Port supports simultaneous Lane Margining on both Lanes.</p> <p>0b: Each Lane is tested separately 1b: Both Lanes can be tested simultaneously</p> |
| 0 | 4:3 | <p>Independent High/Low Voltage Margin</p> <p>This field indicates how voltage margins are returned:</p> <p>00b: The Router returns the minimum between the high and low voltage margins 01b: The Router returns either the high voltage margin or the low voltage margin 10b: The Router returns both the high voltage margin and the low voltage margin 11b: Reserved</p> |
| | 5 | <p>Time Margining</p> <p>This field indicates if Time Margining is supported.</p> <p>0b: Time Margining is not supported 1b: Time Margining is supported</p> |
| | 12:6 | <p>Voltage Margin Steps – Mandatory Range</p> <p>This field contains the number of voltage margin steps that are supported by the USB4 Port. The number of margin steps applies to both the high margin and the low margin separately. A Router shall set this value to a minimum of 25.</p> |
| | 18:13 | <p>Maximum Voltage Offset – Mandatory Range</p> <p>This field contains the maximum supported voltage offset. The maximum voltage offset applies to both the high margin and the low margin separately. A value of n in this field corresponds to a maximum voltage offset of $[74 + n \times 2]$ mV.</p> |
| | 19 | <p>Optional Voltage Offset Range</p> <p>This field indicates if the USB4 Port supports an optional voltage offset range.</p> <p>0b: Optional Voltage Offset Range is not supported 1b: Optional Voltage Offset Range is supported</p> |
| | 26:20 | <p>Voltage Margin Steps – Optional Range</p> <p>This field contains the number of voltage margin steps that are supported by the USB4 Port for the Optional Voltage Offset Range. The number of margin steps applies to both the high margin and the low margin separately. A Router shall set this value to a minimum of 25. This field is reserved if the <i>Optional Voltage Offset Range</i> field is 0b.</p> |
| | 31:27 | Reserved |
| 1 | 7:0 | <p>Maximum Voltage Offset – Optional Range</p> <p>This field contains the maximum supported voltage offset for the optional range. The maximum voltage offset applies to both the high margin and the low margin separately. A value of n in this field corresponds to a maximum voltage offset of $[74 + n \times 2]$ mV. Values larger than A3h are Reserved. This field is reserved if the <i>Optional Voltage Offset Range</i> field is 0b.</p> |
| | 8 | <p>Destructive Time Margining</p> <p>This bit indicates whether the time margining test is destructive or non-destructive:</p> <p>0b: Non-destructive (will not cause actual bit errors on the Link) 1b: Destructive (can cause actual bit errors on the Link)</p> <p>This bit shall be set to 0b if the <i>Time Margining</i> bit is set to 0b.</p> |
| | 10:9 | <p>Independent Left/Right Timing Margin Test</p> <p>This field indicates how timing margins are returned:</p> <p>00b: The Router returns the minimum between the left and right time margins 01b: The Router returns either the left margin or the right time margin 10b: The Router returns both the left margin and the right margin. 11b: Reserved</p> <p>This field shall be set to 0 if the <i>Time Margining</i> bit is set to 0b.</p> |

| DW | Bit(s) | Field Name and Description |
|----|--------|--|
| | 15:11 | Time Margin Steps This field contains the number of time margin steps that are supported by the USB4 Port. The number of margin steps applies to both the left margin and the right margin separately. This field shall be set to 0 if the <i>Time Margining</i> bit is set to 0b. Else, this field shall be set to a value between 07h and 1Fh. |
| 1 | 20:16 | Maximum Time Offset This field contains the maximum time offset supported by the USB4 Port. The maximum time offset applies to both the left margin and the right margin separately. A value of n in this field corresponds to a maximum offset of $[0.2 + 0.01 \times n]$ UI. This field shall be set to 0 if the <i>Time Margining</i> bit is set to 0b. Else, this field shall be set to a value between 0h and 1Eh. |
| | 25:21 | Min BER Level Contour Support The value in this field (n) indicates the minimum BER level contour supported: For even values of n, the minimal BER level contour supported is $1e[-12 + n/2]$. For odd values of n, the minimal BER level contour supported is $3 \times 1e[-12 + (n-1)/2]$. Values larger than 10h are Reserved. This field is Reserved if the <i>Receiver Lane Margining Modes</i> field is 10b. |
| | 30:26 | Max BER Level Contour Support The value in this field (n) indicates the maximum BER level contour supported. For even values n, the maximum BER level contour supported is $1e[-12 + n/2]$. For odd values n, the maximum BER level contour supported is $3 \times 1e[-12 + (n-1)/2]$. Values larger than 10h are Reserved. This field is Reserved if the <i>Receiver Lane Margining Modes</i> field is 10b. |
| | 31 | Reserved |

8.3.2.3.2 RUN_HW_LANE_MARGINING (Conditional)

The RUN_HW_LANE_MARGINING Operation sets the parameters for HW Mode Receiver Lane Margining and executes a HW Mode Receiver Lane Margining test. If the Port Operation completes successfully, the target of the Operation shall set the Completion Metadata listed in Table 8-70.

A USB4 Port shall support the RUN_HW_LANE_MARGINING Port Operation if software margining mode is not supported (see Section 3.8.1). If software margining mode is supported, support for this Port Operation is optional.

Operation Initiation

This Port Operation does not have Operation Data.

Table 8-68. RUN_HW_LANE_MARGINING Operation Metadata

| DW | Bit(s) | Field Name and Description |
|----|--------|--|
| 0 | 2:0 | Lane Select This field indicates which Lane(s) to perform the Receiver Lane Margining test on: 000b: Lane 0 001b: Lane 1 111b: Test runs on all Lanes Other values are Reserved. The Port Operation shall fail if this field is set to 111b and the Router supports Lane Margining on a single Lane only as present in the <i>Two-Lane Margining</i> bit. |

| | | |
|---|-------|--|
| | 3 | <p>Timing Margin Test</p> <p>This bit selects between a voltage margin test and a timing margin test</p> <p>0b: Perform a voltage margin test</p> <p>1b: Perform a timing margin test</p> <p>The Port Operation shall fail if this field is set to 1b and the Router does not support timing margin testing.</p> |
| 0 | 4 | <p>Enable Margin Tests</p> <p>For a voltage margin test, this field enables high margin test or low margin test when the <i>Independent High/Low Voltage Margin</i> field in the READ_LANE_MARGIN_CAP Completion Metadata was set to 01b:</p> <p>0b: Enables low margin test</p> <p>1b: Enables high margin test</p> <p>For a timing margin test, this field enables right margin test or left margin test when the <i>Independent Left/Right Timing Margin Test</i> field in the READ_LANE_MARGIN_CAP Completion Metadata was set to 01b:</p> <p>0b: Enables left margin test</p> <p>1b: Enables right margin test</p> |
| | 9:5 | <p>BER Level Contour</p> <p>This field sets the BER level contour:</p> <p>For even values n, the maximum BER level contour supported is $1e[-12 + n/2]$.</p> <p>For odd values n, the maximum BER level contour supported is $3 \times 1e[-12 + (n-1)/2]$.</p> <p>Values larger than 10h are Reserved.</p> |
| | 10 | <p>Enable Optional Voltage Offset Range</p> <p>For a voltage margin test, this field enables the mandatory voltage margin offset range or the optional extended voltage offset range when the <i>Optional Voltage Offset Range</i> field in the READ_LANE_MARGIN_CAP Completion Metadata was set to 1b:</p> <p>0b: Enables the mandatory Rx voltage offset range</p> <p>1b: Enables the optional Rx voltage offset range</p> |
| | 31:11 | Reserved |

Operation Completion

Completion Data for the RUN_HW_LANE_MARGINING Operation depicted in Table 8-70 contains the results for the HW test. Table 8-69 defines how the contents of each field are selected.

The Completion for this Operation does not have Completion Metadata.

Table 8-69. Contents Selection for RUN_HW_LANE_MARGINING Completion Data

| Field | Lane Select | Independent High/Low Voltage Margin & Independent Left/Right Timing Margin Test | Enable Margin Tests | Contents | |
|------------------------------|--------------|---|--|---|--|
| High / Right Margin (Lane 0) | 001b | x | x | Reserved. | |
| | 000b or 111b | 00b | x | Minimum between the high and low voltage margins / minimum between the left and right time margins. | |
| | | 01b | 0b | 0b | Reserved. |
| | | | 1b | 1b | High voltage margin / right time margin. |
| 10b | x | x | High voltage margin / right time margin. | | |
| Low / Left Margin (Lane 0) | 001b | x | x | Reserved. | |
| | 000b or 111b | 00b | x | Reserved. | |
| | | 01b | 0b | 0b | Low voltage margin / left time margin. |
| | | | 1b | 1b | Reserved. |

| Field | Lane Select | Independent High/Low Voltage Margin & Independent Left/Right Timing Margin Test | Enable Margin Tests | Contents |
|------------------------------|--------------|---|---------------------|---|
| | | 10b | x | Low voltage margin / left time margin. |
| High / Right Margin (Lane 1) | 000b | x | x | Reserved. |
| | 001b or 111b | 00b | x | Minimum between the high and low voltage margins / minimum between the left and right time margins. |
| | | | 0b | Reserved. |
| | | 01b | 1b | High voltage margin / right time margin. |
| | | | x | High voltage margin / right time margin. |
| Low / Left Margin (Lane 1) | 000b | x | x | Reserved. |
| | 001b or 111b | 00b | x | Reserved. |
| | | | 0b | Low voltage margin / left time margin. |
| | | 01b | 1b | Reserved. |
| | | | x | Low voltage margin / left time margin. |

Table 8-70. RUN_HW_LANE_MARGINING Completion Data

| DW | Bit(s) | Field Name and Description |
|----|--------|--|
| 0 | 2:0 | Lane Select This field contains the value of the Lane Select field in the last RUN_HW_LANE_MARGINING Port Operation. |
| | 3 | Timing Margin Test This field contains the value of the Timing Margin Test field in the last RUN_HW_LANE_MARGINING Port Operation. |
| | 4 | Enable Margin Tests This field contains the value of the Enable Margin Tests field in the last RUN_HW_LANE_MARGINING Port Operation. |
| | 9:5 | BER Level Contour This field contains the value of the <i>BER Level Contour</i> field in the last RUN_HW_LANE_MARGINING Port Operation. |
| | 10 | Enable Optional Voltage Offset Range This field contains the value of the <i>Enable Optional Voltage Offset Range</i> field in the last RUN_HW_LANE_MARGINING Port Operation. |
| | 31:11 | Reserved. |
| 1 | 6:0 | High / Right Margin (Lane 0) For a voltage margin test, this field contains the high margin test result for Lane 0 in terms of number of voltage offset steps to the BER Level Contour. If the high margin exceeds the Maximum Voltage Offset this field contains the maximum voltage offset step which was applied during the test. See Note 1. For a time margin test, this field contains the right margin test result for Lane 0 in terms of number of time offset steps to the BER Level Contour. If the right margin exceeds the Maximum Time Offset this field contains the maximum right offset step which was applied during the test. See Note 2. |

| DW | Bit(s) | Field Name and Description |
|----|--------|---|
| 1 | 7 | <p>High / Right Margin Exceeds Maximum Voltage / Time Offset (Lane 0)</p> <p>For a voltage margin test, this field indicates if the high margin exceeds the Maximum Voltage Offset for Lane 0:</p> <p>0b: High margin does not exceed the Maximum Voltage Offset 1b: High margin exceeds the Maximum Voltage Offset</p> <p>For a time margin test, this field indicates if the right margin exceeds the Maximum Time Offset for Lane 0:</p> <p>0b: Right margin does not exceed the Maximum Time Offset 1b: Right margin exceeds the Maximum Time Offset</p> |
| | 14:8 | <p>Low / Left Margin (Lane 0)</p> <p>For a voltage margin test, this field contains the low margin test result for Lane 0 in terms of number of voltage offset steps to the BER Level Contour. If the Low margin exceeds the Maximum Voltage Offset this field contains the maximum voltage offset step which was applied during the test. See Note 3.</p> <p>For a time margin test, this field contains the left margin test result for Lane 0 in terms of number of time offset steps to the BER Level Contour. If the left margin exceeds the Maximum Time Offset this field contains the maximum left offset step which was applied during the test. See Note 4.</p> |
| | 15 | <p>Low / Left Margin Exceeds Maximum Voltage / Time Offset (Lane 0)</p> <p>For a voltage margin test, this field indicates if the low margin exceeds the Maximum Voltage Offset for Lane 0:</p> <p>0b: Low margin does not exceed the Maximum Voltage Offset 1b: Low margin exceeds the Maximum Voltage Offset</p> <p>For a time margin test, this field indicates if the left margin exceeds the Maximum Time Offset for Lane 0:</p> <p>0b: Left margin does not exceed the Maximum Time Offset 1b: Left margin exceeds the Maximum Time Offset</p> |
| | 22:16 | <p>High / Right Margin (Lane 1)</p> <p>For a voltage margin test, this field contains the high margin test result for Lane 1 in terms of number of voltage offset steps to the BER Level Contour. If the high margin exceeds the Maximum Voltage Offset this field contains the maximum voltage offset step which was applied during the test. See Note 1.</p> <p>For a time margin test, this field contains the right margin test result for Lane 1 in terms of number of time offset steps to the BER Level Contour. If the right margin exceeds the Maximum Time Offset this field contains the maximum right offset step which was applied during the test. See Note 2.</p> |
| | 23 | <p>High / Right Margin Exceeds Maximum Voltage / Time Offset (Lane 1)</p> <p>For a voltage margin test, this field indicates if the high margin exceeds the Maximum Voltage Offset for Lane 1:</p> <p>0b: High margin does not exceed the Maximum Voltage Offset 1b: High margin exceeds the Maximum Voltage Offset</p> <p>For a time margin test, this field indicates if the right margin exceeds the Maximum Time Offset for Lane 1:</p> <p>0b: Right margin does not exceed the Maximum Time Offset 1b: Right margin exceeds the Maximum Time Offset</p> |
| | 30:24 | <p>Low / Left Margin (Lane 1)</p> <p>For a voltage margin test, this field contains the low margin test result for Lane 1 in terms of number of voltage offset steps to the BER Level Contour. If the low margin exceeds the Maximum Voltage Offset this field contains the maximum voltage offset step which was applied during the test. See Note 3.</p> <p>For a time margin test, this field contains the left margin test result for Lane 1 in terms of number of time offset steps to the BER Level Contour. If the left margin exceeds the Maximum Time Offset this field contains the maximum left offset step which was applied during the test. See Note 4.</p> |

| DW | Bit(s) | Field Name and Description |
|---|--------|---|
| 1 | 31 | <p>Low / Left Margin Exceeds Maximum Voltage / Time Offset (Lane 1)</p> <p>For a voltage margin test, this field indicates if the low margin exceeds the Maximum Voltage Offset for Lane 1:</p> <p>0b: Low margin does not exceed the Maximum Voltage Offset</p> <p>1b: Low margin exceeds the Maximum Voltage Offset</p> <p>For a time margin test, this field indicates if the left margin exceeds the Maximum Time Offset for Lane 1:</p> <p>0b: Left margin does not exceed the Maximum Time Offset</p> <p>1b: Left margin exceeds the Maximum Time Offset</p> |
| <p>Note 1. The High Margin test result can be reported by application software in units of mV by the following conversion:</p> $\text{High Margin (mV)} = \text{High Margin (steps)} \times \text{Maximum Voltage Offset (mV)} \div \text{Voltage Margin Steps (steps)}$ <p>Note 2. The Right Margin test result can be reported by application software in units of UI by the following conversion:</p> $\text{Right Margin (UI)} = \text{Right Margin (steps)} \times \text{Maximum Time Offset (UI)} \div \text{Time Margin Steps (steps)}$ <p>Note 3. The Low Margin test result can be reported by application software in units of mV by the following conversion:</p> $\text{Low Margin (mV)} = \text{Low Margin (steps)} \times \text{Maximum Voltage Offset (mV)} \div \text{Voltage Margin Steps (steps)}$ <p>Note 4. The Left Margin test result can be reported by application software in units of UI by the following conversion:</p> $\text{Left Margin (UI)} = \text{Left Margin (steps)} \times \text{Maximum Time Offset (UI)} \div \text{Time Margin Steps (steps)}$ | | |



CONNECTION MANAGER NOTE

If the RUN_HW_LANE_MARGINING Operation runs a timing margin test, and the test is destructive, then following the execution of this Operation, the Connection Manager can bring the Router out of compliance mode by transitioning it to the Uninitialized state.

8.3.2.3.3 RUN_SW_LANE_MARGINING (Conditional)

The RUN_SW_LANE_MARGINING sets the parameters for SW Mode Receiver Lane Margining and starts the SW Mode Receiver Lane Margining test.

A USB4 Port shall support the RUN_SW_LANE_MARGINING Port Operation if hardware margining mode is not supported (see Section 3.8.1). If hardware margining mode is supported, support for this Port Operation is optional.

Operation Initiation

This Port Operation does not have Operation Data.

Table 8-71. RUN_SW_LANE_MARGINING Operation Metadata

| DW | Bit(s) | Field Name and Description |
|----|--------|--|
| 0 | 2:0 | <p>Lane Select</p> <p>This field indicates which Lane(s) to perform the Receiver Lane Margining test on:</p> <p>000b: Lane 0 001b: Lane 1 111b: Test runs on all Lanes</p> <p>Other values are Reserved.</p> <p>The Port Operation shall fail if this field is set to 111b and the Router supports Lane Margining on a single Lane only as present in the <i>Two-Lane Margining</i> bit.</p> |
| | 3 | <p>Timing Margin Test</p> <p>This bit selects between a voltage margin test and a timing margin test:</p> <p>0b: Perform a voltage margin test 1b: Perform a timing margin test</p> <p>The Port Operation shall fail if this field is set to 1b and the Router does not support timing margin testing.</p> |
| 0 | 4 | <p>Enable Margin Tests</p> <p>For a voltage margin test, this field enables high margin test or low margin test when the <i>Independent High/Low Voltage Margin</i> field in the READ_LANE_MARGIN_CAP Completion Metadata was set to 01b:</p> <p>0b: Enables low margin test 1b: Enables high margin test</p> <p>For a timing margin test, this field enables right margin test or left margin test when the <i>Independent Left/Right Timing Margin Test</i> field in the READ_LANE_MARGIN_CAP Completion Metadata was set to 01b:</p> <p>0b: Enables left margin test 1b: Enables right margin test</p> |
| | 5 | <p>Enable Optional Voltage Offset Range</p> <p>For a voltage margin test, this field enables the mandatory voltage margin offset range or the optional extended voltage offset range when the <i>Optional Voltage Offset Range</i> field in the READ_LANE_MARGIN_CAP Completion Metadata was set to 1b:</p> <p>0b: Enables the mandatory Rx voltage offset range 1b: Enables the optional Rx voltage offset range</p> |
| | 12:6 | <p>Voltage Offset / Time Offset</p> <p>In a voltage margin test, this field sets the number of voltage offset steps.</p> <p>In a timing margin test, this field sets the number of time offset steps.</p> <p>Values larger than 1Fh are Reserved when performing a timing margin test.</p> |
| | 14:13 | <p>Error Counter</p> <p>00b: NOP – no change in counter setup 01b: CLEAR – set the error counter to 0 and enable counter 10b: START – start counter, continue counting from last value 11b: STOP – stop counter, do not clear value</p> <p>This field affects the Error Counter for the Lanes that perform the test as selected by the <i>Lane Select</i> field.</p> <p>This field is Reserved if the target of the Port Operation supports destructive Time Margining.</p> |
| | 31:15 | Reserved |

Operation Completion

The Completion for this Operation does not have Completion Metadata.

The Completion Data for the RUN_SW_LANE_MARGINING Operation depicted in Table 8-72 contains the configuration for the SW test.

Table 8-72. RUN_SW_LANE_MARGINING Completion Data

| DW | Bit(s) | Field Name and Description |
|----|--------|---|
| 0 | 2:0 | Lane Select This field contains the value of the <i>Lane Select</i> field in the last RUN_SW_LANE_MARGINING Port Operation. |
| | 3 | Timing Margin Test This field contains the value of the <i>Timing Margin Test</i> field in the last RUN_SW_LANE_MARGINING Port Operation. |
| | 4 | Enable Margin Tests This field contains the value of the <i>Enable Margin Tests</i> field in the last RUN_SW_LANE_MARGINING Port Operation. |
| | 5 | Enable Optional Voltage Offset Range This field contains the value of the <i>Enable Optional Voltage Offset Range</i> field in the last RUN_SW_LANE_MARGINING Port Operation. |
| 0 | 12:6 | Voltage Offset / Time Offset This field contains the value of the <i>Voltage Offset / Time Offset</i> field in the last RUN_SW_LANE_MARGINING Port Operation. |
| | 31:13 | Reserved |

8.3.2.3.4 READ_SW_MARGIN_ERR (Conditional)

The READ_SW_MARGIN_ERR Port Operation reads the error indicators of a Receiver SW Lane Margining test.

A USB4 Port shall support the READ_SW_MARGIN_ERR Port Operation if hardware margining mode is not supported (see Section 3.8.1). If hardware margining mode is supported, support for this Port Operation is optional.

If the target of the Port Operation supports Destructive Time Margining, it shall set the Opcode register to a FourCC value of "!CMD".

Operation Initiation

This Port Operation does not have Operation Metadata.

This Port Operation does not have Operation Data.

Operation Completion

If completed successfully, the Port Operation returns Length number of Doublewords of Completion Metadata.

The Completion for this Operation does not have Completion Data.

Table 8-73. READ_SW_MARGIN_ERR Completion Metadata


| DW | Bit(s) | Field Name and Description |
|----|--------|--|
| 0 | 3:0 | Error Counter (Lane 0) This field contains the number of samples with values below the voltage offset measured on Lane 0. The counter value increments from 0 and shall stop counting at 0Fh. |
| | 7:4 | Error Counter (Lane 1) This field contains the number of samples with values below the voltage offset measured on Lane 1. The counter value increments from 0 and shall stop counting at 0Fh. |
| | 31:8 | Reserved |



CONNECTION MANAGER NOTE

If the READ_SW_MARGIN_ERR Operation runs a timing margin test, and the test is destructive, then following the execution of this Operation, the Connection Manager can bring the Router out of compliance mode by transitioning it to the Uninitialized state.

9 USB3 Tunneling



USB4™ Hosts:

A USB4 host shall support USB3 tunneling. The Host Router in a USB4 host shall have one Downstream USB3 Adapter per Downstream Facing Port. Each Downstream USB3 Adapter shall interface to a downstream port on the host controller.

USB4 Hubs:

A USB4 hub shall support USB3 tunneling. A USB4 hub shall incorporate an internal USB SuperSpeed Plus hub per the USB 3.2 Specification with the modifications defined in this chapter.

The Device Router in a USB4 hub shall have:

- An Upstream USB3 Adapter that interfaces with the upstream port of the internal USB 3.2 hub.
- For each Downstream Facing Port, a Downstream USB3 Adapter that interfaces to a downstream port on the internal USB 3.2 hub.

USB4 Devices:

A USB4 device may optionally support USB3 tunneling. A USB4 device that supports USB3 tunneling shall incorporate either an internal USB SuperSpeed Plus hub or an internal USB peripheral device per the USB 3.2 Specification with the modifications defined in this chapter.

The Device Router in a USB4 device that supports USB3 tunneling shall have an Upstream USB3 Adapter that interfaces to the upstream port of the internal USB 3.2 peripheral device or internal USB 3.2 hub.

Note: In this Chapter, the term “internal USB3 device” refers to either an internal USB SuperSpeed Plus hub, internal USB peripheral device, or internal host controller.

9.1 USB3 Adapter Layer

9.1.1 Encapsulation

When a USB3 Adapter Layer receives one of the USB3 constructs listed in this section from the internal USB3 device, it encapsulates the construct in a Tunneled Packet then passes it to the Transport layer for transmission over the USB4 fabric. The USB3 Adapter Layer encapsulates the following USB3 constructs:

- Low Frequency Periodic Signaling (LFPS).
- Ordered Sets.
- Link Commands.
- Link Management Packets (LMP).
- Transaction Packets (TP).
- Isochronous Timestamp Packets (ITP).
- Data Packets (DP).

A USB3 Adapter layer shall follow the rules below when encapsulating a USB3 construct into a Tunneled Packet:

- An LFPS event, a Ordered Set, a Link Command, a Link Management Packet, a Transaction Packet, and an Isochronous Timestamp Packet shall each be encapsulated into a single separate Tunneled Packet.
- A Data Packet shall be encapsulated into one or more Tunneled Packets. A Tunneled Packet shall not contain data from more than one Data Packet. See Section 9.1.1.8.
- Idle Symbols shall be discarded by a USB3 Adapter Layer.
- The byte and bit ordering of an LFPS within a Tunneled Packet shall be as described in Section 9.1.1.1.
- The byte and bit ordering of an Ordered Set within a Tunneled Packet shall be as described in Section 9.1.1.2.
- The bytes and bits in a Tunneled Packet payload, other than LFPS and Ordered Sets, shall be packed in the same order as the original USB3 construct, including the USB3 framing symbols. The first byte (i.e. the least-significant byte of the encapsulated construct) is mapped to B0 in the Tunneled Packet payload (see Figure 4-14). Bits [7:0] in each byte of the encapsulated construct are mapped to bits [7:0], respectively, in the corresponding byte of the Tunneled Packet payload. Section 9.1.1.3 shows how a USB3 Link Command is mapped to a Tunneled Packet.
- A USB3 Adapter shall not discard a Tunneled Packet due to lack of credits in the USB4 Fabric. When insufficient credits are available, the Router shall queue the Tunneled Packet and shall transport it once sufficient credits are available.



IMPLEMENTATION NOTE

The amount of buffering at the USB3 Adapter is implementation specific as it balances the tradeoff between USB3 Tunneling performance and USB3 link latency. It is recommended that implementations make the amount of buffers configurable.

The PDF field in a Tunneled Packet identifies the type of USB3 construct contained therein. Table 9-1 defines the PDF values that shall be used for each type of USB3 construct.

Table 9-1. PDF Values for USB3 Tunneling Packets

| PDF | Type | Payload of Tunneled Packet |
|---------|-------------------|---|
| 0h | LFPS | Indication for an LFPS sequence |
| 1h | Ordered Set | One TS1, TS2, or SDS Ordered Set |
| 2h | Link Command | One USB3 Link Command |
| 3h | LMP / TP / ITP | One LMP, TP, Deferred DPH, or ITP header packet |
| 4h | Start DP Segment | The first segment of a USB3 Data Packet (DPH and DPP) |
| 5h | Middle DP Segment | A segment of a USB3 Data Packet that is not the first segment and is not the last segment |
| 6h | End DP Segment | The last segment of a USB3 Data Packet that is broken into more than one segment |
| 7h - Fh | Rsvd | Reserved |

If a USB3 Adapter receives a Tunneled Packet with a PDF value other than 0 to 6, it shall discard the Tunneled Packet and shall not send any Packets in response.

9.1.1.1 LFPS Encapsulation


An LFPS Tunneled Packet shall contain a single Doubleword of payload as defined in Table 9-2 and shall have the structure shown in Figure 9-1. An LFPS event that is not listed in Table 9-2 shall not be transferred to the Transport Layer.

A USB3 Adapter Layer shall set no more than one of bits [22:17] in an LFPS Tunneled Packet.

Table 9-2. LFPS Tunneled Packet Payload

| Bits | Name | Description |
|-------|-----------------------|--|
| 7:0 | <i>Rsvd</i> | Reserved. |
| 15:8 | <i>LBPM</i> | Shall contain the LBPM PHY Capability byte. Only valid when the LBPM Enable bit is set to 1b. |
| 16 | <i>Rx Term Enable</i> | This bit represents the state of the local low-impedance receiver termination. It shall be set to 1b if the low-impedance receiver termination is on. It shall be set to 0b if the low-impedance receiver termination is off. In any LFPS packet, this bit shall represent the current value of the low-impedance receiver termination. |
| 17 | <i>SCD1</i> | Set to 1b following assertion of SCD1 by the internal USB3 device. Otherwise shall be set to 0b. |
| 18 | <i>SCD2</i> | Set to 1b following assertion of SCD2 by the internal USB3 device. Otherwise shall be set to 0b. |
| 19 | <i>U2 Exit</i> | Set to 1b following assertion of U2 Exit by the internal USB3 device. Otherwise shall be set to 0b. |
| 20 | <i>U3 Wakeup</i> | Set to 1b following assertion of U3 Wakeup by the internal USB3 device. Otherwise shall be set to 0b. |
| 21 | <i>Warm Reset</i> | Set to 1b following assertion of Warm Reset by the internal USB3 device. Otherwise shall be set to 0b. |
| 22 | <i>LBPM Enable</i> | Set to 1b following assertion of LBPM by the internal USB3 device. Otherwise shall be set to 0b. |
| 23 | <i>Rsvd</i> | Reserved. |
| 31:24 | <i>CRC</i> | Cyclic Redundancy Code covering bits [23:0] of the payload. It shall be calculated as defined in Section 9.1.1.1.1. |

Figure 9-1. LFPS Tunneled Packet Format



* SupplD

9.1.1.1.1 CRC

The *CRC* field shall cover bits [23:0] of payload. It shall be calculated as defined in Section 5.1.2.1.1, using the following bit order:

1. Bit 7 to bit 0
2. Bit 15 to bit 8
3. Bit 23 to bit 16

When a USB3 Adapter Layer receives an LFPS Tunneled Packet, it shall verify the *CRC* field value. The USB3 Adapter Layer may correct single-bit errors in the LFPS Tunneled Packet payload. After correcting an error, a Router shall continue on as if the error had never occurred. When an error is detected but not corrected, the packet with the error shall be dropped and no other action taken on its behalf.

See Appendix A for example CRC calculations.

9.1.1.1.2 Rx Term Enable

The USB3 Adapter Layer shall send an LFPS Tunneled Packet to the Transport Layer when the local value of Rx termination changes or when the *Path Enable* bit in the USB3 Adapter Configuration Capability changes from 0b to 1b. The packet shall be sent 3 times.

A Router shall implement a mechanism that allows a USB3 Adapter Layer to indicate far-end receiver termination to the internal USB3 device.

When a Router is not in Sleep state and is not in the process of transitioning into or out of Sleep state, a USB3 Adapter Layer shall indicate far-end receiver termination to the internal USB3 device when all of the following are true:

- Path Setup is complete (i.e. the *Path Enable* bit and the *Valid* bit in the USB3 Adapter Configuration Capability are both set to 1b).
- After Path Setup is complete, the USB3 Adapter Layer received an LFPS Tunneled Packet with the Rx Term Enable bit set to 1b and the USB3 Adapter Layer has not received any subsequent LFPS Tunneled Packets with the Rx Term Enable bit set to 0b.

When a Router transitions to Sleep state, a USB3 Adapter Layer shall maintain the same indicator value as before entry to Sleep state.

When the Router exits Sleep state, the USB3 Adapter Layer shall continue to maintain the indicator value until the *Valid* bit in USB3 Adapter Configuration Capability is set to 1b. After the *Valid* bit is set to 1b, it shall set the indicator as defined above.

9.1.1.1.3 SCD1

The USB3 Adapter Layer shall send an LFPS Tunneled Packet to the Transport Layer with the *SCD1* bit set to 1b when instructed to do so by the internal USB3 device. The packet shall be sent 3 times.

A USB3 Adapter Layer that receives an LFPS Tunneled Packet from the Transport Layer with the *SCD1* bit set to 1b shall indicate reception of SCD1 to the internal USB3 device until it receives a Tunneled Packet with different contents.

9.1.1.1.4 SCD2

The USB3 Adapter Layer shall send an LFPS Tunneled Packet to the Transport Layer with the *SCD2* bit set to 1b when instructed to do so by the internal USB3 device. The packet shall be sent 3 times.

A USB3 Adapter Layer that receives an LFPS Tunneled Packet from the Transport Layer with the *SCD2* bit set to 1b shall indicate reception of SCD2 to the internal USB3 device until it receives a Tunneled Packet with different contents.

9.1.1.1.5 U2 Exit

The USB3 Adapter Layer shall send an LFPS Tunneled Packet to the Transport Layer with the *U2 Exit* bit set to 1b when the internal USB3 device indicates a U2 Exit event. The packet shall be sent 3 times.

A USB3 Adapter Layer that receives an LFPS Tunneled Packet from the Transport Layer with the *U2 Exit* bit set to 1b shall indicate reception of a U2 Exit LFPS to the internal USB3 device until it receives a Tunneled Packet with different contents.

9.1.1.1.6 U3 Wakeup

When the internal USB3 device indicates the start of a U3 Wakeup event and a USB3 Path is enabled, a USB3 Adapter Layer shall send 3 LFPS Tunneled Packets to the Transport Layer. Each LFPS Tunneled Packet shall have the *U3 Wakeup* bit set to 1b.

If a USB3 Adapter Layer detects the beginning of a U3 Wakeup event before a USB3 Path is enabled, it shall ignore the USB3 Wakeup event and shall not send any LFPS Tunneled Packets for that event.

The USB3 Adapter Layer may send an LFPS Tunneled Packet with the *U3 Wakeup* bit set to 0b to the Transport Layer when the internal USB3 device starts to transmit SS signaling.

A USB3 Adapter Layer that receives an LFPS Tunneled Packet from the Transport Layer with the *U3 Wakeup* bit set to 1b shall indicate reception of a U3 Wakeup LFPS to the internal USB3 device until it receives a Tunneled Packet with different contents.

9.1.1.1.7 Warm Reset

When Warm Reset is asserted by the internal USB3 device, the Downstream USB3 Adapter Layer shall:

- Discard any queued Tunneled Packets.

Note: Only Tunneled Packets in the USB3 Adapter Layer are discarded. A Tunneled Packet that is passed to the Transport Layer before Warm Reset is asserted may still be transmitted.

- Send an LFPS Tunneled Packet to the Transport Layer with the *Warm Reset* bit set to 1b. The packet shall be sent 3 times.

While Warm Reset is active, a Downstream USB3 Adapter shall discard any Tunnel Packets it receives and clear any indication of LFPS reception towards the internal USB3 device.

An Upstream USB3 Adapter Layer that receives an LFPS Tunneled Packet with the *Warm Reset* bit set to 1b shall indicate assertion of a Warm Reset to the internal USB3 device until it receives a Tunneled Packet with different contents.

The Downstream USB3 Adapter Layer shall send an LFPS Tunneled Packet to the Transport Layer with the *Warm Reset* bit set to 0b when Warm Reset is de-asserted by the internal USB3 device. The packet shall be sent 3 times.

An Upstream USB3 Adapter Layer that receives an LFPS Tunneled Packet from the Transport Layer with the *Warm Reset* bit set to 0b shall indicate de-assertion of a Warm Reset to the internal USB3 device until it receives a Tunneled Packet with different contents.

After Warm Reset is de-asserted, a USB3 Adapter Layer shall not forward any Ordered Sets, packets, or events to the internal USB3 device that were received before or during Warm Reset assertion.

Note: The duration between the received Tunneled Packets with the Warm Reset bit set to 1b and the received Tunneled Packets with the Warm Reset bit set to 0b can be smaller than t_{Reset} as defined in the USB 3.2 Specification.

9.1.1.1.8 LBPM

The USB3 Adapter Layer shall send an LFPS Tunneled Packet to the Transport Layer with the *LBPM Enable* bit set to 1b when instructed to do so by the internal USB3 device. The *LBPM* field shall contain the USB3 LBPM construct. Bits[7:0] of the USB3 LBPM are mapped to Bits[7:0] of the *LBPM* field, respectively. The LFPS Tunneled Packet shall be sent 3 times.

A USB3 Adapter Layer that receives an LFPS Tunneled Packet with the *LBPM Enable* bit set to 1b shall do the following until it receives a Tunneled Packet with different contents:

- Indicate reception of LBPM to the internal USB3 device.
- Communicate the value in the *LBPM* field to the internal USB3 device.

9.1.1.1.9 LFPS Stop

A USB3 Adapter Layer shall send an LFPS Tunneled Packet to the Transport Layer with bits [22:17] equal 00h for the following cases:

- The Internal USB3 device indicates no LFPS Response Timeout for U3 Exit.
- The Internal USB3 device indicates Warm Reset done.

It is recommended that a USB3 Adapter Layer send an LFPS Tunneled Packet to the Transport Layer with bits [22:17] equal to 00h for the following cases:

- The Internal USB3 device indicates a Polling LFPS Timeout.
- The Internal USB3 device indicates no LFPS Response Timeout for U2 Exit.

The LFPS Tunneled Packet shall be sent 3 times.


9.1.1.2 Ordered Set Encapsulation

An Ordered Set Tunneled Packet shall have the single Doubleword payload defined in Table 9-3 and shall have the structure shown in Figure 9-2.

A USB3 Adapter Layer shall set one and no more than one of bits [2:0] in an Ordered Set Tunneled Packet.

Table 9-3. Ordered Set Tunneled Packet Payload

| Bits | Name | Description |
|-------|--------------------|--|
| 0 | TS1 | Set to 1b after a TS1 Ordered Set or a TSEQ Ordered Set is received from the internal USB3 device. Otherwise shall be set to 0b. |
| 1 | TS2 | Set to 1b after a TS2 Ordered Set is received from the internal USB3 device. Otherwise shall be set to 0b. |
| 2 | SDS | Set to 1b after a SDS Ordered Set is received from the internal USB3 device. Otherwise shall be set to 0b. |
| 7:3 | Rsvd | Reserved. |
| 15:8 | Link Functionality | TS1 or TS2 Ordered Sets – Shall contain the USB3 Link Functionality construct. Bits[7:0] of the USB3 Link Functionality are mapped to Bits[7:0] of the <i>Link Functionality</i> field, respectively. TSEQ Ordered Set – set to 00h by the USB3 Adapter Layer. SDS Ordered Set – set to 00h by the USB3 Adapter Layer. |
| 23:16 | Rsvd | Reserved. |
| 31:24 | CRC | Cyclic Redundancy Code covering bits[23:0] of the payload. It shall be calculated as defined in Section 9.1.1.1.1. |

Figure 9-2. Ordered Set Tunneled Packet Format

* SupplD

When a USB3 Adapter Layer receives a TSEQ, TS1, TS2, or SDS Ordered Set from the internal USB3 device that is not identical to the previous Ordered Set received from the internal USB3 device, it shall send 3 copies of an Ordered Set Tunneled Packet to the Transport Layer, with the respective fields set per Table 9-3.

Note: TSEQ Ordered Sets are used by a USB device for synchronization. They are not needed on USB4. However, a receiving USB device cannot be idle during the time when TSEQ are normally transmitted, so USB4 treats a TSEQ Ordered Set the same as a TS1 Ordered Set.

All other Ordered Sets received from the internal USB3 device shall be discarded by the USB3 Adapter Layer and shall not cause an Ordered Set Tunneled Packet to be sent to the Transport Layer.

A USB3 Adapter Layer shall only send Ordered Sets that target Lane 0 to the Transport Layer. A USB3 Adapter Layer shall discard Ordered Sets that target a non-zero Lane so that they are not sent over the USB4 Fabric.

When a USB3 Adapter Layer receives an Ordered Set Tunneled Packet, it shall verify the *CRC* field value. The USB3 Adapter Layer may correct single-bit errors in the Ordered Set Tunneled Packet payload. After correcting an error, a Router shall continue on as if the error had never occurred. When an error is detected but not corrected, the packet with the error shall be dropped and no other action taken on its behalf.

After verifying the *CRC* field value:

- If the *TS1* bit in the received Ordered Set Tunneled Packet is set to 1b, the USB3 Adapter Layer shall indicate reception of TS1 Ordered Set to the internal USB3 device.
- If the *TS2* bit in the received Ordered Set Tunneled Packet is set to 1b, the USB3 Adapter Layer shall indicate reception of TS2 Ordered Set to the internal USB3 device.
- If the *SDS* bit in the received Ordered Set Tunneled Packet is set to 1b, a USB3 Adapter Layer shall indicate reception of SDS Ordered Set to the internal USB3 device.
- Communicate the value in the *Link Functionality* field to the internal USB3 device.

9.1.1.3 Link Command Encapsulation

A USB3 Adapter Layer that receives a Link Command from the internal USB3 device shall send a Link Command Tunneled Packet to the Transport Layer with a 2-Doubleword payload. The payload shall contain the Link Command.

Upon receiving a Link Command Tunneled Packet from the Transport Layer, a USB3 Adapter Layer shall transfer the Link Command Tunneled Packet to the internal USB3 Port.

Figure 9-3. Link Command Tunneled Packet Format

| | | | | | |
|-----------------------------|---|-----------------------------|-----------------------------|-----------------------------|---|
| 31 | | | | | 0 |
| PDF = 2 | * | HOP ID | Length = 8 | HEC | |
| SLC | | SLC | SLC | EPF | |
| Link Command Word Byte 0 | | Link Command Word Byte 1 | Link Command Word Byte 0 | Link Command Word Byte 1 | |

* SupplID

9.1.1.4 Idle Symbols

A USB3 Adapter Layer shall drop any Idle Symbols received from the internal USB3 device.



IMPLEMENTATION NOTE

A USB3 Adapter Layer can send Idle symbols to an internal USB3 device as needed to maintain the expected behavior of the internal USB3 device. For example:

- *A USB3 Adapter Layer may send Idle Symbols towards the internal USB3 device while in U0 whenever no other Symbols are sent to the internal USB3 device.*
- *A USB3 Adapter Layer may send 16 Idle Symbols towards the internal USB3 device after it sends an SDS Ordered Set to the internal USB3 device.*

9.1.1.5 LMP Encapsulation

A USB3 Adapter Layer that receives a USB3 Link Management Packet (LMP) from the internal USB3 device shall send an LMP Tunneled Packet to the Transport Layer with a 5 Doubleword payload. The payload shall contain the framed (see Section 7.2.1.1.1 in the USB 3.2 Specification) LMP.

Upon receiving an LMP Tunneled Packet from the Transport Layer, a USB3 Adapter Layer shall transfer the LMP to the internal USB3 device.

9.1.1.6 TP Encapsulation

A USB3 Adapter Layer that receives a USB3 Transaction Packet (TP) from the internal USB3 device shall send a TP Tunneled Packet to the Transport Layer with a 5 Doubleword payload. The payload shall contain the framed (see Section 7.2.1.1.1 in the USB 3.2 Specification) TP.

If the USB3 Adapter Layer, receives a Deferred DPH from the internal USB3 device, it shall send a TP Tunneled Packet to the Transport Layer with a 5 Doubleword payload. The payload shall contain the Deferred DPH and its USB3 framing symbols.


Upon receiving a TP Tunneled Packet from the Transport Layer, a USB3 Adapter Layer shall transfer the TP to the internal USB3 device.

9.1.1.7 ITP Encapsulation

A Downstream USB3 Adapter Layer that receives an Isochronous Timestamp Packet (ITP) from the internal USB3 device shall:

1. Set the *TMU Time Capture* field to the nanosecond portion of the current Host Router Time (t_g).
2. Send an ITP Tunneled Packet with the format defined in Figure 9-4 to the Transport Layer. The payload contains the 5 Doublewords of a framed (see Section 7.2.1.1.1 in the USB 3.2 Specification) ITP, followed by 2 Doublewords of *TMU Time Capture* field.

Figure 9-4. Tunneled ITP Packet Format



* SupplID

An Upstream USB3 Adapter Layer that receives an ITP Tunneled Packet from the Transport Layer shall:

1. Update the *Delta* field of the ITP according to the formula below:
 - Updated $\Delta = \Delta + (\text{Host Router time Nanosecond} - \text{TMU Time Capture}) / \text{tIsochTimeStampGranularity}$.

Note: tIsochTimeStampGranularity is defined in the USB 3.2 Specification.

 - A USB3 Adapter Layer shall meet the required accuracy of the *Delta* field in the ITP as defined by Section 8.7 of the USB 3.2 Specification.
2. Update the *CRC-16* field in the ITP.

3. Forward the ITP to the internal USB3 device.

If the *Time Disruption* bit is set to 1b in Router Configuration Space, an Upstream USB3 Adapter Layer shall set the *Delayed* bit in the Link Control Word to 1b.

If the *Time Disruption* bit is set to 0b, an Upstream USB3 Adapter Layer may set the *Delayed* bit to 1b. The decision when to set the *Delayed* bit is implementation specific.

A USB3 Adapter Layer shall recalculate the *CRC5* field within the Link Control Word if it modified the *Delayed* bit.




IMPLEMENTATION NOTE

It is recommended that the TMU in a Router sample time as close as possible to the internal USB3 device.

9.1.1.8 Data Packet (DP) Encapsulation

A USB3 Adapter Layer segments a USB3 Data Packet and its USB3 framing symbols into one or more Tunneled Packets for transport over the USB4 Fabric. Before segmentation, the USB3 Data Packet includes the Data Packet Header (DPH) and the Data Packet Payload (DPP) as depicted in Figure 9-5.

Figure 9-5. Structure of an Unsegmented USB3 Data Packet




If the size of a USB3 Data Packet is 252 bytes or less, then the USB3 Adapter Layer shall send a single Tunneled Packet of type Start DP Segment to the Transport Layer. The Tunneled Packet shall carry the unsegmented USB3 Data Packet as payload and any padding needed for DW alignment.

If the size of a USB3 Data Packet is larger than 252 bytes, the USB3 Data Packet shall be segmented into multiple Tunneled Packets as shown in Figure 9-6. When a USB3 Adapter Layer segments a USB3 Data Packet into Multiple Tunneled Packets, it shall follow the rules below:

- All Tunneled Packets with the possible exception of the last packet shall include 252 bytes of payload.
- The first Tunneled Packet shall be of type *Start DP Segment*.

- Any following Tunneled Packets other than the last Tunneled Packet shall be of type *Middle DP Segment*.
- The last Tunneled Packet shall:
 - Be of type *End DP Segment*.
 - Pad the payload to be Doubleword aligned.
 - Payload size shall not exceed 252 bytes.
- The Tunneled Packets shall be sent to the Transport layer such that the byte ordering of the original USB3 Data Packet is maintained.

Figure 9-6. Segmentation of a USB3 Data Packet



When re-assembling Tunneled Packets into USB3 Data Packets, the USB3 Adapter Layer shall:

- Drop a Tunneled Packet of type *Middle DP Segment* that comes immediately after a Tunneled Packet of type *End DP Segment*.
- Drop a Tunneled Packet of type *End DP Segment* that comes immediately after a Tunneled Packet of type *End DP Segment*.

- Nullify or partially nullify a USB3 Data Packet when both of the following are true:
 - A Tunneled packet of type *Start DP Segment* is received and the length information in the USB3 DPH indicates it does not fit into a single Tunneled Packet.
 - A Tunneled Packet of type *End DP Segment* is not received within *tReassemble* from the reception of the Tunneled Packet of type *Start DP Segment*.

9.1.2 Bandwidth Negotiation

A Host Router uses the *Consumed Upstream Bandwidth*, *Consumed Downstream Bandwidth*, *Allocated Upstream Bandwidth*, and *Allocated Downstream Bandwidth* fields in Adapter Configuration Space to negotiate bandwidth between a Connection Manager and the internal host controller. The Connection Manager sets the *Allocated Upstream Bandwidth* and *Allocated Downstream Bandwidth* fields. The internal host controller sets the *Consumed Upstream Bandwidth* and *Consumed Downstream Bandwidth* fields.

Note: All bandwidth values are USB4 Link bandwidth.

To prevent a race condition between the Connection Manager and the internal host controller, the Host Router implements a locking mechanism that uses the Connection Manager Request and *Host Controller Ack* fields as follows:

- When the *Host Controller Ack* field = 0 (The default value):
 - The internal host controller may read the *Allocated Upstream Bandwidth* and/or *Allocated Downstream Bandwidth* fields.
 - The internal host controller may update the *Consumed Upstream Bandwidth* and/or *Consumed Downstream Bandwidth* fields.
 - The value in the *Consumed Upstream Bandwidth* field shall not exceed the value in the *Allocated Upstream Bandwidth* field.
 - The value in the *Consumed Downstream Bandwidth* field shall not exceed the value in the *Allocated Downstream Bandwidth* field.
- When the *Host Controller Ack* field = 1:
 - The internal host controller shall not read the *Allocated Upstream Bandwidth* or *Allocated Downstream Bandwidth* fields.
 - The internal host controller shall not update *Consumed Upstream Bandwidth* or *Consumed Downstream Bandwidth* fields.


The internal host controller shall negotiate bandwidth as shown in Figure 9-7.

When the USB3 link reaches U0, a USB3 Adapter Layer shall:

- Set the *Actual Link Rate* field to indicate the established link rate.
- Set the *USB3 Link Valid* field to 1b.

When the USB3 link is not in U0, Recovery, U2 or U3, a USB3 Adapter Layer shall set the *USB3 Link Valid* field to 0b.

Figure 9-7. Bandwidth Negotiation by the Internal Host Controller



Request = Connection Manager Request bit in the USB3 Adapter Configuration Capability

Ack = Host Controller Ack bit in the USB3 Adapter Configuration Capability




CONNECTION MANAGER NOTE

A Connection Manager uses the Connection Manager Request and Host Controller Ack fields as follows:

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

The Connection Manager negotiates bandwidth as shown in Figure 9-8.

Figure 9-8. Bandwidth Negotiation by the Connection Manager



Request = Connection Manager Request bit in the USB3 Adapter Configuration Capability

Ack = Host Controller Ack bit in the USB3 Adapter Configuration Capability

9.1.3 Timing Parameters

Table 9-4 lists the timing parameters for a USB3 Adapter Layer.

Table 9-4. USB3 Adapter Timing Parameters

| Parameter | Description | Min | Max | Units |
|-------------|--|-----|-----|-------|
| tReassemble | Time between Tunneled Packets which is used to determine loss of a USB3 Data Packet segment. | 100 | | μs |

9.2 Internal USB3 Device

This section defines the functionality of the internal USB3 device ports that interface with a USB3 Adapter.

A USB3 Physical Layer is not needed in a USB3 port that interfaces with a USB3 Adapter. Therefore, Physical Layer scrambling is not performed, regardless of the value of the *Disable Scrambling* bit in a received TS Ordered Set. SKIP insertion is also not performed.



IMPLEMENTATION NOTE

The internal USB3 device in a USB4 device/hub must preserve standard USB 3.2/USB 2.0 functionality for the cases where the USB4 device/hub is connected to a host port that supports only USB 3.2 functionality. As such, it follows the USB 3.2 specification

requirement to fall back to USB 2.0 mode when unable to detect SuperSpeed terminations after 8 consecutive far-end receiver termination detection attempts.

For a USB4 Link, the USB3 Adapter on the USB4 device/hub does not expose terminations until after the USB3 tunnel is established. An implementation must ensure that its integrated USB3 device does not erroneously fall back to USB2 mode due to the device not detecting far-end receiver terminations after 8 attempts. This can be done through various implementation dependent methods.

9.2.1 Link Layer

An internal USB3 device port that interfaces to a USB3 Adapter Layer shall implement a Link Layer as defined in the USB 3.2 Specification with the modifications, configurations, and parameters described in this section.

The Link Layer shall support Gen 2 Single-Lane (2x1) and may support Gen 2 Dual-Lane (2x2). No other Link capabilities shall be supported.

An Internal USB3 device shall not send a LBPM PHY Capability with a value of 0 in bits [3:2].

Note: The USB3 LBPM handshake between the two internal USB3 devices ports determines the link that will be established.

9.2.1.1 Link Training and Status State Machine (LTSSM)

The Link Layer shall implement a Link Training and Status State Machine (LTSSM) with the following adjustments:

- Loopback state shall not be supported.
- Compliance Mode state shall not be supported.
- U1 state shall not be supported.
 - U1 LGO shall not be sent.
- When the LTSSM of an Upstream Facing Port is in Disabled state, it shall unconditionally transition to the Rx.Detect state.
- TS1 and TS2 shall not be scrambled.

Note: The actual number of generated Ordered Sets and LFPS is not communicated between the two sides of the link.

9.2.1.2 Timers and Timeouts

Table 9-5 lists the timing parameters that have different values than those defined in the USB 3.2 Specification. An internal USB3 device shall use the parameter values in Table 9-5.

Table 9-5. USB3 Timers and Timeout Values

| Parameter | Value |
|------------------------|------------|
| tPollingSCDLFPSTimeout | 60 μ s |
| tPortConfiguration | 40 μ s |
| PENDING_HP_TIMER | 20 μ s |
| PM_LC_TIMER | 20 μ s |
| PM_ENTRY_TIMER | 16 μ s |



IMPLEMENTATION NOTE

It is recommended that implementations make the values in Table 9-5 configurable.

9.2.2 USB3 Protocol Layer

The internal USB3 device shall implement a Protocol Layer as defined in the USB 3.2 Specification with the following adjustments:

- LDM Protocol shall not be supported over ports that interfaces a USB3 Adapter.

9.2.3 Descriptors

The internal USB3 device shall implement Descriptors as defined in the USB 3.2 Specification with the following adjustments:

- If a Router supports entering CL1 or CL2 when the USB3 link between the USB3 Adapter and the internal USB3 device is in U2 state, then the value for wU2DevExitLat shall include the maximum sum of the Router enter and exit time from the supported CLx.

9.3 Paths

A USB3 Adapter Layer shall put HopID=8 in the header of an outgoing Tunneled Packet before handing it off to the Transport Layer for routing.

9.3.1 Path Setup



CONNECTION MANAGER NOTE

To set up a USB3 Path, a Connection Manager shall configure the USB3 Path, then set both the Path Enable bit and the Valid bit in the USB3 Adapter Configuration Capability to 1b. A Connection Manager shall configure a USB3 Path with the Dedicated Flow Control Buffer Allocation scheme.

When the *Path Enable* bit and the *Valid* bit in the USB3 Adapter Configuration Capability are both set to 1b, the USB3 Adapter Layer may:

- Indicate far-end receiver termination to the internal USB3 device as defined in Section 9.1.1.1.2.
- Issue USB3 Tunneled Packets to the Transport Layer as defined in this chapter.



CONNECTION MANAGER NOTE

When tearing down an old USB3 Path before setting up a new USB3 Path, a Connection Manager shall wait at least 500 ms after setting the Path Enable bit to 0b before setting up the new USB3 Path. This is to allow enough time for the internal USB3 device port to reach the Detect state.

9.3.2 Path Teardown



CONNECTION MANAGER NOTE

Before tearing-down a USB3 Path, a Connection Manager needs to disable the Path by setting the Path Enable bit in the USB3 Adapter Configuration Capability to 0b and the Valid bit in the USB3 Adapter Configuration Capability to 1b.

When a Router detects a disconnect on a Downstream Facing Port, it shall set the *Path Enable* bit to 0b in the USB3 Adapter Configuration Capability of the corresponding USB3 Adapter. The corresponding USB3 Adapter is the USB3 Adapter that is associated with the Lane Adapter in the Router's Routing Tables.

When the *Path Enable* bit is set to 0b and the *Valid* bit is set to 1b in the USB3 Adapter Configuration Capability:


- The USB3 Adapter Layer shall:
 - Not issue any Tunneled Packets to the Transport Layer.
 - Remove far-end receiver termination to the internal USB3 device as defined in Section 9.1.1.1.2.
- The internal USB3 device port shall detect a disconnect within 500 ms.
- An integrated Enhanced SuperSpeed Hub within the Device Router shall ensure that SuperSpeed or Enhanced SuperSpeed devices on its downstream-facing ports transition to default state. This can be achieved by either issuing a Warm Reset to the devices or initiating a disconnect/reconnect (either by cycling power or removing terminations).



CONNECTION MANAGER NOTE

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

10 DisplayPort™ Tunneling



This chapter describes DP IN and DP OUT Adapters, which tunnel DP traffic over the USB4™ Fabric. This chapter only applies to USB4 products that implement DisplayPort Source or Sink.

A USB4 host shall support DP tunneling. A Host Router shall contain at least one DP IN Adapter and may optionally contain one or more DP OUT Adapters.

A USB4 hub shall support DP Tunneling. A USB4 Hub shall contain at least one DP OUT Adapter and may optionally contain one or more DP IN Adapters.


Note: The DP OUT Adapter in a USB4 hub is to enable a DisplayPort Sink connected on a DFP (or non-USB display connector if USB4 hub is a USB4-Based Dock), See the USB Type-C Specification for more details.

A USB4 peripheral device may optionally support DP Tunneling. If a USB4 peripheral device supports DP Tunneling, it shall contain at least one DP Adapter.

10.1 DP Adapter Protocol Stack

Figure 10-1 shows the protocol stack layers for a DP Adapter.

Figure 10-1. DP Adapter Protocol Stack Layers



10.1.1 Transport Layer

A DP IN Adapter and a DP OUT Adapter both implement a Transport Layer as described in Chapter 5.

10.1.2 Protocol Adapter Layer

The Protocol Adapter Layer in a DP IN Adapter encapsulates Main-Link data and AUX channel transactions into Tunneled Packets for transmission over the USB4 Fabric. It also reconstructs AUX channel transactions from Tunneled Packets received over the USB4 Fabric.

The Protocol Adapter Layer in a DP OUT Adapter reconstructs Main-Link data and AUX channel transactions from Tunneled Packets received over the USB4 Fabric. It also encapsulates AUX channel transactions into Tunneled Packets for transmission over the USB4 Fabric.

10.1.3 DP Physical Layer

A DP Adapter shall either implement the DisplayPort Physical Layer as defined in the DisplayPort 1.4a Specification or shall implement its functional equivalent (e.g. DP Adapter is connected to a DPRX or a DPTX as part of an SoC).


A DP IN Adapter which does not implement a Physical Layer shall generate a stream of DisplayPort Tunneled packets as if a Physical Layer exists. This is to ensure that a DP OUT Adapter can properly reconstruct the DisplayPort Physical stream.

10.2 DP Adapter States

A DP IN Adapter and a DP OUT Adapter follow the state machine shown in Figure 10-2.

Note: A DP IN Adapter uses upstream device detection as defined in the DisplayPort 1.4a Specification to determine when a Plug, Replug or Unplug event occurs. A DP OUT Adapter uses the hot plug detection method defined in the DisplayPort 1.4a Specification to determine when a Plug, Replug or Unplug event occurs.

Figure 10-2. DP Adapter State Machine



10.2.1 Reset

A DP Adapter enters this state when the Router with the DP Adapter enters the Uninitialized state.

While in the Reset state, a DP Adapter shall set all Configuration Spaces to their default values. A DP IN Adapter shall drive HPD signal low. A DP OUT Adapter shall not apply DP_PWR.

A DP Adapter exits this state when the Router with the DP Adapter transitions to the Enumerated state.

10.2.2 Present

A DP Adapter enters this state upon any of the following conditions:

- The DP Adapter is in the Reset state and the Router with the DP Adapter is in the Enumerated state.
- The DP Adapter is in the Plugged state and an Unplug event occurs as defined in the DisplayPort 1.4a Specification.
- The DP Adapter is in the Paired state and both of the following are true:
 - Either the AUX Path is disabled (AE bit is 0b) or the Video Path is disabled (VE bit is 0b).
 - The DP Adapter is unplugged as defined in the DisplayPort 1.4a Specification.

While in the Present state, a DP Adapter shall set all Configuration Spaces to their default values. A DP IN Adapter shall drive HPD signal low. A DP OUT Adapter shall apply DP_PWR.

A DP Adapter exits this state upon a Plug or Replug event. Plug and Replug events are defined in the DisplayPort 1.4a Specification.

10.2.3 Plugged

A DP Adapter enters this state upon any of the following conditions:

- A Plug or Replug event occurs as defined in the DisplayPort 1.4a Specification.
- The AUX Path is disabled (*AE* bit is 0b), the Video Path is disabled (*VE* bit is 0b), and the DP Adapter is plugged as defined in the DisplayPort 1.4a Specification.

A DP Adapter exits this state upon any of the following conditions:

- An Unplug event as defined in the DisplayPort 1.4a Specification.
- Both the AUX Path and Video Path are enabled (*AE* bit is 1b and *VE* bit is 1b).

While in the Plugged state, a DP IN Adapter shall drive HPD signal low.

10.2.4 Paired

A DP Adapter enters this state when both the AUX Path and Video Path are enabled (*AE* bit is 1b and *VE* bit is 1b).

While in this state, a DP Adapter may send and receive Tunneled Packets over the USB4 Fabric.

- A DP Adapter exits this state after the AUX Path is disabled (*AE* bit is 0b) or the Video Path is disabled (*VE* bit is 0b).

10.3 Interfaces

10.3.1 DisplayPort

A DP Adapter shall support three modes of operation:

- LTTPR Non-Transparent – LT-tunable PHY Repeater (Non-Transparent Mode).
- LTTPR Transparent – LT-tunable PHY Repeater (Transparent Mode).
- Non-LTTPR – Non-LT-tunable PHY Repeater.

DisplayPort Link Training is executed over the tunnel in one of two ways:

- DPTX Managed, Sequential Link Training – In this mode the DPTX is aware and manage the Link Training of the two DisplayPort Links, #1, the link between DPTX and DP In Adapter and #2, the link between DP Out Adapter and DPRX. The Link Training is done sequentially the upstream link (#1) first and the downstream link (#2) second. This mode is used when operating in LTTPR Non-Transparent mode. See Section 10.4.10.1.
- Autonomous, Concurrent Link Training – In this mode the DPTX does not manage the Link Training between DP Out Adapter and DPRX; rather, it manages only the Link Training to the DP IN Adapter. In parallel to the training between the DPTX and DP In Adapter, DP OUT Adapter manages the link training to DPRX. DP In Adapter indicates the upstream Link Training completion only after it is notified of the downstream Link Training completion by DP Out Adapter, ensuring the end-to-end Link Training completion by the time DPTX exits Link Training sequence. This way of Link Training is used when operating in Non-LTTPR and LTTPR Transparent modes. See Section 10.4.10.2.

Aux Transaction handling has two major aspects to attend when executed over the tunnel:

- AUX Transaction initiation operates in two ways:
 - DPTX Only – In this mode only DPTX initiates AUX Transactions. This mode is used when operating in LTTTPR Non-Transparent mode.
 - DPTX and DP OUT Adapter – In this mode both, the DPTX and the DP OUT Adapter, are initiating AUX Transactions. A DP OUT Adapter mainly initiate AUX Transactions as part of the autonomous and concurrent Link Training Process. This mode is used when operating in Non-LTTTPR and LTTTPR Transparent modes, see Section 10.4.4.2.3 and Section 10.4.4.3.3 respectively.
- Timeout and DEFER usage:
 - Allowed – A DP IN Adapter uses timeout to govern the AUX Link and uses AUX DEFER when timeout expires. This mode is used when operating in Non-LTTTPR mode, see Section 10.4.4.2.1.
 - Disallowed – A DP IN Adapter does not activate a timeout for AUX Responses, will not send DEFER and will not gate AUX Responses coming from DPRX. This mode is used when operating in LTTTPR Non-Transparent and LTTTPR Transparent modes, see Section 10.4.4.3.1.

Table 10-1 summarizes the above modes of operations.

Table 10-1. DisplayPort Modes Of Operation Over DisplayPort Tunneling

| Mode | Link Training | AUX Transaction Initiation | AUX DEFER Usage |
|------------------------|--------------------------|----------------------------|-----------------|
| Non-LTTTPR | Autonomous, Concurrent | DPTX and DP OUT Adapter | Allowed |
| LTTTPR Transparent | Autonomous, Concurrent | DPTX and DP OUT Adapter | Disallowed |
| LTTTPR Non-Transparent | DPTX Managed, Sequential | DPTX Only | Disallowed |

After reset, a DP Adapter shall operate in Non-LTTTPR mode. A DP Adapter shall transition between the three modes as described in Section 10.4.6.1. A DP Adapter shall transition to Non-LTTTPR mode upon exit from the Paired state.

Unless otherwise stated, the requirements in this chapter apply to DP Adapters in all three modes.

10.3.1.1 LTTTPR Non-Transparent

A DP IN Adapter shall implement LTTTPR UFP. A DP OUT Adapter shall implement LTTTPR DFP. LTTTPR UFP and LTTTPR DFP are defined in the DisplayPort 1.4a Specification.

10.3.1.2 Non-LTTTPR

A DP IN Adapter shall implement Non-LTTTPR Non-Transparent UFP. A DP OUT Adapter shall implement Non-LTTTPR Non-Transparent DFP. In Non-LTTTPR mode, unlike LTTTPR mode, the DPTX is unaware of the existence of additional DP Links. From the DPTX point of view, each operation appears to be carried out as if the DPRX is directly attached to the DPTX. Non-LTTTPR behavior is defined for AUX Handling in Section 10.4.4.2, Link Training in Section 10.4.10.2, and Connection Manager Discovery in Section 10.3.4.

10.3.1.3 LTTTPR Transparent

Operating in LTTTPR Transparent mode (i.e. handling AUX transactions and performing Link Training) is the same as operating in Non-LTTTPR mode with the following exceptions:

- The timeouts and timers are different
- The DP IN Adapter does not send AUX DEFER as is the case with LTTTPR Non-Transparent mode and does not gate AUX Transactions (see Section 10.4.4.2.2) except during autonomous, concurrent Link Training operation

10.3.2 Programming Model


10.3.2.1 Adapter Configuration Space

A DP Adapter implements an Adapter Configuration Space as defined in Section 8.2.2.

10.3.2.2 Path Configuration Space

A DP Adapter shall implement a Path Configuration Space as defined in Section 8.2.3. As shown in Figure 10-3, a DP Adapter shall support one MAIN-Link Path, one AUX Ingress Path, and one AUX Egress Path. The AUX Ingress Path of the DP IN Adapter corresponds to the AUX Egress Path of the DP OUT Adapter and the AUX Egress Path of the DP IN Adapter corresponds to the AUX Ingress Path of the DP OUT Adapter.

Figure 10-3. DP Adapter Path Directions



CONNECTION MANAGER NOTE

It is recommended that a Connection Manager use the parameters in Table 10-2 for a DP Path.

Table 10-2. Recommended Path Parameters

| | MAIN-Link | | | AUX IN to OUT | | | AUX OUT to IN | | |
|----------------|------------|-------------------------------|-------------|---------------|--------------|-------------|---------------|--------------|------------|
| | IN to USB4 | USB4 to USB4 | USB4 to OUT | IN to USB4 | USB4 to USB4 | USB4 to OUT | OUT to USB4 | USB4 to USB4 | USB4 to IN |
| IFC | NA | 0 | 0 | NA | 1 | 1 | NA | 1 | 1 |
| EFC | 0 | 0 | NA | 1 | 1 | NA | 1 | 1 | NA |
| ISE | NA | 0 | 0 | NA | 0 | 0 | NA | 0 | 0 |
| ESE | 0 | 0 | NA | 0 | 0 | NA | 0 | 0 | NA |
| Credits | NA | 12 (Non-Flow Control buffers) | | NA | 1 | 1 | NA | 1 | 1 |

10.3.3 Hot Plug and Hot Removal Events

10.3.3.1 DP OUT Adapters

When a DP OUT Adapter detects a Plug Event (as defined in the DisplayPort 1.4a Specification), it shall:

- Send a Hot Plug Event Packet with the *UPG* bit set to 0b as described in Section 6.8 within tDPPlug.
- Set the *Plugged* bit to 1b.

When a DP OUT Adapter detects an Unplug Event (as defined in the DisplayPort 1.4a Specification), it shall:

- Send a Hot Plug Event Packet with the *UPG* bit set to 1b as described in Section 6.8 within tDPPlug.
- Set the *Plugged* bit to 0b.

10.3.3.2 DP IN Adapters

A Router shall send a Hot Plug Event Packet as described in Section 6.8 within tDPPlug of when both of the following are true:

- A DP IN Adapter detects a Source (as defined in the DisplayPort 1.4a Specification).
- The DP IN Adapter that detected the Plug Event has sufficient DP stream resources available to support a DP stream.

A Router shall send a Hot Plug Event Packet with the *UPG* bit set to 1b as described in Section 6.8 within tDPPlug of when either of the following are true:


- A DP IN Adapter detects the removal of a Source (as defined in the DisplayPort 1.4a Specification).
- The Router has freed the DP stream resources allocated to the DP IN Adapter such that the DP IN Adapter can no longer support a DP stream.

DP stream resources allocation are managed by the Connection Manager using the DP Stream Resource Allocation Commands defined below in Section 10.3.3.2.1.

10.3.3.2.1 DP Stream Resource Allocation


Figure 10-4 shows three examples of how a DP stream resource can map to a DP IN Adapter within a Router.

Figure 10-4. DP Stream Resource Mapping Examples



Shared DP IN Physical layer

Direct Connection



SoC Solution

A Router shall support the DP Stream Resource Commands listed in Table 10-3.

Table 10-3. DP Stream Resource Allocation Commands

| Name | Purpose | Reference |
|------------------------|---|-------------------|
| QUERY_DP_RESOURCE | Queries a Router to determine whether or not a DP IN Adapter has sufficient DP stream resource availability for DP Tunneling. | Section 8.3.1.1.1 |
| ALLOCATE_DP_RESOURCE | Causes a Router to allocate a DP stream resource to a DP IN Adapter. | Section 8.3.1.1.2 |
| DEALLOCATE_DP_RESOURCE | Causes a Router to free a previously allocated DP stream resource. | Section 8.3.1.1.3 |

10.3.4 DisplayPort Over USB4 Fabric**10.3.4.1 DisplayPort Data Packet Types**

A DP Adapter encapsulates DisplayPort traffic into the Tunneled Packet types defined in this section. The Tunneled Packet types for DP Adapters are defined in Table 10-4 and Table 10-5. The Tunneled Packets types defined in Table 10-4 shall only be used for the AUX Path. The Tunneled Packet Types defined in Table 10-5 shall only be used for the Main-Link Path.

Table 10-4. AUX Path Tunneled Packet Types

| PDF | Type | Reference |
|-------|-------------------|--------------------|
| 0h | AUX Packet | Section 10.3.4.2.1 |
| 1h | HPD Status Packet | Section 10.3.4.2.2 |
| 2h | SET_CONFIG Packet | Section 10.3.4.2.3 |
| 3h | ACK Packet | Section 10.3.4.2.4 |
| 4h-Fh | Reserved | |

Table 10-5. Main-Link Path Tunneled Packet Types

| PDF | Type | Reference |
|-------|----------------------------------|------------------|
| 0h | Reserved | |
| 1h | SST Video Data Packet | Section 10.5.1.1 |
| 2h | SST Blank Start Packet | Section 10.5.1.3 |
| 3h | SST Main Stream Attribute Packet | Section 10.5.1.2 |
| 4h | SST Secondary Stream Packet | Section 10.5.1.4 |
| 5h | DP Clock Sync Packet | Section 10.6.1.3 |
| 6h | Multi Stream Packet | Section 10.5.2.3 |
| 7h | FEC Decode Packet | Section 10.5.3 |
| 8h-Fh | Reserved | |

If a DP Adapter receives a Tunneled Packet on the AUX Path with a PDF value other than 0 to 3, it shall discard the Tunneled Packet and shall not send any Packets in response.

If a DP OUT Adapter receives a Tunneled Packet on the Main-Link Path with a PDF value other than 1 to 7, it shall discard the Tunneled Packet and shall not send any Packets in response.

10.3.4.2 AUX Path Packet


A DP Adapter Layer uses the HopID values in the *AUX Tx HopID* and *AUX Rx HopID* and fields in the Protocol Adapter Configuration Capability to identify AUX Path Packets. When generating an AUX Path Packet, a DP Adapter Layer shall put the value in the *AUX Tx HopID* field into the *HopID* field of the Tunneled Packet header. When a DP Adapter Layer receives a Tunneled Packet with a HopID that is equal to the *AUX Rx HopID* field, it shall treat that packet as an AUX Path Packet.

10.3.4.2.1 AUX Packet

The DisplayPort AUX channel is a half-duplex, bi-directional channel. It is used to carry configuration messages. Configuration messages on the AUX channel use a request/response transaction format. The DisplayPort 1.4a Specification requires each AUX channel transaction to be framed by <SYNC> and <STOP> symbols. Figure 10-5 shows the cases of the minimum and maximum sized AUX transactions.

Figure 10-5. AUX Channel Framing

A) Minimum Size Aux Transaction




B) Maximum Size Aux Transaction




AUX Packets are used to forward AUX channel request and reply commands across the USB4 Fabric. Both command types, I2C-Over-AUX Transaction and DisplayPort Transaction, are supported. An AUX Packet shall have the format shown in Table 10-7.

Figure 10-6. AUX Packet Format

A) Minimum Size Aux Packet



B) Maximum Size Aux Packet



* SuppID

The Tunneled Packet Header for an AUX Packet shall have the *PDF* field set to 0.

AUX Packet payload shall contain the following:

- **CRC:** See Section 10.3.4.2.1.1.
- **AUX Payload:** Shall contain the bytes contained between the <SYNC> and <STOP> framing bytes of a DisplayPort AUX transaction. The number of bytes in this field varies between 1 and 20.
- **Reserved:** Shall be one byte set to 00h.

10.3.4.2.1.1 CRC

The CRC32 computation in an AUX Packet shall be based on the following CRC:

- Width: 32
- Poly: 1EDC 6F41h
- Init: FFFF FFFFh
- RefIn: True
- RefOut: True
- XorOut: FFFF FFFFh

The CRC32 protects the AUX Payload. The CRC32 is computed over 21 bytes regardless of the size of the AUX payload. If the AUX Payload is less than 21 bytes in length, a DP Adapter shall add the required number of zero-padding bytes for the computation of the CRC. The padding bytes shall not be transmitted in the AUX Packet. The CRC32 shall be generated by the DP Adapter that creates the AUX Packet and shall be checked by the DP Adapter that receives the AUX Packet. A DP Adapter that receives an AUX Packet with a CRC error shall drop that packet.

Figure 10-7 shows an example of an AUX Packet with CRC32 calculated over 6 bytes of actual data.

Figure 10-7. AUX Packet Example


AUX ACK - 4 Bytes Read Respond

| | | | | |
|---------|---|----------------|---------------|-----|
| 31 | 0 | | | |
| PDF = 0 | * | HOP ID | Length = 10 | HEC |
| CAh | | 5Ch | 3Bh | C4h |
| 00h | | 14h | 0Ah | C1h |
| 80h | | Reserved = 00h | Padding Bytes | |

* SupplID

10.3.4.2.2 HPD Packets

This packet is used by a DP OUT Adapter to notify a DP IN Adapter of a Plug, Re-Plug, or Unplug HPD event. An HPD Packet shall have the format shown in Figure 10-8.

Figure 10-8. HPD Packet Format

* SupplID

The *PDF* field in the header shall be set to 1 and the *Length* field shall be 4.

HPD Packet payload shall contain the following:

- **ECC [7:0]**: Error correction field that is calculated over bits [31:8] of the HPD Packet payload. See Section 5.1.2.3 for calculation.
- **Reserved [30:8]**: Shall be set to 0.
- **Plug (P) Flag [Bit 31]**: Shall be set to 0 if the HPD signal is low for more than 2 ms. Shall be set to 1 if the HPD signal is high.

A DP OUT Adapter sends an HPD Packet when it detects a Plug, Replug, or Unplug event (see Section 10.4.3). A DP OUT Adapter may also send an HPD Packet to the DP IN Adapter periodically.


When a DP IN Adapter receives an HPD Packet, it shall check the *ECC* field of the packet payload. The DP IN Adapter shall correct single-bit errors in the HPD Tunneled Packet payload. If an uncorrectable error is detected, the HPD Packet shall be dropped. Otherwise the DP IN Adapter shall:

- Generate a Plug/Re-plug HPD event if the *P Flag* in the HPD Packet payload is set to 1b.
- Generate an Unplug HPD event if the *P Flag* in the HPD Packet payload is set to 0b.
- Acknowledge the HPD Packet by sending an ACK Packet to the DP OUT Adapter within *tDPAckResponse* of receiving the HPD Packet.

If a DP OUT Adapter does not receive an ACK Packet with the *Type* field set equal to 8h within *tDPAckTimeout* of sending an HPD Packet, it may perform Link-Init as described in Section 10.4.13.

10.3.4.2.3 SET_CONFIG Packet

This packet is used to transport configuration commands between two DP Adapters. A SET_CONFIG Packet shall have the format shown in Figure 10-9. The different SET_CONFIG Packet MSG types are listed in Table 10-6.

Figure 10-9. SET_CONFIG Packet Format

* SupplID

The *PDF* field in the header shall be set to 2 and the *Length* field shall be 4.

SET_CONFIG Packet payload shall contain the following:

- **ECC [7:0]**: Error correction field that is calculated over bits [31:8] of the SET_CONFIG Packet payload. See Section 5.1.2.3 for calculation.
- **Link Rate 0 (LR0) [8]**: This field is used in combination with the *Link Rate 1 (LR1)* field as defined below.
- **Lane Count (LC) [11:9]**: This field shall specify the selected lane count according to the following encodings:
 - 000b: Link Down
 - 001b: 1 Lane
 - 010b: 2 Lanes
 - 100b: 4 LanesAll other values are reserved.
- **Reserved [12]**: This field shall be set to 1b by sender and ignored by receiver.
- **Link Rate 1 (LR1) [13]**: This field is used in combination with the *Link Rate 0 (LR0)* field where LR = {LR1,LR0}. The LR value shall specify the selected Link rate according to the following encodings:
 - 00b: 1.62GHz
 - 01b: 2.70GHz
 - 10b: 5.40GHz
 - 11b: 8.10GHzThis field is only valid when the *Lane Count* field is greater than 000b.
- **Training Pattern Support (TPS) [15:14]**: This field shall specify the supported TPS which can be used in EQ Phase in Non-LTTPR and LTTPR Transparent link training.
- **TPS3 Support [14]**:
 - 0b: TPS3 is not supported
 - 1b: TPS3 is supported
- **TPS4 Support [15]**:
 - 0b: TPS4 is not supported
 - 1b: TPS4 is supported
- **MSG Type [23:16]**: This field specifies the MSG type of the SET_CONFIG Packet. This field shall carry one of the MSG Type values listed in Table 10-6.
- **MSG Data [31:24]**: This field holds the MSG Data value that is associated with the specific MSG Type. The MSG Data shall match the MSG Type of the packet as set forth in Table 10-6.

A DP IN Adapter shall set the *LC* and *LR* fields in all SET_CONFIG Packets to be the same as the last SET_LINK SET_CONFIG Packet it sent.

A DP OUT Adapter shall set the *LC* and *LR* fields in all SET_CONFIG Packets to be the same as the last SET_LINK SET_CONFIG Packet it received.

Table 10-6. SET_CONFIG Message

| MSG Type | Type Value | Direction | MSG Data | Reference |
|------------------------|------------|-----------|--|------------------------------------|
| SET_LINK | 0x01 | Both | [0] – DP Link Training Mode 0b: Autonomous, concurrent (applicable to Non-LTTPR and LTTPR Transparent) 1b: DPTX-managed, sequential (applicable to LTTPR Non- Transparent) [7:1] – Reserved | Section 10.4.10 Section 10.4.12 |
| STATUS_TRAINING_FAIL | 0x02 | OUT to IN | [0] – LANE0_CHANNEL_EQ_DONE [1] – LANE0_SYMBOL_LOCKED [2] – LANE1_CHANNEL_EQ_DONE [3] – LANE1_SYMBOL_LOCKED [4] – LANE2_CHANNEL_EQ_DONE [5] – LANE2_SYMBOL_LOCKED [6] – LANE3_CHANNEL_EQ_DONE [7] – LANE3_SYMBOL_LOCKED | Section 10.4.10.2 |
| STATUS_LOST_CONNECTION | 0x03 | OUT to IN | [7:0] – Reserved | Section 13.8.2 |
| SET_PHY_TEST_MODE | 0x05 | IN to OUT | [7:0] – Reserved | Section 10.4.14 |
| SET_DOWNSPREAD | 0x07 | IN to OUT | Value written to DPCD 00107h | Section 10.4.7 |
| SET_POWER | 0x08 | IN to OUT | [2:0] – Power state as written/read to/from DPCD 600h [7:3] – Reserved | Section 10.4.11 Section 13.8.5 |
| SET_MFDP | 0x09 | Both | [0] – MFDP Enable 0: MFDP Disabled 1: MFDP Enabled [7:1] – Reserved | Section 10.4.5 |
| SET_FEC_READY | 0x0A | Both | [0] – Ready bit as written to DPCD 120h 0: Disable 1: Enable [7:1] – Reserved | Section 10.4.9 |
| SET_SINK_COUNT | 0x0C | Both | [5:0] – SINK_COUNT The value return by reading DPCD 200h [7:6] – Reserved | Section 10.4.6.3 Section 13.8.4 |
| IRQ | 0x0D | OUT to IN | [7:0] – Reserved | Section 10.4.3.3 |
| SET_STREAM_MODE | 0x0F | Both | [0] – Mode 0: SST Mode 1: MST Mode [7:1] – Reserved | Section 10.4.8 |
| SET_AUX_INIT | 0x12 | Both | [0] – 1b [7:1] – Reserved | Section 10.4.5 |
| SET_CMN_DPRX | 0x13 | IN to OUT | [7:0] – DPRX DPCD Revision 10h: DPCD r1.0 11h: DPCD r1.1 12h: DPCD r1.2 13h: DPCD r1.3 14h: DPCD r1.4a All other values are Reserved | Section 10.4.6.2 |

| MSG Type | Type Value | Direction | MSG Data | Reference |
|-----------------|------------|-----------|---|-------------------|
| SET_LTTTPR_MODE | 0x17 | IN to OUT | [0] – LTTTPR_Mode 0: LTTTPR Transparent 1: LTTTPR Non-Transparent [7:1] – Reserved | Section 10.4.6.1 |
| SET_TRAINING | 0x18 | IN to OUT | [7:0] – Training Stage (TS) 0h: Training Done with DPRX 1h: TPS1 2h: TPS2 3h: TPS3 7h: TPS4 FFh: Training Done with LTTTPR Non-Transparent UFP All other values are Reserved | Section 10.4.10.1 |
| SET_VSPE | 0x19 | IN to OUT | [1:0] – Voltage Swing Level. [2] – Max Swing Reached [4:3] – Pre-emphasis Level. [5] – Max Pre Emphasis Reached [7:6] – Reserved | Section 10.4.10.1 |
| STATUS_CR_DONE | 0x1D | OUT to IN | [0] – CR_DONE0 for Lane 0 [1] – CR_DONE1 for Lane 1 [2] – CR_DONE2 for Lane 2 [3] – CR_DONE3 for Lane 3 [6:4] – Reserved [7] – Phase: 0b: Status at end of CR_DONE Phase 1b: Status at end of EQ Phase | Section 10.4.10.2 |

After a DP Adapter sends a SET_CONFIG Packet, it shall wait for an ACK Packet with the *Type* field equal to 0h. After receiving an ACK Packet with the *Type* field set to 0h, the DP Adapter shall wait tDPSetConfigGap before sending the next SET_CONFIG Packet.

When a DP Adapter receives a SET_CONFIG Packet, it shall check the *ECC* field of the packet payload. The DP Adapter shall correct single-bit errors in the SET_CONFIG Packet payload. If an uncorrectable error is detected, the SET_CONFIG Packet shall be dropped. Otherwise, the DP Adapter shall respond with an ACK Packet with the *Type* field equal to 0h. The ACK Packet shall be sent within tDPAckResponse of receiving the SET_CONFIG Packet. If a DP Adapter does not receive an ACK Packet with the *Type* field equal to 0h within tDPAckTimeout of sending a SET_CONFIG Packet, it may issue Link-Init as described in Section 10.4.13.

A DP Adapter that receives a SET_CONFIG Packet with a value in the *Type* field that is not listed in Table 10-6 shall respond with an ACK packet with the *Type* field equal to 0h.

10.3.4.2.4 ACK Packet

An ACK Packet is used to acknowledge receipt of an HPD or SET_CONFIG Packet. An ACK Packet shall have the format shown in Figure 10-10.

Figure 10-10. ACK Packet Format

* SupplID

The *PDF* field in the header shall be set to 3 and the *Length* field shall be 4.

ACK Packet payload shall contain the following:

- **ECC [7:0]**: Error correction field that is calculated over bits [31:8] of the ACK Packet payload. See Section 5.1.2.3 for calculation.
- **Reserved [27:8]**: Shall be set to 0.
- **Type [31:28]**: Shall be set to 8h to acknowledge the receipt of a HPD Packet. Shall be set to 0h to acknowledge the receipt of a SET_CONFIG Packet. All other values are reserved.

10.3.4.3 Main-Link Path Packet Formats

A DP Adapter Layer uses the HopID values in the *Video HopID* field in the Protocol Adapter Configuration Capability to identify Main-Link Path Packets. When generating a Main-Link Path Packet, a DP IN Adapter Layer shall put the value in the *Video HopID* field into the *HopID* field of the Tunneled Packet header. When a DP OUT Adapter Layer receives a Tunneled Packet with a HopID that is equal to the *Video HopID* field, it shall treat that packet as a Main-Link Path Packet.


Main-Link Path Packets are described in detail in Section 10.5 and Section 10.6.

10.4 System Flows

10.4.1 Connection Manager Discovery

Figure 10-11 shows the sequence of events that take place between a DP Source plug and Hot Plug Detect.

Figure 10-11. Power On to HPD Sequence



CONNECTION MANAGER NOTE

A Connection Manager shall not pair a DP IN Adapter until after all of the following occur:

- The Connection Manager has either received a Hot Plug Event Packet from the DP IN Adapter or the CM has confirmed that the DP IN Adapter is available using a QUERY_DP_RESOURCE command as described in Section 8.3.1.1.1, with the DisplayPort Number parameter equal to the DP IN Adapter number.
- The Connection Manager has allocated a DP stream resource using an ALLOCATE_DP_RESOURCE command as described in Section 8.3.1.1.2, with a DisplayPort Number parameter equal to the DP IN Adapter number.

A Connection Manager shall not pair a DP OUT Adapter until after either the Connection Manager has received a Hot Plug Event Packet from the DP OUT Adapter or the ADP_DP_CS_2.HPD Status is set to 1b in the DP OUT Adapter Configuration Capability Field.

10.4.2 Path Configuration

10.4.2.1 Setup

A DP OUT Adapter shall poll the DP_STATUS.CMHS field and DP_REMOTE_CAP.Protocol Adapter Version field for as long as the values in those fields are both 0. When either DP_STATUS.CMHS = 1 or DP_REMOTE_CAP.Protocol Adapter Version > 0, the DP OUT Adapter shall do the following:

- If DP_REMOTE_CAP.Protocol Adapter Version was set to non-zero value while DP_STATUS_CTRL.CMHS remained zero, a DP OUT Adapter shall conclude it is a TBT3 Connection Manager and shall continue the flow as defined in Section 13.8.3.
- If DP_STATUS_CTRL.CMHS is set to 1 and DP_STATUS_CTRL.UF is zero, a DP OUT Adapter shall conclude it is a TBT3 DP IN Adapter and shall continue the flow as defined in Section 13.8.3.
- If DP_STATUS_CTRL.CMHS is set to 1 and DP_STATUS_CTRL.UF is set to one, a DP OUT Adapter shall reset DP_STATUS_CTRL.CMHS to zero.

A DP Adapter shall set the DP_COMMON_CAP register, to reflect the lowest common capability between DP_LOCAL_CAP and DP_REMOTE_CAP fields.

After Path configuration, a Connection Manager enables the DP Adapters at both ends of the Path by setting the VE (*Video Enable*) and AE (*Aux Enable*) bits in the ADP_DP_CS_0 registers of the Protocol Adapters to 1b.



CONNECTION MANAGER NOTE

Before configuring a Path between DP IN and DP OUT Adapters, the Connection Manager shall perform the following steps:

1. **Available Bandwidth** – A Connection Manager calculates the available bandwidth on each USB4 Link along the entire path, from DP IN Adapter to DP OUT Adapter. Based on the minimum available bandwidth, it concludes what DP Link can be established. Table 10-7 shows the Asynchronous Bandwidth needed for each DP Link lane count and link rate.
2. **Capabilities Exchange** – A Connection Manager does the following:
 - a. Sets the DP OUT Adapter field, DP_STATUS_CTRL.UF, according to the revision supported by the paired DP IN Adapter.
 - b. Sets the DP OUT Adapter DP_STATUS_CTRL.CMHS to 1b.
 - c. Polls the DP OUT Adapter DP_STATUS_CTRL.CMHS until it is reset to zero by the DP OUT Adapter.
 - d. Reads the DP_LOCAL_CAP register of the DP IN and DP OUT Adapters at each end of the Path:
 - The Connection Manager copies 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.
 - The Connection Manager copies 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 available bandwidth is insufficient with respect to the required bandwidth in Table 10-7, the Connection Manager limits the DP Link, lane count and/or link rate parameters accordingly. A Connection Manager 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. The use of this register by the DP IN Adapter is defined in Section 10.4.5.
3. **CM_ID** – A Connection Manager sets ADP_DP_CS_2.CM_ID in the DP IN Adapter to the Connection Manager index
4. **DP BW Allocation Mode** – If the Connection Manager supports DP BW Allocation Mode, it does the following:
 - a. Checks if the DP IN Adapter supports DP BW Allocation Mode by reading DP_LOCAL_CAP.DP_IN_BW_Allocation Mode Support.
 - b. If the DP IN Adapter supports DP BW Allocation Mode, the Connection Manager does the following in the DP IN Adapter Configuration Space:
 - i. Sets ADP_DP_CS_2.CM BW Allocation Mode Support to 1b.
 - ii. If the Path being setup between the two Adapters travels through the same USB4 Links as another DP Path, then it sets ADP_DP_CS_2.Group_ID to a non-zero value for all the Paths which travel through the same USB4 Links, otherwise it sets ADP_DP_CS_2.Group_ID to 0h.

- iii. Sets *ADP_DP_CS_2.NRD_Maximal_Link_Rate* to the *DP_LOCAL_CAP.Maximal_Link_Rate* of the DP IN Adapter or DP OUT Adapter (whichever is lower).
- iv. Sets *ADP_DP_CS_2.NRD_Maximal_Lane_Count* to the *DP_LOCAL_CAP.Maximal_Lane_Count* of the DP IN Adapter or DP OUT Adapter (whichever is lower).
- v. Sets *ADP_DP_CS_2.Estimated BW* according to Section 10.7.
- vi. Sets *ADP_DP_CS_2.Granularity* to any of the valid values.
- vii. Sets *DP_STATUS.Allocated BW* to 0h.
- viii. Sets *ADP_DP_CS_2.CM Ack* to 0b.

A Connection Manager shall configure the Output HopID to be 8 for the segment of an Aux Path that goes from a USB4 Port to a DP IN or DP OUT Adapter. A Connection Manager shall configure the Output HopID to be 9 for the segment of a Main-Link Path that goes from a USB4 Port to a DP OUT Adapter.

After Path configuration, a Connection Manager enables the DP Adapters at both ends of the Path by setting the VE (Video Enable) and AE (Aux Enable) bits in the *ADP_DP_CS_0* registers of the Protocol Adapters to 1b. The Connection Manager needs to use single Write Request to set the VE and AE bits to 1b so that they both are written at the same time.

The Connection Manager shall enable the DP IN Adapter before enabling the DP OUT Adapter.

Table 10-7. DisplayPort Required Bandwidth (Gbps)

| Mode | 4 Lanes | 2 Lanes | 1 Lane |
|-----------------|---------|---------|--------|
| HBR3 (8.1 Gbps) | 26 | 13 | 6.5 |
| HBR2 (5.4 Gbps) | 17.6 | 8.8 | 4.4 |
| HBR (2.7 Gbps) | 8.8 | 4.4 | 2.2 |
| RBR (1.62 Gbps) | 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.

10.4.2.2 Tear-down

When the *ADP_DP_CS_0.AE* bit and the *ADP_DP_CS_0.VE* bit are both set to 0, a DP Adapter shall set all the fields in its DP Adapter Configuration Capability to their default values.



CONNECTION MANAGER NOTE

Before tearing down a Path between two DP Adapters, the Connection Manager shall disable the DP Adapters on both ends of the Path by setting the ADP_DP_CS_0.VE (Video Enable) and ADP_DP_CS_0.AE (Aux Enable) bits to 0b in register of each DP Adapter. The Connection Manager shall use single Write Request to set the VE and AE bits to 0b so that they both are written at the same time.

After tearing down a DP Path, the Connection Manager releases the DP stream resource, using a DEALLOCATE_DP_RESOURCE command as defined in Section 8.3.1.1.3, with a DisplayPort Number parameter equal to the DP IN Adapter number being released.

If only one DP IN Adapter remains in a Group_ID, a Connection Manager sets the DP IN Adapter Group_ID to 0h.

10.4.3 HPD Event Propagation

The DisplayPort 1.4a Specification defines three kinds of Hot Plug Detect (HPD) events:

- IRQ – The Sink device sends an Interrupt Request to the Source device.
- Unplug – The Sink device is no longer attached to the Source device.
- Plug/Re-plug – The Sink device is newly attached to the Source device.

HPD events are transported across the USB4 Fabric from the DP OUT Adapter to the DP IN Adapter. Plug/Re-plug and Unplug events are transported using HPD Packets. IRQ events are transported using SET_CONFIG Packets.

10.4.3.1 HPD Plug

After a Path is setup between a DP OUT Adapter and a DP IN Adapter per 10.4.2.1, the DP OUT Adapter shall send an HPD Packet with the *P Flag* set to 1b. Upon receiving an HPD Packet with the *P Flag* set to 1b, the DP IN Adapter shall respond with an ACK Packet, execute the DP Adapter Init flow as defined in Section 10.4.5, and then drive HPD signal high on the DisplayPort interface.

After the DP IN Adapter drives HPD high, both DP Adapters shall be ready to handle AUX transactions.

10.4.3.2 HPD Unplug

A DP OUT Adapter detects unplug of a DPRX as defined in the DisplayPort 1.4a Specification. Upon unplug detection, the DP OUT Adapter shall send an HPD Packet with the *P Flag* set to 0b to the DP IN Adapter. The DP IN Adapter shall respond with an ACK Packet and drive the HPD signal low.

When the HPD signal is low, a DP OUT Adapter shall monitor the HPD signal (as defined by the DisplayPort 1.4a Specification) to detect a replug. It may also disable its Main-Link transmitters.

When the HPD signal is low, a DP IN Adapter may disable its DP Link receiver.

10.4.3.3 IRQ

A DP OUT Adapter detects an IRQ as defined in the DisplayPort 1.4a Specification. Upon IRQ detection, a DP OUT Adapter shall send a SET_CONFIG Packet of MSG type IRQ to the DP IN Adapter. However, if Link training is in process or it just ended, a DP OUT Adapter shall wait *tIRQDelay* after it sent a SET_CONFIG Packet to the DP IN Adapter (reporting that Link training is completion) before sending the IRQ SET_CONFIG Packet.

A DP IN Adapter that receives a SET_CONFIG Packet of MSG type IRQ shall respond with an ACK Packet and drive the IRQ event (according to the DisplayPort 1.4a Specification) towards the DPTX.

10.4.3.4 HPD Delay Requirements

Section 3.6.6.2 of the DisplayPort 1.4a Specification defines the delay from a detection of an HPD event by the DP OUT Adapter to the start of the event's generation by the DP IN Adapter on the DisplayPort HPD signal. The HPD event detection is defined in the DisplayPort 1.4a Specification – Section 5.1.4. Table 10-8 defines the maximum propagation delay through the DP Adapters that shall be used for HPD Events.

Table 10-8. HPD Event Propagation Delay Requirement

| | Max Delay (μs) |
|--------|----------------|
| DP IN | 90 |
| DP OUT | 90 |

The propagation delay through the DP OUT Adapter shall be measured from when the event is detected by the DP OUT Adapter to when the last bit of the corresponding packet, SET_CONFIG or HPD, is sent over the AUX Path. The propagation delay does not include the time it takes the Connection Manager to establish the AUX Path.

The propagation delay through the DP IN Adapter shall be measured from when the last bit of an HPD Event Packet, SET_CONFIG or HPD, arrives at the DP IN Adapter to when event is driven on the HPD signal.

A DP OUT Adapter shall send HPD packet with the *P Flag* set to 1b within tDPPlug of transitioning to the Paired state.

10.4.3.5 Manual HPD Control

A Connection Manager uses the *HPDC* and *HPDS* fields to control how a DP IN Adapter drives the DisplayPort HPD signal.

- When HPDC transitions from 0b to 1b, a DP IN Adapter shall drive the DisplayPort HPD signal low.
- When HPDS transitions from 0b to 1b, a DP IN Adapter shall drive the DisplayPort HPD signal high.
- The DisplayPort HPD signal level shall be set according to the most recent event, whether it is HPDS set, HPDC set, IRQ SET_CONFIG Packet or HPD Packet.

10.4.4 AUX Request and Response Handling

10.4.4.1 LTTPR Non-Transparent Mode

This section defines the AUX Request and Response handling by DP Adapters operating in LTTPR Non-Transparent mode.

With respect to DisplayPort signals, the DP IN and DP OUT Adapters serve as UFP and DFP respectively. However, in terms of AUX Handling, the DP IN Adapter acts as both UFP and DFP. The AUX handling roles for UFP and DFP are defined in the DisplayPort 1.4a Specification - Sections 3.6.5.3.2 and 3.6.5.3.3.

10.4.4.1.1 DP IN Adapter Requirements

Upon reception of a DisplayPort AUX request, a DP IN Adapter shall send an AUX Packet containing the request over the AUX Path. The AUX request coming from the DPTX shall not be modified by the DP IN Adapter.

A DP IN Adapter handles three types of AUX transactions:

- **Target** – Includes AUX transactions which target the LTTTPR or LT-tunable PHY Repeater DPCD Capability and ID Field (DisplayPort 1.4a Specification - Table 3-34).
 - Figure 10-12 shows an example of a Target transaction that is an AUX read transaction from DPCD address F0020h (TRAINING_AUX_RD_INTERVAL_PHY_REPEATER1). The transaction becomes a Target transaction because the DP Adapters are Repeater1 in this use case example.
- **Snoop** – Includes AUX transactions which hold valuable information for the LTTTPR, but do not target the LTTTPR.
 - Figure 10-13 shows an example of a Snoop transaction that is an AUX write transaction to DPCD address 00600h (SET_POWER). The transaction becomes a Snoop transaction because the power state is of interest to the LTTTPR. The DP IN Adapter follows with a SET_CONFIG Packet of type SET_POWER.
- **Pass-Through** – Includes AUX transactions which are irrelevant to the LTTTPR operation.

For Pass-Through transactions, there is no additional action required by the DP IN Adapter.

For Target transactions, the DP IN Adapter shall set the AUX_PEND flag as defined in the DisplayPort 1.4a Specification – Section 3.6.5.3.2, and take the appropriate action when the request comes back according to the DisplayPort 1.4a Specification – Section 3.6.5.3.2.

When a DP IN Adapter updates the DP OUT Adapter as the results of a Snoop or Target transaction, it shall perform the update only after receiving an AUX ACK response for all cases beside SET_CONFIG Packet of type SET_VSPE. The update shall be done by sending SET_CONFIG Packets to the DP OUT Adapter.

Figure 10-12. Target AUX Transaction Flow





Figure 10-13. Snoop AUX Transaction Flow**10.4.4.1.2 DP OUT Adapter Requirements**

A DP OUT Adapter shall convert an incoming AUX Packet received from the USB4 Fabric into a DisplayPort AUX request. The content of the request shall not be modified by the DP OUT Adapter.

A DP OUT Adapter shall convert an incoming DisplayPort AUX response into an AUX Packet and send it on the AUX Path. The content of the response shall not be changed by the DP OUT Adapter.

10.4.4.2 Non-LTTPR Mode

This section defines how a DP Adapter, operating in Non-LTTPR mode, handles AUX Requests and Responses. A DP IN Adapter shall implement AUX Replier. A DP OUT Adapter shall implement AUX Requester.

10.4.4.2.1 AUX Timeout Timers

The AUX Response Timeout timer in a DP IN Adapter shall be set to 300 μ s.

The AUX Reply Timeout timer in a DP OUT Adapter shall be set to 400 μ s.

10.4.4.2.2 DP IN Adapter Requirements

A DP IN Adapter that receives an AUX Request shall classify the AUX Transaction as one of the three following types:

1. Internal AUX Transaction – AUX Request which targets only DPCD addresses that are defined as internal in Table 10-9.
2. External AUX Transaction – AUX Request which targets only DPCD addresses that are not defined as internal in Table 10-9.
3. Combined AUX Transaction – AUX Request which targets both Internal and External DPCD addresses. The AUX Response is initially generated by the DPRX and altered by the DP IN Adapter.

Note: A DP IN Adapter may assume that an AUX Write Request is not classified as a Combined Aux Transaction.

Table 10-9. DPCD Internal Addresses

| Functionality | Address | Name |
|--|-----------------|------------------------------------|
| Link Training Control | 00100h | LINK_BW_SET |
| | 00101h | LANE_COUNT_SET |
| | 00102h | TRAINING_PATTERN_SET |
| | 00103-6h | TRAINING_LANE _x _SET |
| | 00107h | DOWNSPREAD_CTRL |
| Link Status ¹ | 00202h | LANE0_1_STATUS |
| | 00203h | LANE2_3_STATUS |
| | 00204h | LANE_ALIGN_STATUS_UPDATED |
| | 00206h | ADJUST_REQUEST_LANE0_1 |
| | 00207h | ADJUST_REQUEST_LANE2_3 |
| | 0200Ch | LANE0_1_STATUS_ESI |
| | 0200Dh | LANE2_3_STATUS_ESI |
| | 0200Eh | LANE_ALIGN_STATUS_UPDATED_ESI |
| Link Quality Control | 0010Bh | LINK_QUAL_LANE0_SET |
| | 0010Ch | LINK_QUAL_LANE1_SET |
| | 0010Dh | LINK_QUAL_LANE2_SET |
| | 0010Eh | LINK_QUAL_LANE3_SET |
| DP Tunneling over USB4 | E0000h – E00FFh | DP Tunneling over USB4 Field DPCDs |
| Notes: | | |
| 1. Link Status DPCD Registers are Internal only during Link Training phase. Link Training phase starts when DPTX writes TRAINING_PATTERN_SELECT to a non-zero value and ends when it writes zero to TRAINING_PATTERN_SELECT. | | |

For Internal AUX Transactions, a DP IN Adapter shall not send the AUX Request downstream and shall self-generate the AUX Response.

For External and Combined AUX Transactions, a DP IN Adapter shall send the AUX Request downstream to the DP OUT Adapter. The DP IN Adapter shall generate an AUX DEFER if the Response timer expires while waiting for the AUX Response. Response timer expiration is defined in the DisplayPort 1.4a Specification.

A DP IN Adapter shall not send an AUX Request to a DP OUT Adapter while the AUX Reply to the preceding AUX Request is outstanding.

DP IN Adapter shall increment AUX_REQ_CNTR by 1 on every received AUX request from DPTX and shall reset to zero on transition to IDLE state.

Figure 10-14 shows the Non-LTTPR DP IN Adapter AUX handling state machine.

Figure 10-14. DP IN Adapter AUX Handling State Machine




Table 10-10. DP IN Adapter AUX Handling State Machine

| Number | Current State | Next State | Condition | Action |
|--------|---------------------|---------------------|--|--|
| 1 | IDLE | AUX_REQ_OUTSTANDING | DP IN Adapter received an AUX Request. | If AUX Request is for an Internal AUX Transaction, DP IN Adapter shall process Request. If AUX Request is not for an Internal AUX Transaction, DP IN Adapter send Request downstream. |
| 2 | AUX_REQ_OUTSTANDING | IDLE | DP IN Adapter has an AUX Response available and DP IN Adapter is in Talk Mode as defined in the DisplayPort 1.4a Specification. | DP IN Adapter shall send the AUX Response and reset AUX_REQ_CNTR to zero. |
| 3 | AUX_REQ_OUTSTANDING | IDLE | DP IN Adapter received SET_CONFIG (SET_AUX_INIT) or AUX_REQ_CNTR > 2 while DP IN Adapter either does not have an AUX Response available or is not in Talk Mode as defined in the DisplayPort 1.4a Specification. | DP IN Adapter shall reset AUX_REQ_CNTR to zero. |

| Number | Current State | Next State | Condition | Action |
|--------|---------------------|---------------------|---|--|
| 4 | AUX_REQ_OUTSTANDING | DISCARD_RESPONSE | DP IN Adapter received an AUX Request that is different than the pending AUX Request and AUX_REQ_CNT < 3. | None. |
| 5 | DISCARD_RESPONSE | IDLE | DP IN Adapter received SET_CONFIG(SET_AUX_INIT) or AUX_REQ_CNTR > 2. | DP IN Adapter shall reset AUX_REQ_CNTR to zero. |
| 6 | DISCARD_RESPONSE | AUX_REQ_OUTSTANDING | DP IN Adapter has an AUX Response available. | DP IN Adapter shall discard the AUX Response, set AUX_REQ_CNTR to 1, and send the AUX Request. |

10.4.4.2.3 DP OUT Adapter Requirements

A DP OUT Adapter handles two types of AUX Transactions:

- DPTX initiated.
- DP OUT Adapter initiated.

A DPTX initiated AUX Request has higher priority than a DP OUT Adapter initiated AUX Transaction.

A DP OUT Adapter that receives a DPTX initiated AUX Request while handling a DP OUT Adapter Initiated AUX Transaction, shall send the DPTX initiated AUX Request as soon as it is in Talk Mode.

10.4.4.2.3.1 DP TX Initiated AUX Transactions

A DP OUT Adapter that receives an AUX Request from a DP IN Adapter shall initiate the AUX Request as soon as it is in Talk Mode.

A DP OUT Adapter that receives an AUX Response shall send the AUX Response over the AUX Path to the DP IN Adapter.

If the AUX Reply Timer expires before an AUX Response is received, a DP OUT Adapter shall send a SET_CONFIG of type SET_AUX_INIT and shall not retry the AUX Request.

10.4.4.2.3.2 DP OUT Adapter Initiated AUX Transactions

A DP OUT Adapter initiates AUX Transactions for:

- Autonomous, concurrent Link Training.
- Down-Spread Control access, see Section 10.4.7.

10.4.4.3 LTPR Transparent Mode

This section defines how a DP Adapter, operating in LTPR Transparent mode handles AUX Requests and Responses. A DP IN Adapter shall implement AUX Replier. A DP OUT Adapter shall implement AUX Requester.

10.4.4.3.1 AUX Timeout Timers

The AUX Response Timeout timer in a DP IN Adapter shall not be activated.

The AUX Reply Timeout timer in a DP OUT Adapter shall be set to 3.2ms.

10.4.4.3.2 DP IN Adapter Requirements

A DP IN Adapter that receives an AUX Request shall classify the AUX Transaction as one of the following types:

- Internal AUX Transaction – AUX Request which targets only DPCD addresses that are defined as internal in Table 10-9.
- External AUX Transaction – AUX Request which targets only DPCD addresses that are not defined as internal in Table 10-9.
- Combined AUX Transaction – AUX Request which targets both Internal and External DPCD addresses. The AUX Response is initially generated by the DPRX and altered by the DP IN Adapter.

Note: A DP IN Adapter may assume that an AUX Write Request is not classified as a Combined Aux Transaction.

For External and Combined AUX Transactions, a DP IN Adapter shall send the AUX Request downstream to the DP OUT Adapter.

A DP IN Adapter shall not:

- Generate AUX DEFER.
- Gate any External or Combined AUX Request sent by the DPTX.
- Gate any AUX Response sent by the DPRX.

Note: DPTX enabling LTTPR Transparent mode shall support AUX Reply Timeout of 3.2ms as described in the DisplayPort 1.4a Specification. Once DPTX enables LTTPR Transparent mode, it is the responsibility of DPTX to manage the AUX Requests and AUX Responses, not DP In Adapter.

10.4.4.3.3 DP OUT Adapter Requirements

A DP OUT Adapter handles two types of AUX Transactions:

- DPTX initiated.
- DP OUT Adapter initiated.

A DPTX initiated AUX Request has higher priority than a DP OUT Adapter initiated AUX Transaction.

A DP OUT Adapter that receives a DPTX initiated AUX Request while handling a DP OUT Adapter Initiated AUX Transaction, shall send the DPTX initiated AUX Request as soon as it is in Talk Mode.

10.4.4.3.3.1 DP TX Initiated AUX Transactions

A DP OUT Adapter that receives an AUX Request from a DP IN Adapter shall initiate the AUX Request as soon as it is in Talk Mode.

A DP OUT Adapter that receives an AUX Response shall send the AUX Response over the AUX Path to the DP IN Adapter.

If the AUX Reply Timer expires before an AUX Response is received, a DP OUT Adapter shall send a SET_CONFIG of type SET_AUX_INIT and shall not retry the AUX Request.

10.4.4.3.2 DP OUT Adapter Initiated AUX Transactions

A DP OUT Adapter initiates AUX Transactions for:

- Autonomous, concurrent Link Training.
- Down-Spread Control access, see Section 10.4.7.


10.4.4.4 AUX Delay Requirements

Table 10-11 defines the maximum delay allowed for a DP Adapter, per direction of a transaction. Figure 10-15 correlates the delay requirements to an AUX transaction sequence. The DP-to-USB4 Fabric delay shall be measured from the time the last bit arrives at the DP Adapter from the DisplayPort interface to the time when the first bit is sent to the USB4 Fabric. USB4 Fabric-to-DP delay shall be measured from the time when the last bit of the AUX Packet arrives at the DP Adapter to the time when the first bit is sent to the DisplayPort Interface, assuming the DP Adapter is in Talk Mode.

Table 10-11. AUX Delay Requirements

| Direction of Transport | Max Delay (μ s) |
|-------------------------------------|----------------------|
| DP IN: DP to USB4 Fabric (t_1) | 45 |
| DP OUT: USB4 Fabric to DP (t_2) | 10 |
| DP OUT: DP to USB4 Fabric (t_3) | 10 |
| DP IN: USB4 Fabric to DP (t_4) | 45 |

Figure 10-15. AUX Timing



10.4.4.5 Aggregated DisplayPort Capabilities

A DP IN Adapter shall update its DP_LOCAL_CAP and DP_COMMON_CAP registers if it receives an AUX Read Response that has lower parameter values than the registers currently contain. Before transmitting the AUX Read Response, the DP IN Adapter shall update the AUX Read Response to reflect the aggregated DisplayPort Capabilities as shown in Table 10-12.

Table 10-12. Aggregated DisplayPort Capabilities

| DPCD Register (Address) | DPCD Field Name | DisplayPort Capability |
|--|---------------------------|---|
| DPCD_REV (0000h / 02200h) | Major Revision Number | $\max\{12h, \min [\text{Downstream response}, \text{DP_COMMON_CAP.Maximal DPCD Rev.}]\}$. |
| MAX_LINK_RATE (00001h / 02201h) | MAX_LINK_RATE | Minimum of downstream response and DP_COMMON_CAP.Maximal Link Rate. |
| MAX_LANE_COUNT (00002h / 02202h / F0004h) | MAX_LANE_COUNT | Minimum of downstream response and DP_COMMON_CAP.Maximal Lane Count. If MFDP Mode is set, then Maximum is 2. |
| | POST_LT_ADJ_REQ_SUPPORTED | If in LTTPR Non-Transparent mode, this field is unchanged. Else set to zero. |

| DPCD Register (Address) | DPCD Field Name | DisplayPort Capability |
|--|----------------------------------|---|
| | TPS3_SUPPORTED | <i>Downstream response</i> bit AND DP_COMMON_CAP.TPS3 <i>Capability</i> bit. |
| DP_TUNNELING_MAX_LINK_RATE (E0028h) | DP_TUNNELING_MAX_LINK_RATE | If DP BW Allocation Mode was enabled, set to Minimum of: <ul style="list-style-type: none"> Maximum of DPRX response for MAX_LINK_RATE And ADP_DP_CS_2.NRD_Maximal_Link_Rate Else set to 0h. |
| DP_TUNNELING_MAX_LANE_COUNT (E0029h) | DP_TUNNELING_MAX_LANE_COUNT | If DP BW Allocation Mode was enabled, set to Minimum of: <ul style="list-style-type: none"> Maximum of DPRX response for MAX_LANE_COUNT And ADP_DP_CS_2.NRD_Maximal_Lane_Count Else set to 0h. |
| MAX_DOWNSPREAD (00003h / 02203h) | NO_AUX_TRANSACTION_LINK_TRAINING | Set to zero. |
| | TPS4_SUPPORTED | <i>Downstream response</i> bit AND DP_COMMON_CAP.TPS4 <i>Capability</i> bit. |
| TRAINING_AUX_RD_INTERVAL (0000Eh / 0220Eh) | TRAINING_AUX_RD_INTERVAL | If in LTTPR Non-Transparent mode, this field is unchanged. Else, the downstream response is either unchanged or the DP IN Adapter may increase the value. Note that per DP v1.4a, the maximum value for this field is 4h. |
| MSTM_CAP (00021h) | MSTM_CAP | <i>Downstream response</i> bit AND DP_COMMON_CAP.MSTM <i>Capability</i> bit. |
| DSC_SUPPORT (00060h) | DSC Support | <i>Downstream response</i> bit AND NOT(DP_COMMON_CAP.DSC Not Supported bit). |
| FEC_CAPABILITY (00090h) | FEC_CAPABLE | <i>Downstream response</i> bit AND NOT(DP_COMMON_CAP.FEC Not Supported bit). |
| DPRX_FEATURE_ENUMERATION_LIST (02210h) | GTC_CAP | Set to zero. |
| | SST_SPLIT_SDP_CAP | <i>Downstream response</i> bit AND DP_COMMON_CAP.Secondary Split <i>Capability</i> bit. |
| PHY_REPEATER_CNT (F0002h) | PHY_REPEATER_CNT | If DP_COMMON_CAP.LTTPR Not Supported is set to 1 then set to zero. If DP_COMMON_CAP.LTTPR Not Supported is set to 0 then respond as an LTTPR according to the DisplayPort 1.4a Specification. |

10.4.4.6 DPCD DP Tunneling over USB4

When a DP IN Adapter receives an AUX Request targeting the DP Tunneling over USB4 Field DPCDs, it shall respond with its internal data.

A DP IN Adapter shall set the *DP Tunneling Support* bit to 1b in the DP_TUNNELING_CAPABILITIES DPCD (Address E000Dh bit offset 0).

10.4.5 DP Adapters Init Flow

A DP IN Adapter in the Paired state shall do the following after receiving a first HPD Packet with the *P flag* set to 1b:

- Update MFDP Mode inner variable according to Section 10.4.5.1.
- Send a SET_CONFIG Packet of type SET_AUX_INIT.

The DP IN Adapter shall not drive HPD high on the DisplayPort Interface until after it performs the steps above.

10.4.5.1 Multi-Function DP

If any of the DisplayPort connections are Multi-Function (as defined in the DisplayPort Alt Mode Specification), a DP IN Adapter sets an internal MFDP Mode flag to 1b. Otherwise the internal MFDP Mode flag is set to 0b. Section 10.4.4.5 describes how the MFDP Mode flag is used.

If a DP IN Adapter is not connected as part of a Multi-Function (as defined in the DisplayPort Alt Mode Specification), it shall send a SET_CONFIG Packet of type SET_MFDP with the *MFDP Enable* bit set to 0b. Otherwise it may send a SET_CONFIG Packet of type SET_MFDP with the *MFDP Enable* bit set to 1b.


A DP OUT Adapter which receives a SET_CONFIG Packet of type SET_MFDP shall respond with a SET_CONFIG Packet of type SET_MFDP within tDPInit. If the DP OUT Adapter is connected as part of Multi-Function as defined in the DisplayPort 1.4a Specification, then the *MFDP Enable* bit shall be set to 1b, otherwise it shall be set to 0b.

10.4.6 Source Discovery

According to the DisplayPort 1.4a Specification, upon HPD Plug detection the DPTX executes the following steps:

- LTTTPR recognition.
- DPRX Capabilities read.
- Sink Count read.

Figure 10-16. Example DP Source Discovery Sequence



10.4.6.1 LTTPR Recognition and Modes Change

A DPTX performs an AUX read from the *LT-tunable PHY Repeater DPCD Capability and ID* Field to discover the presence, count and capabilities of any downstream LTTPR. A DP IN Adapter shall modify the resulting AUX read response as defined in Section 10.4.4.5.

This action of AUX read from LT-tunable PHY Repeater DPCD registers by DPTX prompts DP Adapter to transition from Non-LTTPR mode to LTTPR mode. Upon transition to LTTPR mode, DP Adapter is in LTTPR Transparent mode by default. DPTX may prompt the transition to LTTPR Non-Transparent mode and back to LTTPR Transparent mode by setting PHY_REPEATER_MODE register at DPCD F0003h, 55h for LTTPR Transparent mode and AAh for LTTPR Non-Transparent mode.

If DP_COMMON_CAP.LTTPR Not Supported is set to 1b, a DP Adapter shall operate only in Non-LTTPR mode. Otherwise it transitions between the three operation modes according to Table 10-13.

Table 10-13. DP Adapter Operation Mode Transitions

| Mode | DP IN Adapter | DP OUT Adapter |
|-----------------------|---|--|
| Non-LTTPR | Adapter has exited the Paired State | Adapter has exited the Paired State or Adapter received a SET CONFIG Packet of type SET_AUX_INIT |
| LTTPR Transparent | Adapter receives first DPTX access to the <i>LT-tunable PHY Repeater DPCD Capability and ID</i> Field or DPTX changes the PHY_REPEATER_MODE to Transparent mode (55h) | Adapter received a SET CONFIG Packet of type SET_LTTPR_MODE with <i>LTTPR_Mode</i> set to 0b |
| LTTPR Non-Transparent | DPTX changes the PHY_REPEATER_MODE to Non-Transparent mode (AAh) | Adapter received a SET CONFIG Packet of type SET_LTTPR_MODE with <i>LTTPR_Mode</i> set to 1b |

A DP IN Adapter shall do the following before transitioning to LTTPR Transparent mode:

- Complete the AUX Transaction according to the current operation mode.
- Send a SET CONFIG Packet of type SET_LTTTPR_MODE with LTTTPR_Mode set to 0b.

A DP IN Adapter shall do the following before transitioning to LTTTPR Non-Transparent mode:

- Complete the AUX Transaction according to the current operation mode.
- Send a SET CONFIG Packet of type SET_LTTTPR_MODE with LTTTPR_Mode set to 1b.

10.4.6.2 DPRX Capabilities Read

The DPRX Capabilities read is performed by the DPTX. In response to a DPRX Capabilities read, a DP IN Adapter shall:

1. Snoop the read response and record the values of the DPCD_REV, MAX_LINK_RATE, MAX_LANE_COUNT, TPS3_SUPPORTED and TPS4_SUPPORTED fields located at DPCD addresses 00000h, 00001h, 00002h and 00003h respectively.
2. Update the *Maximal DPCD Rev*, *Maximal Link Rate*, *Maximal Lane Count*, *TPS3 Capability* and *TPS4 Capability* fields in the DP_LOCAL_CAP and DP_COMMON_CAP registers to reflect the lowest common capabilities between the existing values of those registers and the recorded values from step 1.
3. Send a SET_CONFIG Packet of type SET_CMN_DPRX with MSG Data equal to the snooped DPCD_REV which reflects the DPRX DPCD_REV.
4. Set the *DPRX Capabilities Read Done* field in the DP_COMMON_CAP register to 1b. Note that this field is set to 1b regardless of whether or not the values in the *Maximal Link Rate* and *Maximal Lane Count* fields in Step 2 were changed.

This process enables the Connection Manager to know when to read the modified capabilities and recalculate the maximal Bandwidth that may be consumed by this DP Link.

10.4.6.3 Sink Count Read

When DPCD addresses 00200h or 02002h are read by the DPTX, a DP IN Adapter shall snoop the read response and record the value of the SINK_COUNT. When the recorded SINK_COUNT value is zero, the DP IN Adapter shall send a SET_CONFIG Packet of type SET_SINK_COUNT, reflecting the recorded value.

When a DP OUT Adapter receives a SET_CONFIG Packet of type SET_SINK_COUNT with the SINK_COUNT value equal to zero, the DP OUT Adapter shall:

- Report an Unplug event as defined in Section 10.3.3.
- Set the ADP_DP_CS_2.HPD Status to 0b in the *DP OUT Adapter Configuration Capability* Field.

While a DP OUT Adapter is Unplugged and has the ADP_DP_CS_2.HPD Status set to 0b in the *DP OUT Adapter Configuration Capability* Field, it shall do the following upon IRQ detection:

- Report a Plug event as defined in Section 10.3.3.
- Set the ADP_DP_CS_2.HPD Status to 1b in the *DP OUT Adapter Configuration Capability* Field.

10.4.7 Down-Spread Control

When DPCD address 00107h is written by a DPTX, a DP IN Adapter, operating in Non-LTTTPR or LTTTPR Transparent modes, shall respond with AUX ACK and shall send a SET_CONFIG Packet of type SET_DOWNSPREAD. The *MSG Data* field in the SET_CONFIG Packet shall be equal to the value written by the DPTX.

A DP OUT Adapter that receives a SET_CONFIG Packet of type SET_DOWNSPREAD shall initiate an AUX write request to DPCD address 00107h with the value received in the MSG Data of the SET_CONFIG Packet.

10.4.8 Stream Mode Set

If the DPTX writes to DPCD address 00111h, a DP IN Adapter shall snoop the write request and record the value of the MST_EN bit. A DP IN Adapter shall send a SET_CONFIG Packet of type SET_STREAM_MODE, reflecting the recorded value of the MST_EN bit when a new recorded MST_EN value is different than the previous value.

The MST_EN default value at the DP Adapters shall be as defined in the DisplayPort 1.4a Specification.

A DP OUT Adapter that receives a SET_CONFIG Packet of type SET_STREAM_MODE shall respond with a SET_CONFIG Packet of type SET_STREAM_MODE, to signify the acknowledgment of the mode change. The value of the *MSG Data* field in the return packet has no meaning and shall be ignored by the DP IN Adapter.

10.4.9 DSC and FEC Enable

DSC capability discover, configuration and enabling flows are carried over AUX transaction, which are Pass-Through transactions for the DP Adapters. Handling the control DSC symbol is described in Section 10.5.1.1.

If the DPTX writes to DPCD address 00120h, a DP IN Adapter shall snoop the write request and record the values of the *FEC_READY*, *FEC_ERROR_COUNT_SEL* and *LANE_SELECT* fields. A DP IN Adapter shall send a SET_CONFIG Packet of type SET_FEC_READY, reflecting the recorded value of the FEC_READY bit when a new recorded FEC_READY value is different than the previous value.

The FEC_READY default value at the DP Adapters shall be as defined in the DisplayPort 1.4a Specification.

A DP OUT Adapter that receives a SET_CONFIG Packet of type SET_FEC_READY shall respond with a SET_CONFIG Packet of type SET_FEC_READY within tDPInit, to signify the acknowledgment of the mode change. The value of the *MSG Data* field in the return packet has no meaning and shall be ignored by the DP IN Adapter.

FEC can be enabled or disabled by the DPTX using FEC_DECODE_EN and FEC_DECODE_DIS symbol sequences as defined in Section 10.5.3.

10.4.10 DP Link Training

10.4.10.1 LTTTPR

The DP IN and DP OUT Adapters shall follow the LTTTPR Non-Transparent link training as defined in the DisplayPort 1.4a Specification while noting the following points:

- DP IN as UFP and DFP – As described in Section 10.4.4, the DP IN Adapter serves as UFP and DFP for AUX handling, therefore it updates the DP OUT Adapter with the different stages of the LTTTPR Non-Transparent link training through SET_CONFIG Packets.
- Training Patterns – Training Patterns are not carried over the USB4 Fabric.

10.4.10.1.1 DP IN Adapter Requirements

- A DP IN Adapter acts as the UFP of LTTTPR and trains its LTTTPR receiver as defined in the DisplayPort 1.4a Specification.
- When a DP Source trains the downstream DP Links, it alternates between all possible training patterns, including TPS1. Therefore, after DP Link training is finished on the UFP, a DP IN Adapter shall maintain symbol lock and lane alignment in its receiver while DPTX trains the rest of downstream DP Links.
- A DP IN Adapter shall send a SET_CONFIG Packet of type SET_LINK after DPTX writes TPS1 to the DP IN Adapter TRAINING_PATTERN_SET_PHY_REPEATERx DPCD register. The SET_CONFIG packet shall have the following values:
 - LC = LANE_COUNT_SET value written by DPTX.
 - LR = LINK_BW_SET value written by the DPTX.

- MSG Data = 1b, representing DP Link Training Mode = LTTTPR Non-Transparent.
- A DP IN Adapter shall send a SET_CONFIG Packet of type SET_TRAINING with TS = 0xFF after DPTX writes 0x0 to the DP IN Adapter TRAINING_PATTERN_SET_PHY_REPEATERx DPCD register.
- A DP IN Adapter shall send a SET_CONFIG Packet of type SET_TRAINING with TS = 0x0 after DPTX writes 0x0 to the DPRX TRAINING_PATTERN_SET DPCD register.
- After DP Link training is finished on the UFP, a DP IN Adapter shall detect a Training pattern on its receiver. After detecting the Training pattern, the DP IN Adapter shall send a single corresponding SET_CONFIG Packet of type SET_TRAINING for every change in Training pattern. The TS field shall be equal to the detected Training pattern as defined in Table 10-6.
- A DP IN Adapter shall send SET_CONFIG Packet of Type SET_VSPE when DPTX writes the TRAINING_LANE0_SET or TRAINING_LANE0_SET_PHY_REPEATERx DPCD register of the next downstream receiver. The MSG Data field shall carry the value in the write request. The SET_VSPE SET_CONFIG Packet shall be sent by the DP IN Adapter before sending the AUX Request Packet.
- A DP Adapter shall have higher priority generating and parsing SET_CONFIG packets over AUX Transaction.
- A DP IN Adapter shall respond to a DPTX that link training has ended successfully only when all the following are true:
 - The DP IN Adapter internal status indicates link training has ended successfully.
 - The DP IN Adapter sent at least nine DP Clock Sync Packets after it sent the SET_CONFIG Packet of type SET_LINK.

10.4.10.1.2 DP OUT Adapter Requirements

- A DP OUT Adapter that receives a SET_CONFIG Packet of Type SET_LINK with the DP Link Training Mode bit set to 1b shall start its internal Symbol clock PLL according to the Link Rate field, and start the Lifetime Counter as defined in Section 10.6.1.2.
- A DP OUT Adapter that receives a SET_CONFIG Packet of Type SET_TRAINING with TS field equal to 1, 2, 3 or 7 shall transmit TPS1, TPS2, TPS3 or TPS4 accordingly.
- A DP OUT Adapter that receives a SET_CONFIG Packet of Type SET_TRAINING with TS field equal to 0 shall transmit IDLE pattern according to the DisplayPort 1.4a Specification.
- A DP OUT Adapter shall set its Voltage Swing (VS) and Pre-Emphasis (PE) levels for all enabled lanes upon receiving a SET_CONFIG Packet of type SET_VSPE. The VS and PE levels shall be according to the MSG Data. The DP OUT Adapter shall transition on the 10-bit symbol boundary when:
 - Transitioning from one training pattern to another training pattern.
 - Transitioning to IDLE sequence after DP Link training is done.

10.4.10.1.3 DP Link Training Example

The DP Link Training process in this section describes a system which, from the DisplayPort point of view, consists of DPTX, DPRX and one LTTTPR in LTTTPR Non-Transparent mode. The DPRX's maximum supported link rate is HBR2, and it does not support TPS4. The DPTX establishes a DP Link of 2 lanes at HBR2 link rate.


10.4.10.1.3.1 LTTTPR – CR_DONE Phase

Figure 10-17 shows a sequence of events that starts from DP Link training initiation and gets the LTTTPR to achieve CR_DONE:

1. DP Link Training Start

- As defined in the DisplayPort 1.4a Specification, the DPTX starts the DP Link training with a write to DPRX's LINK_BW_SET and LANE_COUNT_SET.
 - The DP IN Adapter snoops that transaction to record the Training parameters.
- TPS1 Initiate
 - DPTX starts transmitting TPS1.
 - DPTX writes TPS1 to TRAINING_PATTERN_SET_PHY_REPEATER1 and Level0 to Voltage Swing and Pre-emphasis at TRAINING_LANEEx_SET_PHY_REPEATER1.
 - The DP IN Adapter alters the AUX NACK response to an AUX ACK response.
 - DP OUT Adapter Initiation
 - A DP IN Adapter sends SET_CONFIG Packet of type SET_LINK, carrying the DP Link training parameters and *DP Link Training Mode* = 1b.
 - The DP OUT Adapter starts its Symbol clock PLL according to the received *Link Rate* field and starts the Lifetime Counter.
 - CR Status
 - DPTX reads LTTPR's status from LANE0_1_STATUS & ADJUST_REQUEST_LANE0_1.
 - The DP IN Adapter alters the AUX response according to its internal CR_DONE state. In this example, it achieves CR_DONE on both Lanes.

Figure 10-17. DP Link Training – LTTPR CR_DONE



10.4.10.1.3.2 LTTPR – EQ Phase

Figure 10-18 shows a sequence of events that starts from DPTX Transmit TPS4 and ends as DPTX finishes LTTPR EQ phase:


- TPS4 Initiate
 - DPTX starts transmitting TPS4.
 - DPTX writes TPS4 to TRAINING_PATTERN_SET_PHY_REPEATER1 and Level0 to Voltage Swing and Pre-emphasis at TRAINING_LANEEx_SET_PHY_REPEATER1.
- EQ Status
 - The DP IN Adapter achieves equalization, symbol lock and lanes alignment.

- The DP IN starts its Lifetime Counter.
- DPTX reads LTTTPR's status from LANE0_1_STATUS, LANE_ALIGN_STATUS_UPDATED & ADJUST_REQUEST_LANE0_1.
- The DP IN Adapter alters the AUX response according to:
 - Its internal EQ_DONE, SYMBOL_LOCK & LANE_ALIGNMENT states.
 - It sent at least nine DP Clock Sync Packets.
- In this example, it achieves EQ_DONE, SYMBOL_LOCK & LANE_ALIGNMENT on both Lanes and was able to send at least nine DP Clock Sync Packets.

3. End LTTTPR Training

- DPTX writes 0x0 to TRAINING_PATTERN_SET_PHY_REPEATER1 signifying that DP Link training between the DPTX and the UFP of the LTTTPR is done.
- A DP IN Adapter sends a SET_CONFIG Packet of type SET_TRAINING with a TS field = 0xFF.

Figure 10-18. DP Link Training - LTTTPR - EQ Phase



10.4.10.1.3.3 DPRX - CR_DONE Phase

Figure 10-19 shows a sequence of events which starts from transition to TPS1 across all DP Links, tunes the voltage swing, and ends when DPRX achieves CR_DONE:

1. DP OUT Adapter Transmit enable
 - DPTX switches from transmitting TPS4 to transmitting TPS1.
 - A DP IN Adapter detects TPS1 and sends SET_CONFIG Packet of type SET_TRAINING with a TS field = 0x1.
 - A DP OUT Adapter enables and configures its transmitters according to the received SET_CONFIG Packets. In this example, the DP OUT Adapter enables 2 Lanes at a bit rate of HBR2, transmitting TPS1.
 - DPTX writes TPS1 to TRAINING_PATTERN_SET and Level0 to Voltage Swing and Pre-emphasis at TRAINING_LANE_x_SET to DPRX.
2. DPRX CR Status


2. EQ Status

- DPTX reads DPRX's status from LANE0_1_STATUS, LANE_ALIGN_STATUS_UPDATED & ADJUST_REQUEST_LANE0_1.
- DPRX, in this example, achieves EQ_DONE, SYMBOL_LOCK & LANE_ALIGNMENT on both lanes.

3. DP Link Training End

- DPTX writes 0x0 to TRAINING_PATTERN_SET signifying that DP link training between the DPTX and the DPRX, including the LTTPR, is done.
- A DP IN Adapter sends SET_CONFIG of type SET_TRAINING with TS field = 0x0.
- A DP OUT Adapter switches from transmitting TP3 to transmitting IDLE pattern as defined in the DisplayPort 1.4a Specification.
- DPTX switches from transmitting TPS3 to transmitting IDLE.
- A DP IN Adapter detects the end of the link training and upon receiving SR it starts tunneling the high speed data over the Main-Link Path.
- A DP OUT Adapter switches from the self-generated IDLE pattern to the high speed tunnel data coming over the Main-Link Path.

Figure 10-20. DP Link Training - DPRX - EQ Phase



10.4.10.2 Non-LTTTPR and LTTTPR Transparent

When a DPTX performs link training, it trains the DP IN Adapter receiver – it is unaware of the second DisplayPort link being trained by the DP OUT Adapter. The two DisplayPort links are trained simultaneously. The DP IN Adapter aggregates the status from each link when responding to the DPTX.

A DP Adapter shall perform DisplayPort link training according to the DisplayPort 1.4a Specification with the modifications and requirements defined in Section 10.4.10.2.1 and Section 10.4.10.2.2.

10.4.10.2.1 DP IN Adapter Requirements

A DP IN Adapter shall send a SET_CONFIG Packet of type SET_LINK after DPTX writes TPS1 to the DP RX TRAINING_PATTERN_SET DPCD register. The SET_CONFIG packet shall have the following values:

- LC = LANE_COUNT_SET value written by DPTX.
- LR = LINK_BW_SET value written by the DPTX.
- TPS = Reflects TPS3 and TPS4 support as indicated in the DP_COMMON_CAP register.
- MSG Data = 0b, representing *DP Link Training Mode* = Non-LTTPR and LTTPR Transparent modes.

A DP IN Adapter shall respond to a status read of LANEx_CR_DONE as follows:

- If a SET_CONFIG Packet of type STATUS_CR_DONE was not received since link training started, set the LANEx_CR_DONE bits to 0b.
- If a SET_CONFIG Packet of type STATUS_CR_DONE was received since link training started, set the LANEx_CR_DONE bits to be the internal DP IN Adapter status for a lane ANDed with the relevant bit present in the last received SET_CONFIG MSG Data.

A DP IN Adapter shall respond to a DPTX that link training has ended successfully only when all the following are true:

- The DP IN Adapter internal status indicates link training has ended successfully.
- The DP IN Adapter received a SET_CONFIG Packet of type SET_LINK, carrying the same LC and LR fields that it sent to the DP OUT Adapter when link training was initiated.
- The DP IN Adapter sent at least nine DP Clock Sync Packets after it received a SET_CONFIG Packet of type STATUS_CR_DONE.

While the conditions (as defined in this section) for successful link training are not met, a DP IN Adapter shall respond to a status read of INTERLANE_ALIGN_DONE, LANEx_CHANNEL_EQ_DONE and LANEx_SYMBOL_LOCKED as follows:

- If a SET_CONFIG Packet of type STATUS_TRAINING_FAIL was received since link training started, set the following bits:
 - INTERLANE_ALIGN_DONE shall be set to 0b.
 - LANEx_CHANNEL_EQ_DONE is equal to the DP IN internal status ANDed with LANEx_CHANNEL_EQ_DONE that was received as MSG Data by the STATUS_TRAINING_FAIL.
 - LANEx_SYMBOL_LOCKED is equal to the DP IN internal status ANDed with LANEx_SYMBOL_LOCKED that was received as MSG Data by the STATUS_TRAINING_FAIL.
- Else, a DP IN Adapter shall use one or more of the methods below to indicate to DPTX that link training has not completed successfully yet.
 - Set INTERLANE_ALIGN_DONE to 0b.
 - Set LANEx_CHANNEL_EQ_DONE to 0b for any of the active lanes.
 - Set LANEx_SYMBOL_LOCKED to 0b for any of the active lanes.

Note: Which indication(s) to negate is implementation specific.

10.4.10.2.2 DP OUT Adapter Requirements

A DP OUT Adapter receiving a SET_CONFIG Packet of type SET_LINK, with LC field other than 0h shall:

- Initiate link training with the target Link Rate and Lane Count received from the SET_LINK Packet.
- Link training proceeds according to the DisplayPort 1.4a spec except that a DP OUT Adapter that concludes that it needs to either reduce the Link Rate or Lane Count shall treat it as link training failure and shall not reduce the Link Rate or Lane Count.

A DP OUT Adapter which finishes the Clock Recovery Sequence (as defined in the DisplayPort 1.4a Specification) shall send a SET_CONFIG Packet of type STATUS_CR_DONE, reflecting the LANEx_CR_DONE statuses of the active lanes. The *Phase* field shall be set to 0b.

A DP OUT Adapter in EQ phase which detects that the DP receiver has lost Clock Recovery on one or more of the active lanes shall conclude that link training has failed and shall send a SET_CONFIG Packet of type STATUS_CR_DONE, reflecting the new LANEx_CR_DONE statuses of the active lanes. The *Phase* field shall be set to 1b.

If link training fails for a reason other than lost Clock Recovery, a DP OUT Adapter shall send a SET_CONFIG Packet of type STATUS_TRAINING_FAIL. The MSG Data shall be set as follows:

- LANEx_CHANNEL_EQ_DONE is equal to the value of the last read from LANEx_CHANNEL_EQ_DONE field in DPRX.
- LANEx_SYMBOL_LOCKED is equal to the value of the last read from LANEx_SYMBOL_LOCKED field in DPRX.

If link training finishes successfully, a DP OUT Adapter shall:

- Send a SET_CONFIG Packet of type SET_LINK, with the same *LC* and *LR* fields it received from the DP IN Adapter when link training was initiated.
- Generate IDLE pattern (including SR) for both MST and SST DP Links.

10.4.10.3 Transition to High Speed Tunnel

A DP IN Adapter shall start converting DisplayPort Main-Link Symbols into Tunneled Packets and sending those Packets over the Main-Link Path when all of the following are true:

- Link Training has completed successfully.
- The DP IN Adapter received an SR.

A DP OUT Adapter transitions from sending IDLE pattern to reconstructing the DisplayPort Main-Link symbols according to Section 10.5.4.

Note: The rules for sending Clock Sync Packets are defined in Section 10.6. DP Clock Sync Packets are not gated by receiving an SR from DPTX.

10.4.11 Power States Set

When DPTX writes to DPCD address 00600h, a DP IN Adapter shall snoop the write request and record the value of the *SET_POWER_STATE* field. A DP IN Adapter shall send a SET_CONFIG Packet of type SET_POWER, reflecting the recorded value in the following cases:

- A first DPCD write of address 00600h after an HPD Plug event.
- The new recorded SET_POWER_STATE is different than the previous value.

The SET_POWER_STATE default value at the DP Adapters is defined in the DisplayPort 1.4a Specification.

If the SET_POWER_STATE is 2h or 5h, a DP IN Adapter may choose to initiate the DP Main-Link disable flow defined in Section 10.4.12.

The Low Power state enables both DP Adapters to save power according to the SET_POWER_STATE of the DPRX. The means by which the DP Adapters may save power are implementation specific and outside the scope of this specification.

10.4.12 DP Main-Link Disable

A DP IN Adapter may disable its Main-Link receivers and send a Main-Link disable message in the form of a SET_CONFIG Packet of type SET_LINK with *Link Rate* and *Lane Count* fields set to zero in the following events:

- The DP IN Adapter experiences excessive disparity or symbol error which results in the clearing of LANEx_SYMBOL_LOCKED and INTERLANE_ALIGNED_DONE bits.

A DP OUT Adapter which receives a Main-Link disable message shall disable its transmitters.

Note: It is expected that the system recovers from a DP-Link failure as follows:

1. DPRX experiences a DP Link failure and sends an IRQ.
2. DPTX reads the DP-Link status and discovers the DP-Link failure.
3. DPTX retrains the DP-Link.

10.4.13 Link-Init

A Link-Init action may be triggered by:

- Connection Manager sets the SWLI bit in ADP_DP_CS_2 register to 1b.
- ACK Packet timeout expires.

Upon Link-Init activation, a DP IN Adapter shall turn off its DisplayPort receivers and stop any transmission of Tunneled Packets over the Main-Link Path until the end of the next successful Link training, as defined in Section 10.4.10.3.

Upon Link-Init activation, a DP OUT Adapter shall turn off its DisplayPort transmitters.

Note: It is expected that the system recovers from a DP-Link failure as follows:

1. DPRX experiences a DP-Link failure and sends an IRQ.
2. DPTX reads the DP-Link status and discovers the DP-Link failure.
3. DPTX retrains the DP-Link.

10.4.14 DP PHY Testability

The method for testing the PHY layer in a DP IN Adapter is defined in Section 10.4.14.1. The method for testing the PHY layer DP OUT Adapter is defined in Section 10.4.14.2.

10.4.14.1 DP IN Adapter PHY Layer Testing

The PHY layer of a DP IN Adapter shall be tested as described in the DisplayPort 1.4a PHY CTS with the changes listed below.

- Before entering DP IN PHY Test Mode:
 - Connect a Router with a DP OUT Adapter and a DPRX. Both the DP OUT Adapter and the DPRX need to support the Link Rate and Lane Count required by the test.
 - Verify that a DP Link is established.
- Entering DP IN PHY Test Mode:
 - The DP IN Adapter shall enter DP IN PHY Test Mode when the DPTX writes a non-zero value to LINK_QUAL_LANEx_SET in the DPCD registers.
- While in DP IN PHY Test Mode:

- The DP IN Adapter shall keep the Hot Plug Detect signal high.
- The DP IN Adapter shall respond to all AUX transactions related to the PHY layer testing.
- Exiting DP IN PHY Test Mode:
 - The DP IN Adapter shall exit DP IN PHY Test Mode when the DPTX initiates DP Link Training.

10.4.14.2 DP OUT Adapter PHY Layer Testing

The PHY layer of a DP OUT Adapter shall be tested as described in the DisplayPort 1.4a PHY CTS with the changes listed below.

- Before entering DP OUT PHY Test Mode:
 - Connect a Router with a DP IN Adapter and a DPTX. Both the DP IN Adapter and the DPTX need to support the test required Link Rate and Lane Count.
 - Verify that a DP Link is established.
- Entering DP OUT PHY Test Mode:
 - When the DPTX reads the following sequence, the DP IN Adapter shall send a SET_CONFIG Packet of Type SET_PHY_TEST_MODE and enter DP OUT PHY Test Mode:
 - AUTOMATED_TEST_REQUEST is set to 1b (DPCD 00201h or 02003h bit 1)
 - PHY_TEST_PATTERN is set to 1b (DPCD 00218h bit 3)
 - The DP OUT Adapter shall enter DP OUT PHY Test Mode when it receives a SET_CONFIG Packet of type SET_PHY_TEST_MODE.
- While in DP OUT PHY Test Mode:
 - The DP OUT Adapter shall act as the DPTX under test.
 - The DP OUT Adapter shall send HPD Packet with Plug Flag set to 0b.
 - The DP IN Adapter shall not forward any AUX Transactions to the DP OUT Adapter.
- Exiting DP OUT PHY Test Mode
 - Upon a DPRX HPD signal de-assertion:
 - The DP OUT Adapter shall exit DP OUT PHY Test Mode.
 - The Connection Manager tears down the DP Paths, causing both of the DP Adapters to enter the Present State.

10.5 High Speed Tunneling

A DP IN Adapter converts DisplayPort Main-Link Symbols into Tunneled Packets. A DP OUT Adapter reconstructs the DisplayPort Main-Link Symbols received in Tunneled Packets.

MST and SST share the following base concepts:

- All Main-Link Symbols, Data and Control, generated by the DP OUT Adapter are identical to the Main-Link Symbols received by the DP IN Adapter over the DP Link.

- The Main-Link Symbol clock generated by the DP OUT Adapter corrects any drift with respect to the DP IN Adapter Main-Link Symbol recovered clock. The total number of cycles over time is identical.
- All Main-Link Symbols are de-serialized, 8b/10 ANSI decoded and de-scrambled by the DP IN Adapter and packed into Tunneled Packets as 8-bit data characters.
- All Main-Link Stuffing Symbols are discarded by the DP IN Adapter and reconstructed by the DP OUT Adapter. Stuffing Symbols consists of:
 - SST – Dummy Symbols during horizontal and vertical Blanking and Stuffing Symbols during Active Video.
 - MST – SF and VCPF Symbol Sequences.
- FEC RS parity symbols are not packed into Tunneled Packets. A DP IN Adapter performs FEC Decoding while a DP OUT Adapter performs FEC Encoding.
- HDCP is supported. Encryption and decryption are not performed.
- A DP IN Adapter minimizes the delay between the time a Main-Link Tunneled Packets is ready to the time it is sent to the Transport Layer.


10.5.1 SST Tunneling

The video and audio data carried over a DisplayPort SST Main-Link is organized into frames. Each frame contains several active scan lines followed by a vertical blanking period. Each scan line contains active pixel data followed by a horizontal blanking period. The active pixel data is organized into fixed-size chunks called Transfer Units. Each Transfer Unit (TU) contains active pixel data followed by Stuffing Symbols. Stuffing Symbols are also sent during the horizontal and vertical blanking periods.

When a DisplayPort SST Main-Link is mapped onto USB4, the continuous Main-Link data stream is encapsulated into Tunneled Packets. Before encapsulation, all the Stuffing Symbols (within a TU and during the blanking periods) are discarded by the DP IN Adapter. Stuffing Symbols are recreated by the DP OUT Adapter when the Main-Link data stream is extracted from the Tunneled packets. In order to enable accurate reconstruction of the Stuffing Symbols at the DP OUT Adapter, the DP IN Adapter includes a *Fill Count* field in each Tunneled Packet. The *Fill Count* field specifies the number of Stuffing Symbols that were discarded immediately preceding the packet.

Figure 10-21 shows how a scan line is encapsulated into Tunneled Packets by a DP IN Adapter.

Figure 10-21. Main-Link SST Stream to Tunneled Packets



Note: DisplayPort Tunneling does not support Default Framing mode.

10.5.1.1 Video Data Packet

10.5.1.1.1 Transfer Unit Set

DisplayPort encodes active pixel data from a scan line into a sequence of Transfer Units. A Transfer Unit (TU) can be 32 to 64 link symbols per lane in length. The length of a TU (TU_Size) is fixed by the DisplayPort source based on the selected display mode. Each TU contains active pixel data followed by zero or more Stuffing Symbols. The active pixel data in a TU can range from 1 byte to TU_Size bytes in length. The Stuffing Symbols in a TU vary as follows:

- When the number of active pixel symbols is less than (TU_Size - 2), Stuffing Symbols are delineated by Fill Start (FS) and Fill End (FE) control symbols.
- When the number of active pixel symbols is equal to (TU_Size - 2), the FS and FE control symbols are the only Stuffing Symbols present.
- When the number of active pixel symbols is equal to (TU_Size - 1), the FE control symbol is the only Stuffing Symbol present.
- When the number of active pixel symbols is equal to TU_Size, no Stuffing Symbols are present.

At the end of a scan line, the leftover pixel data symbols constitute the last TU of the scan line. The last TU of the scan line does not contain any Stuffing Symbols. The leftover pixel data symbols are followed by the Blank Start (BS) control symbol.

A DP IN Adapter shall pack the active pixel data of a TU into either one or two TU Sets. A TU Set is created by selecting a symbol from each lane of the Main-Link in a cyclic way, starting with lane 0. The packing of active pixel data into a TU Set for 4-Lane, 2-Lane and 1-Lane configurations is shown in Figure 10-22, Figure 10-23, and Figure 10-24 respectively.

Note: DisplayPort Tunneling does not support the case of a TU that holds no active pixel data.

Figure 10-22. TU Set Packing for a 4-Lane Main-Link




Figure 10-23. TU Set Packing for a 2-Lane Main-Link






Figure 10-24. TU Set Packing for a 1-Lane Main-Link



When EOC control link symbols are present in a TU, a DP IN Adapter shall pack the EOC symbols using the same scheme as for active pixel symbols. The EOC symbols shall be packed in their 8-bit value representation (DCh). Figure 10-25 shows an example of EOC symbol packing for a 2-lane Main-Link with n active video symbols before the EOC symbol and m active video symbols after it. The Video Count value of $n+m+1$ in this example is padded with 2 bytes to make the TU Set 32-bit aligned.

Figure 10-25. EOC Symbol Packing Example



Each TU Set shall be prepended with a TU Set Header. Figure 10-26 shows the two formats for a TU Set Header.

Figure 10-26. TU Set Header Format

A – No EOC (Default)



B – EOC



The header in Figure 10-26(A) shall be used for a TU Set that contains a TU with no EOC Symbol. The header in Figure 10-26(B) shall be used for a TU Set that contains a TU with an EOC symbol. The fields forming the TU Set Header shall be as defined below:

- **ECC [7:0]:** Error correction field that is calculated over bits [31:8] of the TU Set Header. See Section 5.1.2.3 for calculation.
- **Video Count [13:8]:** This field shall contain the number of active video symbols per lane. A value of zero represents 64. The total number of active video symbols in the TU Set is equal to (Video Count * Number of Lanes).
- **(No EOC) Fill Count [27:14]:** This field shall have the value as defined in 10.5.1.5.
- **(EOC) Fill Count [21:14]:** This field shall have the value as defined in 10.5.1.5.

- **(EOC) EOC Index [27:22]:** This field shall contain the index of the EOC symbol inside the TU Set. Symbols within a TU Set are indexed starting with zero (for the first symbol) and ending with (Video Count – 1) for the last symbol.
- **L [Bit 28]:** Last TU Flag. This flag shall be set to 1b if the TU is the last TU Set of a line. Otherwise shall be set to 0b.
- **RSV [30:29]:** Reserved.
- **EOC [Bit 31]:** This field shall be set to 1b if the TU Set contains an EOC symbol. Otherwise shall be set to 0b.

Note: DisplayPort Tunneling does not support the cases where hblank is bigger than 16384 cycles.

10.5.1.1.2 Packet Format


A Video Data Packet has the *PDF* field in the Tunneled Packet header set to 1. A Video Data Packet shall have the format shown in Figure 10-27. As shown in Figure 10-27, the first 4 bytes of the Video Data Packet payload contain the first TU Set header, followed by the video data payload of the first TU set. Additional TU Sets are appended to the payload of the Video Data Packet. It is recommended that a Video Data Packet include as many TU Sets as possible in order to reduce overhead on the USB4 Fabric.

All TU Set headers within the Video Data Packet payload shall be aligned to 32-bits by adding padding bytes at the end of the TU set payload if necessary. The padding bytes may contain any value, however it is recommended that zero padding be used.

A DP IN Adapter shall follow the rules below when constructing a Video Data Packet:

- A Video Data Packet shall contain at least 1 TU Set and no more than 16 TU Sets. If a TU cannot fit within a single Video Data Packet, it shall be split into two TU sets.
- When a TU is split into two TU Sets, the remainder of the active pixel symbols shall be sent in the first TU Set in the next Video Data Packet. The *Fill Count* field in the TU Set Header of the second Video Data Packet for a TU that is split into multiple Video Data Packets shall be set to 0.
- The length of all TU Set Padding is included when calculating the *Length* field in the Tunneled Packet header.

Figure 10-27. Video Data Packet Format




* SupplID

10.5.1.2 Main Stream Attribute Packet

The Main Stream Attribute Packet has the *PDF* field in the Tunneled Packet header set to 3. The Main Stream Attribute Packet is used by a DP IN Adapter to transmit the display mode attributes once per frame during the vertical blanking period. It is identified by the presence of two Secondary Start (SS) control symbols on each lane of the Main-Link.

A Main Stream Attribute Packet shall consist of an MSA header followed by packet payload. The packet payload shall contain the encoding of the 36-byte attribute information following the <SS, SS> control symbol pair. The MSA Header format is shown in Figure 10-28.


Figure 10-28. MSA Header Format



The fields forming an MSA Header shall be as defined below:

- **ECC [7:0]:** Error correction field that is calculated over bits [31:8] of the MSA Header. See Section 5.1.2.3 for calculation.
- **Fill Count [24:8]:** This field shall contain the fill count as defined in Section 10.5.1.5.
- **Reserved [31:25]:** Reserved.

Figure 10-29. MSA Packet Format



Upon receiving a Main Stream Attribute Packet, a DP OUT Adapter shall do the following:

- Verify the *ECC* field in the MSA header. If an uncorrectable error has occurred, the Main Stream Attribute Packet shall be dropped and ignored.
- Send Fill Count number of Stuffing Symbols on all lanes of the DP Main-Link according to Section 10.5.1.5.
- Send the control symbol pair <SS, SS> on all lanes of the DP Main-Link.
- Send the stream attribute information contained in the first 36 bytes of the payload of the Main Stream Attribute Packet by steering bytes from the payload onto the lanes of the DP Main-Link in a round robin fashion starting with lane 0.
- Send the control symbol <SE> on all lanes of the DP Main-Link.

10.5.1.3 Blank Start Packet

The Blank Start Packet has the *PDF* field in the Tunneled Packet header set to 2. It used by a DP IN Adapter to indicate the beginning of a horizontal or vertical blanking period, reset scrambler

and transfer control information at the beginning of a blanking period. A Blank Start Packet shall consist of a Blank Start header followed by packet payload. The packet payload shall contain the encoding of all 4 sets of <VB-ID, Mvid 7:0, Maud 7:0>. The Blank Start header format is shown in Figure 10-30.

Figure 10-30. Blank Start Header Format



The fields forming the Blank Start header shall be as defined below:


- **ECC [7:0]**: Error correction field that is calculated over bits [31:8] of the Blank Start Header. See Section 5.1.2.3 for calculation.
- **Fill Count [24:8]**: This field shall have the value as defined in Section 10.5.1.5.
- **Reserved [29:25]**: Reserved.
- **CP [30]**: Content Protection Flag shall be set to 1b if Blank Start DP Control Link Symbols sequence were <BS,CP,CP,BS> or <SR,CP,CP,SR>.
- **SR [31]**: Scrambler Reset Flag shall be set to 1b if Blank Start DP Control Link Symbols sequence were <SR,BF,BF,SR> or <SR,CP,CP,SR>.

Figure 10-31 shows the 12-byte encoding for the three different lane configurations. The packet payload is created by interleaving the control information carried on the lanes of the DP Main-Link one symbol at a time starting with lane 0 of the DP Main-Link.

Upon receiving a Blank Start Packet, a DP OUT Adapter shall perform the following operations:

1. Verify the *ECC* field at the Blank Start Header. If an uncorrectable error has occurred, the Blank Start Packet shall be dropped.
2. Generate Stuffing Symbols on each lane of the DP Main-Link according to the *Fill Count* field in the Blank Start Header and Section 10.5.1.5.
3. Generate control symbols on each lane of the DP Main-Link marking the start of the blanking period based on the *SR Flag* and *CP Flag* as shown in Table 10-14.
4. Steer the three double-words of Blank Start Packet payload starting with the first double-word onto the Main-Link by interleaving a byte at a time onto the lanes of the DP Main-Link starting with lane 0.

Figure 10-31. Blank Start Packet Format



* SupplID

Table 10-14. Blank Start Control Link Symbols Mapping

| SR | CP | Blank Start Control Link Symbols Sequence |
|----|----|---|
| 0 | 0 | <BS, BF, BF, BS> |
| 0 | 1 | <BS, CP, CP, BS> |
| 1 | 0 | <SR, BF, BF, SR> |
| 1 | 1 | <SR, CP, CP, SR> |

10.5.1.4 Secondary Data Packet

The DisplayPort Secondary-Data Packet (SDP) is used to carry secondary data embedded within the DisplayPort stream during the horizontal or vertical blanking period. The start of an SDP is identified by the presence of a single SS control symbol on each lane of the DP-Main Link. The end of an SDP is identified by the presence of a single SE control symbol.

A DP IN Adapter shall stop packing secondary data into a Secondary TU upon one of the following cases:

- Secondary End Control symbol <SE> is present on the DP Main-link.
- Maximum capacity is reached:
 - For 1-Lane and 2-Lanes links: maximum capacity is reached when the *Secondary Count* field is 64.
 - For 4-Lane links: maximum capacity is reached when the *Secondary Count* field is 62.
- Secondary data stream was split by MSA, BS or Active video as defined in the DisplayPort 1.4a Specification.

Upon receiving a Secondary Data Packet, a DP OUT Adapter shall do the following for each Secondary TU:


1. Verify the *ECC* field in the Secondary TU Header. If an uncorrectable error has occurred, the whole Secondary TU and the subsequent TUs within that packet shall be dropped and ignored.
2. Send Fill Count number of Stuffing Symbols on all lanes of the DP Main-Link according to Section 10.5.1.5.
3. If this Secondary TU is the first Secondary TU to follow a non-Secondary Data Packet and the *NSS* bit is not set in the Secondary TU Header, send the control symbol <SS> on all lanes of the DP Main-Link.
4. Send the secondary data contained in the Secondary TU.
 - a. The number of cycles of data shall be as according to the *Secondary Count* field in the Secondary TU Header.
 - b. The secondary data shall be sent on the DP Main-Link by steering bytes from the payload onto the lanes in a round robin fashion starting with lane 0.
5. If the *L Flag* is set and the *NSE* bit is not set in the Secondary TU header, send the control symbol <SE> on all lanes of the DP Main-Link.

10.5.1.4.2 Packet Format

A Tunneled Secondary Data Packet has the *PDF* field in the Tunneled Packet header set to 4. A Tunneled Secondary Data Packet shall have the format shown in Figure 10-33. As shown in Figure 10-33, the first 4 bytes of the Tunneled Secondary Data Packet payload contain the first Secondary TU Header, followed by the first Secondary TU payload. Additional Secondary TUs can only be added to the payload of the Tunneled Secondary Data Packet if the whole Secondary TU fits in the payload. A Tunneled Secondary Data Packet shall not include Secondary TUs from more than one DisplayPort SDP. It is recommended that a Tunneled Secondary Data Packet include as many Secondary TUs as possible in order to reduce overhead on the USB4 Fabric.

All Secondary TU Headers within the Tunneled Secondary Data Packet payload shall be aligned to 32-bits by adding Secondary TU Padding bytes at the end of the Secondary TU payload if necessary. Secondary TU Padding bytes values are don't care. The length of all Secondary TU Padding is included when calculating the *Length* field in the Tunneled Packet header.

Figure 10-33. Tunneled Secondary Data Path Format




* SupplD

10.5.1.4.3 Secondary Data to Secondary TU Mapping Examples


Secondary data may be packed into several Secondary TUs depending on the secondary data length, the DP Main-Link lane count, and whether or not the secondary data is split by a non-Secondary Data Packet. Figure 10-34 shows several examples of packing the secondary data into Secondary TUs and the corresponding control bits in the Secondary TU header.

Figure 10-34. Secondary Data to Secondary TUs Examples


A. Secondary Data to Single Secondary TU




B. Secondary Data to three Secondary TUs due to Max Capacity reach



C. Secondary Data to three Secondary TUs due to single Split and Max Capacity reach



D. Secondary Data to three Secondary TUs due to two Splits



10.5.1.5 Fill Count

Every Tunneled Packet or TU that carries data from the SST Main-Link Path, contains a *Fill Count* field. The *Fill Count* field represents the number of DP link clock cycles between two adjacent packets over the DisplayPort Main-Link. The *Fill Count* field is an unsigned integer. Setting all bits to 1 in the *Fill Count* field is a special marking for the minimal value of -1.

Note: The term “previous packet” within this section refers to the previous packet or the previous TU within the same Tunneled Packet.

DP IN Adapter shall calculate the *Fill Count* field according to the following formula:

$$\text{Fill Count field} = \text{Act_Fill_Count} + \text{DP_K_Prev_Pkt} - \text{Prev_Factor}$$

where:

- Act_Fill_Count is the count of DP Link clock cycles between the previous packet and the current packet. The following cycles shall be counted as the Act_Fill_Count:

- Stuffing Symbols.
- BE - Blanking End.
- FS - Filling Start.
- FE - Filling End.
- DP_K_Prev_Pkt is the number of DP Link clock cycles which carried DP Control Link Symbols in the previous packet. DP Control Link Symbols which are counted in the Act_Fill_Count are not to be counted for this attribute.
 - EOC Control Link Symbol is treated as data, and not counted for the DP_K_Prev_Pkt attribute.
- Prev_Factor is an adjustment factor which is based on the previous packet type and sub type as appear in Table 10-15.

Table 10-15. Fill Count Prev_Factor

| Previous Packet | Prev_Factor |
|--------------------------------------|-------------|
| Main Stream | 3 |
| Blank Start | 1 |
| Video TU | 0 |
| Secondary(NSS = 0 ; NSE = 0 ; L = 0) | 0 |
| Secondary(NSS = 0 ; NSE = 0 ; L = 1) | 2 |
| Secondary(NSS = 0 ; NSE = 1 ; L = 0) | NA |
| Secondary(NSS = 0 ; NSE = 1 ; L = 1) | 1 |
| Secondary(NSS = 1 ; NSE = 0 ; L = 0) | -1 |
| Secondary(NSS = 1 ; NSE = 0 ; L = 1) | 1 |
| Secondary(NSS = 1 ; NSE = 1 ; L = 0) | NA |
| Secondary(NSS = 1 ; NSE = 1 ; L = 1) | 0 |


A DP OUT Adapter shall use the following formula to calculate the actual number of Stuffing Symbols to be driven over the DP link:

$$\text{Act_Fill_Count} = \text{Fill Count field} + \text{Prev_Factor} - \text{DP_K_Prev_Pkt}$$


Figure 10-35 and Figure 10-36 show examples of *Fill Count* field calculations.

Figure 10-35. Non-Secondary Data Packet Fill Count Examples


A. Previous Packet - Video



B. Previous Packet - Video



D. Previous Packet - Blank Start



E. Previous Packet - MSA






Figure 10-36. Secondary Data Packet Fill Count Examples


A. Secondary TU: L = 1 ; NSE = 0 ; NSS = 0




B. Secondary TU: L = 0 ; NSE = 0 ; NSS = 0



C. Secondary TU: L = 1 ; NSE = 1 ; NSS = 0



D. Secondary TU: L = 0 ; NSE = 0 ; NSS = 1



A DP IN Adapter may put any value in the *Fill Count* field in the first Tunneled Packet sent on the Main-Link Path after link training. A DP OUT Adapter shall ignore the *Fill Count* field in the first Tunneled Packet sent on the Main-Link Path after DP link training.

10.5.2 MST Tunneling

The MST Link is built out of continually transported Multi Stream Transport Packets (MTP). An MTP is 64 link symbol cycles (i.e. 64 time slots) long. An MTP starts with an MTP Header in the first time slot (i.e. Time Slot 0).

MTPs are broken up into Sub-MTP Transfers Units before being encapsulated in Tunneled Packets. A DP IN Adapter converts each MTP into one or more Sub-MTP TU. A DP OUT Adapter recreates MTPs out of the Sub-MTP TUs.

10.5.2.1 Sub-MTP TU

A Sub-MTP TU holds part of an MTP. A Sub-MTP TU consists of a header and optionally Parameter and/or Data bytes. Figure 10-37 shows the possible Sub-MTP TU structures.

Figure 10-37. Sub-MTP TU Structures


A. Header Only




B. Header and Parameter



C. Header and Data



D. Header, Parameter and Data



A Sub-MTP TU Header shall have the format shown in Figure 10-38. The fields forming the Sub-MTP TU Header shall be as defined below:

- **ECC [7:0]**: Error correction field that is calculated over bits [23:8] of the Sub-MTP TU Header. See Section 5.1.2.3 for calculation.
- **Data Count[13:8]**: This field shall contain the number of Data symbols per lane. The total number of Data symbols in the Sub-MTP TU is equal to (*Data Count* * Number of Lanes).
- **Type[17:14]**: This field shall contain the Type encoding of the Sub-MTP TU. The Type encodings are defined in Table 10-16 for slot zero and in Table 10-17 for non-zero slots.
- **Slot Number[23:18]**: This field shall contain the first slot number in the MTP which this Sub-MTP TU Header is describing.

Figure 10-38. Sub-MTP TU Header Format




Table 10-16. Slot Zero Sub-MTP TU Header Types

| Header Name | Type Encoding [17:14] | Data Count | Valid for # Lanes in DP-Main Link |
|-----------------|-----------------------|------------|-----------------------------------|
| MTPH | 0 | 0-63 | 1, 2, 4 |
| SR MTPH | 1 | 0-63 | 1, 2, 4 |
| 1 Data MTPH | 3 | 0-63 | 1, 2, 4 |
| 2 Data MTPH | 4 | 0-63 | 2 |
| 4 Data MTPH | 5 | 0-63 | 4 |
| 1 K-Symbol MTPH | 8 | 0-63 | 1, 2, 4 |

| Header Name | Type Encoding [17:14] | Data Count | Valid for # Lanes in DP-Main Link |
|-----------------|-----------------------|------------|-----------------------------------|
| 2 K-Symbol MTPH | 9 | 0-63 | 2 |
| 4 K-Symbol MTPH | 11 | 0-63 | 4 |
| Corrupt | 12 | 0 | 1, 2, 4 |

Table 10-17. Non-Slot Zero Sub-MTP TU Header Types

| Header Name | Type Encoding [17:14] | Data Count | Valid for # Lanes in DP-Main Link |
|-------------|-----------------------|------------|-----------------------------------|
| Data | 0 | 1-63 | 1, 2, 4 |
| Shifting SR | 1 | 0-63 | 1, 2, 4 |
| BS | 2 | 0-63 | 1, 2, 4 |
| BE | 3 | 0-63 | 1, 2, 4 |
| SS | 4 | 0-63 | 1, 2, 4 |
| SE | 5 | 0-63 | 1, 2, 4 |
| SF | 6 | 0 | 1, 2, 4 |
| VCPF | 7 | 0 | 1, 2, 4 |
| 1 K-Symbol | 8 | 0-63 | 1 |
| 2 K-Symbol | 9 | 0-63 | 1, 2 |
| 3 K-Symbol | 10 | 0-63 | 1 |
| 4 K-Symbol | 11 | 0-63 | 1, 2, 4 |
| Corrupt | 12 | 0 | 1, 2, 4 |
| EOC | 13 | 0-63 | 1, 2, 4 |
| Unallocated | 15 | 0 | 1, 2, 4 |

A DP IN Adapter that receives an MST stream shall perform the DisplayPort PHY layer and De-scrambler functions defined in the DisplayPort 1.4a Specification.

Upon reception of the first SR after link training completion, a DP IN Adapter shall:

- Track the total number of allocated MST slots by snooping any DPCD AUX transactions that configure the VC Payload ID Table.
- Start mapping MTP from the DP Main-Link into Sub-MTP TUs.

Mapping an MTP to Sub-MTP TUs consists of determining the appropriate header, then appending any Parameter and/or Data bytes. For slot zero, Table 10-18 describes when each Sub-MTP TU header type is used and what Parameter bytes (if any) are included. For non-zero slots, Table 10-19 describes when each Sub-MTP TU header type is used and what Parameter bytes (if any) are included.

Table 10-18. Slot Zero Sub-MTP TU Packet Rules

| Header Name | Parameter Bytes | Parameter Contents | Conditions When Used |
|-------------|-----------------|--------------------|--|
| MTPH | 0 | N/A | Data symbol equal 00h is present on all active lanes. |
| SR MTPH | 0 | N/A | SR Control symbol is present on all active lanes. |
| 1 Data MTPH | 1 | Byte 0 = Dx | The same Data symbol (Dx) which is not equal to 00h, is present on all active lanes. |

| Header Name | Parameter Bytes | Parameter Contents | Conditions When Used |
|-----------------|-----------------|--|--|
| 2 Data MTPH | 2 | Byte 0 = Dx0 Byte 1 = Dx1 | Two different Data symbols, Dx0 and Dx1 are present on lanes 0 and 1 respectively. |
| 4 Data MTPH | 4 | Byte 0 = Dx0 Byte 1 = Dx1 Byte 2 = Dx2 Byte 3 = Dx3 | One or more Data symbols are not the same across all 4 lanes (Dx0, Dx1, Dx2, Dx3 are present on lanes 0, 1, 2 and 3 respectively). |
| 1 K-Symbol MTPH | 1 | lower nibble = x upper nibble = 0h | The same K-Symbol Index (Cx) is present on all active lanes. |
| 2 K-Symbol MTPH | 1 | lower nibble = x0 upper nibble = x1 | Two different K-Symbols Indexes, Cx0 and Cx1 are present on lanes 0 and 1 respectively. |
| 4 K-Symbol MTPH | 2 | Byte 0: lower nibble = x0 upper nibble = x1 Byte 1: lower nibble = x2 upper nibble = x3 | One or more K-Symbol Indexes are not the same across all 4 lanes (Cx0, Cx1, Cx2 and Cx3 are present on lanes 0, 1, 2 and 3 respectively). |
| Corrupt | 0 | N/A | Either Data and K-Symbols are present for the same slot, or an error indication from the DP PHY Layer is present that cannot be resolved using majority voting as defined in the DisplayPort 1.4a Specification. |

Table 10-19. Non- Zero Slot Sub-MTP TU Packet Rules

| Header Name | Parameter Bytes | Parameter Contents | Conditions When Used |
|-------------|-----------------|--|--|
| Data | 0 | N/A | The slot is allocated and has Data Symbols on all active lanes. |
| Shifting SR | 0 | N/A | An SR Control symbol is present on all active lanes for the 4th time as described in Section 10.5.2.2.2. |
| BS | 0 | N/A | The BS Control Link Symbol Sequence C0-C0-C0-C0 is present in full. |
| BE | 0 | N/A | The BE Control Link Symbol Sequence C1-C1-C1-C1 is present in full. |
| SS | 0 | N/A | The SS Control Link Symbol Sequence C3-C3-C3-C3 is present in full. |
| SE | 0 | N/A | The SE Control Link Symbol Sequence C6-C6-C6-C6 is present in full. |
| SF | 0 | N/A | The SF Control Link Symbol Sequence C4-C4-C4-C4 is present in full and the last Sub-MTP TU Header sent was VCPF. |
| VCPF | 0 | N/A | The VCPF Control Link Symbol Sequence C0-C1-C2-C3 is present in full. |
| 1 K-Symbol | 1 | lower nibble = x upper nibble = 0h | A K-Symbol Index (Cx) is followed by either a non-K-Symbol ¹ or slot zero. |
| 2 K-Symbol | 1 | lower nibble = x0 upper nibble = x1 | Two K-Symbols Indexes (Cx0 and Cx1) are followed by either a non K-Symbol ¹ or slot zero. |

| Header Name | Parameter Bytes | Parameter Contents | Conditions When Used |
|--|-----------------|--|--|
| 3 K-Symbol | 2 | Byte 0: lower nibble = x0 upper nibble = x1 Byte 1: lower nibble = x2 upper nibble = 0h | Three K-Symbols Indexes (Cx0, Cx1 and Cx2) are followed by either a non-K-Symbol ¹ or slot zero. |
| 4 K-Symbol | 2 | Byte 0: lower nibble = x0 upper nibble = x1 Byte 1: lower nibble = x2 upper nibble = x3 | Four K-Symbol Indexes (Cx0, Cx1, Cx2 and Cx3) are present that do not match Types 2 to 7 or Type 13. |
| Corrupt | 0 | N/A | Either Data and K-Symbols are present for the same slot or an error indication from the DP PHY Layer is present and cannot be resolved using majority voting as defined in the DisplayPort 1.4a Specification. |
| EOC | 0 | N/A | The EOC Control Link Symbol Sequence (C2-C2-C2) is present in full. |
| Unallocated | 0 | N/A | Slot is the first Unallocated slot for this MTP. |
| Notes: | | | |
| 1. An unallocated slot is considered a "non-K-Symbol". | | | |

When a Parameter includes a Data byte, the DP IN Adapter shall append the de-scrambled Data byte as the Parameter. For a K-Symbol, the nibble describing the index is equal to the de-scrambled Index in case of C0-C7 Symbols and equal to 8 in case of an SR, as described in Table 10-20.

Table 10-20. K-Code Index Nibble in Parameter Byte

| Symbol | Index Nibble |
|--------|--------------|
| C0 | 0 |
| C1 | 1 |
| C2 | 2 |
| C3 | 3 |
| C4 | 4 |
| C5 | 5 |
| C6 | 6 |
| C7 | 7 |

A DP IN Adapter shall pack the Data bytes by selecting a byte from each Lane of the Main-Link in a cyclic way, starting with lane 0. The packing of Data bytes into a Sub-MTP TU for a 4-Lane, 2-Lane and 1-Lane configuration are shown in Figure 10-39, Figure 10-40, and Figure 10-41 respectively. The figures use an MTP with 1 cycle of blank end followed by n slots of video data followed by unallocated slots to demonstrate packing.

A DP IN Adapter shall follow the rules below when constructing a Sub-MTP TU:

- A Sub-MTP TU is byte-aligned.
- The total length of a Sub-MTP TU (Header + Parameter Bytes + Data Bytes) does not exceed 252 Bytes.

- Slot 0 always starts a new Sub-MTP TU.
- A Sub-MTP TU includes data from one MTP only.
- A DP IN Adapter shall map an MTP into the minimum possible number of Sub-MTP TU
- A Sub-MTP TU shall be split into 2 Sub-MTP TUs if it does not fit into an MTP packet according to Section 10.5.2.3.

Figure 10-39. Sub-MTP TU 4-Lane Mapping




Figure 10-40. Sub-MTP TU 2-Lane Mapping





Figure 10-41. Sub-MTP TU 1-Lane Mapping




10.5.2.2 MTP to Sub-MTP TU Examples

10.5.2.2.1 No Allocation

Figure 10-42 shows the 1-lane case where an MTP that has no allocated slots is mapped into two Sub-MTP TUs for both default MTPH and SR MTPH.


Figure 10-42. Unallocated Sequence, 1-Lane



10.5.2.2.2 Shifting SR

The MST Link Timing Robustness section in the DisplayPort 1.4a Specification requires that a DP IN Adapter ignore the appearance of SR at an unexpected time slot. Upon detecting a sequence of four consecutive SR with 2^{16} time-slot intervals, a DP IN Adapter shall switch to the new SR location. The first three SR that are not at Slot Zero's original location shall be mapped to a non-zero Slot Type 8 (1 K-Symbol), for the case of 1-lane, carrying Parameter byte = 8 (as defined in Table 10-19). The fourth SR shall be mapped to a non-zero Slot Type 1 (Shifting SR) forcing the next slot to be slot number 1. Figure 10-43 shows a 3rd and 4th SR in a 1-lane case of unallocated slots. For the case of 2-Lane DP links, the first three SR shall be mapped to Type 9 (2 K-Symbols). For the case of and 4-Lane DP links, the first three SR shall be mapped to Type 11 (4 K-Symbols).


Figure 10-43. Shifting SR, 1-Lane



10.5.2.2.3 ACT

Figure 10-44 shows the ACT sequence for a 1-lane case, where originally none of the slots are allocated.

Figure 10-44. ACT Sequence, 1-Lane



10.5.2.2.4 SF and VCPF

A DP IN Adapter shall not map a SF sequence into a Sub-MTP TU unless the SF sequence comes immediately after a VCPF sequence. Figure 10-45A shows the case where a SF sequence is present after Data symbols, which results in no SF Sub-MTP TU. Figure 10-45B shows the case where a VCPF sequence is present after Data symbols, which results in a one-time VCPF Sub-MTP TU. Figure 10-45C shows the case where a VCPF sequence is followed by SF sequence, which results a one-time VCPF Sub-MTP TU followed by one-time SF Sub-MTP TU.

Figure 10-45. SF and VCPF Sequence 4-Lane

A. SF only results No SF Sub-MTP TU

| Slots | 0 | 1 | 2 | 3 | .. | n | n+1 | .. | 63 |
|--------|----|-----------------|-----------------|----|-----|----|-----|-----|----|
| Lane 0 | 00 | V ₁₀ | V ₂₀ | C4 | ... | C4 | | ... | |
| Lane 1 | 00 | V ₁₁ | V ₂₁ | C4 | ... | C4 | | ... | |
| Lane 2 | 00 | V ₁₂ | V ₂₂ | C4 | ... | C4 | | ... | |
| Lane 3 | 00 | V ₁₃ | V ₂₃ | C4 | ... | C4 | | ... | |

| | |
|------------------------|-----------------|
| Slot=0;Type=0; DC=2 | V ₁₀ |
| V ₁₁ | V ₁₂ |
| V ₁₃ | V ₂₀ |
| V ₂₁ | V ₂₂ |
| V ₂₃ | |

| |
|---------------------------|
| Slot=n+1; Type=15;DC=0 |
|---------------------------|

B. VCPF only results VCPF Sub-MTP TU

| Slots | 0 | 1 | 2 | 3 | .. | n | n+1 | .. | 63 |
|--------|----|-----------------|-----------------|----|-----|----|-----|-----|----|
| Lane 0 | 00 | V ₁₀ | V ₂₀ | C0 | ... | C0 | | ... | |
| Lane 1 | 00 | V ₁₁ | V ₂₁ | C1 | ... | C1 | | ... | |
| Lane 2 | 00 | V ₁₂ | V ₂₂ | C2 | ... | C2 | | ... | |
| Lane 3 | 00 | V ₁₃ | V ₂₃ | C3 | ... | C3 | | ... | |

| | |
|------------------------|-----------------|
| Slot=0;Type=0; DC=2 | V ₁₀ |
| V ₁₁ | V ₁₂ |
| V ₁₃ | V ₂₀ |
| V ₂₁ | V ₂₂ |
| V ₂₃ | |

| |
|------------------------|
| Slot=3;Type=7; DC=0 |
|------------------------|

| |
|---------------------------|
| Slot=n+1; Type=15;DC=0 |
|---------------------------|

C. VCPF followed by SF – VCPF and SF Sub-MTP TUs

| Slots | 0 | 1 | 2 | 3 | .. | m | m+1 | .. | n | n+1 | .. | 63 |
|--------|----|-----------------|-----------------|----|-----|----|-----|-----|----|-----|-----|----|
| Lane 0 | 00 | V ₁₀ | V ₂₀ | C0 | ... | C0 | C4 | ... | C4 | | ... | |
| Lane 1 | 00 | V ₁₁ | V ₂₁ | C1 | ... | C1 | C4 | ... | C4 | | ... | |
| Lane 2 | 00 | V ₁₂ | V ₂₂ | C2 | ... | C2 | C4 | ... | C4 | | ... | |
| Lane 3 | 00 | V ₁₃ | V ₂₃ | C3 | ... | C3 | C4 | ... | C4 | | ... | |

| | |
|------------------------|-----------------|
| Slot=0;Type=0; DC=2 | V ₁₀ |
| V ₁₁ | V ₁₂ |
| V ₁₃ | V ₂₀ |
| V ₂₁ | V ₂₂ |
| V ₂₃ | |

| |
|------------------------|
| Slot=3;Type=7; DC=0 |
|------------------------|

| |
|--------------------------|
| Slot=m+1; Type=6;DC=0 |
|--------------------------|

| |
|---------------------------|
| Slot=n+1; Type=15;DC=0 |
|---------------------------|


10.5.2.3 MST Packet Format

A DP IN Adapter uses an MST Packet to transport one or more Sub-MTP TUs. An MST Packet has the *PDF* field in the Tunneled Packet header set to 6. An MST Packet shall have the format shown in Figure 10-46.

A DP IN Adapter shall follow the rules below when constructing an MST Packet:

- The first 3 bytes of the MST Packet payload contains the first Sub-MTP TU Header.
- When concatenating two Sub-MTP TUs, the first TU is not be padded.
- The maximum number of MTPs packed into one MST Packet does not exceed 17.
- The *Length* field in the Tunneled Packet Header does not include padding bytes.
- The Payload of the Tunneled Packet shall be between 230 and 252 Bytes (inclusive), unless the Payload contains 17 MTPs.

Figure 10-46. MST Packet Format



* SuppID

10.5.2.4 MST Packets to DP MTP

A DP OUT Adapter receiving MST Packets from the USB4 Fabric recreates the DP MTP stream based on the Sub-MTP TUs included in the MST Packets and the slot number. A DP OUT Adapter is unaware of the number of allocated slots and each Sub-MTP TU parsing and handling is independent of the previous one. A DP OUT Adapter shall analyze each Sub-MTP TU Header and recreate the MTP K-Code and data bytes according to Table 10-16 (for slot zero) or Table 10-17 (for non-zero slots). If a DP OUT Adapter has a slot for which it lacks sufficient information to recreate the MTP K-Code and/or data bytes, it shall follow the rules below:

- If the last Sub-MTP TU Header was VCPF, insert VCPF.
- If the last Sub-MTP TU Header was Unallocated, insert unallocated (data bytes equal zero).
- For all other cases, insert SF.

After creating the MTPs, the DP OUT Adapter shall follow all the PHY Layer functions required functions by the DisplayPort 1.4a Specification.

10.5.3 FEC

This section describes the methods by which a DP Adapter supports FEC functionality while ensuring that all Main-Link Symbols, Data and Control, generated by the DP OUT Adapter are identical to the Main-Link Symbols received by the DP IN Adapter over the DP Link.

10.5.3.1 SR Count

A DP Adapter shall implement the SR Count counter, which counts the number of cycles that have elapsed since the last SR. A DP IN Adapter uses this counter to measure and report to the DP OUT Adapter, the number of cycles from the last SR it received to the first cycle of the FEC_DECODE_EN or FEC_DECODE_DIS sequences. A DP OUT Adapter uses this counter to count the number of cycles which elapsed since the last SR it transmitted in order to recreate the

FEC_DECODE_EN or FEC_DECODE_DIS sequences at the same cycle, and after the same SR, as the DP IN Adapter has received.

A DP IN Adapter shall initiate the SR Count at the first cycle after receiving an SR. A DP OUT Adapter shall initiate the SR Count at the first cycle after transmitting an SR. For SST, the count starts after receiving or transmitting all four Enhanced Framing Mode Symbols. SR Count counts any link clock cycle, including cycles which carry FEC-related symbols.

10.5.3.2 DP IN Adapter Requirements

A DP IN Adapter shall:

- Implement FEC Decoding as defined in the DisplayPort 1.4a Specification.
- Construct an FEC_DECODE Packet as defined in Section 10.5.3.4 upon FEC_DECODE_EN or FEC_DECODE_DIS sequence detection.
- The Adapter Layer shall prioritize the FEC_DECODE Packet over all other Main-Link Path packets when pass it to the Transport Layer. If the Adapter Layer is constructing a Main-Link Path Packet when the FEC_DECODE_EN or FEC_DECODE_DIS sequence is received over the DP Main-Link, the Adapter Layer may first finish constructing the Main-Link Path Packet and pass it to the Transport layer before passing the FEC_DECODE Packet.

A DP IN Adapter shall not:

- Tunnel any FEC-related symbols including FEC_PARITY_MARKER, FEC_DECODE & FEC_PARITY_PH.
- Count the Link cycles of FEC Symbols for fill count purposes.

10.5.3.3 DP OUT Adapter Requirements

A DP OUT Adapter shall:

- Implement FEC Encoding as defined in the DisplayPort 1.4a Specification.
- Apply majority voting for the repeated fields with in the FEC_DECODE Packet.
 - SR Count.
 - FEN.
 - FDS.
- Generate FEC_DECODE_EN and FEC_DECODE_DIS upon reception of a FEC_DECODE Packet.
 - FEC_DECODE_EN sequence shall be generated if *FEN* field in the FEC_DECODE Packet is 1b.
 - FEC_DECODE_DIS sequence shall be generated if *FDS* fields in the FEC_DECODE Packet is 1b.
 - The first symbol of the FEC_DECODE_EN and FEC_DECODE_DIS sequences shall be transmitted according to the SR Count field of the FEC_DECODE Packet, i.e. the FEC_DECODE_EN and FEC_DECODE_DIS sequence shall be transmitted *SR Count* link clock cycles after the most recently transmitted SR.
 - When a FEC_DECODE Packet is received, a DP OUT Adapter compares the Packet SR Count and the Counter SR Count as follows:
 - If Packet SR Count > Counter SR Count the DP OUT Adapter waits for the Counter SR Count to be equal to Packet SR Count then generate the FEC sequence.
 - Else, a DP OUT Adapter waits for next SR to be transmitted, then waits for the Counter SR Count to be equal to Packet SR Count then generate the FEC sequence.


A DP OUT Adapter shall not:

- Count the Link cycles of FEC Symbols for fill count purposes.

10.5.3.4 FEC_DECODE Packet

A FEC_DECODE Packet is sent over the Main-Link Path from DP IN Adapter to DP OUT Adapter. A FEC_DECODE Packet shall have the format shown in Figure 10-47.


Figure 10-47. FEC_DECODE Packet Format



* SupplD

The *PDF* field in the header shall be set to 7 and the *Length* field shall be 14h. An FEC Command shall have the format defined in Figure 10-48. The three FEC Commands in an FEC_DECODE Packet shall be identical to each other.

Figure 10-48. FEC Command Format



The fields in an FEC Command shall contain the following:

- **SR Count [29:0]:** This field contains the number of DP Link clock cycles between the last received SR and the first FEC_DECODE_EN or FEC_DECODE_DIS sequences. The minimum value for this field is 1h (occurs when SR is immediately followed by FEC_DECODE sequence).
- **FEC DISABLE (FDS) [30]:** This field shall be set to 1b if a FEC_DECODE_DIS sequence was detected. In all other cases it shall be set to 0b.
- **FEC ENABLE (FEN) [31]:** This field shall be set to 1b if a FEC_DECODE_EN sequence was detected. In all other cases it shall be set to 0b.

10.5.4 DP OUT Adapter Buffer

In order to reconstruct DP Main-Link traffic without any interruptions, a DP OUT Adapter shall implement a buffer that can be used to compensate for the jitter in the latency of the received Tunneled Packets.

After the DP link is successfully trained, a DP OUT Adapter shall start generating idle pattern as described in the DisplayPort 1.4a Specification. The DP OUT Adapter transitions from sending self-generated idle pattern to reconstructing the DP Main-Link from the Tunneled Packets after completing the following steps:


1. The DP OUT Adapter shall adjust the PLL frequency at least once as a result of an Adjust PLL event as described in Section 10.6.
2. The DP OUT Adapter ensures, in an implementation specific manner, that within eight PLL frequency adjustments the link symbol clock frequency difference between its own and the DPTX is such that buffer overflow and buffer underrun is avoided.
3. The DP OUT Adapter shall wait to receive an SR from the DP IN Adapter.
 - All Main-Link Path Tunneled Packets, besides DP Clock Sync Packets, are dropped by the DP OUT Adapter until reception of the SR.
 - After the next step is completed, the received SR, shall be the first symbol driven by the DP OUT Adapter as the reconstructed Main-link.
4. The DP OUT Adapter delays sending the SR in Step 3 for a number of Accumulation Cycles (see Section 10.5.4.2).
 - During the delay, the DP OUT Adapter shall accumulate the DP Main-Link traffic from the DP IN Adapter.

10.5.4.1 Buffer Operation

During Active Video periods, the buffer is filled to some extent due to the Accumulation Cycles that the DP OUT Adapter waits when switching from self-generated idle pattern to reconstruction of DP Main-Link.

During Blanking periods and idle pattern, where the maximal distance between two BS could be 8K cycles, the buffer might become empty. If the buffer becomes empty, a DP OUT Adapter shall continue to drive Dummy Symbols on the DP Main Link, assuming that the next Main-Link Path Tunneled Packet holds a Fill Count value larger than the driven Dummy Cycles. Figure 10-49 shows the process of switching from Active Video to Blanking.

Figure 10-49. Active Video to Blanking



10.5.4.2 Accumulation Cycles

Accumulation Cycles are counted in Link Symbol clock cycles, where one Accumulation Cycle equals one Link Symbol clock cycle. The number of Accumulation Cycles that a DP OUT Adapter waits is a function of the following:

- DP IN Jitter – the Jitter between Tunneled Packets from the DP IN Adapter.
 - The delay from when a DP construct is received over the DP Main-Link until a DP IN Adapter generates the Main-Link Tunnel Packet can vary where:
 - Minimal delay is for MSA and BS.
 - Maximal delay is:

- MST: $17 \text{ MTPs} + \text{Max FEC Cycles}^1 = 1088 + 37 = 1125 \text{ LS Cycles}$.
- SST: $16 \text{ Video TUs} + \text{Max FEC Cycles} = 1024 + 37 = 1061 \text{ LS Cycles}$.
- USB4 Delay Jitter – the Jitter cause by delays over the USB4 Fabric.
 - The following can create delays at the Egress Adapter for each Hop:
 - Time Sync Packets.
 - Credit Grant Packets.
 - Control Packets.
 - Other Tunneled Packets (on both same priority and lower priority Paths).
 - Packet forwarding jitter within a Router.
 - The delay depends on the number Credit Grant Records, the number of additional DP Paths, the USB4 Link speed, and the assumptions regarding how many delay factors on average affect each Hop.

A DP OUT Adapter shall report the Maximum Accumulation Cycles it performs. The maximum is needed when operating at the maximum DisplayPort Link Rate.

The following is an example of a DP OUT Adapter that supports the maximum DP Link Rate of HBR3, supports MST, and assumes that a DP Tunneled Packet is delayed at every Hop by two full size Packets:

- Worst case Accumulating Cycles = (DP IN Jitter) + (USB4 Delay Jitter)
 - = $(17\text{MTPs} + \text{Max FEC Cycles}) + (((5 \text{ Hops} * 2 * 260 \text{ Bytes} * 8) * \text{USB4 UI}) / (10 * \text{HBR3 UI}) + (6 \text{ Hops} * t_{\text{TunneledPacketJitter}} / (10 * \text{HBR3 UI})))$
 - = $(1088 + 37) + ((20800 \text{ Bits} * 0.1\text{ns} / 1.23\text{ns}) + (6 * 100\text{ns} / 1.23\text{ns})) = 3305 \text{ Link Symbol Cycles}$

10.5.5 HDCP

High-bandwidth Digital Content Protection (HDCP) is performed in the DisplayPort Link Layer. A DisplayPort tunnel is transparent to HDCP. A DP Adapter is not aware of whether or not HDCP is used and does not participate in the HDCP process.

A DP IN Adapter shall not perform HDCP decryption. It shall not drop or modify an AUX Request or AUX Response associated with HDCP functionality.

A DP OUT Adapter shall not perform HDCP encryption.

Note: When HDCP is active, the Stuffing Symbols generated by a DP OUT Adapter do not have the same values as generated by the DPTX.

10.6 DP Link Clock Sync

A DisplayPort Sink device recovers the link symbol clock from the Main-Link and uses the M and N values embedded within the video stream to derive the stream clock (pixel clock) from the recovered link symbol clock. When a DisplayPort link is tunneled across a USB4 Fabric, the DP IN and DP OUT Adapters are in different clock domains, Link Symbol clock synchronization is needed so that:

- The DisplayPort Sink is able to recreate the pixel clock.

¹ Max FEC Cycles = $3x \text{FEC_PARITY_PH} + 1x \text{FEC_PM} = 3x12 + 1x1 = 37 \text{ Cycles}$

- Symbols received by a DP IN Adapter are identical to the symbols sent by the DP OUT Adapter.

This section describes the method by which the DP OUT Adapter generates the DP-Main Link Symbol clock at the same frequency as the DP IN Adapter recovers it.

10.6.1 Synchronization Method

Clock Synchronization starts as soon as the Link Symbol clock is stable. It is achieved by periodically following the steps below:

1. Both the DP IN Adapter and DP OUT Adapter measure the number of link symbol clocks within a synchronized measurement window.
2. The DP IN Adapter sends its measurement to the DP OUT Adapter.
3. The DP OUT Adapter computes the difference between the frequencies and adjusts its Link symbol clock frequency.

10.6.1.1 Events

10.6.1.1.1 Measuring Events

The measuring events at a DP Adapter are derived from the Host Router time, which guarantees high accuracy, synchronized events. The two measuring events used for the DP Main-Link Symbol Clock Synchronization scheme are:


- Measure Window Event – occurs every 2^{21} ns (i.e. when bit 21 of the *Nanosecond* field of the Host Router time changes polarity).
- Update Counter Event – occurs every 2^{12} ns (i.e. when bit 12 of the *Nanosecond* field of the Host Router time changes polarity).

10.6.1.1.2 Adjust PLL Event

An Adjust PLL event is used by a DP OUT Adapter to initiate the PLL frequency adjustment. An Adjust PLL Event occurs when there have been 128 Update Counter Events since the last Measure Window Event.

Figure 10-50 illustrates when Adjust PLL Events occur.

Figure 10-50: Adjust PLL Event Occurrence



10.6.1.2 Lifetime Counter

The Lifetime Counter (LC) is a free running 64-bit counter that advances in proportion to the DP Main-Link Symbol clock. It increases by 1 for each Symbol clock when DP Main-Link operates in RBR or HBR, and advances by 1 for every 2 Symbol clocks when DP Main-Link operates in HBR2 or HBR3. The LC is reset when DP Main-Link is down and starts counting when DP Main-Link Symbol clock is stable.

A DP OUT Adapter shall start counting as soon as Link Symbol clock is stable for starting link training with DPRX.

A DP IN Adapter shall start counting as soon as it completed its equalization process.

In order to filter out the variation introduced by spread-spectrum modulation, the LC shall be filtered using a first order IIR filter. Therefore, the value of the LC is sampled upon an Update Counter Event, converted to fractional representation, and then used for calculating the Filtered Lifetime Counter (FLC), according to formula:

$$\begin{aligned}
 LC_Frac[n] &= \{LC[n], 00000000b\} \\
 FLC[n] &= LC_Frac[n] \gg 5 + FLC[n-1] - FLC[n-1] \gg 5 ; n > 0 \\
 FLC[n] &= LC_Frac[n] \gg 5 ; n = 0
 \end{aligned}$$

where:

LC[n] - is the current sample of LC.

LC_Frac[n] - is the current sample of LC represented in fractional format.

FLC[n-1] - is the previous output of the filter.


FLC[n] - is the new output of the filter.

The filtering operation shall be done with 8-bit truncation at the fraction part to assure reproducible result. All operands of IIR filter shall have the same format of 64 bits of integer followed by 8 bits of fraction.


The *Window Count* field, LC Counter, LC_Frac operand, and FLC Counter structures are shown in Figure 10-51. An illustration of the computation logic is shown in Figure 10-52.

Figure 10-51. Lifetime Counter Format


A. Window Count



B. Lifetime Counter (LC)



C. Lifetime Counter – Fraction representation (LC_Frac)



D. Filtered Lifetime Counter (FLC)




Figure 10-52. Filtered Lifetime Counter Logic Concept




Table 10-21 shows an example of five FLC calculations, starting from $n=0$.

Table 10-21. FLC Calculation Examples

| n | LC | FLC |
|---|--------|---------|
| 0 | 0xBC | 0x5E0 |
| 1 | 0x50E | 0x2E21 |
| 2 | 0x95F | 0x77A8 |
| 3 | 0xDB1 | 0xE173 |
| 4 | 0x1203 | 0x16A80 |

10.6.1.3 DP Clock Sync Packet

A DP Clock Sync Packet has the *PDF* field in the Tunneled Packet header set to 5. It is sent over the Main-Link Path from a DP IN Adapter to a DP OUT Adapter. A DP Clock Sync Packet shall have the format shown in Figure 10-53.

Figure 10-53. DP Clock Sync Packet Format

| | | | | | |
|-----|--------------|---|--------|--------------|-----|
| | 31 | | | | 0 |
| DW0 | PDF = 5 | * | HOP ID | Length = 20h | HEC |
| DW1 | Reserved | | | | |
| DW2 | Reserved | | | | |
| DW3 | Window Count | | | | |
| DW4 | FLC [31:0] | | | | |
| DW5 | FLC [63:32] | | | | |
| DW6 | Reserved | | | FLC [71:64] | |
| DW7 | Reserved | | | | |
| DW8 | CRC32 | | | | |

* SupplID

The fields forming a DP Clock Sync Packet shall be as defined below:

- **Reserved:** This field is reserved and shall be set to 0.
- **Window Count:** This field is defined in Section 10.6.2.1. The *Window Count* field structure is shown in Figure 10-51(A).
- **FLC:** This field contains the snapshot of the Filtered Lifetime Counter at the time the Window Measured Event occurred.
- **CRC32:** This field contains a CRC32 computed over the entire payload using the following DW order: DW1, DW3, DW2, DW7, DW6, DW5, DW4. The following CRC shall be used:
 - Width: 32
 - Poly: 1EDC 6F41h
 - Init: FFFF FFFFh
 - RefIn: True
 - RefOut: True
 - XorOut: FFFF FFFFh

Figure 10-54 shows an example of a DP Clock Sync Packet.

Figure 10-54. DP Clock Sync Packet Example

| | | | | |
|--------------|-----|--------|--------------|-----|
| 31 | 0 | | | |
| PDF = 5 | * | HOP ID | Length = 20h | HEC |
| 00_00_00_00h | | | | |
| 00_00_00_00h | | | | |
| 09h | 73h | 2Fh | C6h | |
| CAh | FDh | 40h | 02h | |
| 34h | 56h | 78h | 9Ah | |
| 00_00_00h | | | | 12h |
| 00_00_00_00h | | | | |
| 47h | 64h | 79h | 0Ah | |

* SupplD

10.6.2 DP Adapter Requirements**10.6.2.1 DP IN Adapter Requirements**

A DP IN Adapter shall:

- Implement a Lifetime Counter as described in Section 10.6.1.1.2.
- Implement the logic to perform the LC filtering.
- Update FLC upon an Update Counter Event.
- Upon the first Measure Window Event:
 - Capture the current FLC.
 - Store the current captured FLC to be used as previous captured FLC at the next Measure Window Even.
- Upon each subsequent Measure Window Event:
 - Capture the current FLC.
 - Compute the Window Count by calculating the current captured FLC minus the FLC that was captured at the previous Measure Event.
 - Construct a DP Clock Sync Packet and send it over the Main-Link Path within tDPClockSync after the Measure Window Event.
 - Store the current captured FLC to be used as previous captured FLC at the next Measure Window Event.

Note: A DP IN Adapter may send a DP Clock Sync Packet before DP Link Training is completed.

10.6.2.2 DP OUT Adapter Requirements

A DP OUT Adapter shall:

- Implement a Lifetime Counter as described in Section 10.6.1.2.
- Implement the logic to perform the LC filtering.

- Update FLC upon an Update Counter Event.
- Upon a Measure Window Event:
 - Capture the current FLC.
 - Compute the Window Count as described in Section 10.6.1.3.
- Upon receiving a DP Clock Sync Packet after the Measure Window Event and before the PLL Adjust Event, compute the PLL frequency adjustment. The method for computing the PLL frequency adjustment is outside the scope of this specification.
- Upon an Adjust PLL Event, adjust the PLL frequency based on the computation performed at the Measure Window Event.

If a DP OUT Adapter receives a DP Clock Sync Packet after an Adjust PLL Event but before the next Measure Window Event, it shall not adjust the PLL and shall silently discard the packet.

If a DP OUT Adapter receives a DP Clock Sync Packet before it computed the first Window Count, it shall not adjust the PLL and shall silently discard the packet.

When a DP OUT Adapter changes the DisplayPort Main-Link transmitter frequency as a result of adjusting the PLL frequency, it shall adhere to the DisplayPort 1.4a Specification. The PLL frequency adjustment shall be completed within tDPPLLAdjust after the Adjust PLL Event.



IMPLEMENTATION NOTE

A DP OUT Adapter needs to take measures to ensure that the number of Link Symbol clock cycles between it and the DP IN Adapter don't drift over time. This includes the following:

- *The Lifetime Counters in the DP IN Adapter and the DP OUT Adapter start counting at different points in time. This causes an initial offset between the FLC values in the DP IN and DP OUT Adapter, which a DP OUT Adapter can calculate at the first Measure Window Event where it has FLC values from the DP IN Adapter (as provided in a DP Clock Sync Packet). In subsequent PLL frequency adjustments, a DP OUT Adapter should maintain a distance equal to the initial offset when adjusting the PLL frequency.*
- *A DP OUT Adapter uses two data points when computing the PLL frequency adjustment: Window Count, and FLC. Including the FLC as a data point, helps eliminate drift.*

10.7 DP BW Allocation Mode

A DP IN Adapter may optionally support DP BW Allocation Mode. If a DP IN Adapter supports DP BW Allocation Mode, it shall do so as defined in this section.

When DP BW Allocation Mode is enabled (see Section 10.7.1), the DPTX and the Connection Manager communicate through the DP IN Adapter. The interaction between a DP IN Adapter and a DPTX is described in Section 10.7.2. The interaction between a DP IN Adapter and a Connection Manager is described in Section 10.7.3.

10.7.1 DP BW Allocation Mode Enablement

If a DP IN Adapter supports DP BW Allocation Mode, it shall:

- Set the DP_LOCAL_CAP.DP_IN_BW_Allocation Mode Support bit to 1b.
- Update the AUX Response for DPCD DP TUNNELING and PANEL REPLAY OPTIMIZATION SUPPORT. DP_IN_BW_Allocation_Mode_Support (E000Dh, bit 7) to 1b.
- Update the AUX Response for DPCD USB4_DRIVER_BW_CAPABILITY.USB4_Driver_BW_Allocation_Mode_Support (E0020h, bit 7) to have the same value as ADP_DP_CS_2.CM BW Allocation Mode Support.

When a Connection Manager changes ADP_DP_CS_2.CM BW Allocation Mode Support bit, a DP IN Adapter shall:

- Set the *BW_Allocation_Capability_Changed* field to 1b in DP_TUNNELING_STATUS DPCD register.
- Set DP_TUNNELING_IRQ bit (Bit 5 of LINK_SERVICE_IRQ_VECTOR_ESI0 register at DPCD 02005h)
- If Unmask_BW_Allocation_IRQ is 1b, generate an IRQ_HPD.

DP BW Allocation Mode enablement is initiated by the DPTX only if both, the DP IN Adapter and the Connection Manager, support the DP BW Allocation Mode.

When DPTX sets DPTX_BW_ALLOCATION_MODE_CONTROL.DP_Display_Driver_BW_Allocation_Mode_Enable (*E0030h, bit 7*) to 1b, a DP IN Adapter shall:

- Enable DP BW Allocation Mode
- Send the Connection Manager a Notification Packet with Event Code = DP_BW as defined in Table 6-11.

10.7.2 Interaction with DPTX

Table 10-22 lists the fields in the DPCD registers that are used by DPTX to negotiate bandwidth allocation with a Connection Manager. Table 10-23 describes how a DP IN Adapter maps the DPCD fields to Adapter Configuration Space.

Table 10-22. DPCD Bandwidth Allocation Registers

| DPCD Register | Register Address | Bits | DPCD Field Type | DPCD Field Name |
|----------------------------------|------------------|-------|--|---|
| DP_TUNNELING_CAPABILITIES | E000Dh | [7] | Read Only | DP_IN_BW_Allocation_Mode_Support |
| USB4_DRIVER_ID | E000Fh | [3:0] | Read Only | USB4_Driver_ID |
| USB4_DRIVER_BW_CAPABILITY | E0020h | [7] | Read Only | USB4_Driver_BW_Allocation_Mode_Support |
| DP_IN_ADAPTER_TUNNEL_INFORMATION | E0021h | [2:0] | Read Only | Group_ID |
| DP_BW_GRANULARITY | E0022h | [1:0] | Read Only | Granularity |
| ESTIMATED_BW | E0023h | [7:0] | Read Only | Estimated BW |
| ALLOCATED_BW | E0024h | [7:0] | Read Only | Allocated BW |
| DP_TUNNELING_STATUS | E0025h | [0] | Clearable, Read Only (Bit is cleared when 1 is written by way of an AUX_CH write transaction) | BW Request Failed |
| | | [1] | | BW Request Succeeded |
| | | [2] | | Estimated BW Changed |
| | | [3] | | BW_Allocation_Capability_Changed |
| DP_TUNNELING_MAX_LINK_RATE | E0028h | [7:0] | Read Only | DP_Tunneling_Max_Link_Rate |
| DP_TUNNELING_MAX_LANE_COUNT | E0029h | [7:0] | Read Only | DP_Tunneling_Max_Lane_Count |
| DPTX_BW_ALLOCATION_MODE_CONTROL | E0030h | [6] | Read/Write | Unmask_BW_Allocation_IRQ |
| | | [7] | Read/Write | DP_Display_Driver_BW_Allocation_Mode_Enable |
| REQUESTED_BW | E0031h | [7:0] | Read/Write | Requested BW |

Table 10-23. DP IN Adapter Configuration Space Mapping

| DPCD Field Name | DPCD Field Type | DP IN Adapter Register | DP IN Field |
|---|-----------------|------------------------|---------------------------------------|
| USB4_Driver_BW Allocation Mode Support | Read Only | ADP_DP_CS_2 | <i>CM BW Allocation Mode Support</i> |
| USB4_Driver_ID | Read Only | ADP_DP_CS_2 | <i>CM_ID</i> |
| Group_ID | Read Only | ADP_DP_CS_2 | <i>Group_ID</i> |
| Granularity | Read Only | ADP_DP_CS_2 | <i>Granularity</i> |
| Estimated BW | Read Only | ADP_DP_CS_2 | <i>Estimated BW</i> |
| Allocated BW | Read Only | DP_STATUS | <i>Allocated BW</i> |
| DP_Display_Driver_BW_Allocation_Mode_Enable | Read/Write | ADP_DP_CS_8 | <i>DPTX BW Allocation Mode Enable</i> |
| Requested BW | Read/Write | ADP_DP_CS_8 | <i>Requested BW</i> |

When a DP IN Adapter receives a DPCD AUX Write transaction that targets a DPCD register within Table 10-23, and the targeted DPCD field Type is Read/Write, it shall update the corresponding field in Adapter configuration space with the value of the write transaction.

When a DP IN Adapter receives a DPCD AUX Read transaction that targets a DPCD register within Table 10-23, and the targeted DPCD field Type is Read Only, it shall update the read transaction with the value in the corresponding field in Adapter Configuration Space.

When DPTX sends a DPCD AUX write transaction that targets the REQUESTED_BW register, a DP IN Adapter shall:

- Store the current Allocated BW in an internal variable.
- If the recovery timer is advancing, stop and reset it.
- Initiate a bandwidth request handshake with the Connection Manager as defined in Section 10.7.3.

When a Connection Manager writes a value to the Allocated BW field that is equal to or greater than the Requested BW, a DP IN Adapter shall:

- Set the BW Request Succeeded field to 1b in DP_TUNNELING_STATUS DPCD register.
- Set DP_TUNNELING_IRQ bit (Bit 5 of LINK_SERVICE_IRQ_VECTOR_ESI0 register at DPCD 02005h).
- If Unmask_BW_Allocation_IRQ is 1b, generate an IRQ_HPD.
- If the ESTIMATED_BW field was locked for updates due to bandwidth request failure, unlock it.

When a Connection Manager writes a value to the Allocated BW field that is smaller than the Requested BW, a DP IN Adapter shall:


- Set the ESTIMATED_BW field to the Allocated BW, and lock its value (i.e. ignore any changes in the Estimated_BW field).
- Set the BW Request Failed to 1b in DP_TUNNELING_STATUS DPCD register.
- Set DP_TUNNELING_IRQ bit (Bit 5 of LINK_SERVICE_IRQ_VECTOR_ESI0 register at DPCD 02005h).

- If Unmask_BW_Allocation_IRQ is 1b, generate an IRQ_HPD.
- Start the recovery timer.
- If the recovery timer has reached tDPBWRecoveryTimeout, the DP IN Adapter initiates a bandwidth allocation request. The DP BW that a DP IN Adapter requests shall be the same bandwidth as before the failed bandwidth allocation (i.e. the same value as in the Allocated BW field before the DPTX last updated the DPCD Requested BW register).

Note: When a DPTX writes to the REQUESTED_BW DPCD register, the DP IN Adapter should store the current Allocated BW in an internal variable. Then, in case of a tDPBWRecoveryTimeout timeout, the DP IN Adapter knows what Requested BW value to request.

A DP IN Adapter interacts with DPTX according to Figure 10-55.

Figure 10-55: DP IN Adapter Interaction with DPTX During DP BW Allocation



10.7.2.1 Estimated Bandwidth

Whenever the bandwidth allocation changes, the Connection Manager updates the available bandwidth by writing to ADP_DP_CS_2.Estimated BW field. Note that the Estimated BW field includes the bandwidth that is already allocated to the DP IN Adapter.

Upon a change in the Estimated BW field, a DP IN adapter shall:

- Set Estimated BW Changed bit in the DPCD DP_TUNNELING_STATUS register to 1b.
- Set DP_TUNNELING_IRQ bit (Bit 5 of LINK_SERVICE_IRQ_VECTOR_ESI0 register at DPCD 02005h).
- If Unmask_BW_Allocation_IRQ is 1b, generate an IRQ_HPD.



CONNECTION MANAGER NOTE

When there is a change in bandwidth across its Domain, the Connection Manager recalculates the estimated bandwidth for each DP IN Adapter in the Domain. The estimated bandwidth for a DP IN Adapter includes:

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

After the estimated bandwidth is recalculated, a Connection Manager updates the ADP_DP_CS_2.Estimated BW field.

10.7.3 Interaction with the Connection Manager

When DPTX sends a DPCD AUX write transaction that targets the REQUESTED_BW field, a DP IN Adapter shall:

1. Set the ADP_DP_CS_8.DPTX Req field to 1b.
2. Send the Connection Manager a Notification Packet with Event Code = DP_BW as defined in Table 6-11.
3. Wait for the Connection Manager to set the ADP_DP_CS_2.CM Ack bit to 1b.
4. Set the ADP_DP_CS_8.DPTX Req field to 0b.




CONNECTION MANAGER NOTE

When a Connection Manager receives a DP_BW Notification Packet, it completes the following steps within 100ms:

1. *Reads ADP_DP_CS_8.Requested BW.*
2. *Allocates bandwidth:*
 - a. *If the Request bandwidth is less than or equal to the currently allocated bandwidth, then the new allocation will be equal to the new requested bandwidth.*
 - *If the DP IN Adapter has a Group_ID different than zero, then for the following tDPBWRebalance, the Connection Manager will reserve the released bandwidth to be allocated to DP IN Adapters with the same Group_ID.*
 - b. *Else, the Connection Manager tries to allocate bandwidth from the available bandwidth and by freeing bandwidth from other clients.*
3. *Writes the new allocated bandwidth to the DP_STATUS.Allocated BW field.*
4. *Sets the ADP_DP_CS_2.CM Ack bit to 1b.*
5. *Waits for the DP IN Adapter to set the ADP_DP_CS_8.DPTX Req bit to 0b.*
6. *Sets the ADP_DP_CS_2.CM Ack bit to 0b.*

A DP IN Adapter interacts with a Connection Manager according to Figure 10-56.

Figure 10-56: DP BW Allocation Interaction with Connection Manager




10.8 Timing Parameters

Table 10-24 lists the timing parameters for a DP Adapter.

Table 10-24. DP Adapter Timing Parameters

| Parameter | Description | Min | Max | Units |
|----------------------|---|-----|-----|-------|
| tDPAckResponse | Maximum time between receiving a SET_CONFIG or HPD Packet and sending the corresponding ACK. | -- | 4 | μs |
| tDPAckTimeout | Maximum time after sending a SET_CONFIG or HPD Packet that a DP Adapter waits for the corresponding ACK Packet. | -- | 400 | μs |
| tDPSetConfigGap | Minimum time between receiving an ACK Packet and sending a SET_CONFIG Packet. | 50 | -- | μs |
| tDPPlug | Maximum time a Router can wait after detection of a Source/Sink Connect/Disconnect before sending the corresponding Hot Plug/Hot Removal Event. Maximum time a DP OUT Adapter can wait after transitioning to the Paired state before sending an HPD Packet. | -- | 150 | μs |
| tDPInit | Maximum time between receiving a SET_CONFIG Packet as part of DP Adapters Init Flow and responding with a SET_CONFIG Packet. | -- | 1 | ms |
| tDPPLLAdjust | Maximum time between detecting an Adjust PLL Event and completing the PLL adjustment. | -- | 100 | μs |
| tDPClockSync | Maximum time to send a DP Clock Sync Packet. | -- | 100 | μs |
| tIRQDelay | The time a DP OUT Adapter waits after Link training is complete, before it sends SET_CONFIG IRQ. | 200 | 300 | μs |
| tDPBWRecoveryTimeout | The time a DP IN Adapter waits for DPTX to request bandwidth in case of an allocation failure. | 10 | 10 | sec |
| tDPBWRebalance | The time a Connection Manager reserves released bandwidth to be used by other DP IN Adapters with the same Group_ID. | 10 | 10 | sec |

11 PCI Express Tunneling



USB4™ Hosts:

A USB4 host may optionally support PCIe Tunneling. A USB4 host that tunnels PCIe traffic shall do one of the following:

- Incorporate an internal PCIe Switch.
- Connect to a Root Complex via Root Ports.
- Connect to a Root Complex via other means that meet the PCIe Specification.

The Host Router in a USB4 host that supports PCIe Tunneling shall have one Downstream PCIe Adapter per Downstream Facing Port. Each Downstream PCIe Adapter shall interface to a downstream port of the internal PCIe Switch or Root Complex.

USB4 Hubs:

A USB4 hub shall support PCIe Tunneling. A USB4 hub shall contain an internal PCIe Switch.

The Device Router in a USB4 hub shall have:

- An Upstream PCIe Adapter that interfaces to the internal PCIe Switch.
- For each Downstream Facing Port, a Downstream PCIe Adapter that interfaces to a downstream port of the internal PCIe Switch.

USB4 Devices:

A USB4 device may optionally support PCIe Tunneling. A USB4 device that tunnels PCIe traffic shall contain either an internal PCIe Switch or an internal PCIe Endpoint.

The Device Router in a USB4 device that supports PCIe tunneling shall have an Upstream PCIe Adapter that interfaces to the internal PCIe Switch or Endpoint.

Note: In this Chapter, the term “internal PCIe Port” refers to either an internal PCIe Switch Port, an internal PCIe Endpoint upstream port, or a PCIe root complex downstream port.

A native PCIe device within a USB4 hub or USB4 device shall implement LTR.

11.1 PCIe Adapter Layer

11.1.1 Encapsulation

The PCIe Adapter Layer shall encapsulate the following PCIe constructs in Tunneled Packets:

- Transaction Layer Packets (TLP).
- Data Link Layer Packets (DLLP).
- Ordered Sets.
- Out-of-band events.

A PCIe Adapter Layer shall not encapsulate Idle data Symbols into Tunneled Packets.

A PCIe Adapter layer shall follow the rules below when encapsulating a PCIe construct into a Tunneled Packet:

- A TLP or a DLLP shall fit into a single Tunneled Packet.
- A Tunneled Packet may contain a single TLP or DLLP, or it may contain a DLLP followed by a TLP or DLLP. It shall not contain any other TLP/DLLP combinations.
- Each Ordered Set shall be encapsulated into a separate Tunneled Packet.
- Each PCIe out-of-band event shall be encapsulated into a separate Tunneled Packet.
- The order of bytes and bits in the Tunneled Packet shall be identical to the original PCIe construct. The least-significant byte of the encapsulated construct is mapped to B0 in the Tunneled Packet Payload (see Figure 4-14). Bit 7 in each byte of the encapsulated construct is mapped to bit 7 in respective byte of the Tunneled Packet Payload.

The *PDF* field in a Tunneled Packet identifies the type of PCIe construct contained therein. Table 11-1 defines the PDF values that shall be used for each type of PCIe construct.

Table 11-1. PDF Values for PCIe Tunneled Packets

| PDF | Type | Contents of Tunneled Packet |
|---------|-----------------------|---|
| 0h | Rsvd | Reserved. |
| 1h | Ordered Set | One PCIe EIOS, TS1, or TS2 Ordered Set. |
| 2h | Electrical Idle State | Indication of an Electrical Idle state on PCIe. |
| 3h | TLP/DLLP | PCIe TLPs and/or DLLPs. |
| 4h | Rsvd | Reserved. |
| 5h | PERST Active | PCIe Reset (PERST) in active state. |
| 6h | PERST Inactive | PCIe Reset (PERST) in inactive state. |
| 7h – Fh | Rsvd | Reserved. |

If a PCIe Adapter receives a Tunneled Packet with a Rsvd PDF value, it shall discard the Tunneled Packet and shall not send any Packets in response.

11.1.1.1 PCIe TLP and DLLP

A PCIe Adapter shall not discard a PCIe TLP or DLLP due to lack of credits in the USB4 Fabric. When credits are not available, the Router shall queue the PCIe TLP and shall transport it once sufficient credits become available. When credits are not available, a Router may hold off creation of new DLLPs.

When a PCIe Adapter Layer receives a Tunneled Packet with a pre-header that does not match any of the pre-headers defined in this section, it shall discard the packet and report the mismatch as a PCIe Receiver Error to the internal PCIe Port.



IMPLEMENTATION NOTE

The amount of buffering at the PCIe Adapter is implementation specific as it balances the tradeoff between PCIe tunneling performance and PCIe link latency. It is recommended that implementations make the amount of buffers configurable.

11.1.1.1.1 TLP Encapsulation

A PCIe Adapter shall not send a Tunneled Packet for the following:

- A Nullified TLP.
- A Truncated TLP.
- A Truncated DLLP.

As shown in Figure 11-1, before encapsulation into a Tunneled Packet, a PCIe Adapter Layer shall:

1. Truncate a TLP by removing the STP Symbol, four leading Reserved bits, and the END Symbol.
2. Add the TLP pre-header defined in Table 11-2.

TLP bytes are mapped into the payload of a Tunneled Packet as shown in Figure 11-1. The TLP pre-header is mapped into byte B0 starting with the *Type* field. The truncated TLP immediately follows the TLP pre-header.

Figure 11-1. Tunneled PCIe TLP




Table 11-2. TLP Pre-Header

| Field | Size | Description |
|-------|--------|---|
| TYPE | 4 bit | Identifies the PCIe construct. Set to 1111b for TLP. |
| LEN | 11 bit | Length of TLP starting after the TLP <i>Sequence Number</i> field. Counted in DWs with a fix offset of 4. Value of 0 represents the minimum case of 4 DWs (TLP header size of 3 DW plus 1 DW of LCRC). Value of 38 represents the maximum case of 42 DWs (32 DW of TLP, plus 4 DW of maximal TLP header, plus 1 DW of LCRC, plus 1 DW of TLP Digest, plus 4 DW of End-to-End TLP Prefix). |
| CHK | 5 bit | Error detection code on the vector, $V[14:0] = \{TYPE, LEN\}$. CHK[3:0] are calculated as follows: <ul style="list-style-type: none"> CHK[0] = $V[14] \wedge V[13] \wedge V[12] \wedge V[10] \wedge V[8] \wedge V[7] \wedge V[4] \wedge V[0]$ CHK[1] = $V[12] \wedge V[11] \wedge V[10] \wedge V[9] \wedge V[7] \wedge V[5] \wedge V[4] \wedge V[1]$ CHK[2] = $V[13] \wedge V[12] \wedge V[11] \wedge V[10] \wedge V[8] \wedge V[6] \wedge V[5] \wedge V[2]$ CHK[3] = $V[14] \wedge V[13] \wedge V[12] \wedge V[11] \wedge V[9] \wedge V[7] \wedge V[6] \wedge V[3]$ CHK[4] = XORing of [14:0] bits. A PCIe Adapter Layer that receives a TLP Tunneled Packet with a mismatch in this field shall report the mismatch as a PCIe Receiver Error and the TLP shall be discarded. |

Figure 11-2 shows an example of a single TLP, carrying the PTM ResponseD message, encapsulated into a Tunneled Packet.

Figure 11-2. Tunneled PTM Example

| 31 | | | | | | 0 | |
|-----------------------------|---|-----------------------------|--|---------------------------------|--|-----------------------------|--|
| PDF = 3 | * | HOP ID | | Length = 1Ch | | HEC | |
| {TYPE, LEN[10:7]} = F0h | | {LEN[6:0], CHK[4]} = 05h | | {CHK[3:0], TLP Seq[11:8]} = 70h | | TLP Seq[7:0] = 1Bh | |
| Header 0 = 74h | | Header 1 = 00h | | Header 2 = 00h | | Header 3 = 01h | |
| Header 4 = 56h | | Header 5 = 20h | | Header 6 = 00h | | Header 7 = 53h | |
| TMU_to_PT_M_B [63:56] = F2h | | TMU_to_PT_M_B [55:48] = 44h | | TMU_to_PT_M_B [47:40] = E9h | | TMU_to_PT_M_B [39:32] = 09h | |
| TMU_to_PT_M_B [31:24] = FDh | | TMU_to_PT_M_B [23:16] = 97h | | TMU_to_PT_M_B [15:8] = 7Eh | | TMU_to_PT_M_B [7:0] = 03h | |
| TMU_to_PT_M_A [7:0] = 5Dh | | TMU_to_PT_M_A [15:8] = F5h | | TMU_to_PT_M_A [23:16] = 1Ah | | TMU_to_PT_M_A [31:24] = 80h | |
| LCRC[31:24] = EBh | | LCRC[23:16] = 5Dh | | LCRC[15:8] = E2h | | LCRC[7:0] = 5Ah | |

* SupplD


11.1.1.1.2 DLLP Encapsulation

Before encapsulating a DLLP into a Tunneled Packet, a PCIe Adapter Layer shall:

1. Truncate a DLLP by removing the SDP Symbol and the END Symbol.
2. Add a DLLP pre-header of 0FACH to a DLLP as shown in Figure 11-3.

DLLP bytes are mapped into the payload of the Tunneled Packet as in Figure 11-3. The DLLP pre-header is mapped into bytes B0 and B1. The truncated DLLP immediately follows the DLLP pre-header.

Figure 11-3. Tunneled PCIe DLLP




11.1.1.1.3 Mixed DLLP/TLP Encapsulation


When a Tunneled Packet contains a DLLP followed by a TLP or DLLP, each DLLP or TLP shall contain its own pre-header.

Figure 11-4. PCIe DLLP and TLP Tunneled Packet Payload

DLLP followed by DLLP



DLLP followed by TLP



11.1.1.2 PCIe Ordered Sets

A PCIe Adapter Layer that receives a TS Ordered Set (TS1 or TS2) or an Electrical Idle Ordered Set (EIOS) from the internal PCIe Port shall transfer the Ordered Set to the Transport Layer as defined in the sections below. All other Ordered Sets shall be discarded and shall not be transferred to the Transport Layer.

11.1.1.2.1 Training Sequence (TS) Ordered Sets

Transmitter rules:

- When a PCIe Adapter Layer receives one or more identical TS1 Ordered sets in a row from the internal PCIe Port, it shall transmit 16 copies of the TS1 Ordered set before it transmits another Tunneled PCIe Packet. The PCIe Adapter Layer may transmit more than 16 TS1 Ordered sets before transmitting a Tunneled PCIe Packet.

Note: If the Internal PCIe Port sends more than 16 TS1 Ordered Sets, the PCIe Adapter Layer is not required to send more than 16 TS1 Ordered Sets.

- When a PCIe Adapter Layer receives one or more identical TS2 Ordered sets in a row from the internal PCIe Port, it shall transmit 16 copies of the TS2 Ordered set before it transmits another Tunneled PCIe Packet. The PCIe Adapter Layer may transmit more than 16 TS2 Ordered sets before transmitting a Tunneled PCIe Packet.

Note: If the Internal PCIe Port sends more than 16 TS2 Ordered Sets, the PCIe Adapter Layer is not required to send more than 16 TS2 Ordered Sets.

- A PCIe Adapter Layer shall only transmit TS Ordered Sets that target Lane 0. TS Ordered Sets that target a non-zero Lane shall be discarded and shall not be send over the USB4 Fabric.
- A PCIe Adapter Layer shall modify a TS Ordered Set as defined in Table 11-3.

Table 11-3. TS Ordered Sets

| PCIe Symbol Number | USB4 Byte Number | TS1 Contents | TS2 Contents |
|--------------------|------------------|--|--|
| 0 | 0 | A PCIe Adapter Layer shall replace the COM Symbol with the value BCh. | Same as for TS1. |
| 1 | 1 | Link Number A PCIe Adapter Layer shall replace the PAD symbol with the value F7h. All other values are unchanged. | Same as for TS1. |
| 2 | 2 | Lane Number The PCIe Adapter Layer shall replace a PAD Symbol with the value F7h. Lane Number = 0 is unchanged. | Same as for TS1. |
| 3 | 3 | N_FTS The byte value is unchanged by the PCIe Adapter Layer. | Same as for TS1. |
| 4 | 4 | Data Rate Identifier The byte value is unchanged by the PCIe Adapter Layer. | Same as for TS1. |
| 5 | 5 | Training Control The byte value is unchanged by the PCIe Adapter Layer. Note that loopback request (Bit 2) is not supported. | Same as for TS1. |
| 15:6 | 15:6 | A PCIe Adapter Layer shall replace the D10.2 Symbol with the value 4Ah. | A PCIe Adapter Layer shall replace the D5.2 Symbol with the value 45h. |

Receiver rules:

Upon receiving a TS Ordered Set from the Transport Layer, a PCIe Adapter Layer shall transfer the TS Ordered Set to the internal PCIe Port.

**IMPLEMENTATION NOTE**

A PCIe Adapter Layer may send back-to-back identical copies of the TS Ordered Set to the PCIe Physical Layer Logical sub-block until either a different Transport Layer Packet is received or the LTSSM in the PCIe Physical Layer Logical sub-block enters a Configuration.Idle or Recovery.Idle state.

11.1.1.2.2 Electrical Idle Ordered Sets (EIOS)

When a PCIe Adapter Layer receives an Electrical Idle Ordered Set (EIOS) from the internal PCIe Port, it shall send two consecutive EIOS Tunneled Packets over the USB4 Fabric. An EIOS shall be encapsulated in a Tunneled Packet with payload as defined in Table 11-4.

Table 11-4. Electrical Idle Ordered Sets

| Symbol | Contents |
|--------|--------------------------------|
| 0 | COM Symbol with the value BCh. |
| 1 – 3 | IDL Symbol with the value 7Ch. |

A PCIe Adapter Layer that receives an EIOS Tunneled Packet shall:

1. Transfer the EIOS to the internal PCIe Port.
2. Indicate Rx Electrical Idle to the internal PCIe Port.

11.1.1.3 Electrical Idle State

Upon detecting that a PCIe Physical Layer Logical sub-block is in Electrical Idle state, a PCIe Adapter Layer shall send at least 3 Electrical Idle State Tunneled Packets. The payload for an Electrical Idle State Tunneled Packet shall consist of one DW that contains a value of 0000 0000h.

A PCIe Adapter Layer that receives an Electrical Idle State Tunneled Packet shall indicate Rx Electrical Idle to the internal PCIe Port.

When a PCIe Adapter Layer receives a Tunneled Packet that is not an Electrical Idle State Tunneled Packet or an EIOS Tunneled Packet, it shall stop indicating Rx Electrical Idle to the internal PCIe Port.

11.1.1.4 PERST

11.1.1.4.1 PERST Tunneled Packets

PERST Active and PERST Inactive Tunneled Packets are used to propagate PCIe PERST# assertion and de-assertion. A Router shall send PERST Active and PERST Inactive packets only from a Downstream PCIe Adapter. The payload for a PERST Active Tunneled Packet and a PERST Inactive Tunneled Packet shall consist of one DW that contains a value of 0000 0000h.

11.1.1.4.2 PERST Propagation

By default, a Device Router asserts PERST# on all physical PCIe ports and internal PCIe ports.

11.1.1.4.2.1 PERST Activation

Upon detecting an assertion of PERST#, a Host Router shall:

1. Discard any queued Tunneled Packet in the PCIe Adapter Layer.
2. Send at least three PERST Active Tunneled Packets on all Downstream PCIe Adapters that have the *Path Enable* bit set to 1b.

When a Device Router receives a PERST Active Tunneled Packet it shall:

1. Discard any Tunneled Packets that are queued in a PCIe Adapter Layer.
2. Send at least 3 PERST Active Tunneled Packets on all Downstream PCIe Adapters that have the *Path Enable* bit set to 1b. The Router shall also assert PERST# on all physical PCIe ports and to the internal PCIe Port.

While PERST# is asserted, a Downstream PCIe Adapter Layer shall discard any received PCIe Tunneled Packets. The Adapter Layer shall not send any PCIe tunneled Packets except for the PERST Active Tunneled Packets.

11.1.1.4.2.2 PERST Inactivation

Upon detecting a de-assertion of PERST#, a Host Router shall send at least 3 PERST Inactive Tunneled Packets on all Downstream PCIe Adapters that have the *Path Enable* bit set to 1b.

After receiving a PERST Active Tunneled Packet, if a Device Router receives any PCIe Tunneled Packet other than a PERST Active Tunneled Packet on its Upstream PCIe Adapter, it shall:

1. Send at least 3 PERST Inactive Tunneled Packets on all Downstream PCIe Adapters that have the *Path Enable* bit set to 1b.
2. De-assert PERST# on all PCIe physical ports and to the internal PCIe Port.

After PERST# is de-asserted, a Downstream PCIe Adapter Layer shall not forward to the internal PCIe Port any Ordered Sets, packets, or events that were received before or during PERST# assertion.

11.1.2 USB4 Hot-Plug

When a Device Router is hot-plugged, the Upstream PCIe Adapter in the hot-plugged Router shall maintain the internal PCIe Port in PCIe Warm Reset. The Device Router shall maintain the Warm Reset until either a PERST Inactive Tunneled Packet or a PCIe TS Ordered Set is received by the Upstream PCIe Adapter.

11.2 Internal PCIe Ports

This section defines the functionality of an internal PCIe port that interfaces to a PCIe Adapter. Each internal PCIe port that interfaces to a PCIe Adapter shall implement a Physical Layer Logical sub-block, a Data Link Layer, and a Transaction Layer as defined in the PCIe Specification with the modifications, configurations, and parameters described in this section.

The Logical sub-block only supports PCIe Gen 1. The PCIe link width can be any value supported by the PCIe Specification and may be different at the Upstream Adapter and Downstream Adapter.

11.2.1 PCIe Physical Layer Logical Sub-block

The Logical sub-block shall update the PCIe configuration registers with the following characteristics:

- PCIe Gen 1 protocol behavior.
- *Max Link Speed* field in the Link Capabilities Register set to 0001b (data rate of 2.5 GT/s only).

Note: These settings do not represent actual throughput. Throughput is implementation dependent and based on the USB4 Fabric performance.

11.2.1.1 Encoding

The Physical Layer Logical sub-block shall:

- Not scramble the bits it delivers to the PCIe Adapter Layer.
- Not de-scramble the bits it receives from the PCIe Adapter Layer, regardless of the Disable Scrambling bit received in the TS Ordered Set.

11.2.1.2 Link Training and Status State Machine (LTSSM)

The Physical Layer Logical sub-block shall implement a Link Training and Status State Machine (LTSSM). The LTSSM shall support the L1 state. The LTSSM shall not support the following:

- Loopback state.
- L0s state.
- L1 PM substates.
- Changes in PCIe link speed.
- Lane-to-Lane de-skew.
- Inferring Electrical Idle in states other than L0.
- Inferring Electrical Idle in L0 state by absence of Flow Control Update DLLPs is optional.

Note: Negotiation and Reduction of PCIe link width is not applicable as the only TS which traverse over the tunnel are TS for Lane 0.

The following changes shall also apply to the LTSSM:

- A PCIe Upstream port in Recovery.idle shall transition to L0 state when it receives a TLP or DLLP. If the port does not receive any TLP or DLLP, it shall transition to the L0 state tRecovery time after it entered the Recovery.idle state.

Note: The tRecovery wait period is needed to detect a possible transition to Disabled state or to Hot Reset State.

A PCIe downstream port, shall immediately send UpdateFC after transitioning from L1 to L0. It is recommended that a PCIe port transmit at least 16 TS1 after receiving the first TS1 during Recovery.RcvrLock.



IMPLEMENTATION NOTE

Idle data Symbols are not sent over the USB4 Fabric. Therefore, a receiving internal PCIe port cannot depend on the reception of Idle data Symbols in its LTSSM. One possible implementation is that the PCIe Adapter Layer generates the Idle data Symbols towards the internal PCIe port before transitioning to L0 state. Another possible implementation is for the LTSSM to proceed from either Configuration.Idle state or Recovery.Idle state to L0 state without reception of any Idle data Symbols. Note that TS2 Ordered Sets received after the transition to L0 state are ignored until a Tunneled Packet containing a construct other than a TS2 Ordered Set is received.

11.2.1.3 ASPM L1 Entry

The following rules shall be implemented in order to avoid ambiguity in ASPM L1 handshake termination when the PCIe downstream port responds with the PM_Active_State_Nak message:

- The PCIe upstream port shall send no more than 10 additional PM_Active_State_Request DLLPs after sending an ACK DLLP for the received PM_Active_State_Nak message.
- After receiving the ACK DLLP for the PM_Active_State_Nak message, the PCIe downstream port shall wait 9.5 microseconds as described in the PCIe Specification and shall drop the first 10 PM_Active_State_Request DLLPs it receives. If it receives a TLP or DLLP which is not a PM_Active_State_Request DLLP, it may not drop the next PM_Active_State_Request DLLPs.

11.2.1.4 Clock Tolerance Compensation

Clock tolerance compensation shall not be performed.

11.2.1.5 Compliance Mode

Compliance Mode shall not be supported.

11.2.1.6 Clock Power Management

Clock Power Management shall not be supported.

11.2.1.7 L2 State

A Router does not support propagation of PCIe WAKE unless it is in the Sleep state. It is up to system software to ensure that the PCIe link only enters the L2 state when a Router is in the Sleep state and the PCIe device is enabled for wake from L2 state.

11.2.2 PCIe Data Link Layer

An internal PCIe Port shall implement a PCIe Data Link Layer as defined in the PCIe Specification.

11.2.3 PCIe Transaction Layer

The Transaction Layer shall support:

- A Max Payload Size of 128B.
- The PCIe Latency Tolerance Reporting (LTR) mechanism. The Latency Tolerance Reporting LTR Capability shall be implemented in the PCIe upstream port.

- PCIe hot-add and hot-removal (“hot-plug”).

In a USB4 Hub, the Transaction Layer shall additionally support:

- Access Control Services (ACS) as described in the PCIe Specification with ACS P2P Egress Control required. Note that support for ACS P2P Egress Control is optional in the PCIe Specification.
- Flattening Portal Bridge (FPB).

It is recommended that a Transaction Layer also support Downstream Port Containment (DPC) Extended Capability.

The Transaction Layer shall not support local TLP Prefix and Protocol MUXing.

A USB4 host needs to be hardened against malicious devices and malformed requests. The Transaction Layer in an Internal PCIe Port in a USB4 host, in conjunction with the System Software, needs to be able to provide appropriate protection against requests from rogue endpoints. The mechanism to provide such protection is implementation specific, but System Software needs certain functionality to be provided by the hardware.

In a USB4 Host, the Transaction Layer shall additionally provide functionality to:

- Ensure that a transaction received on the PCIe Root Port is appropriate for the Requester ID. This can be done using ACS Source Validation or by an implementation-specific mechanism that is more appropriate for the architecture of the Host.
- Allow blocking of Upstream Memory Requests with the AT field set to a value other than the default. This can be done using ACS Translation Blocking or by an implementation-specific mechanism that is more appropriate for the architecture of the Host. When Upstream Memory Requests are not blocked the following applies:
 - A Host supporting ATS Translation Services shall check and filter any request marked “Translated” to ensure that the transaction is from the provided SourceID and the PCIe Function is permitted to access the specific address range needed by the transaction.
 - A Host that is not supporting ATS and does not block Upstream Memory Requests delivered with non-default AT field values, may use a proprietary method for implementing external address translation, with equivalent (or improved) functionality to ATS, as long as it can authenticate the source, and guarantee that it is privileged for address translation.
 - A Host that is neither supporting ATS nor an equivalent scheme shall never agree to process Upstream Memory Requests without qualifying the source and translating the address accordingly.
- Provide per-source address filtering (and optionally translation). This can be done with an IOMMU or other implementation-specific structures. It is recommended that a USB4 Host use a mechanism that allows System Software to revoke mappings with low performance overhead.

It is recommended that the USB4 Host also provide:

- Flattening Portal Bridge or equivalent to improve the scalability and runtime reallocation of Routing ID and Memory Space resources.

11.2.4 PCIe Link Timers (Informative)

As the delay through PCIe tunnel may vary, it is recommended that a Router implement PCIe link timer ranges listed in Table 11-5.

Table 11-5. PCIe Link Timer Ranges


| Timer | Min | Max | PCIe Specification Reference |
|--------------|----------|----------|---|
| UpdateFC | 30 [μs] | 100 [μs] | 2.6.1.2 FC Information Tracked by Receiver |
| Retrain Link | 200 [μs] | 1 [ms] | 2.6.1.2 FC Information Tracked by Receiver |
| Replay | 10 [μs] | 256 [μs] | 3.5.2.1 LCRC and Sequence Number Rules (TLP Transmitter) |
| L1 Reject | 9.5 [μs] | 128 [μs] | 5.4.1.2.1. Entry into the L1 State See the PCIe Specification, Implementation Note titled <i>ASPM L1 Accept/Reject Considerations for the Upstream Component</i> |

11.2.5 Precision Time Measurement (PTM) Mechanism

A USB4 Hub shall support PTM as defined in the PCIe Specification with the modifications in this section. It is recommended that a USB4 Host also support PTM. If a USB4 Host supports PTM, it shall do so as defined in this section.

The modified PTM link protocol runs between a PCIe Adapter Layer and an Internal PCIe Port within a Router. When PCIe traffic is tunneled through more than one Router, the PTM link protocol runs on each Router that the PCIe traffic traverses. Physical PCIe Ports perform the PCIe PTM link protocol without any modifications.

Figure 11-5: Example of PTM Relationships



A PTM link protocol that runs across a USB4 Fabric needs to be adjusted for the additional latency added by the Fabric. This is done by calculating a mathematical relationship between the PTM Master Time and the TMU Host Router Time at one Router, then using that relationship to reconstruct the PTM Master Time from TMU Host Router Time samples at downstream Device Routers.

There are two parameters that characterize the relationship between the PTM Master Time and the TMU Host Router Time. These parameters are called $TMU_to_PTM_A$ and $TMU_to_PTM_B$. Together, they are referred to as the TMU_to_PTM Parameters.

There is one Router that calculates the TMU_to_PTM Parameters. This Router is referred to as the Parameter Generator. A Router that reconstructs PTM Master Time is referred to as a Parameter Consumer. The TMU_to_PTM Parameters are propagated throughout a USB4 Fabric via Responded messages.

Note: USB4 buffering and arbitration methods can introduce random and varying delays in a USB4 Fabric. Furthermore, the propagation delay between PTM messages is not symmetrical. Because of this the PTM link protocol described in the PCIe specification does not give sufficient accuracy without modification.

11.2.5.1 Parameter Generator

The following Routers shall act as a Parameter Generator:

- A Host Router with its PTM function enabled.
- A Device Router with its PTM function enabled and the *Root Select* bit in its PTM Control Register set to 1b.

A Parameter Generator shall calculate the TMU_to_PTM Parameters as defined in Section 11.2.5.3.1. It is recommended that a Parameter Generator calculate the TMU_to_PTM Parameters periodically, similar to PTM periodic calculation. Periodic calculations increase accuracy.

Instead of the *PTM Master Time* and *Propagation Delay* fields defined in the PCIe Specification, a Parameter Generator shall include the most recent TMU_to_PTM Parameters in the PTM ResponseD Message as depicted in Figure 11-6.

Figure 11-6: PTM ResponseD Message




Figure 11-2 shows an example of how a PTM ResponseD Message is encapsulated into a Tunneler Packet.

If the *Time Disruption* bit is set to 1b in Router Configuration Space, a Parameter Generator shall not reply with a PTM ResponseD Message to a PTM Request.

11.2.5.2 Parameter Consumer

A Device Router shall act as Parameter Consumer if its PTM function is enabled and the *Root Select* bit in its PTM Control Register is set to 0b.

The Parameter Consumer shall not reply with a PTM ResponseD Message to a PTM Request if any of the following conditions are true:

- The *Time Disruption* bit is set to 1b in Router Configuration Space.
- TMU_to_PTM Parameters were not received after the *Time Disruption* bit was cleared.
- A rule in the PCIe Specification forbids sending a PTM ResponseD Message.

When sending a PTM ResponseD Message through a Downstream PCIe Adapter, a Parameter Consumer shall use the PTM ResponseD Message format defined in Figure 11-6. The TMU_to_PTM Parameters in the ResponseD Message shall be the same as the last TMU_to_PTM Parameters received on the Upstream PCIe Adapter.

When sending a PTM ResponseD Message to a Native Downstream PCIe Port, a Parameter Consumer shall use the PTM ResponseD Message format defined in the PCIe Specification. The Parameter Consumer shall calculate the PTM Master Time for the PTM ResponseD Message as defined in Section 11.2.5.3.2.

11.2.5.3 PTM Calculations

11.2.5.3.1 TMU_to_PTM Parameters

The TMU_to_PTM Parameters represent the linear relationship between the PTM Master Time and the TMU Host Router Time. TMU_to_PTM_A is a 32-bit number that represents the slope of the linear line. TMU_to_PTM_B is a 64-bit number that represents the PTM Master Time intercept of the linear line. The TMU_to_PTM Parameters are formatted as follows:

- TMU_to_PTM_A
 - TMU_to_PTM_A is a 32-bit fixed-point value, with the binary point between bit 31 and bit 30.
 - TMU_to_PTM_A[31] is 0b for values in the range (0,1) (excluding 0 and 1) and is 1b for values in the range [1,2) (including 1 and excluding 2).
 - TMU_to_PTM_A[30:0] is a binary representation of the fraction part of TMU_to_PTM_A.
 - For example, the value 0.5 is represented by 400...0h and the value 1.5 is represented by C00...0h.
- TMU_to_PTM_B
 - TMU_to_PTM_B[63:0] value is in nanoseconds and is in 2's complement notation

The TMU_to_PTM Parameters are calculated from two samples of PTM Master Time and TMU Host Router Time ((tmu0, ptm0) and (tmu1, ptm1)). In general the TMU_to_PTM Parameters can be calculated by solving the following set of equations:

$$ptm0 = TMU_to_PTM_A * tmu0 + TMU_to_PTM_B$$

$$ptm1 = TMU_to_PTM_A * tmu1 + TMU_to_PTM_B$$

However, quantization error can occur when using the equations above because the TMU_to_PTM_A is only 32 bits wide. To compensate for any quantization error, a Parameter Generator shall calculate the TMU_to_PTM Parameters as follows:

1. Calculate the ideal 64-bit parameters:
 - $A_ideal[63:0] = (ptm1 - ptm0) / (tmu1 - tmu0)$
 - where A_ideal [63:0] is a 64-bit fixed-point value with the binary point between bit 63 and bit 62.
 - $B_ideal[63:0] = ptm1 - A_ideal[63:0] * tmu1$
 - where B_ideal is a 64-bit signed integer in 2's complement notation.
2. Calculate the ptm_error at ptm1 point using the 32 bits representation of the slope:
 - $ptm_error = ptm1 - (A_ideal[63:32] * tmu1 + B_ideal[63:0])$


- where `ptm_error` is a 64-bit signed integer in 2's complement notation.

3. Assign the TMU to PTM Parameters:

- $TMU_to_PTM_A[31:0] = A_ideal[63:32]$
- $TMU_to_PTM_B[63:0] = B_ideal[63:0] + ptm_error$

Figure 11-7 shows the relationships between the terms used in the calculation above.

Figure 11-7: TMU to PTM Parameters Illustration



IMPLEMENTATION NOTE

It is recommended that a Parameter Generator apply filtering to any measured values before calculating the TMU_to_PTM Parameters.

11.2.5.3.2 PTM Master Time Reconstruction

A Parameter Consumer shall reconstruct the PTM Master Time as follows:

$$PTM_Master_Time(t) = TMU_to_PTM_A * TMU_time(t) + TMU_to_PTM_B$$

where:

$TMU_time(t)$ is the TMU Host Router Time.

$TMU_to_PTM_A$ and $TMU_to_PTM_B$ are the most recent TMU_to_PTM Parameters received on the Upstream Facing Port.

11.2.6 Timing Parameters

Table 11-6 lists the timing parameters for an internal PCIe Port.

Table 11-6. PCIe Adapter Timing Parameters

| Parameter | Description | Min | Max | Units |
|-----------|---|-----|-----|-------|
| tRecovery | Time in the Recovery.Idle state before transitioning to L0 state, in the absence of DLLP or TLP. Applicable only for a PCIe upstream port. | 200 | - | μs |

11.3 Paths

A PCIe Adapter Layer shall put HopID = 8 in the header of an outgoing Tunneled Packet before handing it off to the Transport Layer for routing.

**CONNECTION MANAGER NOTE**

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

11.3.1 Path Set-Up

Before setting up a Path, the Router does not indicate in-band presence to the Internal PCIe Port. The internal PCIe Port LTSSM is in the Detect state.

When the *Path Enable* bit is set to 1b, a Router shall indicate in-band presence to the internal PCIe Port. The PCIe Adapter Layer shall also enable sending of PCIe Tunneled Packets to the Transport Layer.

When an internal PCIe Port detects an in-band presence, it shall:

- Transition its LTSSM to the Polling state.
- Generate a PCIe interrupt to indicate a hot plug event to software.

**CONNECTION MANAGER NOTE**

A Connection Manager shall configure all of the Path Configuration Spaces for a PCIe Path before setting the Path Enable bit to 1b in the PCIe Adapters at either end of the Path.

The Connection Manager shall wait for the following before it enables a PCIe Path in an Adapter:

- *The previous PCIe Path set for the Adapter (if any) completed its tear-down flow.*
- *The LTSSM in the PCIe Port is in the Detect state, as reflected in PCIe Adapter Configuration Space. This allows stabilization of the PCIe link after the previous tunnel teardown, and to facilitate fast and reliable reconnect.*

11.3.2 Path Tear-Down

When a Router detects a disconnect on a Downstream Facing Port, it shall set the *Path Enable* bit to 0b in the PCIe Adapter Configuration Capability of the PCIe Adapter associated with the Lane Adapter via the Router's Routing Tables.

When the *Path Enable* bit is set to 0b in its PCIe Adapter Configuration Capability, a PCIe Adapter shall:

- Disable sending of PCIe Tunneled Packets from the PCIe Adapter Layer to the Transport Layer.

- If the PCIe Adapter is a Downstream PCIe Adapter, clear the In-band Presence indication to the internal PCIe port.
- If the PCIe Adapter is an Upstream PCIe Adapter, drive PERST# to default as defined in Section 11.1.1.4.2.

When an internal PCIe Port detects that the in-band presence indicator is cleared, it shall:


- Transition its LTSSM the Detect state. It is recommended that the LTSSM not wait the full amount of time defined in the PCIe Specification before transitioning to the Detect state.
- Generate a PCIe interrupt to indicate a hot removal event to software.



CONNECTION MANAGER NOTE

Before tearing down a PCIe Path, a Connection Manager shall set the Path Enable bit in the PCIe Adapter Configuration Capability to 0b.

12 Host Interface



This chapter defines the Host Interface Adapter Layer within a Host Interface Adapter.

A Host Router shall implement a single Host Interface Adapter Layer as defined in Sections 12.1 through 12.4 and Section 12.7. A PCIe Host Router shall also implement the PCIe-specific Host Interface Adapter Layer functionality defined in Sections 12.5 and 12.6. A non-PCIe Host Router shall implement functionality equivalent to that defined in Sections 12.5 and 12.6. The method to do so is system specific.

A Host Interface Adapter Layer shall implement N Transmit Descriptor Rings and N Receive Descriptor Rings, where $N \geq 2$. For any given HopID n (where $n = 0$ through $N - 1$), the Transmit Descriptor Ring is referred to as “Transmit Ring n ” and the Receive Descriptor Ring is referred to as “Receive Ring n ”. For example, the Transmit Descriptor Ring for HopID 0 is referred to as Transmit Ring 0 and the Receive Descriptor Ring for HopID 0 is referred to as Receive Ring 0.

A Host Router shall support operation with a Connection Manager that interfaces to the Router via Transmit Ring 0 and Receive Ring 0. The Connection Manager may be implemented in Host system software (an “external Connected Manager”), or it may be implemented in Host Router firmware or hardware (an “embedded Connection Manager”). If a Host Router has an embedded Connection Manager, it shall implement a mechanism to disable the embedded Connection Manager.

Note: A Host Router needs to operate with an external Connection Manager so that compliance testing can be performed. Transmit Ring 1 and Receive Ring 1 are used for compliance testing, but can be used for other purposes as well.

A Device Router shall not contain a Host Interface Adapter.

12.1 Descriptor Ring Mode

Transmit Ring 0 and Receive Ring 0 shall operate in Raw Mode only. All other Descriptor Rings shall support operation in both Raw Mode and Frame Mode. For a PCIe Host Interface Adapter Layer, the mode of operation for a Descriptor Ring is determined by the *Raw Mode* bit.

12.1.1 DW, Byte, and Bit Order

Data Buffers, Transmit Descriptors, and Receive Descriptors shall be stored in host memory in the following manner:

- For a Data Buffer: the lowest-addressed byte of data shall be mapped to bits [31:24] of the first payload DW depicted in Figure 5-1. The following addressed byte shall be mapped to bits [23:16] of the first payload DW depicted in Figure 5-1, and so on.
- For a Transmit Descriptor: the lowest-addressed byte of descriptor data in memory shall be mapped to bits [7:0] of the first DW of the Transmit Descriptor depicted in Figure 12-4. The following addressed byte shall be mapped to bits [15:8] of the first DW of the Transmit Descriptor depicted in Figure 12-4, and so on.
- For a Receive Descriptor: the lowest-addressed byte of descriptor data in memory shall be mapped to bits [7:0] of the first DW of the Receive Descriptor depicted in Figure 12-5 and Figure 12-6. The following addressed byte shall be mapped to bits [15:8] of the first DW of the first DW of the Receive Descriptor depicted in Figure 12-5 and Figure 12-6, and so on.
- Within each byte, bit[i] in host memory shall be mapped to bit[i] in the corresponding byte of the Data Buffer, Transmit Descriptor, or Receive Descriptor.

A PCIe Host Interface Adapter Layer shall map the payload of a PCIe TLP into a Transport Layer Packet payload, Transmit Descriptor, or a Receive Descriptor in the following manner:

- For a Transport Layer Packet: Data Byte 0 in the PCIe TLP payload, shall be mapped to bits [31:24] of the first payload DW depicted in Figure 5-1. Data Byte 1 in the PCIe TLP payload, shall be mapped to bits [23:16] of the first payload DW depicted in Figure 5-1, and so on.
- For a Transmit Descriptor: Data Byte 0 in the PCIe TLP payload shall be mapped to bits [7:0] of the first DW of the Transmit Descriptor depicted in Figure 12-4. Data Byte 1 in the PCIe TLP payload shall be mapped to bits [15:8] of the first DW of the Transmit Descriptor depicted in Figure 12-4, and so on.
- For a Receive Descriptor: Data Byte 0 in the PCIe TLP payload shall be mapped to bits [7:0] of the first DW of the Receive Descriptor depicted in Figure 12-5 and Figure 12-6. Data Byte 1 in the PCIe TLP payload shall be mapped to bits [15:8] of the first DW of the first DW of the Receive Descriptor depicted in Figure 12-5 and Figure 12-6, and so on.
- Within each byte, bit[i] in the PCIe TLP payload shall be mapped to bit[i] in the corresponding byte of the Transmit Descriptor, Receive Descriptor, or Transport Layer Packet payload.

12.1.2 Raw Mode

In Raw Mode, the Host sends bytes over the USB4™ Fabric by forming the payload of a Transport Layer Packet (either a Control Packet for Ring 0 or a Tunneled Packet for non-zero Rings) and posting it in a Data Buffer in Host Memory. The Host Router returns bytes to the Host by posting the payload of a Transport Layer Packet to a Data Buffer in Host Memory. Each Data Buffer holds the payload of one Transport Layer Packet.

12.1.3 Frame Mode

In Frame Mode, a Host groups bytes to be sent over the USB4 Fabric into Frames. A Frame shall be between 1 byte and 4096 bytes in length.

A Host Interface Adapter Layer segments a Frame into one or more Tunneled Packets as shown in Figure 12-1. The Tunneled Packet types used by a Host Interface Adapter in Frame Mode are defined in Table 12-1.

Figure 12-1. Segmentation of a Frame




Table 12-1. Frame Mode Tunneled Packet Format

| Type | Segment | PDF |
|------------------------|---|--|
| <i>Start of Frame</i> | Payload shall contain the first segment of a Frame. This Tunneled Packet type shall only be used when the Frame is segmented into more than one Tunneled Packet. | PDF value shall be taken from the <i>SOF PDF</i> field in the first Transmit Descriptor of the Frame (see Section 12.3.1). |
| <i>Middle of Frame</i> | Payload shall contain an intermediate segment of a Frame. This Tunneled Packet type shall only be used when the Frame is segmented into more than two Tunneled Packet. | 0h |

| Type | Segment | PDF |
|---------------------|--|---|
| <i>End of Frame</i> | <p>Payload shall contain the last segment of a Frame when the Frame is segmented into more than one Tunneled Packet.</p> <p>Payload shall contain the full Frame when the Frame is not segmented into more than one Tunneled Packet.</p> <p>Payload shall include the complete 32-bit CRC.</p> | PDF value shall be taken from the <i>EOF PDF</i> field in the last Transmit Descriptor of the Frame (see Section 12.3.1). |

12.2 End-to-End (E2E) Flow Control

The Host Interface Adapter Layer employs an end-to-end (E2E), credit-based flow control mechanism to prevent overflow of a Receive Descriptor Ring. Flow control is managed individually for each Receive Descriptor Ring and for each Transmit Descriptor Ring.

A Host Interface Layer shall not use E2E flow control for Transmit Ring 0 or Receive Ring 0.



CONNECTION MANAGER NOTE

The Connection Managers at either end of a Host-to-Host Path shall ensure that both ends of the Path to have the same flow control scheme.

12.2.1 E2E Flow Control Packets

12.2.1.1 E2E Credit Grant Packet

This packet is used to transmit credit information from a Receive Descriptor Ring to a Transmit Descriptor Ring. An E2E Credit Grant Packet shall include the header in Table 12-2 followed by the payload defined in Table 12-3.

Figure 12-2 shows an example of how an E2E Credit Grant Packet is sent from a Receiving Host Interface (Host 1) to a Transmitting Host Interface (Host 2). The Receiving Host is the Host that received data in a Receive Descriptor Ring. The Transmitting Host is the Host that sends data from a Transmit Descriptor Ring. In the example, the following steps take place:

1. The Host Interface Layer in Host 1 generates an E2E Credit Grant Packet for Receive Ring 7. The HopID in the E2E Credit Grant Packet is equal to the *Transmit E2E HopID* field for Receive Ring 7.
2. The HopID in the E2E Credit Grant Packet is updated according to the Routing Table in the Ingress Adapter. HopID x is set by the Connection Manager during Path Setup and can be any valid HopID.
3. The Ingress Adapter in Host 1 forwards the E2E Credit Grant Packet to the Egress Adapter indicated in its Routing Table. The Egress Adapter forwards the Packet to Host 2.
4. The E2E Credit Grant Packet is forwarded along the Inter-Domain Path between Transmit Ring 5 and Receive Ring 6. The Ingress Adapter in Host 2 receives the E2E Credit Grant Packet with HopID x.
5. The Ingress Adapter in Host 2 updates the HopID of the E2E Credit Grant Packet according to its Routing Table. The Connection Manager configures the Routing Table during Path Setup so that the new HopID is 6. After updating the HopID, the Ingress Adapter forwards the Packet to the Egress Adapter.
6. When the Host Interface Layer in Host 2 gets the E2E Credit Grant Packet, it looks at the HopID to see which Transmit Ring the Packet applies to. Because the HopID is 4, the Host Interface Adapter Layer applies the Packet to Transmit Ring 4.

Figure 12-2. Example of Forwarding an E2E Credit Grant Packet





Table 12-2. E2E Credit Grant Packet Header

| Bits | Field | Value |
|-------|---------------|---|
| 7:0 | <i>HEC</i> | See Section 5.1.2.1.1 |
| 15:8 | <i>Length</i> | 04h |
| 26:16 | <i>HopID</i> | Contains the HopID of a Path that terminates at the destination of this Credit Grant Packet and corresponds to the target Transmit Descriptor Ring. See Figure 12-2. For a PCIe Host Interface Adapter Layer, the inserted HopID value shall be the same value as in the <i>Transmit E2E HopID</i> field (Section 12.6.3.2.5). |
| 27 | <i>SuppID</i> | 0b |
| 31:28 | <i>PDF</i> | Fh |

Table 12-3. E2E Credit Grant Packet Payload

| Bits | Field | Value |
|-------|----------------|---|
| 7:0 | <i>ECC</i> | Error Correction – Error correction field that is calculated over bits [31:8] of the E2E Credit Grant Packet payload. See Section 5.1.2.3 for calculation. |
| 20:8 | <i>Credits</i> | Credits – Indicates the total number (modulo 8192) of credits granted to the target Transmit Descriptor Ring since initialization. |
| 30:21 | <i>Rsvd</i> | Reserved |
| 31 | <i>CS</i> | Credit Sync – Shall be set to 0b. |

Figure 12-3. E2E Credit Grant / Sync Packet Format



**CONNECTION MANAGER NOTE**

A Connection Manager shall set the Ring Valid bit in a Transmit Ring to 1b before using the Transmit Ring's Path to send an E2E Credit Grant Packet. For Example, in Figure 12-2 Host 1 needs to set the Ring Valid bit for Transmit Ring 5 to 1b before sending E2E Credit Grant Packets with HopID = 5.

A Connection Manager shall set the Ring Valid bit in a Receive Ring to 1b before using the Receive Ring's Path to receive an E2E Credit Grant Packet. For example, in Figure 12-2 Host 2 needs to set the Ring Valid bit for Receive Ring 6 to 1b before forwarding E2E Credit Grant Packets to Transmit Ring 4.

12.2.1.2 E2E Credit Sync Packet

This packet synchronizes the credit count between a Transmit Descriptor Ring and a Receive Descriptor Ring. An E2E Credit Sync Packet shall include the header in Table 12-4 followed by payload defined in Table 12-5.

Table 12-4. E2E Credit Sync Packet Header

| Bits | Field | Value |
|-------|---------------|---|
| 7:0 | <i>HEC</i> | See Section 5.1.2.1.1 |
| 15:8 | <i>Length</i> | 04h |
| 26:16 | <i>HopID</i> | Contains the HopID that corresponds to the Transmit Descriptor Ring that the packet applies to. |
| 27 | <i>SuppID</i> | 0b |
| 31:28 | <i>PDF</i> | Fh |

Table 12-5. E2E Credit Sync Packet Payload

| Bits | Field | Value |
|-------|----------------|--|
| 7:0 | <i>ECC</i> | Error Correction – Error correction field that is calculated over bits [31:8] of the E2E Credit Sync Packet payload. See Section 5.1.2.3 for calculation. |
| 20:8 | <i>Credits</i> | Credits Raw Mode: Represents the number (modulo 8192) of credits consumed by the Transmit Descriptor Ring up to (and including) the last packet transmitted prior to this E2E Credit Sync Packet. Frame Mode: Represents the number (modulo 8192) of credits consumed by the Transmit Descriptor Ring. The number of credits includes all Frames sent prior to the E2E Credit Sync Packet. It does not include any partially sent Frames. |
| 30:21 | <i>Rsvd</i> | Reserved |
| 31 | <i>CS</i> | Credit Sync – Shall be set to 1b. |

12.2.2 Flow Control Rules**12.2.2.1 Credit Update**

Credits are used to track the number of Receive Descriptors allocated for a Receive Descriptor Ring. Credits shall be given in units of Receive Descriptors where one credit corresponds to one Receive Descriptor.

12.2.2.2 Credit Counter Synchronization

If an E2E Credit Grant Packet is lost for any reason, the next successfully received E2E Credit Grant Packet corrects credit tracking in the transmitting Host Interface.

If an E2E flow controlled Tunneled Packet sent by a Host Interface is dropped or otherwise lost for any reason, the Receiving Host Interface credit counter loses synchronization with the Transmitting Host Interface. In order to re-establish synchronization, a Transmitting Host Interface sends E2E Credit Sync Packets periodically to the Receiving Host Interface.

12.2.2.3 Transmitting Host Interface Rules

The Host Interface Adapter Layer determines if sufficient credits have been advertised to permit the transmission of a Tunneled Packet from the Transmitting Host Interface:

- If E2E flow control is disabled for the Transmit Descriptor Ring, the Host Interface Adapter Layer shall not require any credits to be available before transmitting a Tunneled Packet from this Transmit Descriptor Ring.
- If E2E flow control is enabled for the Transmit Descriptor Ring and the Ring is in Raw Mode, the Host Interface Adapter Layer shall not transmit a Tunneled Packet unless at least one credit for the corresponding Transmit Descriptor Ring is available.
- If E2E flow control is enabled for the Transmit Descriptor Ring and the Ring is in Frame Mode, the Host Interface Adapter Layer shall not transmit a Tunneled Packet unless at least one credit for the corresponding Transmit Descriptor is available for the Frame.

A Transmitting Host Interface Adapter Layer shall not send E2E Credit Sync Packets for a Transmit Descriptor Ring with E2E flow control disabled or when the egress Link is in a low-power state.

When a Transmitting Host Interface Adapter Layer receives a Credit Grant Packet, it shall verify the *ECC* field value in the Credit Grant Record.

- The Adapter Layer shall correct any single-bit errors. After correcting an error, the Adapter Layer shall continue on as if the error had never occurred.
- If an uncorrectable error is detected, the Credit Grant Packet shall be dropped.

The following is a recommended scheme using state variables to manage credits for each Transmit Descriptor Ring enabled with E2E flow control. The state variables are cleared to their default values when the Transmit Ring is enabled. Other implementations are possible as long as the requirements in Section 12.2 are met:

- *Credits Consumed*
 - Contains a count (modulo 8192) of the total number of flow control credits consumed by Tunneled Packets sent by the Transmitting Host Interface since the Ring was initialized.
 - Default value is 0.
 - In Raw Mode, incremented each time a Tunneled Packet is sent by the Transmitting Host Interface: $\text{Credits Consumed} = (\text{Credits Consumed} + 1) \bmod 8192$.
 - In Frame Mode, incremented each time a last Tunneled Packet is sent for a Frame by the Transmitting Host Interface: $\text{Credits Consumed} = (\text{Credits Consumed} + 1) \bmod 8192$.
 - Included in the *Credits* field of an E2E Credit Sync Packet each time an E2E Credit Sync Packet is sent.
- *Credit Limit*
 - Contains the most recent number of flow control credits received from the Receiving Host Interface at the other end.
 - Default value is 0.

- Updated each time an E2E Credit Grant Packet is received for the Transmit Descriptor Ring by overwriting with the value contained in the *Credits* field of the E2E Credit Grant Packet.

12.2.2.4 Receiving Host Interface Rules

A Receiving Host Interface Adapter Layer shall track credits individually for each Receive Descriptor Ring enabled with E2E flow control.

A Receiving Host Interface Adapter Layer shall send an E2E Credit Grant Packet every $tE2ERate$ for each Receive Descriptor Ring with E2E flow control enabled. It shall also send an E2E Credit Grant Packet each time additional Receive Descriptors are made available to the Receive Descriptor Ring in Host Memory. An E2E Credit Grant Packet shall carry the most recent credit count for the Descriptor Ring.

A Host Interface Adapter Layer shall not send E2E Credit Grant Packets for a Receive Descriptor Ring with E2E flow control disabled.

If the Adapter indicated by the *Transmit E2E HopID* field for a Receive Descriptor Ring is in a Low Power state, a Host Interface Adapter Layer shall only send an E2E Credit Grant Packet for the Receive Descriptor Ring when the credit count for the Receive Ring changes.

The number of available credits advertised for a Receive Descriptor Ring shall not exceed 8192.

When a Receiving Host Interface Adapter Layer receives a Credit Sync Packet, it shall verify the *ECC* field value in the Credit Sync Packet payload.

- The Adapter Layer shall correct any single-bit errors. After correcting an error, the Adapter Layer shall continue on as if the error had never occurred.
- If an uncorrectable error is detected, the Credit Sync Packet shall be dropped.

The following is a recommended scheme using state variables to manage credits for each Receive Descriptor Ring enabled with E2E flow control. The state variables are cleared to their default values when the Receive Ring is enabled. Other implementations are possible as long as the requirements in Section 12.2 are met:

- *Credits Allocated*
 - Contains the total number (modulo 8192) of flow control credits allocated to the Receive Descriptor Ring since initialization.
 - On initialization, set to the number of Receive Descriptors provided by the Host in the Receive Descriptor Ring.
 - Included in the *Credits* field of an E2E Credit Grant Record each time an E2E Credit Grant Packet is sent.
 - Incremented (modulo 8192) by the number of Receive Descriptors provided each time a Host provides additional Receive Descriptors.
 - Each time an E2E Credit Sync Packet is consumed by the Host Interface Adapter Layer, the *Credits Allocated* counter is updated as follows:
$$Credits\ Allocated = (Credits\ Allocated + (Credits\ Consumed - Credits\ Received)) \bmod 8192$$
where *Credits Consumed* is the value in the *Credits* field in the E2E Credit Sync Packet.
- *Credits Received*
 - Contains the total number (modulo 8192) of flow control credits used by the Receive Descriptor Ring since initialization.
 - Set to 0 on initialization.

- Incremented (modulo 8192) each time the Host Interface Adapter Layer receives a Tunneled Packet (in Raw Mode) or the last Tunneled Packet of a Frame (in Frame Mode).
- Each time an E2E Credit Sync Packet is consumed by the Host Interface Adapter Layer, the *Credits Received* variable is updated after the *Credits Allocated* variable has been updated:

$$Credits\ Received = \text{value in the } Credits \text{ field in the E2E Credit Sync Packet}$$

12.3 Transmit Interface

A Host Interface Adapter Layer shall implement a Transmit Descriptor Ring for each supported Path.

12.3.1 Transmit Descriptor Structure

The Host Interface Adapter Layer shall fetch Transmit Descriptors from Host Memory. A Transmit Descriptor shall have the format depicted in Figure 12-4 and listed in Table 12-6.

Figure 12-4. Transmit Descriptor Structure




Table 12-6. Transmit Descriptor Contents

| DW | Bits | Name | Function |
|----|-------|---------------------|---|
| 0 | 31:0 | <i>Address Low</i> | Address Low – Lower 32 bits of the physical address of a Data Buffer in Host Memory. The least significant 2-bits of the <i>Address Low</i> field shall be 00b. This field is written by the Host. An Adapter Layer shall not modify this field. |
| 1 | 31:0 | <i>Address High</i> | Address High – Upper 32 bits of the physical address of a Data Buffer in Host Memory. This field is written by the Host. An Adapter Layer shall not modify this field. |
| 2 | 11:0 | <i>Data Length</i> | Data Length – Number of bytes to be transmitted from this Data Buffer. A value of 0 indicates that 4096 bytes shall be fetched. This field is written by the Host. An Adapter Layer shall not modify this field. |
| | 15:12 | <i>EOF PDF</i> | EOF PDF Raw Mode: Contains the PDF value for the Transport Layer Packet carrying data from a Transmit Descriptor. Shall be set to a value between 0h and Eh. Frame Mode: Contains the PDF value for the last Tunneled Packet for a Transmit Descriptor. If the Transmit Descriptor is the last (or only) Transmit Descriptor for a Frame, shall be set to a value between 1h and Eh. Else, shall be set to 0h. This field is written by the Host. An Adapter Layer shall not modify this field. |

| DW | Bits | Name | Function |
|----|-------|-------------------------|---|
| 2 | 19:16 | <i>SOF PDF</i> | SOF PDF Raw Mode: A Router shall ignore this field. Frame Mode: Contains the PDF value for the first Tunneled Packet for a Transmit Descriptor that is segmented into more than one Tunneled Packet. If the Transmit Descriptor is the first (or only) Transmit Descriptor for a Frame, shall be set to a value between 1h and Eh. Else, shall be set to 0h. The SOF PDF shall be different than the EOF PDF for the Frame. This field is written by the Host. An Adapter Layer shall not modify this field. |
| | 20 | <i>Rsvd</i> | Reserved |
| | 21 | <i>Descriptor Done</i> | Descriptor Done (DD) – This bit is set to 0b by the Host when posting a Data Buffer to be transmitted. If the <i>Request Status</i> bit is set to 1b, the Host Interface Adapter Layer shall set this bit to 1b after the last byte of the Data Buffer is sent to the Transport Layer. If the <i>Request Status</i> bit is set to 0b, the Host Interface Adapter Layer shall not write to this bit. |
| | 22 | <i>Request Status</i> | Request Status (RS) – This bit determines whether or not the Host Interface Adapter Layer updates transmission status in the <i>DD</i> bit. This field is written by the Host. An Adapter Layer shall not modify this field. |
| | 23 | <i>Interrupt Enable</i> | Interrupt Enable (IE) – If this bit is set to 1b, then the Host Interface Adapter Layer shall issue an interrupt indicating the completion of the Data Buffer after setting the <i>DD</i> bit to 1b. This field is written by the Host. An Adapter Layer shall not modify this field. |
| | 31:24 | <i>Offset</i> | Offset – The offset in bytes of the beginning of data to be transmitted. Offset is relative to the beginning of the Data Buffer as defined in the <i>Address Low</i> and <i>Address High</i> fields. This field is written by the Host. An Adapter Layer shall not modify this field. |
| 3 | 31:0 | <i>Rsvd</i> | Reserved |

12.3.2 Transmit Flow

A Host Interface Adapter Layer shall only process a Transmit Descriptor Ring when both the Transmit Descriptor Ring is enabled and at least one Transmit Descriptor is pending. For a PCIe Host Interface Adapter Layer, Transmit Descriptors are pending when the *Producer Index* field for the Transmit Descriptor Ring has a different value than the *Consumer Index* field (see Section 12.6.3.2.3).

Section 12.3.2.1 describes how a Host Interface Adapter Layer processes Transmit Descriptor Ring in Frame Mode. Section 12.3.2.2 describes how a Host Interface Adapter Layer processes Transmit Descriptor Ring in Raw Mode.



CONNECTION MANAGER NOTE

A Connection Manager shall only post Control Packets to Transmit Ring 0. It shall not post Control Packets to any other Transmit Ring.

12.3.2.1 Frame Mode

A Frame shall reside in one or more Data Buffers. A Data Buffer shall not contain data from more than one Frame at a time. A Frame may span multiple Data Buffers.

A Host Interface Adapter Layer shall segment a Frame into one or more Tunneled Packets as follows:

- The Host Interface Adapter takes the Data Buffers of one or more Transmit Descriptors to form a Tunneled Packet payload. The Host Interface Adapter Layer prepends a Transport Layer Packet Header to the payload to generate a Tunneled Packet.
- If a Tunneled Packet contains multiple Data Buffers as payload, all Data Buffers shall contain data from the same Frame.
- If a Frame is sent in a single Tunneled Packet, the Tunneled Packet shall be of type *End of Frame*.
- If a Frame is sent in multiple Tunneled Packets:
 - The Tunneled Packet containing the first segment of the Frame shall be of type *Start of Frame* and shall be sent first.
 - Any following Tunneled Packets other than the last Tunneled Packet shall be of type *Middle of Frame*.
 - The Tunneled Packet containing the last segment of the Frame shall be of type *End of Frame* and shall be sent last.

The Router shall append a 32-bit CRC to each Frame. The CRC shall cover the entire Frame. The Router shall calculate the CRC as defined in Table 6-1. The CRC shall be calculated in increasing DW order. Within each DW, CRC shall be calculated from bit[31] to bit[0].

- The CRC shall be placed in the payload of an *End of Frame* Tunneled Packet immediately after the Frame bytes. See Appendix A.7 for an example of a Frame CRC.
 - Bits [7:0] of the calculated CRC shall be placed in bits [31:24] of the *CRC* field.
 - Bits [15:8] of the calculated CRC shall be placed in bits [23:16] of the *CRC* field.
 - Bits [23:16] of the calculated CRC shall be placed in bits [15:8] of the *CRC* field.
 - Bits [31:24] of the calculated CRC shall be placed in bits [7:0] of the *CRC* field.
- An *End of Frame* Tunneled Packet may contain just the 32-bit CRC without any preceding Frame bytes.
- If padding is added to any Tunneled Packets, it shall be added after the Frame CRC is calculated and the Frame and CRC are segmented into Tunneled Packet payload.

The Router shall set the *HopID* value of each Tunneled Packet to *n*, where *n* is the HopID associated with the Transmit Descriptor Ring.

If the *RS* bit in the Transmit Descriptor is set to 1b, a Host Interface Adapter Layer shall do the following after the last byte of the Data Buffer is sent to the Transport Layer:

- Set the *DD* bit to 1b in the Transmit Descriptor in Host Memory.
- If the *IE* bit is set to 1b, issue an interrupt indicating the completion of the Data Buffer. See Section 12.5 for interrupt issuance and handling by a PCIe Host Interface Adapter Layer. Interrupt handling in non-PCIe-based systems is implementation specific.

After fetching a Data Buffer from Host Memory, a PCIe Host Interface Adapter Layer shall increment the *Consumer Index* for the Transmit Descriptor Ring by 1 (with wraparound to 0 when *Consumer Index* = *Ring Size*).

12.3.2.2 Raw Mode

A Host Interface Adapter Layer shall fetch the payload for a Transport Layer Packet from the Data Buffer referenced in the next available Transmit Descriptor.

The Host Interface Adapter Layer is responsible for prepending the Transport Layer Packet Header to the payload fetched from the Data Buffer.

The *PDF* field in the Transport Layer Packet shall match the *EOF PDF* field in the Transmit Descriptor.

If the *RS* bit in the Transmit Descriptor is set to 1b, a Host Interface Adapter Layer shall do the following after the last byte of the Data Buffer is sent to the Transport Layer:

- Set the *DD* bit to 1b in the Transmit Descriptor in Host Memory.
- If the *IE* bit is set to 1b, issue an interrupt indicating the completion of the Data Buffer. See Section 12.5 for interrupt issuance and handling by a PCIe Host Interface Adapter Layer. Interrupt handling in non-PCIe-based systems is implementation specific.

After fetching a Data Buffer from Host Memory, a PCIe Host Interface Adapter Layer shall increment the *Consumer Index* for the Transmit Descriptor Ring by 1 (with wraparound to 0 when *Consumer Index = Ring Size*).

12.4 Receive Interface

A Host Interface Adapter Layer shall implement a Receive Descriptor Ring for each supported Path.

12.4.1 Receive Descriptor Structure

A Receive Descriptor that is fetched from Host Memory shall have the format depicted in Figure 12-5 and listed in Table 12-7. A Receive Descriptor that is posted to Host Memory shall have the format depicted in Figure 12-6 and listed in Table 12-8.

Figure 12-5. Receive Descriptor Structure (Posted by Host)




Table 12-7. Receive Descriptor Contents (Posted by Host)

| DW | Bits | Name | Function |
|----|-------|------------------------------|---|
| 0 | 31:0 | <i>Address Low</i> | Address Low – Lower 32 bits of the physical address of a Data Buffer in Host Memory. The least significant 2-bits of the <i>Address Low</i> field shall be 00b. |
| 1 | 31:0 | <i>Address High</i> | Address High – Upper 32 bits of the physical address of a Data Buffer in Host Memory. |
| 2 | 20:0 | <i>Reserved</i> | Reserved |
| | 21 | <i>Descriptor Done (DD)</i> | Descriptor Done – This bit shall be set to 0b. |
| | 22 | <i>Request Status (RS)</i> | Request Status – This bit is set to 1b to enable descriptor write-back by the Host Interface. Host shall always set this bit to 1b. |
| | 23 | <i>Interrupt Enable (IE)</i> | Interrupt Enable – This bit is set to 1b to enable interrupts indicating completion of the Data Buffer after the <i>DD</i> bit is set to 1b. |
| | 31:24 | <i>Offset</i> | Offset – The offset in bytes where the first byte of data shall be written to the Data Buffer. The offset is relative to the beginning of the Data Buffer as defined in the <i>Address Low</i> and <i>Address High</i> fields. |
| 3 | 31:0 | <i>Reserved</i> | Reserved |

Figure 12-6. Receive Descriptor Structure (Posted by Host Interface Adapter Layer)




Table 12-8. Receive Descriptor Contents (Posted by Host Interface Adapter Layer)

| DW | Bits | Name | Function |
|----|-------|-----------------------------|--|
| 0 | 31:0 | <i>Address Low</i> | Address Low – Shall contain the same value as posted by the Host. |
| 1 | 31:0 | <i>Address High</i> | Address High – Shall contain the same value as posted by the Host. |
| 2 | 11:0 | <i>Data Length</i> | Data Length – Number of bytes posted to the Data Buffer. A value of 0 indicates that 4096 bytes were posted. |
| | 15:12 | <i>EOF PDF</i> | EOF PDF Raw Mode: Shall be set to the value in the <i>PDF</i> field of the Transport Layer Packet. Frame Mode: Shall be set to the value in the <i>PDF</i> field of the End of Frame Transport Layer Packet carrying the last (or only) segment of the Frame. |
| | 19:16 | <i>SOF PDF</i> | SOF PDF Raw Mode: Shall be set to 0h. Frame Mode: Shall be set to 0h if Frame is sent in a single Tunneled Packet. Else, shall be set to the value in the <i>PDF</i> field of the <i>Start of Frame</i> Transport Layer Packet carrying the first segment of the Frame. |
| | 20 | <i>CRC Error</i> | CRC Error – indicates if there is an error in the Frame CRC Raw Mode: This field shall be set to 0b. Frame Mode: This bit shall be set to 1b if the CRC check failed for the reassembled Frame. Otherwise, this bit shall be 0b. |
| | 21 | <i>Descriptor Done (DD)</i> | Descriptor Done If the <i>RS</i> bit is set to 1b when the Receive Descriptor is read from Host Memory, this bit shall be set to 1b after the last byte has been written to the Data Buffer. |
| | 22 | <i>Buffer Overflow (BO)</i> | Buffer Overflow Raw Mode: This bit shall be set to 1b if the size of the received Transport Layer Packet exceeds the available space in the Data Buffer. It shall be set to 0b otherwise. Frame Mode: This bit shall be set to 1b if the size of the reassembled Frame exceeds the available space in the Data Buffer. It shall be set to 0b otherwise. |

| | | | |
|---|-------|------------------------------|--|
| 2 | 23 | <i>Interrupt Enable (IE)</i> | Interrupt Enable Shall contain the same value as posted by the Host. |
| | 31:24 | <i>Offset</i> | Offset Shall contain the same value as posted by the Host. |
| 3 | 31:0 | <i>Reserved</i> | Reserved |

12.4.2 Receive Flow

This section defines how a Host Interface Adapter Layer processes a Transport Layer Packet received from the Transport Layer. When a Host Interface Adapter Layer receives a Transport Layer Packet from the Transport Layer it shall:

1. Select the Receive Descriptor Ring corresponding to the *HopID* field in the received Transport Layer Packet.
2. If the Receive Descriptor Ring is disabled, discard the Transport Layer Packet.
3. Else if the Receive Descriptor Ring operates in Frame Mode, process the Transport Layer Packet according to Section 12.4.2.1.
4. Else if the Receive Descriptor Ring operates in Raw Mode, process the Transport Layer Packet according to Section 12.4.2.2.

12.4.2.1 Frame Mode

A Host Interface Adapter Layer in Frame Mode shall process Tunneled Packets received for a Receive Descriptor Ring as follows:

- All Tunneled Packets for a Frame shall be posted into the same Data Buffer.
- A Frame shall be posted using the next available Receive Descriptor. For a PCIe Host Interface Adapter Layer, the *Producer Index* field points to the next available Receive Descriptor.
 - If a Receive Descriptor is not available in Host Memory, discard the packet.
 - If a Receive Descriptor is not available for posting a *Start of Frame* packet, and potentially subsequent *Middle of Frame* packets, but is available for posting later packets for the same Frame, then subsequent Packets may be discarded.
- If the size of the packet payload exceeds the remaining available size of the Data Buffer:
 - Optionally write to the Data Buffer the part of the packet that fits into the Data Buffer. Any further received Tunneled Packets that belong to the same Frame shall be dropped and shall not be written to Host Memory.
 - Set the *Buffer Overflow* bit in the Receive Descriptor to 1b.
 - The *CRC Error* bit, the *SOF PDF* field, and the *EOF PDF* field shall represent the Frame as received, including packets not written to Host Memory.
- If a packet is a *Middle of Frame* packet and either no *Start of Frame* packet has been received for the Frame or an *End of Frame* packet was previously received for the Frame:
 - Post packet payload to a Data Buffer using the next available Receive Descriptor.
 - Set the *SOF PDF* field to 0b and the *CRC Error* bit to 1b in the Receive Descriptor after the Receive Descriptor is written back to Host Memory.
- If a packet is an *End of Frame* packet and either no *Start of Frame* packet has been received for the Frame or an *End of Frame* packet was previously received for the Frame:
 - Post packet payload to a Data Buffer using the next available Receive Descriptor.
 - Set the *SOF PDF* field to 0b.

- If a packet is a *Start of Frame* packet and the previous packet received was a *Start of Frame* packet or a *Middle of Frame* packet, post the packet to the beginning of the Data Buffer (overwriting any previous packets).

Note: This applies even if the overwritten Frame overflowed the Data Buffer.

- Packet payload for a Frame shall be posted in the order received. If the Packet payload is from a *Start of Frame* packet, it shall be posted at the offset in the *Offset* field of the Receive Descriptor. Otherwise, payload shall be posted immediately after the previously received payload for the Frame.
 - The 32-bit CRC received with a Frame shall not be posted to the Data Buffer.
- If a packet is an *End of Frame* packet:
 - If the *RS* bit in the Receive Descriptor is set to 1b, then:
 - Write back the Receive Descriptor with the contents defined in Table 12-8.
 - If the *IE* bit is set to 1b, issue an interrupt indicating the completion of the Data Buffer.
 - For a PCIe Host Interface Adapter Layer, increment the *Producer Index* by 1 with wraparound to 0 when *Producer Index* = *Ring Size*.

12.4.2.2 Raw Mode

A Host Interface Adapter Layer in Raw Mode shall process Transport Layer Packets received for a Receive Descriptor Ring as follows:

- Each Transport Layer Packet payload shall be posted into a separate, single, Data Buffer. It shall be posted at the offset in the *Offset* field of the Receive Descriptor.
- A Transport Layer Packet shall be posted using the next available Receive Descriptor. For a PCIe Host Interface Adapter Layer, the *Producer Index* field points to the Receive Descriptor.
 - If a Receive Descriptor is not available in host Memory, discard the packet.
- If the size of payload of the received Transport Layer Packet exceeds the available size of the Data Buffer:
 - Optionally, write the part of the packet that fits into the Data Buffer.
 - Set the *Buffer Overflow* bit in the Receive Descriptor to 1b.
 - If the *RS* bit in the Receive Descriptor is set to 1b, then:
 - Write back the Receive Descriptor with the contents defined in Table 12-8.
 - If the *IE* bit is set to 1b, issue an interrupt indicating the completion of the Data Buffer.
 - For a PCIe Host Interface Adapter Layer, increment the *Producer Index* by 1 with wraparound to 0 when *Producer Index* = *Ring Size*.
- When posting packet payload to a Data Buffer:
 - If the *RS* bit in the Receive Descriptor is set to 1b, then:
 - Write back the Receive Descriptor with the contents defined in Table 12-8.
 - If the *IE* bit is set to 1b, issue an interrupt indicating the completion of the Data Buffer.
 - For a PCIe Host Interface Adapter Layer, increment the *Producer Index* by 1 with wraparound to 0 when *Producer Index* = *Ring Size*.

12.5 Interrupts

A PCIe Host Router shall support INTx Interrupt Signaling and Message Signaled Interrupt (MSI and MSI-X).

12.5.1 Interrupt Causes

If the *Interrupt Enable* bit is set to 1b in a Transmit Descriptor, then after the Host Interface Adapter Layer writes back to that Transmit Descriptor, it shall set to 1b the *Transmit Data Buffer Interrupt* bit in the Interrupt Status Registers that corresponds to this Transmit Descriptor Ring.

If the *Interrupt Enable* bit is set to 1b in a Receive Descriptor, then after the Host Interface Adapter Layer has written back a Receive Descriptor, it shall set to 1b the *Receive Data Buffer Interrupt* bit in the Interrupt Status Registers that corresponds to this Receive Descriptor Ring.

The Host Interface Adapter Layer shall monitor how many unused Receive Descriptors are in each Receive Descriptor Ring. A Receive Descriptor is “unused” when the Host has indicated that the Receive Descriptor is available, but a Host Interface Adapter layer has not yet posted a packet into its Data Buffer. For each Receive Descriptor Ring, a Host Interface Adapter Layer shall set to 1b the corresponding bit in the Receive Ring Vacancy Status register to 1b when either:

- The *Receive Ring Vacancy Control* field for that Receive Descriptor Ring is zero and Receive Descriptor Ring has no unused Receive Descriptors.
- The number of unused Receive Descriptors for a Receive Descriptor Ring is less than or equal to the value indicated by the *Receive Ring Vacancy Control* field for that Receive Descriptor Ring.

If a *Receive Ring Vacancy Status* bit is set to 1b, the Host Interface Adapter Layer shall set the corresponding bit to 1b in the Interrupt Status Registers.

12.5.2 Interrupt Masks

When the Host Interface Adapter Layer sets a bit in the *Interrupt Status* Registers to 1b, it shall initiate an interrupt request to the interrupt moderation function if the corresponding bit in the *Interrupt Mask* Registers is set to 1b. It shall not initiate an interrupt if the corresponding bit in the *Interrupt Mask* Registers is set to 0b.

12.5.3 Interrupt Vectors

In MSI mode of operation, the Host Interface Adapter Layer shall request a single interrupt vector. All interrupt causes are mapped to the single interrupt vector.

In MSI-X mode of operation, the Host Interface Adapter Layer shall support up to 16 interrupt vectors. When a bit is set to 1b in the *Interrupt Status* Registers and the corresponding bit in the *Interrupt Mask* Registers is set to 1b, the Host Interface Adapter Layer shall issue the interrupt vector associated with the interrupt cause in the *Interrupt Vector Allocation* Registers (IVAR).

12.5.4 Interrupt Moderation


Interrupt moderation limits the rate of interrupts issued to the Host by delaying the time when an interrupt is issued to the Host. Each interrupt vector is moderated independently, where *Interrupt Throttling Rate* Register *i* (ITR[*i*]) moderates interrupt vector *i* (*i*=0,...,15). In INTx or MSI modes of operation, interrupt moderation only uses the ITR[0] register.

If the *Interval* field in the ITR[*i*] register (*i*=0,...,15) is set to 0, an interrupt of interrupt vector *i* shall not be moderated. An interrupt message shall be sent when the bit in the *Interrupt Status* Registers that is associated with interrupt vector *i* is set to 1b.

If the *Interval* field in the ITR[*i*] register (*i*=0,...,15) is set to a non-zero value, an interrupt message shall be sent according to the flow defined in Figure 12-7.

Section 12.6.3.4.7 defines the time interval for interrupt moderation.

Figure 12-7. Interrupt Moderation



* An interrupt is pending when a bit in the Interrupt Status Registers that is associated with interrupt vector *i* is set to 1b

12.6 Programming Interface

A PCIe Host Interface Adapter Layer shall support the programming interface defined in this section.

This section defines the Configuration Space of the Host Interface Adapter Layer. The Configuration Space resides in a single PCIe Memory BAR.

The Host Interface Adapter Layer shall discard any write request to a non-implemented register within the memory BAR.

The returned value by the Host Interface Adapter Layer to a read request for a non-implemented register within the memory BAR, is implementation-specific.

The Host Interface Adapter Layer shall support aligned 32-bit accesses (with all Byte Enables set to 1b) to the Memory BAR.

The Host Interface Adapter Layer shall expose a value of 0C 03 40h in the *Class Code* field of the PCI Configuration Space.



CONNECTION MANAGER NOTE

A Connection Manager shall limit the size of a read and write accesses from the PCIe Memory BAR to a single DW.

12.6.1 Access Types

Table 12-9 defines the access types that are possible for a particular configuration field.

Table 12-9. Access Types

| Access Type | Description |
|-------------|--|
| R/W | Read/Write. A field with this access type shall be capable of both read and write operations. The value read from this field shall reflect the last value written to it unless the field was reset in the interim. |
| R/W S | Read/Write Status. A field with this access type shall be capable of both read and write operations. The value read from this field may or may not reflect the last value written. |
| RO | Read Only. A write to a field with this access type shall have no effect. A read shall return a meaningful value. |
| R/Clr | Read Clear. A field with this access type shall be cleared to 0 after it is read. A write to a field with this attribute shall have no effect on its value. |
| W/Clr | Write Clear. A field with this access type shall be cleared to 0 after it is written to. A read shall return a meaningful value. |
| R/W SC | Read/Write Self Clearing. When set to 1b a field with this access type causes an action to be initiated. A field with this attribute shall read as 0b after the action is complete. The value returned prior to completion of the action is vendor defined. |
| WO | Write Only. A field with this access type shall be capable of write operations. Reading the field returns a vendor-defined value. |
| Rsvd | Reserved. Reserved for future implementation. A write to this field shall have no effect. A read shall return 0. |
| RsvdZ | Reserved and Zero. Reserved for future implementation. A read shall return 0. |

**CONNECTION MANAGER NOTE**

A Connection Manager shall only write 0 to a field marked as RsvdZ field.

12.6.2 Registers Summary

The Memory BAR contains the 32-bit registers listed in Table 12-10.

Certain fields have one instance per Transmit or Receive Descriptor Ring. The offset of such fields is defined as function of n, where n equals the HopID of the Transmit or Receive Descriptor Ring. N is equal to the value in the *Total Paths* field and represents the number of Transmit or Receive Descriptor Rings.

Table 12-10. Summary of Memory BAR Registers

| Offset | Register Name | Section |
|----------------------------------|-----------------------------|------------|
| Host Interface Control | | |
| 39640h | Host Interface Capabilities | 12.6.3.1.1 |
| 39858h | Host Interface Reset | 12.6.3.1.2 |
| 39864h | Host Interface Control | 12.6.3.1.3 |
| 39880h | Host Interface CL1 Enable | 12.6.3.1.4 |
| 39884h | Host Interface CL2 Enable | 12.6.3.1.5 |
| Transmit Descriptor Rings | | |
| 00000h + n*10h | Base Address Low | 12.6.3.2.1 |
| 00004h + n*10h | Base Address High | 12.6.3.2.2 |
| 00008h + n*10h | Producer & Consumer Indexes | 12.6.3.2.3 |

| Offset | Register Name | Section |
|--|------------------------------------|-------------|
| 0000Ch + n*10h | Ring Size | 12.6.3.2.4 |
| 19800h + n*20h | Ring Control | 12.6.3.2.5 |
| Receive Descriptor Rings | | |
| 08000h + n*10h | Base Address Low | 12.6.3.3.1 |
| 08004h + n*10h | Base Address High | 12.6.3.3.2 |
| 08008h + n*10h | Producer & Consumer Indexes | 12.6.3.3.3 |
| 0800Ch + n*10h | Ring & Buffer Size | 12.6.3.3.4 |
| 29800h + n*20h | Ring Control | 12.6.3.3.5 |
| 29804h + n*20h | PDF Bit Masks | 12.6.3.3.6 |
| Interrupts | | |
| 37800h : 37800h + (4 * ceiling(3N/32) - 1) | Interrupt Status (ISR) | 12.6.3.4.1 |
| 37808h : 37808h + (4 * ceiling(3N/32) - 1) | Interrupt Status Clear (ISC) | 12.6.3.4.2 |
| 37810h : 37810h + (4 * ceiling(3N/32) - 1) | Interrupt Status Set (ISS) | 12.6.3.4.3 |
| 38200h : 38200h + (4 * ceiling(3N/32) - 1) | Interrupt Mask (IMR) | 12.6.3.4.4 |
| 38208h : 38208h + (4 * ceiling(3N/32) - 1) | Interrupt Mask Clear (IMC) | 12.6.3.4.5 |
| 38210h : 38210h + (4 * ceiling(3N/32) - 1) | Interrupt Mask Set (IMS) | 12.6.3.4.6 |
| 38C00h : 38C3Ch | Interrupt Throttling Rate (ITR) | 12.6.3.4.7 |
| 38C40h : 38C40h + (4 * ceiling(3N/8) - 1) | Interrupt Vector Allocation (IVAR) | 12.6.3.4.8 |
| 18C00h : 18C00h + (4 * ceiling(N/8) - 1) | Receive Ring Vacancy Control | 12.6.3.4.9 |
| 19400h : 19400h + (4 * ceiling(N/32) - 1) | Receive Ring Vacancy Status | 12.6.3.4.10 |

12.6.3 Registers Description

12.6.3.1 Host Interface Control

12.6.3.1.1 Host Interface Capabilities

The Host Interface Capabilities register specifies the parameters supported by the Host Interface Adapter Layer.

Table 12-11. Host Interface Capabilities Register

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|----------------|
| 10:0 | Total Paths The total number of transmit/receive Paths supported by the Host Interface Adapter Layer. <i>Note: The maximum supported value for Total Paths is 21.</i> | RO | Vendor defined |
| 31:11 | Reserved | Rsvd | 0 |

12.6.3.1.2 Host Interface Reset

The Host Interface Reset register resets the Host Interface Adapter Layer.

Table 12-12. Host Interface Reset Register

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|--------|---------------|
| 0 | RST When set to 1b, shall reset the Host Interface registers and the E2E flow control counters to their default values. The Host Interface Adapter Layer shall complete the reset within tHIReset time. | R/W SC | 0b |
| 31:1 | Reserved | Rsvd | 0 |

12.6.3.1.3 Host Interface Control**Table 12-13. Host Interface Control Register**

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|----------------|
| 16:0 | Reserved | Rsvd | Vendor Defined |
| 17 | Disable ISR Auto-Clear This bit controls the clearing of the Interrupt Status Registers. See Section 12.6.3.4.1. | R/W | 0b |
| 31:18 | Reserved | Rsvd | Vendor Defined |

12.6.3.1.4 Host Interface CL1 Enable

The Host Interface CL1 Enable Register contains one bit per Transmit Descriptor Ring.

Table 12-14. Host Interface CL1 Enable

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|---------------|
| 31:0 | Host Interface CL1 Enable bits When bit n ($n=0,\dots,Total\ Paths - 1$) is set to 0b, it prevents the Link associated with Transmit Descriptor Ring n from entering into CL1 state. When bit n is set to 1b it enables the Link associated with Transmit Descriptor Ring n to enter CL1 state. | R/W | 0 |

12.6.3.1.5 Host Interface CL2 Enable

The Host Interface CL2 Enable Register contains one bit per Transmit Descriptor Ring.

Table 12-15. Host Interface CL2 Enable

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|---------------|
| 31:0 | Host Interface CL2 Enable bits When bit n ($n=0,\dots,Total\ Paths - 1$) is set to 0b, it prevents the Link associated with Transmit Descriptor Ring n from entering into CL2 state. When bit n is set to 1b it enables the Link associated with Transmit Descriptor Ring n to enter CL2 state. | R/W | 0 |

12.6.3.2 Transmit Descriptor Rings**12.6.3.2.1 Base Address Low****Table 12-16. Base Address Low Register**

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|---------------|
| 31:0 | <p>Base Address Low</p> <p>Lower 32 bits of the physical address of the corresponding Descriptor Ring in Host Memory.</p> <p>The Base Address is aligned to 16 bytes, so that the least significant 4-bits of the <i>Ring Base Address Low</i> field shall be 0h.</p> <p>Writing to this register sets the <i>Producer Index</i> field and the <i>Consumer Index</i> field to their default values.</p> | R/W | 0 |

12.6.3.2.2 Base Address High**Table 12-17. Base Address High Register**

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|---|------|---------------|
| 31:0 | <p>Base Address High</p> <p>Upper 32 bits of the physical address of the corresponding Descriptor Ring in Host Memory.</p> <p>Writing to this register sets the <i>Producer Index</i> field and the <i>Consumer Index</i> field to their default values.</p> | R/W | 0 |

12.6.3.2.3 Producer and Consumer Indexes**Table 12-18. Producer and Consumer Indexes Register**

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|---|------|---------------|
| 15:0 | <p>Consumer Index</p> <p>Index of the next Transmit Descriptor to be processed by the Host Interface Adapter Layer.</p> <p>The <i>Consumer Index</i> field counts in units of Descriptors. A value of 0 refers to the first Transmit Descriptor in the Descriptor Ring.</p> | RO | 0 |
| 31:16 | <p>Producer Index</p> <p>Index of the next Transmit Descriptor that Host writes to.</p> <p>The <i>Producer Index</i> field counts in units of Descriptors. A value of 0 refers to the first Transmit Descriptor in the Descriptor Ring.</p> <p>Router behavior is undefined if a Host writes a value larger than (<i>Ring Size</i> - 1).</p> <p>A Connection Manager shall not increment this field if the new value points to the location of the <i>Consumer Index</i> field (leaving one vacant space in the Ring).</p> | R/W | 0 |

12.6.3.2.4 Ring Size

Table 12-19. Ring Size Register

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|---|------|---------------|
| 15:0 | Ring Size The number of Transmit Descriptors in the Ring. Writing to this register sets the <i>Producer Index</i> field and the <i>Consumer Index</i> field to their default values. Ring Size shall not exceed 4096 descriptors. | R/W | 0 |
| 31:16 | Reserved | Rsvd | 0 |



CONNECTION MANAGER NOTE

A Connection Manager shall write a non-zero value to the Ring Size field before setting the Ring Valid bit to 1b.

12.6.3.2.5 Ring Control**Table 12-20. Ring Control Register**

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|---|------|---------------|
| 27:0 | Reserved | Rsvd | 0 |
| 28 | E2E Flow Control Enable When set to 0b, end-to-end flow control is disabled for this Transmit Descriptor Ring. When set to 1b, end-to-end flow control is enabled for this Transmit Descriptor Ring. | R/W | 0b |
| 29 | No-snoop flag (NS) Determines the value to be set by the Host Interface Adapter Layer in the No Snoop attribute of a PCIe TLP associated with this Descriptor Ring. | R/W | 0b |
| 30 | Raw Mode (RAW) When set to 0b, the Descriptor Ring shall operate in Frame Mode. When set to 1b, the Descriptor Ring shall operate in Raw Mode. | R/W | 0b |
| 31 | Ring Valid When set to 0b, the Descriptor Ring is disabled. When set to 1b, the Descriptor Ring is enabled. | R/W | 0b |

12.6.3.3 Receive Descriptor Rings**12.6.3.3.1 Base Address Low****Table 12-21. Base Address Low Register**

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|---------------|
| 31:0 | Base Address Low Lower 32 bits of the physical address of the corresponding Descriptor Ring in Host Memory. The Base Address is aligned to 16 bytes, so that the least significant 4-bits of the <i>Ring Base Address Low</i> field shall be 0. Writing to this register sets the <i>Producer Index</i> field and the <i>Consumer Index</i> field to their default values. | R/W | 0 |

12.6.3.3.2 Base Address High**Table 12-22. Base Address High Register**

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|---------------|
| 31:0 | Base Address High Upper 32 bits of the physical address of the corresponding Descriptor Ring in Host Memory. Writing to this register sets the <i>Producer Index</i> field and the <i>Consumer Index</i> field to their default values. | R/W | 0 |

12.6.3.3.3 Producer and Consumer Indexes**Table 12-23. Producer and Consumer Indexes Register**

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|---------------|
| 15:0 | Consumer Index Index of the next Receive Descriptor that the Host provides to the Host Interface Adapter Layer. The <i>Consumer Index</i> field counts in units of Descriptors. A Host shall not write a value larger than (<i>Ring Size</i> - 1). A Connection Manager shall not increment this field if the new value points to the location of the <i>Producer Index</i> field (leaving one vacant space in the Ring). | R/W | 0 |
| 31:16 | Producer Index Index of the next Receive Descriptor that the Host Interface Adapter Layer writes to. The <i>Producer Index</i> field counts in units of Descriptors. | RO | 0 |

**CONNECTION MANAGER NOTE**

A Connection Manager shall set the Consumer Index for a Ring to its default value after it sets the Ring Valid bit to 0b.

12.6.3.3.4 Ring and Buffer Size**Table 12-24. Ring Size Register**

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|---|------|---------------|
| 15:0 | Ring Size Size of the Descriptor Ring in multiples of Receive Descriptors. Writing to this register sets the <i>Producer Index</i> field and the <i>Consumer Index</i> field to their default values. Ring Size shall not exceed 4096 descriptors. | R/W | 0 |
| 27:16 | Data Buffer Size A value of 0 indicates that 4096 bytes were posted. The Host Interface Adapter Layer shall not write to the <i>Data Length</i> field in a Receive Descriptor a value larger than <i>Data Buffer Size</i> minus the value of the <i>Offset</i> field provided by the Host in the Receive Descriptor. | R/W | 0 |
| 31:28 | Reserved | Rsvd | 0 |

**CONNECTION MANAGER NOTE**

A Connection Manager shall write a non-zero value to the Ring Size field before setting the Ring Valid bit to 1b.

12.6.3.3.5 Ring Control**Table 12-25. Ring Control Register**

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|---|------|---------------|
| 11:0 | Reserved | Rsvd | 0 |
| 22:12 | <p>Transmit E2E HopID</p> <p>For a Receive Descriptor Ring sending an E2E Credit Grant Packet:</p> <p>This field specifies the value to be inserted into the HopID field of an E2E Credit Grant Packet for this Receive Descriptor Ring. See Table 12-2.</p> <p>For a Receive Descriptor Ring receiving an E2E Credit Grant Packet:</p> <p>This field specifies the Transmit Descriptor Ring that is the target of E2E Credit Grant Packets received by this Receive Descriptor Ring. See Section 12.2.1.1.</p> <p>This field is only valid if the <i>E2E Flow Control Enable</i> bit in this register is set to 1b.</p> <p><i>Note: This field is only valid when the Ring Valid bit is 1b.</i></p> | R/W | 0 |
| 27:23 | Reserved | Rsvd | 0 |
| 28 | <p>E2E Flow Control Enable</p> <p>When set to 0b, end-to-end flow control is disabled for this Receive Descriptor Ring.</p> <p>When set to 1b, end-to-end flow control is enabled for this Receive Descriptor Ring.</p> | R/W | 0b |
| 29 | <p>No-snoop flag (NS)</p> <p>Determines the value to be set by the Host Interface Adapter Layer in the No Snoop attribute of a PCIe TLP associated with this Descriptor Ring.</p> | R/W | 0b |
| 30 | <p>Raw Mode (RAW)</p> <p>When set to 0b, the Descriptor Ring shall operate in Frame Mode.</p> <p>When set to 1b, the Descriptor Ring shall operate in Raw Mode.</p> | R/W | 0b |
| 31 | <p>Ring Valid</p> <p>When set to 0b, the Descriptor Ring is disabled.</p> <p>When set to 1b, the Descriptor Ring is enabled.</p> | R/W | 0b |

12.6.3.3.6 PDF Bit Masks

Table 12-26. PDF Bit Masks Register

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|---------------|
| 15:0 | <p>EOF PDF Bitmask</p> <p>This field specifies the <i>PDF</i> values that shall be interpreted as <i>End of Frame PDF</i>.</p> <p>If the <i>Raw Mode</i> bit is set to 0b, then the Host Interface Adapter Layer shall treat a received Transport Layer Packet as an End of Frame packet if bit <i>i</i> ($i=1,\dots,14$) in this register is set to 1b, where <i>i</i> is the <i>PDF</i> value of the received Transport Layer Packet.</p> <p>If the <i>Raw Mode</i> bit is set to 1b, this field shall have no effect.</p> <p>A Host shall not set the same <i>PDF</i> value as both <i>EOF PDF</i> and <i>SOF PDF</i>. A Host shall always set bit 0 and bit 15 to 0b.</p> | R/W | 0 |
| 31:16 | <p>SOF PDF Bitmask</p> <p>This field specifies the <i>PDF</i> values that shall be interpreted as <i>Start of Frame PDF</i>.</p> <p>If the <i>Raw Mode</i> bit is set to 0b, then the Host Interface Adapter Layer shall treat a received Transport Layer Packet as a Start of Frame Packet if bit <i>i</i> ($i=17,\dots,30$) in this register is set to 1b, where ($i-16$) is the <i>PDF</i> value of the received Transport Layer Packet.</p> <p>If the <i>Raw Mode</i> bit is set to 1b, this field shall have no effect.</p> <p>A Host shall not set the same <i>PDF</i> value as both <i>SOF PDF</i> and <i>EOF PDF</i>.</p> <p>A Host shall always set bit 16 and bit 31 to 0b.</p> | R/W | 0 |

12.6.3.4 Interrupts

12.6.3.4.1 Interrupt Status

The Interrupt Status Registers illustrated in Figure 12-8, contain an *Interrupt Status* bit for each interrupt cause. The *Interrupt Status* bits are organized as a packed array of Doublewords. The last Doubleword may not necessarily be fully populated.

The first *N* bits are allocated to transmit Data Buffer interrupts, where each Transmit Descriptor Ring has its own interrupt. The second *N* bits are allocated to receive Data Buffer interrupts where each Receive Descriptor Ring has its own interrupt. The last *N* bits are allocated for Receive Descriptor Ring Vacancy interrupts, where each Receive Descriptor Ring has its own interrupt.

Figure 12-8. Structure of the Interrupt Status Registers




Table 12-27. Interrupt Status

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|-----------------|---------------|
| 31:0 | <p>Interrupt Status bits</p> <p>Setting a bit in this register indicates that an interrupt is pending for the corresponding interrupt cause.</p> <p>If the <i>Disable ISR Auto-Clear</i> bit is set to 0b, then a read operation returns the current value and then clears the register to 0. If the <i>Disable ISR Auto-Clear</i> bit is set to 1b, then a read operation returns the current value and does not change the register contents.</p> <p>When a bit is set to 1b in the Interrupt Status Clear Registers, the corresponding bit in the Interrupt Status Registers shall be set to 0b.</p> <p>When a bit is set to 1b in the Interrupt Status Set Registers, the corresponding bit in the Interrupt Status Registers shall be set to 1b.</p> <p>A Host shall not write to this register.</p> | See Description | 0 |

12.6.3.4.2 Interrupt Status Clear

The Interrupt Status Clear Registers have the same structure as the Interrupt Status Registers. Each bit in the Interrupt Status Clear Registers corresponds to a bit at the same relative location in the Interrupt Status Registers.

Table 12-28. Interrupt Status Clear

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|----------------|
| 31:0 | <p>Interrupt Status Clear bits</p> <p>When a bit is set to 1b, the corresponding bit in the Interrupt Status Registers is set to 0b.</p> <p>Writing 0b to a bit in this register has no effect.</p> | WO | Vendor Defined |

12.6.3.4.3 Interrupt Status Set

The Interrupt Status Set Registers have the same structure as the Interrupt Status Registers. Each bit in the Interrupt Status Set Registers corresponds to a bit at the same relative location in the Interrupt Status Registers.

Table 12-29. Interrupt Status Set

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|----------------|
| 31:0 | <p>Interrupt Status Set bits</p> <p>When a bit is set to 1b, the corresponding bit in the Interrupt Status Registers is set to 1b.</p> <p>Writing 0b to a bit in this register has no effect.</p> | WO | Vendor Defined |

12.6.3.4.4 Interrupt Mask

The Interrupt Mask Registers have the same structure as the Interrupt Status Registers. Each bit in the Interrupt Mask Registers corresponds to a bit at the same relative location in the Interrupt Status Registers.

Table 12-30. Interrupt Mask

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|---------------|
| 31:0 | Interrupt Mask bits When a bit is set to 0b, the corresponding interrupt cause is masked and does not generate an interrupt. When a bit is set to 1b, the corresponding interrupt cause generates an interrupt. | R/W | 0 |

12.6.3.4.5 Interrupt Mask Clear

The Interrupt Mask Clear Registers have the same structure as the Interrupt Mask Registers. Each bit in the Interrupt Mask Clear Registers corresponds to a bit at the same relative location in the Interrupt Mask Registers.

Table 12-31. Interrupt Mask Clear

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|----------------|
| 31:0 | Interrupt Mask Clear bits Writing 1b to a bit in this register sets the corresponding bit in the Interrupt Mask Registers to 0b. Writing 0b to a bit in this register does not have any effect. | WO | Vendor Defined |

12.6.3.4.6 Interrupt Mask Set

The Interrupt Mask Set Registers have the same structure as the Interrupt Mask Registers. Each bit in the Interrupt Mask Set Registers corresponds to a bit at the same relative location in the Interrupt Mask Registers.

Table 12-32. Interrupt Mask Set

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|----------------|
| 31:0 | Interrupt Mask Set bits Writing 1b to a bit in this register sets the corresponding bit in the Interrupt Mask Registers to 1b. Writing 0b to a bit in this register does not have any effect. | WO | Vendor Defined |

12.6.3.4.7 Interrupt Throttling Rate (ITR)

Each of the ITR registers defines one MSI-X vector. ITR register n at offset $38C00h + 4*n$ corresponds to MSI-X vector n ($n=0, \dots, 15$).

Table 12-33. Interrupt Throttling Rate (ITR)


| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|---------------|
| 15:0 | Interval Defines the initial value of the <i>Counter</i> field in increments of 256 ns. When set to 0, the interrupt throttling mechanism is disabled. | R/W | 0 |
| 31:16 | Counter A decrementing counter for interrupt throttling. When initialized following the issue of an interrupt, it starts counting with the value written to the <i>Interval</i> field and stops at zero. When written to, starts counting from the value written. | R/W | 0 |

12.6.3.4.8 Interrupt Vector Allocation (IVAR)

The Interrupt Vector Allocation Registers, illustrated in Figure 12-9, contain a 4-bit MSI-X vector number per each interrupt cause. The MSI-X vector values are organized as a packed array of Doublewords. The last Doubleword may not necessarily be fully populated.

Each *MSI-X Vector* field correspond to one interrupt cause. The entry for interrupt cause *n* is located in DW $\lfloor n / 8 \rfloor$ starting at bit $4 * [n - 8 * \lfloor n / 8 \rfloor]$, where the interrupt causes are ordered as in the Interrupt Status Registers.

Figure 12-9. Structure of the Interrupt Vector Allocation Registers (IVAR)



$$m = \text{ceiling}(3N/8) - 1$$


Table 12-34. Interrupt Vector Allocation (IVAR)

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|---|------|---------------|
| 3:0 | Interrupt Vector Allocation (IVAR) Identifies the MSI-X vector issued when the corresponding interrupt cause is asserted. | R/W | 0 |
| ... | ... | ... | ... |
| 31:28 | Interrupt Vector Allocation (IVAR) Identifies the MSI-X vector issued when the corresponding interrupt cause is asserted. | R/W | 0 |

12.6.3.4.9 Receive Ring Vacancy Control

The Receive Ring Vacancy Control Registers, illustrated in Figure 12-10, contain a 4-bit threshold value for each Receive Descriptor Ring. The values are organized as a packed array of Doublewords. The last Doubleword may not necessarily be fully populated.

Each *Threshold* field corresponds to one Receive Descriptor Ring. The entry for Receive Descriptor Ring *n* is located in DW $\lfloor n / 8 \rfloor$ starting at bit $4 * [n - 8 * \lfloor n / 8 \rfloor]$.

Figure 12-10. Structure of the Receive Ring Vacancy Control Register

$$m = \text{ceiling}(N/8) - 1$$

Table 12-35. Receive Ring Vacancy Control

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|---------------|
| 3:0 | <p>Threshold</p> <p>Identifies the vacancy threshold. An interrupt is issued when the corresponding Receive Descriptor Ring vacancy is less than or equal to the threshold (see Section 12.5.1).</p> <p>A value of 0 means the Receive Descriptor Ring does not have any available Receive Descriptors.</p> <p>A value of j ($j > 0$) corresponds to $2^{(j-1)}$ available Receive Descriptors.</p> | R/W | 0 |
| ... | ... | ... | ... |
| 31:28 | <p>Threshold</p> <p>Identifies the vacancy threshold. An interrupt is issued when the corresponding Receive Descriptor Ring vacancy is less than or equal to the threshold (see Section 12.5.1).</p> <p>A value of 0 means the Receive Descriptor Ring does not have any available Receive Descriptors.</p> <p>A value of j ($j > 0$) corresponds to $2^{(j-1)}$ available Receive Descriptors.</p> | R/W | 0 |

12.6.3.4.10 Receive Ring Vacancy Status

The Receive Ring Vacancy Status Registers contain a bit for each Receive Descriptor Ring. The *Receive Ring Vacancy Status* bits are organized as a packed array of Doublewords. The last Doubleword may not necessarily be fully populated.

Each *Receive Ring Vacancy Status* bit corresponds to one Receive Descriptor Ring. Bit n corresponds to Receive Descriptor Ring n .

Table 12-36. Receive Ring Vacancy Status

| Bit(s) | Field Name and Description | Type | Default Value |
|--------|--|------|---------------|
| 31:0 | <p>Receive Ring Vacancy Status bits</p> <p>Setting a bit in this register indicates that the number of unused Receive Descriptors for a Receive Descriptor Ring is less than or equal to the threshold defined by the respective field in the Receive Ring Vacancy Control registers. See Section 12.5.1.</p> | RO | 0 |

12.7 Timing Parameters

Table 12-37 lists the timing parameters for the Host Interface Adapter Layer.

Table 12-37. Host Interface Timing Parameters

| Parameter | Description | Min | Max | Units |
|-----------|--|-----|------|-------|
| tE2ERate | The time interval between periodic E2E Credit Grant Packets. | 1 | 1000 | ms |
| tE2ESync | The time interval between E2E Credit Sync Packets. | 10 | 20 | sec |
| tHIReset | Time to complete a Host Interface Reset. | - | 10 | ms |

13 Interoperability with Thunderbolt™ 3 (TBT3) Systems

This chapter defines requirements for a USB4 Product to be TBT3-Compatible. A USB4 Product that is TBT3-Compatible can operate in a USB4™ Fabric that includes any combination of the following:

- A Thunderbolt 3 Router.
- A Thunderbolt 3 Connection Manager.
- A Thunderbolt 3 Active Cable.

A USB4 host and USB4 peripheral device may optionally support TBT3-Compatibility. If a USB4 host or USB4 peripheral device supports TBT3-Compatibility, it shall do so as defined in this chapter.

A USB4 hub shall support TBT3-Compatibility as defined in this chapter. A USB4 hub shall support TBT3-Compatibility on all of its DFP. If the USB4 hub is a USB4-Based Dock, it shall support TBT3-Compatibility on its UFP in addition to all its DFP.

A Router uses the *TBT3 Not Supported* bit to indicate whether or not it is TBT3-Compatible. This chapter does not apply to a Router with the *TBT3 Not Supported* bit set to 1b.

13.1 Electrical Layer

When TBT3 Mode is established, an Adapter shall run at a TBT3-Compatible speed.

A Router Assembly shall support TBT3-Compatible Gen 2 Speed (10.3125Gbps). A Router Assembly may also optionally support TBT3-Compatible Gen 3 Speed (20.625Gbps).

A Router Assembly may transmit with either USB4 SSC or TBT3 SSC. A Router Assembly receiver must tolerate both USB4 SSC and TBT3 SSC.

A Router Assembly shall meet the specifications described in chapter 3, except for the following set of parameters that shall be used instead of the values specified in chapter 3:

Table 13-1. Thunderbolt 3 Parameters

| Parameter Name | Min | Max | Units |
|------------------------------|---------------|---------|-------|
| USB4_SSC_DOWN_SPREAD_RATE_TX | See Table 3-3 | | KHz |
| USB4_SSC_PHASE_DEVIATION_TX | See Table 3-3 | | ns |
| TBT3_SSC_DOWN_SPREAD_RATE_TX | 35 | 37 | KHz |
| TBT3_SSC_PHASE_DEVIATION_TX | -- | 18.5 | ns |
| SSC_DOWN_SPREAD_RATE_RX | 30 | 37 | KHz |
| SSC_PHASE_DEVIATION_RX | 2.5 | 22 | ns |
| UI (Gen 2 Speed) | 96.9406 | 96.9988 | ps |
| UI (Gen 3 Speed) | 48.4703 | 48.4994 | ps |

The transmit average UI (measured over windows at the size of one SSC cycle) shall be at the range of 97.1348ps to 97.2420ps.

The receiver shall tolerate input signals with maximum SSC down spreading of 5800ppm.

Additional receiver “Case 2a” test setup shall be supported, addressing optical interconnects with limiting modules. Case 2a setup is identical to Case 2 setup described in section 3.5.2, except that the passive cable is replaced with worst-case limiting optical cable (corresponding to the USB Type-C Specification). A receiver shall operate at BER of 1E-12 or lower with neither Forward Error Correction nor Pre-Coding applied when a stressed signal is driven at its input.

13.2 Logical Layer

13.2.1 Sideband Channel


This section applies to a USB4 Port that is TBT3-Compatible.

13.2.1.1 Bidirectional Re-timer

A TBT3 Active Cable can contain two bidirectional Re-timers. A bidirectional Re-timer is a TBT3 Cable Re-timer that sends AT Responses on its SBRX wire. A Router shall implement a bidirectional Sideband Channel when attached directly to a bidirectional Cable Re-timer (i.e. when there are no On-Board Re-timers between the Router and the bidirectional Re-timer). A Router shall implement a unidirectional Sideband Channel when not attached directly to a bidirectional Cable Re-timer.

Figure 13-1 depicts SBTX and SBRX connectivity between two Routers that are connected via an Active Cable with bidirectional Re-timers. The Sideband Channel between Router A and Re-timer A in Figure 13-1 allows Router A to receive transactions from Re-timer A on SBTX and receive transactions from Router B on SBRX.

Figure 13-1. Bidirectional Re-timer Topology



A Router that is connected directly to a bidirectional Re-timer shall support concurrent reception of Transactions on SBTX and on SBRX.

A Router shall drive its SBTX for up to 2 bit times after the last Stop bit of an AT Command.

13.2.1.2 Transactions

This section applies to a TBT3-Compatible Sideband Channel.

13.2.1.2.1 LT Transactions

A Router shall support the additional LT Transaction types defined in Table 13-2.

Table 13-2. TBT3 LT Transaction Types

| Bit(s) | Name | Function |
|--------|------------------|--|
| [3:0] | <i>LSESymbol</i> | This field defines the LT Transaction Type. Undefined Values are Rsvd. 0001b: LT_Gen_2 (Gen 2 speed selected) 0101b: LT_Gen_3 (Gen 3 speed selected) 0110b: LT_Resume2 (Cable Re-timer started USB4 transmission in the direction of the cable) |

13.2.1.2.2 AT Transactions

The structure of the STX Symbol within an AT Transaction shall be as defined in Table 13-3.

Table 13-3. STX Symbol

| Bits | Name | Function |
|-------|---------------------|---|
| [7:6] | <i>StartAT</i> | See Table 4-4. |
| [5] | <i>Rsvd</i> | See Table 4-4. |
| [4] | <i>Responder</i> | Indicates that the transaction originated from a responding Re-timer in an active cable with bidirectional Re-timers. Shall be set to 1b by a Re-timer in an AT Response when operating in an active cable with bidirectional Re-timers. Shall be set to 0b in all other cases. |
| [3] | <i>Bounce</i> | Together with the <i>ReturnBounce</i> bit, serves to forward an AT Transaction from a Router to a Cable Re-timer at the remote end of a cable that can only be accessed by its Link Partner. See Section 13.2.1.2.2.1 for more details. |
| [2] | <i>Recipient</i> | Identifies the intended final recipient of the transaction. For an AT Command: shall be 0b if a Cable Re-timer is the intended final recipient or 1b if a Router is the intended final recipient. For an AT Response: shall be set to 1b. |
| [1] | <i>ReturnBounce</i> | Together with the <i>Bounce</i> bit, serves to forward an AT Transaction from a Router to a Cable Re-timer at the remote end of a cable that can only be accessed by its Link Partner. See Section 13.2.1.2.2.1 for more details. |
| [0] | <i>CmdNotResp</i> | See Table 4-4. |

A Router shall not send an AT Command that targets Register 13 of a Re-timer or Router SB Register Space (see Table 4-17) unless the Re-timer or Router is directly attached to it.

13.2.1.2.2.1 Bounce Mechanism

The Bounce mechanism is used when a Router needs to access the Register Space of a Cable Re-timer that can only be accessed by its Link Partner. A Router shall support the Bounce Mechanism. The Bounce Mechanism consists of the following rules:


- A Router shall set the *Bounce* bit to 1b and the *ReturnBounce* bit to 1b to target a Cable Re-timer that is adjacent to the Router's Link Partner.
- A Router that receives an AT Transaction with the *Bounce* bit set to 1b and the *ReturnBounce* bit to 1b shall set the *Bounce* bit to 0b, then forward the AT Transaction towards its adjacent Cable Re-timer.

Note: A Cable Re-timer responds to an AT Transaction with the *Bounce* bit set to 0b and the *ReturnBounce* bit to 1b. The AT Response from the Cable Re-timer has the *Bounce* Bit set to 1b and the *ReturnBounce* bit set to 0b.

- A Router that receives an AT Response with the *Bounce* bit set to 1b and the *ReturnBounce* bit to 0b shall set the *Bounce* bit to 0b, then forward the AT Response to its Link Partner.

An example of the Bounce Mechanism is shown in Figure 13-2 where Router A is accessing the configuration space of Cable Re-timer B.

Figure 13-2. Bounce Mechanism



13.2.1.2.3 RT Transactions

Byte 2 in a Broadcast RT Transaction shall have the format in Table 4-7 with the changes in Table 13-4.

Table 13-4. Contents of Byte 2 in a Broadcast RT Transaction

| Bits | Name | Function |
|------|-----------------------------|--|
| 0 | <i>USB4</i> | When sending a Broadcast RT Transaction on a USB4 Sideband Channel, the <i>USB4</i> bit shall be set to 1b. When sending a Broadcast RT Transaction on a TBT3-Compatible Sideband Channel, the <i>USB4</i> bit shall be set to 0b. The <i>USB4</i> bit shall match the <i>USB4-Compatible Sideband Channel</i> attribute defined in Table 13-7. |
| 4 | <i>TBT3-CompatibleSpeed</i> | Set to 1b. |

13.2.1.3 SB Register Space

A Router shall support the additional registers and register fields in Table 13-5 and Table 13-6.

Table 13-5. SB Registers

| Register | Size (Bytes) | Name | Description |
|----------|--------------|------|--|
| 3 | 4 | Mode | A FourCC string indicating the operating mode of the Router. |
| 4 | 4 | Type | A FourCC string indicating the type of USB4 device. |
| 10 | 1 | Rsvd | Reserved. |

Table 13-6. SB Registers Fields

| Register | Name | Byte | Bits | Description | Type |
|----------|------|------|------|--|------|
| 3 | Mode | 0 | 7:0 | Contains the ASCII value "A" = 41h | RO |
| 3 | Mode | 1 | 7:0 | Contains the ASCII value "P" = 50h | RO |
| 3 | Mode | 2 | 7:0 | Contains the ASCII value "P" = 50h | RO |
| 3 | Mode | 3 | 7:0 | Contains the ASCII value "Space" = 20h | RO |
| 4 | Type | 0 | 7:0 | Contains the ASCII value "E" = 45h | RO |
| 4 | Type | 1 | 7:0 | Contains the ASCII value "M" = 4Dh | RO |
| 4 | Type | 2 | 7:0 | Contains the ASCII value "Space" = 20h | RO |
| 4 | Type | 3 | 7:0 | Contains the ASCII value "Space" = 20h | RO |
| 10 | Rsvd | 1 | 7:0 | Contains the value 00h | RO |

13.2.1.4 Lane Initialization**13.2.1.4.1 Phase 1 – Determination of Initial Conditions**

During phase 1, Router A determines whether or not TBT3 Mode is established on the Link and whether or not the Cable connected supports CLx states. See the USB PD Specification and the USB Type-C Specification for how to determine if TBT3 Mode is established and if the Cable supports CLx states.

A Router shall not continue on to Phase 2 until it has obtained the connection information described in this section and in Section 4.1.2.1. If TBT3 Mode is established on the Link, a Router shall proceed with Lane Initialization as defined in Section 4.1.2 with the changes defined in this chapter.

13.2.1.4.2 Phase 3 – Determination of USB4 Port Characteristics

A Router shall decide the Lane attributes using the decision criteria in Table 4-18 with the changes defined in Table 13-7.

Table 13-7. Lane Attributes

| Attribute | Action |
|-----------|--|
| RS-FEC | <p>Gen 2 speed:</p> <p>Router A shall enable RS-FEC if at least one side of the Link requests it (i.e. the <i>RS-FEC Request (Gen 2)</i> bit is set to 1b in the SB Register Space of the local USB4 Port and/or its Link Partner). Otherwise, RS-FEC shall not be enabled.</p> <p>Router A shall set the <i>RS-FEC Enabled (Gen 2)</i> bit in the USB4 Port Capability to reflect whether it is operating with RS-FEC.</p> |

| Attribute | Action |
|----------------------------------|--|
| USB4-Compatible Sideband Channel | <p>The Sideband Channel of Router A shall operate as a USB4 Sideband Channel (as defined in Section 4.1) if all of the following are true:</p> <ul style="list-style-type: none"> • The <i>USB4 Sideband Channel Support</i> bit in the SB Register Space of both Routers is 1b. • The Link is over either a Passive Cable or an Active Cable with unidirectional Re-timers. <p>Else, the Sideband Channel of Router A shall operate as a TBT3-Compatible Sideband Channel (as defined in Section 13.2).</p> |

13.2.1.4.3 Phase 4 – Lane Parameters Synchronization and Transmit Start

For a TBT3-Compatible Sideband Channel, this section replaces Step 1 in Section 4.1.2.4.

1. Router A shall do the following for each enabled Lane to indicate that it is ready to start transmission on a given Lane:
 - If the Router Assembly for Router A does not include any On-Board Re-timers, and if operating at Gen 2 speed, Router A shall send an LT_Gen_2 Transaction for each enabled Lane every tLaneParams. Router A shall continue sending LT_Gen_2 Transactions until all of the following are true, then continue to Step 2:
 - At least tLTPHase4 time has passed from completion of Phase 2.
 - Router A has sent LT_Gen_2 Transactions at least twice.
 - Router A has received an LT_Gen_2 Transaction from Router B.
 - If the Router Assembly for Router A does not include any On-Board Re-timers, and if operating at Gen 3 speed, Router A shall send an LT_Gen_3 Transaction for each enabled Lane every tLaneParams. Router A shall continue sending LT_Gen_3 Transactions until all of the following are true, then continue to Step 2:
 - At least tLTPHase4 time has passed from completion of Phase 2.
 - Router A has sent LT_Gen_3 Transactions at least twice.
 - Router A has received an LT_Gen_3 Transaction from Router B.
 - If the Router Assembly for Router A includes one or more On-Board Re-timers, and if operating at Gen 2 speed, Router A shall send a Broadcast RT Transaction every tLaneParams. The Broadcast RT Transaction shall have the parameter values in Table 4-18. Router A shall also send an LT_Gen_2 Transaction for each enabled Lane.
 - Router A shall continue sending the Broadcast RT and LT_Gen_2 Transactions until all of the following conditions are true:
 - At least tLTPHase4 time has passed from completion of Phase 2.
 - Router A has sent LT_Gen_2 Transactions at least twice.
 - Router A has received an LT_Gen_2 Transaction from Router B.
 - If the Router Assembly for Router A includes one or more On-Board Re-timers, and if operating at Gen 3 speed, Router A shall send a Broadcast RT Transaction every tLaneParams. The Broadcast RT Transaction shall have the parameter values in Table 4-18. Router A shall also send an LT_Gen_3 Transaction for each enabled Lane.
 - Router A shall continue sending the Transactions until all of the following conditions are true:
 - At least tLTPHase4 time has passed from completion of Phase 2.
 - Router A has sent LT_Gen_3 Transactions at least twice.
 - Router A has received an LT_Gen_3 Transaction from Router B.

13.2.1.4.4 Phase 5 – Link Equalization

For a TBT3-Compatible Sideband Channel, this section replaces Section 4.1.2.5.

The fifth phase of Lane Initialization consists of negotiating transmitter Feed-Forward Equalization (TxFFE) parameters between each Router or Re-timer and a Router or Re-timer adjacent to it. During TxFFE negotiation, the receiver cycles through one or more of the potential preset numbers defined in Table 3-5. The receiver examines its behavior for each preset and selects one preset value to use. The receiver may follow any order while cycling through the preset numbers.

The TxFFE flow that Router A's transmitter performs depends on whether Router A is connected to an On-Board Re-timer, a Cable Re-timer, or another Router:

- If Router A connects to an On-Board Re-timer in the same Router Assembly, then Router A's transmitter shall perform the transmitter flow in the symmetric equalization flow defined in Section 4.1.2.5.1. The Router and Re-timer use Addressed RT Transactions (with the *Index* field set to 0) to access each other's SB Register Space.
- If Router A connects to a Cable Re-timer:
 - Router A's transmitter shall perform the Primary Partner flow in the Asymmetric TxFFE Parameter Negotiation with a transmitting Primary Partner defined in Section 13.2.1.4.4.1. The Router uses AT Transactions (with the *Recipient* bit set to 0b) to access the Re-timer's SB Register Space.
 - Once a transmitter completes TxFFE negotiation with the Cable Re-timer's receiver, Router A shall send an LT_Resume2 Transaction on the USB4 Port that completed negotiation with the *LSELane* field matching the Lane number that completed negotiation.
- If Router A either connects directly to Router B or connects to an On-Board Re-timer in the Router Assembly of Router B (through a Passive Cable or Cable with re-driver), then Router A's transmitter shall perform the transmitter flow in the symmetric equalization flow defined in Section 4.1.2.5.1. Router A uses AT Transactions (with the *Recipient* bit set to 1b) to access the TxFFE Register of the adjacent USB4 Port.

The TxFFE flow that Router B's receiver performs depends on whether Router B is connected to an On-Board Re-timer, a Cable Re-timer, or another Router. Router B polls the *Tx Active* bit in the SB Register Space of the adjacent Router or Re-timer to identify when the transmitter of the adjacent Router or Re-timer is on.

- If Router B connects to an On-Board Re-timer, then Router B's receiver shall perform the receiver flow in the symmetric equalization flow defined in Section 4.1.2.5.1. The Router and Re-timer use Addressed RT Transactions (with the *Index* bit set to 0b) to access each other's SB Register Space.
 - When a receiver's equalization flow is complete on a Lane, the Router shall set the Lane's *Clock Switch Done* bit to 1b.
- If Router B connects to a Cable Re-timer, then Router B's receiver shall perform the Primary Partner flow in the Asymmetric TxFFE Parameter Negotiation with a Receiving Primary Partner defined in Section 13.2.1.4.4.2. The Router uses AT Transactions (with the *Recipient* bit set to 0b) to access the Re-timer's SB Register Space.
- If Router B either connects directly to Router A or connects to an On-Board Re-timer in the Router Assembly of Router A, then Router B's receiver shall perform the receiver flow in the symmetric equalization flow defined in Section 4.1.2.5.1. Router B uses AT Transactions (with the *Recipient* bit set to 1b) to access the SB Register Space of the adjacent USB4 Port.

The equalization flow for Re-timers is defined in the USB4 Re-Timer Specification.

13.2.1.4.4.1 Asymmetric TxFFE Parameter Negotiation with a Transmitting Primary Partner

During TxFFE negotiation, the transmitter negotiates TxFFE parameters with the receiver at the other end. The transmitting end is defined as the Primary Partner and it is capable of issuing AT Commands. The receiving end is defined as the Subordinate Partner and it is capable of issuing AT Responses.

Transmitting Primary Partner flow:

The steps that the transmitter shall perform are listed below:

1. The transmitter shall start with the *Tx Active* bit set to 1b (default value) in the *Tx Status* byte of the TxFFE register.
2. The transmitter shall send an AT Transaction with a write Command to the receiver that sets the *Tx Active* bit to 1b in the *Partner Tx Status* byte in the TxFFE register.
3. The transmitter shall read the local *Rx Status & TxFFE Request* byte from the receiver. To do this, the transmitter sends an AT Transaction with a read Command targeting the TxFFE register of the receiver.
4. On reception of an AT Response from the receiver, the transmitter shall copy the transaction contents into its *Rx Status & TxFFE Request* byte and take the following actions:
 - If *Rx Locked* = 1b, then negotiation is complete and no further TxFFE negotiation steps shall be taken.
 - Else if *New Request* = 0b and *TxFFE Request* is the same as the previous *TxFFE Request*, the receiver has not provided a new request yet. The Router shall go to Step 3. The Router shall perform Step 3 within tPollTXFFE of receiving the AT Response.
 - Else, this is a new request to update TxFFE parameters. Continue on to Step 5.
5. The transmitter shall update its transmitter parameters based on the new parameters received in the AT Response. To do this, the transmitter loads one of 16 predefined TxFFE configurations that matches the *TxFFE Request* field in the *Rx Status & TxFFE Request* byte.
 - If both Lane Adapters in the Port are enabled and have not yet completed TxFFE negotiation, both transmitters must complete Step 5 before continuing to Step 6. A Lane Adapter shall not continue to Step 6 until one of the following occurs:
 - i. The other Lane Adapter completes Step 5.
 - ii. tTxFFETimeout has passed since starting Step 5.
6. The transmitter shall inform the receiver that it has updated to new parameters by sending an AT Transaction with a write Command to the receiver targeting its *Partner Tx Status* byte in the TxFFE register with the following contents:
 - *Tx Active* = 1b.
 - *TxFFE Setting* = value received in Step 4.
7. The transmitter shall read the local *Rx Status & TxFFE Request* byte from the receiver. To do this, the transmitter sends an AT Transaction with a read Command targeting the TxFFE register of the receiver.
8. On reception of an AT Response from the receiver, the transmitter shall copy the transaction contents into its *Rx Status & TxFFE Request* byte and take the following actions:
 - If *New Request* = 1b and *TxFFE Request* is the same as the previous *TxFFE Request*, the Router shall return to and perform Step 7 within tPollTXFFE of receiving the AT Response.
 - Else, go to Step 3.

Note: The Partner Tx Status byte is defined in the USB4 Re-Timer Specification.

13.2.1.4.4.2 Asymmetric TxFFE Parameter Negotiation with a Receiving Primary Partner

During TxFFE negotiation, the receiver negotiates TxFFE parameters with the transmitter at the other end. The receiving end is defined as the Primary Partner and can issue AT Commands. The transmitting end is defined as the Subordinate Partner and can issue AT Responses.

Receiving Primary Partner flow:

The receiver shall execute the following flow with the transmitter:

1. The receiver shall start with the following default values in the *Rx Status & TxFFE Request* byte of the TxFFE register:
 - *Rx Locked* = 0b.
 - *New Request* bit = 0b.
 - *Rx Active* bit = 0b.
2. The receiver shall read the *Local Tx Status* byte of the transmitter.
 - The receiver sends an AT Command with a read Command to the *Local Tx Status* byte of the TxFFE register of the transmitter.
 - On reception of an AT Response from the transmitter, the receiver shall do the following:
 - If *Tx Active* = 1b (i.e. the transmitter is transmitting), then enable the receiver, set *Rx Active* to 1b, and go to Step 3.
 - Else, repeat Step 2 within tPollTXFFE of receiving the AT Response.
3. The receiver shall evaluate its receiver behavior. If equalization is complete, the receiver shall set the *Rx Locked* field to 1b.
4. The receiver shall do the following:
 - If *Rx Locked* = 1, then go to Step 5.
 - Else, go to Step 6.
5. TxFFE negotiation is complete.
6. The receiver shall select a new set of TxFFE parameters.
7. The receiver shall write the *Partner Rx Status & TxFFE Request* byte at the transmitter as follows:
 - If the Transmitting Primary Partner finished TxFFE on both Lanes, then the receiver sends an AT Command with a write Command targeting the *Partner Rx Status & TxFFE Request* byte of the TxFFE register of the transmitter. Else, the receiver will wait for the next AT Command with a write Command to the TxFFE Register from the Transmitting Primary Partner and use it to write the *Partner Rx Status & TxFFE Request* byte of the TxFFE register of the transmitter. The AT command shall write to following values to the following fields:
 - *New Request* = 1b.
 - *Rx Active* = 1b.
 - *TxFFE Request* = index of selected set of TxFFE parameters.
8. The receiver shall wait for a write Response indicating the transmitter is using the new requested TxFFE settings.
9. The receiver shall evaluate its receiver behavior. If equalization is complete, the receiver shall set the *Rx Locked* field to 1b.
10. The receiver shall write the *Partner Rx Status & TxFFE Request* byte at the transmitter as follows:

- If the Transmitting Primary Partner finished TxFFE on both Lanes, then the receiver sends an AT Command with a write Command targeting the *Partner Rx Status & TxFFE Request* byte of the TxFFE register of the transmitter.
- Else, the receiver will wait for the next AT Command with a write Command to the TxFFE Register from the Transmitting Primary Partner and use it to write the *Partner Rx Status & TxFFE Request* byte of the TxFFE register of the transmitter. The AT Command shall write to following values to the following fields:
 - *New Request* = 0b.
 - *Rx Active* = 1b.

11. The receiver shall do the following:

- If *Rx Locked* = 1, then go to Step 5.
- Else, go to Step 6.

Note: The Transmitting Primary Partner and the Receiving Primary Partner are both on the same Router.

Note: The Partner Rx Status & TxFFE Request byte is defined in the USB4 Re-Timer Specification.

13.2.2 Logical Layer State Machine

13.2.2.1 CLd State

13.2.2.1.1 Behavior in State

A USB4 Port that is TBT3-Compatible shall support the behavior defined in this section in addition to the behavior described in Section 4.2.1.2.2.

A Lane Adapter (that is not the Upstream Adapter) that enters this state due to reception of an LT_Fall Transaction shall start Lane Initialization when it receives a Broadcast RT Transaction, an LT_Gen_2 Transaction, or an LT_Gen_3 Transaction.

- A Lane 0 Adapter shall not start Lane Initialization until it receives one of the following:
 - A Broadcast RT Transaction with the *Lane0Enabled* bit set to 1b.
 - An LT_Gen_2 Transaction with the *LSELane* bit set to 0b.
 - An LT_Gen_3 Transaction with the *LSELane* bit set to 0b.
- A Lane 1 Adapter shall not start Lane Initialization until it receives one of the following:
 - A Broadcast RT Transaction with the *Lane1Enabled* bit set to 1b.
 - An LT_Gen_2 Transaction with the *LSELane* bit set to 1b.
 - An LT_Gen_3 Transaction with the *LSELane* bit set to 1b.

The Lane Adapter shall start Lane Initialization from Phase 4. The Lane Adapter shall maintain any state acquired in Phases 1 through 3 of previous Lane Initialization.

13.2.2.2 TS1 and TS2 Ordered Sets

When operating in TBT3 mode, TS1 and TS2 Ordered Sets shall have the format shown in Table 4-25 with the changes in Table 13-8.

Table 13-8. TS1 and TS2 Ordered Set Structure

| Bits | Name | Description |
|-------|------------------------------|--|
| 28:26 | <i>Lane Bonding Target 2</i> | Lane Bonding Target 2. Transmitter shall either set this value to match the <i>Lane Bonding Target</i> field or shall set this value to 001b. A Receiver shall ignore this field. |

13.2.2.3 Low Power (CL0s, CL1, and CL2)

If the *CLx Protocol Support* bit in the USB4 Port Capability is set to 1b, the Port shall support the low power protocol as defined in Section 4.2.1.6. Else, the Port does not support the low power protocol and shall not initiate low power transactions.

13.2.3 USB4 Link Operation**13.2.3.1 USB4 Link Transitions**

When TBT3 Mode is established on the Link, a USB4 Port shall support the transitions described in Section 4.2.2 with the following changes:

- For a Device Router that supports TBT3 Mode on its Upstream Facing Port, all USB4 Ports shall support operation with two independent Single-Lane Links. Unlike Section 4.2.2, this configuration is not just a transient state between Link Initialization and Lane Bonding. In TBT3 Mode, both Links may operate as independent Single-Lane Links and are configured and managed separately by the Connection Manager.

**CONNECTION MANAGER NOTE**

A Connection Manager shall not set the Lane Bonding bit to 1b if either of the following are true:

- *A Path other than Path 0 is enabled across the Lanes being bonded.*
- *A Lane 1 Adapter is the Upstream Adapter.*

13.2.3.2 Pre-Coding

In addition to the conditions defined in Section 4.3.6.2, Pre-Coding shall be off when the Link uses a TBT3-Compatible Sideband Channel.

13.2.4 Sleep and Wake

This section only applies to a Device Router that supports TBT3-Compatibility on its UFP.

If bits 15:12 in the Connection Manager *USB4 Version* field in the Router Configuration Space are 0b (indicating a TBT3 Connection Manager), a Router shall support Sleep and Wake per Section 4.5 with the changes defined in this section.

13.2.4.1 Entry to Sleep

After the *Enter Sleep* bit is set to 1b in all Ports, a Device Router shall do the following for each USB4 Port:

- Transition the USB4 Adapters to CLd state.
- If any of the following conditions apply, the USB4 Port shall go through disconnect:
 - For Lane 0 in a USB4 Port:
 - The *Lane 0 is Inter-Domain* bit is 0b and the *Lane 0 Configured* bit is 0b.
 - The *Lane 0 is Inter-Domain* bit is 1b and the *Inter-Domain Disconnect on Sleep* bit is set to 1b.

- For Lane 1 in a USB4 Port:
 - The *Lane 1 is Inter-Domain* bit is 0b and the *Lane 1 Configured* bit is 0b.
 - The *Lane 1 is Inter-Domain* bit is 1b and the *Inter-Domain Disconnect on Sleep* bit is set to 1b.

13.2.4.2 Behavior in Sleep State

A Device Router shall retain a copy of the state information listed in Table 13-9 separate from Configuration Space.

Table 13-9. Router State Retained During Sleep

| Field Copied | Configuration Space / Capability |
|----------------------------------|--|
| Upstream Adapter | Router Configuration Space. |
| Lane 0 is Inter-Domain | USB4 Port Region in Vendor Specific Extended 6 Capability of the Router Configuration Space. |
| Lane 1 Configured | USB4 Port Region in Vendor Specific Extended 6 Capability of the Router Configuration Space. |
| Lane 0 is Inter-Domain | USB4 Port Region in Vendor Specific Extended 6 Capability of the Router Configuration Space. |
| Lane 1 is Inter-Domain | USB4 Port Region in Vendor Specific Extended 6 Capability of the Router Configuration Space. |
| Inter-Domain Disconnect on Sleep | USB4 Port Region in Vendor Specific Extended 6 Capability of the Router Configuration Space. |
| Enable Wake on Inter-Domain | USB4 Port Region in Vendor Specific Extended 6 Capability of the Router Configuration Space. |
| Enable Wake Events | USB4 Port Region in Vendor Specific Extended 6 Capability of the Router Configuration Space. |

If Lane 0 of a USB4 Port is disconnected while in Sleep state, then the internal *Lane 0 is Inter-Domain* state and *Lane 0 Configured* state listed in Table 13-9 shall both transition to 0b.

If Lane 1 of a USB4 Port is disconnected while in Sleep state, then the internal *Lane 1 is Inter-Domain* state and *Lane1 Configured* state listed in Table 13-9 shall both transition to 0b.

13.2.4.3 Wake Events

A Device Router shall support all of the wake events listed in the *Enable Wake Events* field of the USB4 Port Region in the Vendor Specific Extended 6 Capability of the Router Configuration Space.

13.2.4.4 Exit from Sleep

A Device Router shall not start Lane Initialization for Lane 0 until the *Start Link Initialization* bit is set to 1b, if either of the following is true on exit from sleep:

- The *Lane 0 Configured* state is set to 0b.
- The *Lane 0 is Inter-Domain* bit is 1b.

A Device Router shall not start Lane Initialization for Lane 1 until the *Start Link Initialization* bit is set to 1b, if either of the following is true on exit from sleep:

- The *Lane 1 Configured* state is set to 0b.
- The *Lane 1 is Inter-Domain* bit is 1b.

13.2.5 Timing Parameters

Table 13-10 lists changes in timing parameters for the Logical Layer.

Table 13-10. Logical Layer Timing Parameters

| Parameter | Description | Min | Max | Units |
|-----------------|--|-----|-----|-------|
| tLTPhase4 | The amount of time that Broadcast RT Transactions, LT_Gen_2 Transactions, or LT_Gen_3 Transactions are sent after completion of Lane Initialization phase 2. | 25 | | ms |
| tLaneParams | The time interval between transmissions of LT_Gen_2 Transactions, between the transmissions of LT_Gen_3 Transactions, or between the transmissions of Broadcast RT Transactions. | 1 | 5 | ms |
| tTxFFETimeout | The amount of time that a Lane Adapter waits for the other Lane Adapter to update its TxFFE parameters before continuing negotiation. | 80 | | ms |
| tTrainingAbort1 | The amount of time in Training state following Lane Initialization. | 2 | -- | sec |

13.3 Transport Layer**13.3.1 Adapter Numbering Rules**

This section only applies to a Device Router that supports TBT3-Compatibility on its UFP.

If bits 15:12 in the *Connection Manager USB4 Version* field in the Router Configuration Space Basic Attributes are 0b (indicating a TBT3 Connection Manager), a Device Router shall expose either one or two USB4 Ports. If the Device Router supports PCIe Tunneling, it shall only expose the PCIe Adapters that are related to the exposed USB4 Ports. The method to signify that an Adapter is invalid is defined in Section 5.2.1. The Lane Adapters in the exposed USB4 Ports shall be assigned consecutive Adapter numbers, starting from 1.

If any of bits 15:12 in the *Connection Manager USB4 Version* field are 1b (indicating a Connection Manager that supports USB4 Version 1.0 or higher) a Device Router may expose additional USB4 Ports and/or additional PCIe Adapters. A Router that exposes additional USB4 Ports and/or additional PCIe Adapters shall do so immediately when the *Connection Manager USB4 Version* field is set to a non-zero value.

13.3.2 Maximum HopID

This section applies to Host Routers and Device Routers.

The *Max Input HopID* and *Max Output HopID* fields in a Lane Adapter shall be at least 15 for USB4 hosts and USB4 hubs.

The *Max Input HopID* and *Max Output HopID* fields in a Lane Adapter of a USB4 device supporting one DisplayPort tunneled stream shall be at least 11.

The *Max Input HopID* and *Max Output HopID* fields in a Lane Adapter of a USB4 device supporting two DisplayPort tunneled streams shall be at least 14.

13.3.3 Connectivity Rules

This section only applies to a Device Router that supports TBT3-Compatibility on its UFP.

A Device Router shall support the Connectivity rules defined in Section 5.2.5 for both the Lane 0 Adapter and Lane 1 Adapter in a USB4 Port. For example, where Section 5.2.5 requires that "A Router shall be able to forward a Control Packet received on any Lane 0 Adapter to the Control Adapter" this section also requires that a Device Router shall be able to forward a Control Packet received on any Lane 1 Adapter to the Control Adapter.

13.3.4 Buffer Allocation

This section only applies to a Device Router that supports TBT3-Compatibility on its UFP.

A TBT3 Connection Manager sets the buffer allocation in a static manner, according to Table 13-11.

Table 13-11. Buffer Allocation by TBT3 Connection Manager

| | Dual Lane Link | Single Lane Link |
|---|---|---|
| Control Path | Assumes 2 buffers | |
| PCIe Path | 32 buffers | 16 buffers |
| Host-to-Host Path | 14 buffers | 6 buffers |
| DP Aux Path | 1 buffer | |
| DP Main-Link Path | The remaining buffers = Total Buffers – 50 ¹ | The remaining buffers = Total Buffers – 26 ¹ |
| <i>Note 1: A TBT3 CM assumes two DP Tunnels through a port, therefore it will keep 2 buffers for DP Aux Path.</i> | | |

13.4 Configuration Layer

This section only applies to Device Routers.

13.4.1 Router Enumeration

On transition to the Uninitialized state, a Router shall:

- Expose USB4 Ports and PCIe Adapters as defined in Section 13.3.1.
- Set its sleep and wake behavior as defined in Section 13.2.4.
- Expose the additional registers defined in Section 13.6.

The Router is enumerated when the *TopologyID Valid* bit is set to 1b.

Following enumeration, if the *Connection Manager USB4 Version* field is set to indicate a TBT3 Connection Manager, then the Router shall maintain the configuration listed above. The Router may set the *Router Ready* bit to 1b.

Following enumeration, if the *Connection Manager USB4 Version* field is set to indicate USB4 Rev. 1.0 support, then the Router shall follow Section 6.7.

13.4.2 Notification Packet

This section only applies to a Device Router that supports TBT3-Compatibility on its UFP.

If bits 15:12 in the *Connection Manager USB4 Version* field in the Router Configuration Space are 0b (indicating a TBT3 Connection Manager), a Hot Plug Acknowledgment Packet may have the *PG* bit set to 00b.

13.4.3 Bit Banging Interface

This section only applies to a Device Router that supports TBT3-Compatibility on its UFP.

A Router shall support the “bit banging” interface defined in Vendor Specific 1 Capability (see Section 13.6.1.1).

A Router shall return the value 00000080h when the 32 bits at addresses [78h:75h] are read.

A Router shall return the value 01h when the byte at address 0148h is read.

A Router shall return the value 00000111h when the 32 bits at addresses [1A7h:1A4h] are read.



CONNECTION MANAGER NOTE

A Connection Manager reads the contents of a DROM attached to a TBT3 Router using the "bit banging" interface defined in Vendor Specific 1 Capability (see Section 13.6.1.1). The DROM Base Address field contains the DROM base address.

13.4.4 Control Packet Routing

13.4.4.1 Downstream-Bound Packets

A Router that receives a Control Packet with the *CM* bit set to 0b, shall route the packet according to the following rules:

- If the packet arrived on the Upstream Adapter, then:
 - If the Router is a Host Router and the *TopologyID Valid* bit in Router Configuration Space is set to 0b, then the Router shall process the packet using the Uninitialized Router Flow in Section 13.4.4.2.
 - Else the Router shall extract the Egress Adapter number from the Route String that corresponds to the Router's depth in the Spanning Tree (as present in the *Depth* field in the Router Configuration Space). For example, given a Route String with a TopologyID of 0,0,0,0,3,1,8 and a Router with a depth of 2, the Egress Adapter number is 3.
 - If the extracted Adapter number is 0, the Control Adapter of the Router shall consume the packet. The Router shall process the packet using the Enumerated Router Flow in Section 6.4.3.2.1.
 - If the extracted Adapter number refers to a Protocol Adapter, the packet shall be dropped and the Router shall send the Connection Manager a Notification Packet with Event Code = ERR_ADP as defined in Table 6-11.
 - If the extracted Adapter number refers to a disconnected or disabled Adapter, the Router shall drop the packet and shall send the Connection Manager a Notification Packet with Event Code = ERR_CONN as defined in Table 6-11.
 - If the extracted Adapter number refers to a connected Adapter and the *Lock* bit in the Adapter Configuration Space is set to 1b, the Router shall drop the packet and shall send the Connection Manager a Notification Packet with Event Code = ERR_LOCK as defined in Table 6-11.
 - Else, the Router shall forward the packet over the Egress Adapter that matches the extracted Adapter number.
- If the packet arrived on an Adapter that is not the Upstream Adapter, then:
 - If the packet is an Inter-Domain Request or an Inter-Domain Response, then the Router shall modify the packet as follows, and then send the packet over the Upstream Adapter.
 - Replace the Route String in the packet with the Route String of the receiving Router within the receiving Domain, then add the Ingress Adapter number of the Adapter connected to the inter-Domain Link (see Figure 6-15).
 - Set the *CM* bit to 1b.

- If the packet is a Read Request or a Write Request and the *TopologyID Valid* bit in the Router Configuration Space is set to 1b, then the Router shall drop the packet and shall send the Adapter that originated the Request a Notification Packet with Event Code = ERR_ENUM as defined in Table 6-11.
- If the Router is a Host Router and the packet is a Read Request or a Write Request and the *TopologyID Valid* bit in the Router Configuration Space is set to 0b, then the Router shall drop the packet and shall send the Adapter that originated the Request a Notification Packet with Event Code = ERR_NUA as defined in Table 6-11.
- Else, Router shall drop the packet and shall not send any packets in response.

13.4.4.2 Uninitialized Router Flow

If the packet is a Read Request or a Write Request that targets Router Configuration Space, the Router shall process the packet as described in Section 6.4.3.3.

Else, Router shall drop the packet and shall not send any packets in response.

13.5 Time Synchronization

This section only applies to a Device Router that supports TBT3-Compatibility on its UFP.

When a USB4 Port is operating as two Single-Lane Links in TBT3 mode, the Time Sync Handshake shall occur on Lane 0 when the Lane 0 Adapter is in CL0 state. If the Lane 0 Adapter is not in the CL0 state but Lane 1 Adapter is in the CL0 state, the Time Sync Handshake may occur on Lane 1.

A Device Router shall implement an additional *Time Disruption* bit as defined in Section 13.7.2.2. This bit is used by the TBT3 Connection Manager to indicate to the Router when a time disruption event is occurring. A Router shall OR the two *Time Disruption* bits to conclude if a time disruption event is occurring.

A Device Router shall implement an additional set of Time Posting registers as defined in Section 13.7.2.2. Those registers are used by the TBT3 Connection Manager to post the time to a Router. A Router shall react to Time Posting if one of the register sets are used. Both register sets have the same effect on LocalTime.

13.6 Configuration Spaces

This section only applies to a Device Router that supports TBT3-Compatibility on its UFP.

This section defines the additional registers and register fields that are used by a TBT3 Connection Manager. If bits 15:12 in the *Connection Manager USB4 Version* field in Router Configuration Space are 0b (indicating a TBT3 Connection Manager), the register definitions in this section replace the register definitions in Chapter 8.



CONNECTION MANAGER NOTE

Only a TBT3 Connection Manager should use the registers defined in this section. A Connection Manager that supports USB4 Version 1.0 or higher shall not use the registers in this section.

Table 13-12 defines additional access types that are allowed for a particular configuration register field.

Table 13-12. Configuration Register Fields Access Types

| Access Type | Description |
|-------------|-------------|
|-------------|-------------|

| | |
|------|--|
| TBT3 | TBT3-Compatible. TBT3 Connection Managers uses this field. A write to this field shall have no effect. A read returns a vendor-defined value. |
|------|--|

13.6.1 Router Configuration Space



CONNECTION MANAGER NOTE

A Connection Manager shall not change the value in the Upstream Adapter field of a Host Router.

Table 13-13 lists the TBT3-Compatible Capabilities in Router Configuration Space. A Capability listed as “Required” shall be present in Router Configuration Space.

Table 13-13. List of TBT3-Compatible Router Configuration Capabilities

| Capability | Required / Optional | Capability ID |
|--|---------------------|---------------|
| <i>TMU Router Configuration</i> | Required | 03h |
| <i>Vendor Specific 1 Capability</i> | Required | 05h |
| <i>Vendor Specific 3 Capability</i> | Required | 05h |
| <i>Vendor Specific 4 Capability</i> | Optional | 05h |
| <i>Vendor Specific Extended 6 Capability</i> | Required | 05h |

13.6.1.1 Vendor Specific 1 Capability

A Vendor Specific 1 Capability shall have the structure depicted in Figure 13-3 and the fields defined in Section 13.6.1.4.1. The Absolute address of the VSC_CS_0 register shall be 0x28.

Figure 13-3. Structure of the Vendor Specific 1 Capability




Table 13-14. Vendor Specific 1 Capability Fields

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|--|----------|----------------|
| 0 | VSC_CS_0 | 7:0 | Next Capability Pointer This field contains the Doubleword index of the next Capability in the Router Configuration Space. It shall be set to 00h if the Vendor Specific Capability is the final Capability in the linked list of Capabilities in the Router Configuration Space. | RO | Vendor Defined |
| | | 15:8 | Capability ID This field shall contain the value 05h indicating this is the start of a Vendor Specific Capability. | RO | 05h |
| | | 23:16 | VSC ID This field shall contain the value 01h indicating this is a Vendor Specific 1 Capability. | RO | 01h |
| | | 31:24 | VSC Length This field shall contain the total number of Doublewords in the VSC structure including Doubleword 0 and the Vendor Specific Doublewords that follow. | RO | 30 |
| 1 | VSC_CS_1 | 2:0 | TBT3-Compatible | TBT3 | 0 |
| | | 6:3 | Plug Event Disable This field is used by a Connection Manager to stop a Router from sending Hot Plug Event Packets for a specific type of Adapter. When a bit in this field is set to 1b, a Router shall not send a Hot Plug Event Packet when a Hot Plug or a Hot Unplug takes place on an Adapter with the Adapter Type specified by the bit. Bit definitions within this field are: Bit 3: Lane Adapter Bit 4: DP OUT Adapter Bit 5: DP IN Adapter (lowest numbered DP IN Adapter) Bit 6: DP IN Adapter (next numbered DP IN Adapter) | R/W | 0 |
| | | 31:7 | TBT3-Compatible | TBT3 | 0 |
| 2 | VSC_CS_2 | 31:0 | TBT3-Compatible | TBT3 | 0 |
| 3 | VSC_CS_3 | 0 | Link Errors – Adapter A* A Router shall set this field to 1b when the <i>Link Errors Enable – Adapter A</i> bit is 1b and one of the bits in the <i>Logical Layer Errors of Adapter A</i> is set to 1b. This bit is not valid if the <i>Link Errors Enable – Adapter A</i> bit is 0b. * Adapter A is the lowest-numbered Lane Adapter. | R/W C | 0b |
| | | 1 | HEC Error – Adapter A A Router shall set this field to 1b when the <i>Link Errors Enable – Adapter A</i> bit is 1b and a Transport Layer Packet is received on Adapter A with an uncorrectable HEC error in the header. This bit is not valid if the <i>HEC Error Enable – Adapter A</i> bit is 0b. | R/W C | 0b |
| | | 2 | Flow Control Error – Adapter A A Router shall set this field to 1b when the <i>Link Errors Enable – Adapter A</i> bit is 1b and a Transport Layer Packet is received on Adapter A for a flow controlled Path where the appropriate buffer (dedicated or shared) has no space for the Packet or is not enabled. This bit is not valid if the <i>Flow Control Error Enable – Adapter A</i> bit is 0b. | R/W C | 0b |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|---|----------|---------------|
| | | 3 | Reserved | Rsvd | 0b |
| | | 4 | Link Errors – Adapter B* A Router shall set this field to 1b when the <i>Link Errors Enable – Adapter B</i> bit is 1b and one of the bits in the <i>Logical Layer Errors of Adapter B</i> is set to 1b. This bit is not valid if the <i>Link Errors Enable – Adapter B</i> bit is 0b. * Adapter B is the second lowest-numbered Lane Adapter | R/W C | 0b |
| | | 5 | HEC Error – Adapter B A Router shall set this field to 1b when the <i>Link Errors Enable – Adapter B</i> bit is 1b and a Transport Layer Packet is received on Adapter B with an uncorrectable HEC error in the header. This bit is not valid if the <i>HEC Error Enable – Adapter B</i> bit is 0b. | R/W C | 0b |
| 3 | VSC_CS_3 | 6 | Flow Control Error – Adapter B A Router shall set this field to 1b when the <i>Link Errors Enable – Adapter B</i> bit is 1b and a Transport Layer Packet is received on Adapter B on a flow controlled Path and the appropriate buffer (dedicated or shared) has no space for the Packet or is not enabled. This bit is not valid if the <i>Flow Control Error Enable – Adapter B</i> bit is 0b. | R/W C | 0b |
| | | 7 | Reserved | Rsvd | 0b |
| | | 8 | Link Errors – Adapter C* A Router shall set this field to 1b when the <i>Link Errors Enable – Adapter C</i> bit is 1b and one of the bits in the <i>Logical Layer Errors of Adapter C</i> is set to 1b. This bit is not valid if the <i>Link Errors Enable – Adapter C</i> bit is 0b. * Adapter C is the third lowest-numbered Lane Adapter | R/W C | 0b |
| | | 9 | HEC Error – Adapter C A Router shall set this field to 1b when the <i>Link Errors Enable – Adapter C</i> bit is 1b and a Transport Layer Packet is received on Adapter C with an uncorrectable HEC error in the header. This bit is not valid if the <i>HEC Error Enable – Adapter C</i> bit is 0b. | R/W C | 0b |
| | | 10 | Flow Control Error – Adapter C A Router shall set this field to 1b when the <i>Link Errors Enable – Adapter C</i> bit is 1b and a Transport Layer Packet is received on Adapter C on a flow controlled Path and the appropriate buffer (dedicated or shared) has no space for the Packet or is not enabled. This bit is not valid if the <i>Flow Control Error Enable – Adapter C</i> bit is 0b. | R/W C | 0b |
| | | 11 | Reserved | Rsvd | 0b |
| | | 12 | Link Errors – Adapter D* A Router shall set this field to 1b when the <i>Link Errors Enable – Adapter D</i> bit is 1b and one of the bits in the <i>Logical Layer Errors of Adapter D</i> is set to 1b. This bit is not valid if the <i>Link Errors Enable – Adapter D</i> bit is 0b. * Adapter D is the fourth lowest-numbered Lane Adapter | R/W C | 0b |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---|--------|---|----------|---------------|
| | | 13 | HEC Error – Adapter D A Router shall set this field to 1b when the <i>Link Errors Enable – Adapter D</i> bit is 1b and a Transport Layer Packet is received on Adapter D with an uncorrectable HEC error in the header. This bit is not valid if the <i>HEC Error Enable – Adapter D</i> bit is 0b. | R/W C | 0b |
| | | 14 | Flow Control Error – Adapter D A Router shall set this field to 1b when the <i>Link Errors Enable – Adapter D</i> bit is 1b and a Transport Layer Packet is received on Adapter D on a flow controlled Path and the appropriate buffer (dedicated or shared) has no space for the Packet or is not enabled. This bit is not valid if the <i>Flow Control Error Enable – Adapter D</i> bit is 0b. | R/W C | 0b |
| | | 15 | Reserved | Rsvd | 0b |
| 3 | VSC_CS_3 | 16 | Link Errors Enable – Adapter A A Connection Manager uses this field to enable Link error reporting for Adapter A. | R/W | 0b |
| | | 17 | HEC Error Enable – Adapter A A Connection Manager uses this field to enable HEC error reporting for Adapter A. | R/W | 0b |
| | | 18 | Flow Control Error Enable – Adapter A A Connection Manager uses this field to enable Flow Control error reporting for Adapter A. | R/W | 0b |
| | | 19 | Reserved | Rsvd | 0b |
| | | 20 | Link Errors Enable – Adapter B A Connection Manager uses this field to enable Link error reporting for Adapter B. | R/W | 0b |
| | | 21 | HEC Error Enable – Adapter B A Connection Manager uses this field to enable HEC error reporting for Adapter B. | R/W | 0b |
| | | 22 | Flow Control Error Enable – Adapter B A Connection Manager uses this field to enable Flow Control error reporting for Adapter B. | R/W | 0b |
| | | 23 | Reserved | Rsvd | 0b |
| | | 24 | Link Errors Enable – Adapter C A Connection Manager uses this field to enable Link error reporting for Adapter C. | R/W | 0b |
| | | 25 | HEC Error Enable – Adapter C A Connection Manager uses this field to enable HEC error reporting for Adapter C. | R/W | 0b |
| | | 26 | Flow Control Error Enable – Adapter C A Connection Manager uses this field to enable Flow Control error reporting for Adapter C. | R/W | 0b |
| | | 27 | Reserved | Rsvd | 0b |
| 28 | Link Errors Enable – Adapter D A Connection Manager uses this field to enable Link error reporting for Adapter D. | R/W | 0b | | |
| 29 | HEC Error Enable – Adapter D A Connection Manager uses this field to enable HEC error reporting for Adapter D. | R/W | 0b | | |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----------|---------------|--------|--|------|----------------|
| | | 30 | Flow Control Error Enable - Adapter D A Connection Manager uses this field to enable Flow Control error reporting for Adapter D. | R/W | 0b |
| | | 31 | Reserved | Rsvd | 0b |
| 4 | VSC_CS_4 | 0 | FL_SK When the <i>Bit Banging Enable</i> bit is set to 1b, a Router shall drive the value of this bit to the clock pin of the Flash memory device. | R/W | 0b |
| | | 1 | FL_CS When the <i>Bit Banging Enable</i> bit is set to 1b, a Router shall drive the value of this bit to the chip select pin of the Flash memory device. | R/W | 0b |
| | | 2 | FL_DI When the <i>Bit Banging Enable</i> bit is set to 1b, a Router shall drive the value of this bit to the data input pin of the Flash memory device. | R/W | 0b |
| 4 | VSC_CS_4 | 3 | FL_DO When the <i>Bit Banging Enable</i> bit is set to 1b, a Router shall set the value of this bit to reflect the data output pin of the Flash memory device. | RO | 0b |
| | | 4 | Bit Banging Enable A Connection Manager sets this bit to 1b to map the FL_SK, FL_CS, FL_DI, and FL_DO bits to the respective pins on the Flash memory interface. When this bit is set to 0b, the values of the FL_SK, FL_CS, FL_DI, and FL_DO bits are not valid. | R/W | 0b |
| | | 5 | Invalid Flash Memory A Router shall set this bit to 0b if it has a Flash Memory that can be accessed via bit banging. Otherwise, this bit shall be set to 1b. | RO | Vendor Defined |
| | | 31:6 | TBT3-Compatible | TBT3 | 0 |
| 5 to 11 | | 31:0 | TBT3-Compatible | TBT3 | 0 |
| 12 | VSC_CS_12 | 31:0 | DROM Base Address This field shall contain the base address (in bytes) of the DROM within the Flash Memory address space. This field shall be between 0000_0000h and 0000_FFFFh. | RO | Vendor Defined |
| 13 to 29 | | 31:0 | TBT3-Compatible | TBT3 | 0 |

13.6.1.2 Vendor Specific 3 Capability

A Vendor Specific 3 Capability shall have the structure depicted in Figure 13-4 and the fields defined in Table 13-15.

Figure 13-4. Structure of the Vendor Specific 3 Capability






Table 13-15. Vendor Specific 3 Capability Fields

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|--------|---------------|--------|--|------|----------------|
| 0 | VSC_CS_0 | 7:0 | Next Capability Pointer This field contains the Doubleword index of the next Capability in the Router Configuration Space. It shall be set to 00h if the Vendor Specific Capability is the final Capability in the linked list of Capabilities in the Router Configuration Space. | RO | Vendor Defined |
| | | 15:8 | Capability ID This field shall contain the value 05h indicating this is the start of a Vendor Specific Capability. | RO | 05h |
| | | 23:16 | VSC ID This field shall contain the value 03h indicating this is a Vendor Specific 3 Capability. | RO | 03h |
| | | 31:24 | VSC Length This field shall contain the total number of Doublewords in the VSC structure including Doubleword 0 and the Vendor Specific Doublewords that follow. | RO | 50 |
| 1 to 3 | | 31:0 | TBT3-Compatible | TBT3 | 0 |
| 4 | VSC_CS_4 | 31:0 | Post Local Time Low This field contains the value that is used to update the least significant 32 bits of the <i>nanoseconds</i> field (i.e. bits 47 to 16) of the LocalTime register when the Post Time (VSC_CS_6/VSC_CS_7) is less or equal to the Host Router Time. | R/W | 0 |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----------|---------------|--------|---|------|---------------|
| 5 | VSC_CS_5 | 31:0 | Post Local Time High This field contains the value that is used to update the most significant 32 bits of the <i>nanoseconds</i> (i.e. bits 79 to 48) field of the LocalTime register when the Post Time (VSC_CS_6/VSC_CS_7) less or equal to the Host Router Time. | R/W | 0 |
| 6 | VSC_CS_6 | 31:0 | Post Time Low This field contains the least significant 32 bits of the <i>Nanoseconds</i> field of the Host Router Time register at which the software updates to the LocalTime register are applied. The update occurs when the Post Time is greater than zero and is less or equal to the Host Router Time. When the Local Time is updated, this register is nullified. | R/W | 0 |
| 7 | VSC_CS_7 | 31:0 | Post Time High This field contains the most significant 32 bits of the <i>Nanoseconds</i> field of the Host Router Time register at which the software updates to the LocalTime register are applied. The update occurs when the Post Time is greater than zero and is less or equal to the Host Router Time. To activate Time Posting, this register should be the last to be written. When the Local Time is updated, this register is nullified. | R/W | 0 |
| 8 | VSC_CS_8 | 31:0 | TBT3-Compatible | TBT3 | 0 |
| 9 | VSC_CS_9 | 15:0 | TBT3-Compatible | TBT3 | 0 |
| | | 17:16 | Reserved | RO | 0b |
| | | 31:18 | TBT3-Compatible | TBT3 | 0 |
| 10 to 25 | | 31:0 | TBT3-Compatible | TBT3 | 0 |
| 26 | VSC_CS_26 | 21:0 | TBT3-Compatible | TBT3 | 0 |
| | | 22 | Time Disruption (TD) The Connection Manager sets this bit to 1b before any of the following time disruptions: <ul style="list-style-type: none"> • TMU mode changes (TSPacketInterval, Direction, Filters, Frequency Measurement Window) • Inter-Domain time sync is enabled • Time Posting is applied After the time disruption has passed, the Connection Manager sets this bit to 0b. | R/W | 0b |
| | | 31:23 | TBT3-Compatible | TBT3 | 0 |
| 27 to 49 | | 31:0 | TBT3-Compatible | TBT3 | 0 |

13.6.1.3 Vendor Specific 4 Capability

If a Router implements Vendor Specific 4 Capability, the first 11 Doublewords shall have the structure depicted in Figure 13-5 and the fields defined in Table 13-16.


Figure 13-5. Structure of the Vendor Specific 4 Capability**Table 13-16. Vendor Specific 4 Capability Fields**

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|---------|---------------|--------|--|------|----------------|
| 0 | VSC_CS_0 | 7:0 | Next Capability Pointer This field contains the Doubleword index of the next Capability in the Router Configuration Space. It shall be set to 00h if the Vendor Specific Capability is the final Capability in the linked list of Capabilities in the Router Configuration Space. | RO | Vendor Defined |
| | | 15:8 | Capability ID This field shall contain the value 05h indicating this is the start of a Vendor Specific Capability. | RO | 05h |
| | | 23:16 | VSC ID This field shall contain the value 04h indicating this is a Vendor Specific 4 Capability. | RO | 04h |
| | | 31:24 | VSC Length This field shall contain the total number of Doublewords in the VSC structure including Doubleword 0 and the Vendor Specific Doublewords that follow. | RO | Vendor Defined |
| 1 to 10 | | 31:0 | TBT3-Compatible | TBT3 | 0 |

13.6.1.4 Vendor Specific Extended 6 Capability


A Vendor Specific Extended 6 Capability shall have the structure depicted in Figure 13-6 and the fields defined in Section 13.6.1.4.1 and Section 13.6.1.4.2. A USB4 Port Region shall exist for each USB4 Port.

Figure 13-6. Structure of the Vendor Specific Extended 6 Capability



A USB4 Port Region shall exist for each USB4 Port. The first USB4 Port Region (USB4 Port Region 0) shall contain information about the USB4 Port with the lowest Adapter Number. Each subsequent USB4 Port Region shall contain information about the USB4 Port with the next highest Adapter Number. For example, Figure 13-7 shows the Vendor Specific Extended 6 Capability for a Router with three USB4 Ports (referred to as Port A, Port B, and Port C) where Port A contains Lane Adapters 1 and 2, Port B contains Lane Adapters 3 and 4, and Port C contains Lane Adapters 5 and 6.

Figure 13-7. Example Vendor Specific Extended 6 Capability



13.6.1.4.1 Common Region

A Common Region shall have the structure depicted in Figure 13-8 and the fields defined in Table 13-17.

Figure 13-8. Structure of the Common Region




Table 13-17. Common Region Fields

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|--|------|---------------|
| 0 | VSEC_CS_0 | 7:0 | Reserved | Rsvd | 00h |
| | | 15:8 | Capability ID This field shall contain the value 05h indicating this is the start of a Vendor Specific Capability. | RO | 05h |

| | | | | | |
|---------|------------|-------|--|------|---------------------|
| | | 23:16 | VSEC ID This field shall contain the value 06h indicating this is a Vendor Specific Extended 6 Capability. | RO | 06h |
| | | 31:24 | VSEC Header This field shall be set to 00h to indicate that the Capability is an Extended Capability. | RO | 00h |
| 1 | VSEC_CS_1 | 15:0 | Next Capability Pointer This field contains the Doubleword index of the next Capability in the Router Configuration Space. It shall be set to 00h if the Vendor Specific Capability is the final Capability in the linked list of Capabilities in the Router Configuration Space. | RO | Vendor Defined |
| | | 31:16 | VSEC Length This field shall contain the total number of Doublewords in the VSEC structure including Doubleword 0, Doubleword 1, and the Vendor-Specific Doublewords that follow. | RO | 256*(USB4 Ports)+20 |
| 2 | CAP_STRUCT | 3:0 | USB4 Ports This field shall contain the number of USB4 Ports supported by the Router. | RO | Vendor Defined |
| | | 7:4 | TBT3-Compatible | TBT3 | 0 |
| | | 15:8 | Common Region Length This field shall contain the size (in Doublewords) of the Common Region. | RO | 14h |
| | | 27:16 | USB4 Port Region Length This field shall contain the size (in Doublewords) of a single USB4 Port Region. | RO | 100h |
| | | 31:28 | TBT3-Compatible | TBT3 | 0 |
| 3 to 19 | | 31:0 | TBT3-Compatible | TBT3 | 0 |

13.6.1.4.2 USB4 Port Regions

A USB4 Port Region shall have the structure depicted in Figure 13-9 and the fields defined in Table 13-18.

When a field in the USB4 Port Region has the same name as a register field in Configuration Space (as defined in Chapter 8), the field in the USB4 Port Region is used instead of the field in Configuration Space. For example, when operating with a TBT3 Connection Manager, a Device Router uses the *Enter Sleep* bit in the USB4 Port Region below instead of the *Enter Sleep* bit in Router Configuration Space to determine Sleep and Wake (see Section 4.5 and Section 13.2.4 for use of *Enter Sleep* bit).

Figure 13-9. Structure of a USB4 Port Region




Table 13-18. USB4 Port Region Fields

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----------|---------------|--------|---|------|---------------|
| 0 to 37 | - | 31:0 | TBT3-Compatible | TBT3 | 0 |
| 38 | PORT_MODE | 5:0 | TBT3-Compatible | TBT3 | 0 |
| | | 6 | Downstream Port Reset (DPR) For a Downstream Facing Port: Setting this bit to 1b initiates a Downstream Port Reset. Setting this bit to 0b brings the Port out of a Downstream Port Reset. For an Upstream Facing Port: A read or write to this bit shall have no effect. | R/W | 0b |
| | | 31:7 | TBT3-Compatible | TBT3 | 0 |
| 39 to 42 | - | 31:0 | TBT3-Compatible | TBT3 | 0 |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|------------|---------------|--------|--|------|----------------|
| 43 | ID_WAKE | 4:0 | TBT3-Compatible | TBT3 | 0 |
| | | 5 | Enable Wake on Inter-Domain (EWI) A Connection Manager sets this bit to 1b to cause a Router to exit from Sleep state when it detects an Inter-Domain Wake on the USB4 Port. The Connection Manager sets the <i>Inter-Domain Disconnect on Sleep</i> bit to 1b when this bit is 0b. | R/W | Vendor Defined |
| | | 31:6 | TBT3-Compatible | TBT3 | 0 |
| 44 to 60 | - | 31:0 | TBT3-Compatible | TBT3 | 0 |
| 61 | ID_SLEEP | 4:0 | TBT3-Compatible | TBT3 | 0 |
| | | 5 | Inter-Domain Disconnect on Sleep (IDD) A Connection Manager sets this bit to 1b to cause a disconnect on the USB4 Port if the USB4 Port is connected to an Inter-Domain Link when the Router enters Sleep state. When this field is 0b, the Connection Manager disables wake on a USB4 Wake in the <i>Enable Wake Events</i> field and sets the <i>Enable Wake on Inter-Domain</i> bit to 1b. | R/W | 0b |
| | | 31:6 | TBT3-Compatible | TBT3 | 0 |
| 62 to 140 | - | 31:0 | TBT3-Compatible | TBT3 | 0 |
| 141 | PORT_ATTR | 11:0 | TBT3-Compatible | TBT3 | 0 |
| | | 12 | Bonding Enabled (BE) An Adapter shall set this bit to 1b when the conditions for Lane bonding are met. See Section 4.1.2.3. An Adapter shall set this bit to 0b when the conditions for Lane bonding are not met. | RO | 0 |
| | | 31:13 | TBT3-Compatible | TBT3 | 0 |
| 142 to 149 | - | 31:0 | TBT3-Compatible | TBT3 | 0 |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|-----|---------------|--------|---|------|---------------------------------------|
| 150 | LC_SX_CTRL | 10:0 | <p>Enable Wake Events</p> <p>A Connection Manager uses this field to enable the various types of wake events.</p> <p>When a bit is set to 1b, the corresponding event shall cause a Router to exit from sleep. When a bit is set to 0b, the corresponding event shall not cause a Router to exit from sleep.</p> <ul style="list-style-type: none"> bit 0: vendor specific event bit 1: wake on connect of a USB4 Router to the USB Type-C® connector bit 2: wake on disconnect of a USB4 Router to the USB Type-C connector bit 3: wake on connect of a DisplayPort device to the USB Type-C connector <p>A Router may exit sleep on connect of a DisplayPort device to any DP connector if this bit is set in one of the USB4 Ports</p> <ul style="list-style-type: none"> bit 4: wake on disconnect of a DisplayPort device to the USB Type-C connector bit 5: wake on a USB4 Wake detected on this Port bit 6: wake on a PCIe Wake indication bit 7: vendor specific event bit 8: vendor specific event bit 9: wake on connect of a USB device to the USB Type-C connector bit 10: wake on disconnect of a USB device to the USB Type-C connector | R/W | Host Router: 180h Else: 140h |
| | | 15:11 | TBT3-Compatible | TBT3 | 0 |
| 150 | LC_SX_CTRL | 16 | <p>Lane 0 Configured (LOC)</p> <p>The Connection Manager sets this bit to 1b to indicate that the Router connected to Lane 0 of the USB4 Port is configured and that entry to Sleep State and exit from Sleep State shall be supported on the Lane.</p> <p>The Connection Manager sets this bit to 0b to indicate that the Router connected to Lane 0 of the USB4 Port is not configured and therefore the USB4 Port is disconnected on entry to Sleep State.</p> | R/W | 0b |
| | | 17 | <p>Lane 0 is Inter-Domain (LOD)</p> <p>A Connection Manager sets this bit to 0b when Lane 0 is not part of an Inter-Domain Link.</p> <p>A Connection Manager sets this bit to 1b when Lane 0 is part of an Inter-Domain Link.</p> | R/W | 0b |
| | | 19:18 | TBT3-Compatible | TBT3 | 0 |
| | | 20 | <p>Lane 1 Configured (L1C)</p> <p>The Connection Manager sets this bit to 1b to indicate that the Router connected to Lane 1 of the USB4 Port is configured and that entry to Sleep State and exit from Sleep State shall be supported on the Lane.</p> <p>The Connection Manager sets this bit to 0b to indicate that the Router connected to Lane 1 of the USB4 Port is not configured and therefore the USB4 Port is disconnected on entry to Sleep State.</p> | R/W | 0b |
| | | 21 | <p>Lane 1 is Inter-Domain (L1D)</p> <p>A Connection Manager sets this bit to 0b when Lane 1 is not part of an Inter-Domain Link.</p> <p>A Connection Manager sets this bit to 1b when Lane 1 is part of an Inter-Domain Link.</p> | R/W | 0b |
| | | 28:22 | TBT3-Compatible | TBT3 | 0 |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|------------|---------------|--------|--|------|---------------|
| | | 29 | Start Link Initialization (SLI) A Connection Manager uses this bit to delay Lane Initialization on a USB4 Port when a Router exits Sleep state (See Section 13.2.4.4) 1b - The USB4 Port shall start Lane Initialization 0b - The USB4 Port shall not start Lane Initialization | R/W | 0b |
| | | 30 | TBT3-Compatible | TBT3 | 0 |
| | | 31 | Enter Sleep (SLP) A Connection Manager sets this bit to 1b to let the Router know that the system is preparing for entry to sleep state. | R/W | 0b |
| 151 | LINK_ATTR | 9:0 | TBT3-Compatible | TBT3 | 0 |
| | | 10 | RS-FEC Enabled (Gen 2) (RE2) An Adapter shall set this bit to 1b when the USB4 Port is operating at Gen 2 speed and RS-FEC is enabled. This bit is set to 0b otherwise. | RO | 0 |
| | | 11 | RS-FEC Enabled (Gen 3) (RE3) An Adapter shall set this bit to 1b when the USB4 Port is operating at Gen 3 speed and RS-FEC is enabled. This bit is set to 0b otherwise. | RO | 0 |
| | | 16:12 | TBT3-Compatible | TBT3 | 0 |
| | | 17 | TBT3-Compatible Mode (TCM) An Adapter shall set this bit to 1b when the Link is operating in TBT3-Compatible Mode. This bit is set to 0b otherwise. | RO | 0 |
| 151 | LINK_ATTR | 18 | CLx Protocol Support (CPS) This bit indicates that the Router and the Cable support the Low Power protocol. When a Router supports the Low Power protocol, it either accepts or rejects entry to a CLx state as defined in Section 4.2.1.6. A Router shall set this bit to 1b if the Sideband Channel operates as a USB4 Sideband Channel and the Cable supports CLx states. Otherwise, Router shall set this bit to 0b. | RO | 0 |
| | | 31:19 | TBT3-Compatible | TBT3 | 0 |
| 152 to 172 | - | 31:0 | TBT3-Compatible | TBT3 | 0 |
| 173 | LINK_REQ | 0 | Request RS-FEC Gen 2 (RS2) A Connection Manager sets this bit to 1b to enable RS-FEC encoding at 10G speeds. If this bit is set to 1b, the Router shall enable RS-FEC encoding at 10G speeds on the Links of this USB4 Port during the next Lane Initialization. If this bit is set to 0b, then RS-FEC enabling at 10G speeds is determined by the setting of the Link Partner. | R/W | 0b |

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|------------|---------------|--------|--|------|---------------|
| | | 1 | Request RS-FEC Gen 3 (RRS3) A Connection Manager sets this bit to 1b to enable RS-FEC encoding at 20G speeds. If this bit is set to 0b, the Router shall enable RS-FEC encoding at 20G speeds on the Links of this USB4 Port during the next Lane Initialization. If this bit is set to 1b, then RS-FEC enabling at 20G speeds is determined by the setting of the Link Partner. | R/W | 1b |
| | | 31:2 | TBT3-Compatible | TBT3 | 0 |
| 174 to 255 | - | 31:0 | TBT3-Compatible | TBT3 | 0 |

13.6.2 Adapter Configuration Space

The Absolute address of the ADP_DP_CS_0 register in a DP Adapter Configuration Capability shall be 0x39.

A Device Router shall ignore an attempt to modify bit 8 in ADP_DP_CS_3 register of a DP OUT Adapter.

A TBT3 Connection Manager assumes that a Router has implemented a register at absolute address 0x10 in the DP OUT Adapter Configuration Space. When a DP OUT Adapter receives a Write Request that targets address 0x10, it shall send a Write Response. A DP OUT Adapter shall not implement a Capability Register at address 0x10 in its Adapter Configuration space.

A DP IN Adapter shall not have a Vendor Specific Capability with VSC ID = 0 or VSC ID = 1.

A DP OUT Adapter shall not have a Vendor Specific Capability with VSC ID = 1.

A TBT3 Connection Manager may initiate a Read Request to the Basic Configuration Registers in the Adapter Configuration Space of the Control Adapter. A Device Router shall respond to such a request in one of the following two ways:

1. Send a Notification Packet with Event Code = ERR_ADP, as defined in Table 6-11.
2. Send a Read Response with the following values:
 - For ADP_CS_2, respond with value 0x01000001.
 - For any other address respond with any value.

13.6.2.1 Basic Attributes

An Adapter shall support the Adapter Configuration Space Basic Attributes in Table 13-19. A Lane 1 Adapter may optionally implement the *HEC Error*, *Flow Control Error*, and/or *Shared Buffering Capable* bits.

Table 13-19. Adapter Configuration Space Basic Attributes

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|--|------|----------------|
| 4 | ADP_CS_4 | 29:20 | Max Credits This field shall be equal to the <i>Total Buffers</i> Field. | RO | Vendor Defined |

| | | | | | |
|--|--|----|---|-----|----|
| | | 31 | <p>Lock (LCK)</p> <p>This bit controls whether or not a Connection Manager can access a Router that is downstream of the Adapter. This bit is only used for the Adapters in a Downstream Facing Port.</p> <p>When the bit is 1b, the Adapter is “locked”, which means that Control Packets are not forwarded to the downstream Router. When the bit is 0b, the Adapter is “unlocked”, and Control Packets can be forwarded to the downstream Router.</p> <p>If the value in the <i>Connection Manager USB4 Version</i> field is greater than 0, an Adapter shall set this bit to 1b after the Adapter goes through a disconnect.</p> <p>If the value in the <i>Connection Manager USB4 Version</i> field is 0, an Adapter shall keep this bit set to 0b.</p> <p>An Adapter may ignore a write to this bit if the Adapter does not have a Router connected downstream.</p> <p><i>Note: A TBT3 Connection Manager does not use this bit and the Adapter remains “unlocked” by default.</i></p> | R/W | 0b |
|--|--|----|---|-----|----|

13.6.2.2 USB4 Port Capability

An Adapter shall support the USB4 Port Capability fields in Table 13-20.

Table 13-20. USB4 Port Capability Fields

| DW | Register Name | Bit(s) | Field Name and Description | Type | Default Value |
|----|---------------|--------|--|------|---------------|
| 1 | PORT_CS_1 | 18:16 | <p>Target</p> <p>This field defines which SB Register Space to access:</p> <ul style="list-style-type: none"> 110b – Near-end Cable Re-timer (via AT Transaction) 111b – Far-end Cable Re-timer (via AT Transaction and using the Bounce mechanism) | R/W | 0 |
| 19 | PORT_CS_19 | 1 | <p>Request RS-FEC Gen 2 (RS2)</p> <p>A Connection Manager uses this bit to request enabling of RS-FEC encoding at Gen 2 speeds. If a Link is Active, the Link must be re-initialized before RS-FEC can be enabled.</p> <p>If this bit is set to 1b, the USB4 Port shall enable RS-FEC at Gen 2 speeds during the next Lane Initialization.</p> <p>If this bit is set to 0b, then the Link Partner response determines whether or not RS-FEC is enabled at Gen 2 speeds.</p> | R/W | 1b |

13.7 PCI Express Tunneling

A TBT3-Compatible Host Router shall support PCIe Tunneling.

13.7.1 PCIe Power Management

13.7.1.1 L1

If TBT3 Mode is established on a link and the *USB4 Sideband Channel Support* bit at the Link Partner is set to 0b, then ASPM L1 operation over PCIe Tunneling is not supported. When ASPM L1 operation is not supported:

- If TBT3 Mode is established on a Downstream Facing Port, the state of the downstream port of the internal PCIe Switch that is connected to that Downstream Facing Port may be ignored when determining the conditions to enter L1 state on the upstream port of the internal PCIe Switch.
- If TBT3 Mode is established on the Upstream Facing Port, the upstream port of the internal PCIe Switch that is connected to that Upstream Facing Port shall disable L1.

If TBT3 Mode is established on a link and the *USB4 Sideband Channel Support* bit at the Link Partner is set to 1b, then ASPM L1 operation over PCIe Tunneling is supported.

Note: Software L1 is always supported.

Note: The internal upstream PCIe port of a TBT3 device, implements an L1 handshake timeout of 10 μ s for ASPM L1 and 2 μ s for Software L1, starting from sending EIOS and ending when an EIOS is received. If timeout expires, the internal upstream PCIe port enters Recovery state.

13.7.1.2 L2

Note: The internal upstream PCIe port of a TBT3 device implements an L2 handshake timeout of 2 μ s, starting from sending EIOS and ending when an EIOS is received. If timeout expires, the internal upstream PCIe port enters either the Detect or Recovery states.

13.8 DisplayPort Tunneling

A DP Adapter shall operate in Non-LTTPR (TBT3-Compatible) mode when the *DP_COMMON_CAP.Protocol Adapter Version* value is less than 4. TBT3-Compatible mode is Non-LTTPR mode as defined in Chapter 10 with the modifications defined in this section.

13.8.1 AUX Handling

13.8.1.1 DP IN Adapter Requirements

A DP IN Adapter operating in TBT3-Compatible mode follows the requirements listed in Section 10.4.4.2.1 except the following:

- The Link Status DPCD registers, as defined in Table 10-9, shall be mapped statically as Internal registers.
- DPCD address 00600h is mapped as Internal register.

A DP IN Adapter operating in TBT3-compatible mode follows the requirements listed in Section 10.4.4.5 except the following:

- DSC Support field in the DSC SUPPORT DPCD register shall always be set to 0b.
- FEC_CAPABLE field in FEC_CAPABILITY DPCD register shall always be set to 0b.
- PHY_REPEATER_CNT field in PHY_REPEATER_CNT DPCD register shall always be set to 0b.

13.8.1.2 DP OUT Adapter Requirements

A DP OUT Adapter operating in TBT3-Compatible mode follows the requirements listed in Section 10.4.4.2.3 except that initiation of an AUX Transaction occurs for the following processes:

- DPRX Capability Discovery - Defined in Section 13.8.3.2.
- IRQ Handling - Defined in Section 13.8.2.

Note: A DP OUT Adapter operating in TBT3-Compatible mode could receive the Reserved byte at the end of the AUX Tunneled Packet with value other than zero.

13.8.2 IRQ Handling

When a DP OUT Adapter operating in TBT3-Compatible Mode detects an IRQ, it shall determine if the IRQ was issued to signify a DP Link loss. If the IRQ was initiated to signify a DP Link loss, the TBT3-Compatible DP OUT Adapter shall:

- Disable its Main-Link transmitters.
- Send a SET_CONFIG Packet of MSG type STATUS_LOST_CONNECTION.
- Transition to IDLE state as described in Section 13.8.6.2.

If the IRQ was not initiated to signify a DP Link loss, the DP OUT Adapter shall send a SET_CONFIG Packet of MSG type IRQ.

When a DP IN Adapter receives a SET_CONFIG Packet of MSG type STATUS_LOST_CONNECTION, it shall:

- Disable its Main-Link receivers and issue an IRQ to the DPTX.
- Transition to IDLE state as described in Section 13.8.6.1.

Note: It is expected that the system recovers from a DP Link failure as follows:

1. DPTX receives an IRQ, reads the DP-Link status and discovers the DP-Link failure.
2. DPTX retrains the DP-Link.

13.8.3 Connection Manager Discovery

13.8.3.1 TBT3 Connection Manager

The methods for detecting whether or not a TBT3 Connection Manager is present in a system is defined in Section 10.4.2.1. After determining that a system contains a TBT3 Connection Manager, a TBT3-Compatible DP OUT Adapter shall perform the following steps:

1. Send a Hot Removal Event Packet to the TBT3 Connection Manager.
2. Read DPRX capabilities by initiating an AUX read Request.
3. Update each field in the DP_LOCAL_CAP register to reflect the lowest common capabilities of:
 - DP OUT Adapter capabilities.
 - DPRX capabilities.
 - The capabilities defined in the DisplayPort Standard, Version 1.2.
4. Wait at least 150 ms.
5. Send a Hot Plug Event Packet to the TBT3 Connection Manager.

13.8.3.2 TBT3 Router Discovery

The methods for detecting whether or not a TBT3 Router is present in a system is defined in Section 10.4.2.1. After determining that the DP IN Adapter at the end of a Path is part of a TBT3 Router, a TBT3-Compatible DP OUT Adapter shall perform the following steps:

1. Read DPRX capabilities by initiating, at minimum, AUX read Requests to the following DPCD addresses.
 - 0000h-0000Eh
 - 00021h
 - 00200h
2. Update each field in the DP_LOCAL_CAP register to reflect the lowest common capabilities of:
 - DP OUT Adapter capabilities

- DPRX capabilities
3. Reset DP_STATUS_CTRL.CMHS to zero.



CONNECTION MANAGER NOTE

When a Connection Manager detects a TBT3 Router with a DP Adapter, it shall do the following:

- *Set the Video HopID field in the DP Adapter Configuration Capability to 9.*
- *Set the AUX Tx HopID field in the DP Adapter Configuration Capability to 8.*
- *Set the AUX Rx HopID field in the DP Adapter Configuration Capability to 8.*

Note that in a TBT3 Router, the Video HopID, AUX Tx HopID, and AUX Rx HopID fields are R/W and default to 0.

13.8.4 Sink Count Read

A DP OUT Adapter, operating in TBT3-Compatible mode, shall read DPCD address 00200h as part of DPRX capability discovery. If the SINK_COUNT value read from the DPRX is 0b, the DP OUT Adapter shall follow the steps defined in Section 10.4.6.3, and shall not wait for a SINK_COUNT SET_CONFIG to be sent by the DP IN Adapter.

A DP IN Adapter, operating in TBT3-Compatible mode, may send a SET_CONFIG Packet of type SET_SINK_COUNT to query the value of the SINK COUNT.

A DP OUT Adapter, receiving a SET_CONFIG Packet of type SET_SINK_COUNT, shall respond to the DP IN Adapter with in tDPInit with a SET_CONFIG Packet of type SET_SINK_COUNT with the SINK_COUNT value equal to the value that was read from DPCD address 00200h.

13.8.5 Power States Set

When DPTX either reads or writes to DPCD address 00600h, a DP IN Adapter, operating in TBT3-Compatible mode shall:

- In case it is a write, respond with an AUX ACK and update its internal power state according to the SET_POWER_STATE of the write request.
- In case it is a read, respond with an AUX Response carrying its internal power state.
- Send a SET_CONFIG Packet of type SET_POWER, carrying the internal power state.

A DP OUT Adapter in TBT3-Compatible mode, receiving a SET_CONFIG Packet of type SET_POWER, shall initiate an AUX Write Request to DPCD address 00600h with the value written equal to the *MSG Data* received by the SET_CONFIG Packet.

13.8.6 DisplayPort Link Training

The DisplayPort link training process in TBT3-Compatible is similar to the Non-LTTPr mode in that the DPTX is not aware of the second DisplayPort link being trained. However, it is different in that the two DisplayPort links are trained sequentially; the DP OUT Adapter starts its DisplayPort link training only when the DP IN Adapter has internally finished training its DisplayPort link.

A DP Adapter in TBT3-Compatible mode shall perform DisplayPort link training according to the DisplayPort 1.4a Specification with the modifications and requirements defined in Section 13.8.6.1 and Section 13.8.6.2.

13.8.6.1 DP IN Adapter Requirements

DP IN Adapter shall follow the state machine described in Figure 13-10 and the transition table described in Table 13-21.

Figure 13-10. DP IN Adapter Link Training State Machine




Table 13-21. DP IN Adapter Link Training State Machine Transition Table

| Current State | Next State | Condition | Action |
|---------------|------------|--|--|
| IDLE | DP_IN_LT | DPTX initiates DisplayPort link training as defined in the DisplayPort 1.4a Specification. | DP IN Adapter performs DisplayPort link training. |
| DP_IN_LT | IDLE | DisplayPort link training between DPTX and the DP IN Adapter fails. | None. |
| DP_IN_LT | DP_OUT_LT | DP IN Adapter achieves LANE _x _CHANNEL_EQ_DONE, LANE _x _SYMBOL_LOCK and INTERNAL_LANE_ALIGN on all active lanes. | DP IN Adapter sends a SET_CONFIG Packet of type SET_LINK with LC and LR that matches the DisplayPort link training parameters set by the DPTX. |
| DP_OUT_LT | FAIL_LT | DP IN Adapter receives a SET_CONFIG Packet of type SET_LINK with LC and LR equal to zero or the DisplayPort link between DPTX and the DP IN Adapter fails. | None. |
| DP_OUT_LT | LINK_UP | DP IN Adapter receives a SET_CONFIG Packet of type SET_LINK with LC and LR not equal to zero or DPTX issues its last AUX status read attempt. | None. |
| LINK_UP | IDLE | Either: DP IN Adapter receives a SET_CONFIG Packet of type SET_LINK with LC and LR equal to zero. or: DP IN Adapter receives a SET_CONFIG Packet of type STATUS_LOST_CONNECTION. or: The DisplayPort link between DPTX and the DP IN Adapter fails. | DP IN Adapter disables its Main-Link receivers, then sends an IRQ. If the transition is due to a failure of the DisplayPort link between DPTX and the DP IN Adapter, the DP IN Adapter sends a SET_CONFIG Packet of type SET_LINK with LC and LR equal to zero. |

| | | | |
|---------|------|---|-------|
| FAIL_LT | IDLE | DisplayPort link training between DPTX and the DP IN Adapter fails. | None. |
|---------|------|---|-------|

A DP IN Adapter shall respond to an AUX status read with LANEx_CHANNEL_EQ_DONE = 1b, LANEx_SYMBOL_LOCKED = 1b and INTERLANE_ALIGN_DONE = 1b on all active lanes is while in LINK_UP state. It shall not respond to an AUX status read with LANEx_CHANNEL_EQ_DONE = 1b, LANEx_SYMBOL_LOCKED = 1b and INTERLANE_ALIGN_DONE = 1b in any other states.

While in the DP_OUT_LT state or FAIL_LT state, a DP IN Adapter shall respond to AUX status read with LANEx_CHANNEL_EQ_DONE = 0 for all lanes.

13.8.6.2 DP OUT Adapter Requirements

A TBT3-Compatible DP OUT Adapter shall follow the state machine described in Figure 13-11 and the transition table described in Table 13-22.

Figure 13-11. DP OUT Adapter Link Training State Machine




Table 13-22. DP OUT Adapter Link Training State Machine Transition Table

| Current State | Next State | Condition | Action |
|---------------|------------|--|--|
| IDLE | DP_OUT_LT | DP OUT Adapter receives a SET_CONFIG Packet of type SET_LINK with LC and LR not equal to zero. | DP OUT Adapter performs DisplayPort link training. |
| DP_OUT_LT | IDLE | Either: DisplayPort link training between DPRX and the DP OUT Adapter fails to train the DisplayPort Link to the target LC and LR. or: DP OUT Adapter receives a SET_CONFIG Packet of type SET_LINK with LC and LR equal to zero. | If the transition is due to failure of the DisplayPort link between DPRX and the DP OUT Adapter, The DP OUT Adapter sends a SET_CONFIG Packet of type SET_LINK with LC and LR equal to zero. |
| DP_OUT_LT | LINK_UP | DPRX and DP OUT Adapter successfully train the DisplayPort link. | DP OUT Adapter sends a SET_CONFIG Packet of type SET_LINK with LC and LR that matches the Link training parameters which were sent by the DP IN Adapter. |

| | | | |
|---------|------|---|--|
| LINK_UP | IDLE | DP OUT Adapter receives a SET_CONFIG Packet of type SET_LINK with LC and LR equal to zero or DPRX reports link failure. | DP OUT Adapter disables its high speed transmitters. If the transition is due to DPRX reporting DisplayPort link failure, the DP OUT Adapter sends a SET_CONFIG packet of type SET_LINK with LC and LR equal to zero. |
|---------|------|---|--|


While performing DisplayPort link training, DP OUT Adapter shall use the target Link Rate and Lane Count received from the DP IN Adapter. If a DP OUT Adapter fails to train the DisplayPort link to the target parameters, it shall report DisplayPort link failure to the DP IN Adapter. It shall not attempt to apply the fallback mechanisms defined in the DisplayPort 1.4a Specification.

13.9 USB3 Functionality

This section only applies to USB4-Based Docks and USB4 hubs that support TBT3-Compatibility on their UFP.

A USB4-Based Dock or USB4 hub shall incorporate an internal host controller to support USB3 functionality. Figure 13-12 shows an example of a USB4-Based Dock that incorporates an internal host controller.

Figure 13-12. Example of a USB4-Based Dock with an Internal Host Controller



The Device Router in a USB4-Based Dock or USB4 hub shall set the *HCI* field to 1b.

During Router enumeration and initialization, after the Device Router sets the *Router Ready* bit to 1b and the Connection Manager sets the *Configuration Valid* bit to 1b, a Device Router in a

USB4-Based Dock or USB4 hub shall do the following in addition to the requirements defined in Section 6.7 before setting the *Configuration Ready* bit to 1b:

- If the *Internal Host Controller On* bit is set to 1b, establish tunneling functionality to/from the internal host controller using PCIe tunneling.



IMPLEMENTATION NOTE

A Device Router may decide to establish tunneling functionality to/from the internal host controller using PCIe tunneling when its Upstream Facing Port operates in TBT3 Mode.



CONNECTION MANAGER NOTE

If a Connection Manager intends to use the internal host controller, then before setting up a PCIe Tunneling Path, the Connection Manager shall set the internal host controller On field to 1b and the Configuration Valid bit to 1b in the Router Configuration Space Basic Attributes in the Device Router in the USB4-Based Dock or USB4 hub. It then polls the Configuration Ready field until it is set to 1b by the Device Router.

When a Connection Manager is going to tunnel USB3 traffic through a USB4-Based Dock or USB4 hub, it can either enable USB3 functionality using the internal host controller (as described in this section) or USB3 Tunneling using the internal USB3 device (as described in Chapter 9). It cannot do both.

13.10 Host-to-Host Tunneling

No additional requirements.


A Verification of CRC, Scrambling, and FEC Calculations

This Appendix gives several examples of CRC and FEC calculations.


A.1 Transport Layer Packet HEC

Figure A-1 provides examples of HEC calculations for the Transport Layer Packet header. In each example, the second line contains the values for the header depicted in the line above it.


Figure A-1. Examples of Transport Layer Packet HEC Calculation



Example 1



Example 2



Example 3

* = SupplD

A.2 Control Packet CRC

Table A-1 provides examples of CRC calculation for Control Packets.

Table A-1. Examples of Control Packet CRC Calculation

| | Route String High | Route String Low | Payload DW 0 | Payload DW 1 | CRC DW |
|---------------------------|-------------------|------------------|--------------|--------------|------------|
| Example 1 – Read Request | 0000 0000h | 0000 0000h | 0400 2000h | - | 2FC9 67ACh |
| Example 2 – Read Response | 8000 0000h | 0000 0000h | 0428 2000h | 1578 8086h | 5A9B 0181h |

A.3 Sideband Channel AT Transaction CRC

Table A-2 provides an example of a CRC calculation for an AT Transaction that reads a single byte from register address 0Ah.

Table A-2. Example of a Read Command

| | Byte | Value |
|---|--------|-------|
| 0 | DLE | FEh |
| 1 | STX | 05h |
| 2 | Data 0 | 0Ah |

| | Byte | Value |
|---|-----------------------|-------|
| 3 | Data 1 | 01h |
| 4 | CRC (low order byte) | A6h |
| 5 | CRC (high order byte) | A1h |
| 6 | DLE | FEh |
| 7 | ETX | 40h |

Table A-3 provides an example of a CRC calculation for an AT Transaction that writes 5 bytes to register address 62h. Note that Byte 8 is a duplicate symbol inserted to distinguish Byte 9 from a DLE symbol. It is not part of the symbols protected by the CRC and is thus not included in the CRC calculation.

Table A-3. Example of a Write Command

| | Byte | Value |
|----|-----------------------|-------|
| 0 | DLE | FEh |
| 1 | STX | 05h |
| 2 | Data 0 | 62h |
| 3 | Data 1 | 85h |
| 4 | Data 2 | 43h |
| 5 | Data 3 | 0Ah |
| 6 | Data 4 | 80h |
| 7 | Data 5 | 51h |
| 8 | Data 6 | FEh |
| 9 | Data 7 | FEh |
| 10 | CRC (low order byte) | 27h |
| 11 | CRC (high order byte) | 88h |
| 12 | DLE | FEh |
| 13 | ETX | 40h |

A.4 Scrambler

Each row in Table A-4 contains the payload of a 64-bit Symbol scrambled by a transmitter into a 64 bit output. The 64-bit input is listed in its order at the Symbol Encoding stage, prior to bit swap. The 64-bit output is listed in the same order as the input.

Table A-4. Examples of Scrambler Computations

| Symbol | Seed | Input | Output |
|--------|----------|----------------------|----------------------|
| Data | 17 8225h | 1001 1CB6 8000 26F9h | 742E 57A6 CFBF FAACH |
| Data | 7F 0A42h | 0000 020E 000B 0782h | 33A5 96AB 5625 6764h |
| Data | 25 B41Bh | 000D 0DCA 000E 05CDh | 3ADA 28D9 392D 6F77h |
| Data | 74 9989h | 0011 097D 0013 016Fh | FB4B D87E A250 5964h |
| OS | 1F EEDDh | 0100 0110 0400 98F2h | CDC6 B616 282A 68F2h |
| OS | 70 A001h | 0100 0110 0400 98F2h | AA11 4439 2C23 3CF2h |
| OS | 27 8D0Dh | 0100 0110 0400 98F2h | 9B7B 542E 1ED7 88F2h |
| OS | 23 E363h | 0100 0110 0400 98F2h | CB81 DDDE 8CA3 58F2h |

A.5 Logical Layer RS-FEC

The following examples contain an RS-FEC block as it appears on the wire. Each table entry represents a byte in decimal format. Serial byte 0 is the first to be sent, and serial byte 197 is the last. Within each byte, the most significant bit is sent first. The examples were generated with scrambling disabled.

- Example 1 (Table A-5) – a sequence of 192 bytes with increasing values {1 (oldest),2,...,192 (newest)}, followed by 2 bytes of Sync Bits with value 0333h (indicating these are all data bytes), followed by 4 bytes of redundancy bits.
- Example 2 (Table A-6) – a sequence of 191 null bytes (value = 0), followed by a single byte of value 1, followed by 2 bytes of Sync Bits with value 0333h (indicating these are all data bytes), followed by 4 bytes of redundancy bits.
- Example 3 (Table A-7) – a sequence of 192 bytes of data with decreasing values {192 (oldest),191,...,1(newest)}, with the 7th 16-byte Symbol payload replaced by a SKIP Ordered Set. These are followed by 2 bytes of Sync Bits, followed by 4 bytes of redundancy bits.
- Example 4 (Table A-8) – a sequence of Idle Packets in the first 192 bytes, followed by 2 bytes of Sync Bits with value 0333h (indicating these are all data bytes), followed by 4 bytes of redundancy bits.

Table A-5. Example 1 – RS-FEC Block

| Symbol | Value | Symbol | Value | Symbol | Value | Symbol | Value |
|--------|-------|--------|-------|--------|-------|--------|-------|
| 0 | 1 | 1 | 2 | 2 | 3 | 3 | 4 |
| 4 | 5 | 5 | 6 | 6 | 7 | 7 | 8 |
| 8 | 9 | 9 | 10 | 10 | 11 | 11 | 12 |
| 12 | 13 | 13 | 14 | 14 | 15 | 15 | 16 |
| 16 | 17 | 17 | 18 | 18 | 19 | 19 | 20 |
| 20 | 21 | 21 | 22 | 22 | 23 | 23 | 24 |
| 24 | 25 | 25 | 26 | 26 | 27 | 27 | 28 |
| 28 | 29 | 29 | 30 | 30 | 31 | 31 | 32 |
| 32 | 33 | 33 | 34 | 34 | 35 | 35 | 36 |
| 36 | 37 | 37 | 38 | 38 | 39 | 39 | 40 |
| 40 | 41 | 41 | 42 | 42 | 43 | 43 | 44 |
| 44 | 45 | 45 | 46 | 46 | 47 | 47 | 48 |
| 48 | 49 | 49 | 50 | 50 | 51 | 51 | 52 |
| 52 | 53 | 53 | 54 | 54 | 55 | 55 | 56 |
| 56 | 57 | 57 | 58 | 58 | 59 | 59 | 60 |
| 60 | 61 | 61 | 62 | 62 | 63 | 63 | 64 |
| 64 | 65 | 65 | 66 | 66 | 67 | 67 | 68 |
| 68 | 69 | 69 | 70 | 70 | 71 | 71 | 72 |
| 72 | 73 | 73 | 74 | 74 | 75 | 75 | 76 |
| 76 | 77 | 77 | 78 | 78 | 79 | 79 | 80 |
| 80 | 81 | 81 | 82 | 82 | 83 | 83 | 84 |
| 84 | 85 | 85 | 86 | 86 | 87 | 87 | 88 |
| 88 | 89 | 89 | 90 | 90 | 91 | 91 | 92 |
| 92 | 93 | 93 | 94 | 94 | 95 | 95 | 96 |
| 96 | 97 | 97 | 98 | 98 | 99 | 99 | 100 |

| Symbol | Value | Symbol | Value | Symbol | Value | Symbol | Value |
|--------|-------|--------|-------|--------|-------|--------|-------|
| 100 | 101 | 101 | 102 | 102 | 103 | 103 | 104 |
| 104 | 105 | 105 | 106 | 106 | 107 | 107 | 108 |
| 108 | 109 | 109 | 110 | 110 | 111 | 111 | 112 |
| 112 | 113 | 113 | 114 | 114 | 115 | 115 | 116 |
| 116 | 117 | 117 | 118 | 118 | 119 | 119 | 120 |
| 120 | 121 | 121 | 122 | 122 | 123 | 123 | 124 |
| 124 | 125 | 125 | 126 | 126 | 127 | 127 | 128 |
| 128 | 129 | 129 | 130 | 130 | 131 | 131 | 132 |
| 132 | 133 | 133 | 134 | 134 | 135 | 135 | 136 |
| 136 | 137 | 137 | 138 | 138 | 139 | 139 | 140 |
| 140 | 141 | 141 | 142 | 142 | 143 | 143 | 144 |
| 144 | 145 | 145 | 146 | 146 | 147 | 147 | 148 |
| 148 | 149 | 149 | 150 | 150 | 151 | 151 | 152 |
| 152 | 153 | 153 | 154 | 154 | 155 | 155 | 156 |
| 156 | 157 | 157 | 158 | 158 | 159 | 159 | 160 |
| 160 | 161 | 161 | 162 | 162 | 163 | 163 | 164 |
| 164 | 165 | 165 | 166 | 166 | 167 | 167 | 168 |
| 168 | 169 | 169 | 170 | 170 | 171 | 171 | 172 |
| 172 | 173 | 173 | 174 | 174 | 175 | 175 | 176 |
| 176 | 177 | 177 | 178 | 178 | 179 | 179 | 180 |
| 180 | 181 | 181 | 182 | 182 | 183 | 183 | 184 |
| 184 | 185 | 185 | 186 | 186 | 187 | 187 | 188 |
| 188 | 189 | 189 | 190 | 190 | 191 | 191 | 192 |
| 192 | 03 | 193 | 51 | 194 | 196 | 195 | 215 |
| 196 | 142 | 197 | 109 | | | | |

Table A-6. Example 2 - RS-FEC Block

| Symbol | Value | Symbol | Value | Symbol | Value | Symbol | Value |
|--------|-------|--------|-------|--------|-------|--------|-------|
| 0 | 0 | 1 | 0 | 2 | 0 | 3 | 0 |
| 4 | 0 | 5 | 0 | 6 | 0 | 7 | 0 |
| 8 | 0 | 9 | 0 | 10 | 0 | 11 | 0 |
| 12 | 0 | 13 | 0 | 14 | 0 | 15 | 0 |
| 16 | 0 | 17 | 0 | 18 | 0 | 19 | 0 |
| 20 | 0 | 21 | 0 | 22 | 0 | 23 | 0 |
| 24 | 0 | 25 | 0 | 26 | 0 | 27 | 0 |
| 28 | 0 | 29 | 0 | 30 | 0 | 31 | 0 |
| 32 | 0 | 33 | 0 | 34 | 0 | 35 | 0 |
| 36 | 0 | 37 | 0 | 38 | 0 | 39 | 0 |
| 40 | 0 | 41 | 0 | 42 | 0 | 43 | 0 |
| 44 | 0 | 45 | 0 | 46 | 0 | 47 | 0 |
| 48 | 0 | 49 | 0 | 50 | 0 | 51 | 0 |
| 52 | 0 | 53 | 0 | 54 | 0 | 55 | 0 |

| Symbol | Value | Symbol | Value | Symbol | Value | Symbol | Value |
|--------|-------|--------|-------|--------|-------|--------|-------|
| 56 | 0 | 57 | 0 | 58 | 0 | 59 | 0 |
| 60 | 0 | 61 | 0 | 62 | 0 | 63 | 0 |
| 64 | 0 | 65 | 0 | 66 | 0 | 67 | 0 |
| 68 | 0 | 69 | 0 | 70 | 0 | 71 | 0 |
| 72 | 0 | 73 | 0 | 74 | 0 | 75 | 0 |
| 76 | 0 | 77 | 0 | 78 | 0 | 79 | 0 |
| 80 | 0 | 81 | 0 | 82 | 0 | 83 | 0 |
| 84 | 0 | 85 | 0 | 86 | 0 | 87 | 0 |
| 88 | 0 | 89 | 0 | 90 | 0 | 91 | 0 |
| 92 | 0 | 93 | 0 | 94 | 0 | 95 | 0 |
| 96 | 0 | 97 | 0 | 98 | 0 | 99 | 0 |
| 100 | 0 | 101 | 0 | 102 | 0 | 103 | 0 |
| 104 | 0 | 105 | 0 | 106 | 0 | 107 | 0 |
| 108 | 0 | 109 | 0 | 110 | 0 | 111 | 0 |
| 112 | 0 | 113 | 0 | 114 | 0 | 115 | 0 |
| 116 | 0 | 117 | 0 | 118 | 0 | 119 | 0 |
| 120 | 0 | 121 | 0 | 122 | 0 | 123 | 0 |
| 124 | 0 | 125 | 0 | 126 | 0 | 127 | 0 |
| 128 | 0 | 129 | 0 | 130 | 0 | 131 | 0 |
| 132 | 0 | 133 | 0 | 134 | 0 | 135 | 0 |
| 136 | 0 | 137 | 0 | 138 | 0 | 139 | 0 |
| 140 | 0 | 141 | 0 | 142 | 0 | 143 | 0 |
| 144 | 0 | 145 | 0 | 146 | 0 | 147 | 0 |
| 148 | 0 | 149 | 0 | 150 | 0 | 151 | 0 |
| 152 | 0 | 153 | 0 | 154 | 0 | 155 | 0 |
| 156 | 0 | 157 | 0 | 158 | 0 | 159 | 0 |
| 160 | 0 | 161 | 0 | 162 | 0 | 163 | 0 |
| 164 | 0 | 165 | 0 | 166 | 0 | 167 | 0 |
| 168 | 0 | 169 | 0 | 170 | 0 | 171 | 0 |
| 172 | 0 | 173 | 0 | 174 | 0 | 175 | 0 |
| 176 | 0 | 177 | 0 | 178 | 0 | 179 | 0 |
| 180 | 0 | 181 | 0 | 182 | 0 | 183 | 0 |
| 184 | 0 | 185 | 0 | 186 | 0 | 187 | 0 |
| 188 | 0 | 189 | 0 | 190 | 0 | 191 | 1 |
| 192 | 03 | 193 | 51 | 194 | 229 | 195 | 109 |
| 196 | 52 | 197 | 141 | | | | |

Table A-7. Example 3 - RS-FEC Block

| Symbol | Value | Symbol | Value | Symbol | Value | Symbol | Value |
|--------|-------|--------|-------|--------|-------|--------|-------|
| 0 | 192 | 1 | 191 | 2 | 190 | 3 | 189 |
| 4 | 188 | 5 | 187 | 6 | 186 | 7 | 185 |
| 8 | 184 | 9 | 183 | 10 | 182 | 11 | 181 |

| Symbol | Value | Symbol | Value | Symbol | Value | Symbol | Value |
|--------|-------|--------|-------|--------|-------|--------|-------|
| 12 | 180 | 13 | 179 | 14 | 178 | 15 | 177 |
| 16 | 176 | 17 | 175 | 18 | 174 | 19 | 173 |
| 20 | 172 | 21 | 171 | 22 | 170 | 23 | 169 |
| 24 | 168 | 25 | 167 | 26 | 166 | 27 | 165 |
| 28 | 164 | 29 | 163 | 30 | 162 | 31 | 161 |
| 32 | 160 | 33 | 159 | 34 | 158 | 35 | 157 |
| 36 | 156 | 37 | 155 | 38 | 154 | 39 | 153 |
| 40 | 152 | 41 | 151 | 42 | 150 | 43 | 149 |
| 44 | 148 | 45 | 147 | 46 | 146 | 47 | 145 |
| 48 | 144 | 49 | 143 | 50 | 142 | 51 | 141 |
| 52 | 140 | 53 | 139 | 54 | 138 | 55 | 137 |
| 56 | 136 | 57 | 135 | 58 | 134 | 59 | 133 |
| 60 | 132 | 61 | 131 | 62 | 130 | 63 | 129 |
| 64 | 128 | 65 | 127 | 66 | 126 | 67 | 125 |
| 68 | 124 | 69 | 123 | 70 | 122 | 71 | 121 |
| 72 | 120 | 73 | 119 | 74 | 118 | 75 | 117 |
| 76 | 116 | 77 | 115 | 78 | 114 | 79 | 113 |
| 80 | 112 | 81 | 111 | 82 | 110 | 83 | 109 |
| 84 | 108 | 85 | 107 | 86 | 106 | 87 | 105 |
| 88 | 104 | 89 | 103 | 90 | 102 | 91 | 101 |
| 92 | 100 | 93 | 99 | 94 | 98 | 95 | 97 |
| 96 | 212 | 97 | 212 | 98 | 212 | 99 | 212 |
| 100 | 212 | 101 | 212 | 102 | 206 | 103 | 176 |
| 104 | 212 | 105 | 212 | 106 | 212 | 107 | 212 |
| 108 | 212 | 109 | 212 | 110 | 206 | 111 | 176 |
| 112 | 80 | 113 | 79 | 114 | 78 | 115 | 77 |
| 116 | 76 | 117 | 75 | 118 | 74 | 119 | 73 |
| 120 | 72 | 121 | 71 | 122 | 70 | 123 | 69 |
| 124 | 68 | 125 | 67 | 126 | 66 | 127 | 65 |
| 128 | 64 | 129 | 63 | 130 | 62 | 131 | 61 |
| 132 | 60 | 133 | 59 | 134 | 58 | 135 | 57 |
| 136 | 56 | 137 | 55 | 138 | 54 | 139 | 53 |
| 140 | 52 | 141 | 51 | 142 | 50 | 143 | 49 |
| 144 | 48 | 145 | 47 | 146 | 46 | 147 | 45 |
| 148 | 44 | 149 | 43 | 150 | 42 | 151 | 41 |
| 152 | 40 | 153 | 39 | 154 | 38 | 155 | 37 |
| 156 | 36 | 157 | 35 | 158 | 34 | 159 | 33 |
| 160 | 32 | 161 | 31 | 162 | 30 | 163 | 29 |
| 164 | 28 | 165 | 27 | 166 | 26 | 167 | 25 |
| 168 | 24 | 169 | 23 | 170 | 22 | 171 | 21 |
| 172 | 20 | 173 | 19 | 174 | 18 | 175 | 17 |
| 176 | 16 | 177 | 15 | 178 | 14 | 179 | 13 |

| Symbol | Value | Symbol | Value | Symbol | Value | Symbol | Value |
|--------|-------|--------|-------|--------|-------|--------|-------|
| 180 | 12 | 181 | 11 | 182 | 10 | 183 | 9 |
| 184 | 8 | 185 | 7 | 186 | 6 | 187 | 5 |
| 188 | 4 | 189 | 3 | 190 | 2 | 191 | 1 |
| 192 | 03 | 193 | 115 | 194 | 25 | 195 | 38 |
| 196 | 17 | 197 | 174 | | | | |

Table A-8. Example 4 - RS-FEC Block

| Symbol | Value | Symbol | Value | Symbol | Value | Symbol | Value |
|--------|-------|--------|-------|--------|-------|--------|-------|
| 0 | 16 | 1 | 128 | 2 | 0 | 3 | 136 |
| 4 | 16 | 5 | 128 | 6 | 0 | 7 | 136 |
| 8 | 16 | 9 | 128 | 10 | 0 | 11 | 136 |
| 12 | 16 | 13 | 128 | 14 | 0 | 15 | 136 |
| 16 | 16 | 17 | 128 | 18 | 0 | 19 | 136 |
| 20 | 16 | 21 | 128 | 22 | 0 | 23 | 136 |
| 24 | 16 | 25 | 128 | 26 | 0 | 27 | 136 |
| 28 | 16 | 29 | 128 | 30 | 0 | 31 | 136 |
| 32 | 16 | 33 | 128 | 34 | 0 | 35 | 136 |
| 36 | 16 | 37 | 128 | 38 | 0 | 39 | 136 |
| 40 | 16 | 41 | 128 | 42 | 0 | 43 | 136 |
| 44 | 16 | 45 | 128 | 46 | 0 | 47 | 136 |
| 48 | 16 | 49 | 128 | 50 | 0 | 51 | 136 |
| 52 | 16 | 53 | 128 | 54 | 0 | 55 | 136 |
| 56 | 16 | 57 | 128 | 58 | 0 | 59 | 136 |
| 60 | 16 | 61 | 128 | 62 | 0 | 63 | 136 |
| 64 | 16 | 65 | 128 | 66 | 0 | 67 | 136 |
| 68 | 16 | 69 | 128 | 70 | 0 | 71 | 136 |
| 72 | 16 | 73 | 128 | 74 | 0 | 75 | 136 |
| 76 | 16 | 77 | 128 | 78 | 0 | 79 | 136 |
| 80 | 16 | 81 | 128 | 82 | 0 | 83 | 136 |
| 84 | 16 | 85 | 128 | 86 | 0 | 87 | 136 |
| 88 | 16 | 89 | 128 | 90 | 0 | 91 | 136 |
| 92 | 16 | 93 | 128 | 94 | 0 | 95 | 136 |
| 96 | 16 | 97 | 128 | 98 | 0 | 99 | 136 |
| 100 | 16 | 101 | 128 | 102 | 0 | 103 | 136 |
| 104 | 16 | 105 | 128 | 106 | 0 | 107 | 136 |
| 108 | 16 | 109 | 128 | 110 | 0 | 111 | 136 |
| 112 | 16 | 113 | 128 | 114 | 0 | 115 | 136 |
| 116 | 16 | 117 | 128 | 118 | 0 | 119 | 136 |
| 120 | 16 | 121 | 128 | 122 | 0 | 123 | 136 |
| 124 | 16 | 125 | 128 | 126 | 0 | 127 | 136 |
| 128 | 16 | 129 | 128 | 130 | 0 | 131 | 136 |
| 132 | 16 | 133 | 128 | 134 | 0 | 135 | 136 |

| Symbol | Value | Symbol | Value | Symbol | Value | Symbol | Value |
|--------|-------|--------|-------|--------|-------|--------|-------|
| 136 | 16 | 137 | 128 | 138 | 0 | 139 | 136 |
| 140 | 16 | 141 | 128 | 142 | 0 | 143 | 136 |
| 144 | 16 | 145 | 128 | 146 | 0 | 147 | 136 |
| 148 | 16 | 149 | 128 | 150 | 0 | 151 | 136 |
| 152 | 16 | 153 | 128 | 154 | 0 | 155 | 136 |
| 156 | 16 | 157 | 128 | 158 | 0 | 159 | 136 |
| 160 | 16 | 161 | 128 | 162 | 0 | 163 | 136 |
| 164 | 16 | 165 | 128 | 166 | 0 | 167 | 136 |
| 168 | 16 | 169 | 128 | 170 | 0 | 171 | 136 |
| 172 | 16 | 173 | 128 | 174 | 0 | 175 | 136 |
| 176 | 16 | 177 | 128 | 178 | 0 | 179 | 136 |
| 180 | 16 | 181 | 128 | 182 | 0 | 183 | 136 |
| 184 | 16 | 185 | 128 | 186 | 0 | 187 | 136 |
| 188 | 16 | 189 | 128 | 190 | 0 | 191 | 136 |
| 192 | 03 | 193 | 51 | 194 | 162 | 195 | 243 |
| 196 | 190 | 197 | 223 | | | | |

A.6 USB3 Tunneling CRC

Figure A-2 provides examples of CRC calculations for USB3 Tunneled Packets containing LFPS and Ordered Sets.

Figure A-2. Examples of USB3 Tunneling Calculations

| | | | | | | | | | | |
|------|------|------------|--------|----------|----------|------------------|------------------|-----------------------|------|------|
| 31 | 24 | 16 | 8 | 0 | | | | | | |
| CRC | R | LBPM En | W R | U3 Wa | U2 Ex | S C D 2 | S C D 1 | R X T E n | LBPM | Rsvd |
| 0x22 | 0x11 | | | | | | | 0x0 | 0x0 | |

Example 1 – USB LFPS Tunnel payload (U3 Wake = 1 ; Rx Term En = 1)

| | | | | | | | | | | |
|------|------|------------|--------|----------|----------|------------------|------------------|-----------------------|------|------|
| 31 | 24 | 16 | 8 | 0 | | | | | | |
| CRC | R | LBPM En | W R | U3 Wa | U2 Ex | S C D 2 | S C D 1 | R X T E n | LBPM | Rsvd |
| 0xC1 | 0x41 | | | | | | | 0x04 | 0x0 | |

Example 2 – USB LFPS Tunnel payload (LBPM En = 1 ; Rx Term En = 1 ; LBPM PHY_CAP_10G = 1)

| | | | | | | |
|------|------|--------------------|------|-------------|-------------|-------------|
| 31 | 24 | 16 | 8 | 0 | | |
| CRC | Rsvd | Link Functionality | Rsvd | S D S | T S 2 | T S 1 |
| 0x83 | 0x0 | 0x0 | 0x2 | | | |

Example 3 – USB Ordered Set Tunnel payload (TS2 = 1)

| | | | | | | |
|------|------|--------------------|------|-------------|-------------|-------------|
| 31 | 24 | 16 | 8 | 0 | | |
| CRC | Rsvd | Link Functionality | Rsvd | S D S | T S 2 | T S 1 |
| 0xFE | 0x0 | 0x0 | 0x4 | | | |

Example 4 – USB Ordered Set Tunnel payload (SDS = 1)

A.7 Host Interface Frame CRC

This section provides an example of CRC calculation for a Host Interface Frame. The Frame payload contains 553 bytes, listed in DW order as delivered to the Transport Layer. Each DW is represented as 4 bytes, each in hexadecimal notation, with bits [31:24] at the left and bits [7:0] at the right.

| | | | |
|----|----|----|----|
| a6 | 25 | 07 | 03 |
| f3 | 05 | 82 | 49 |
| be | 34 | db | 0d |
| 21 | 87 | 2f | 14 |
| 2f | 51 | 4b | 48 |
| 56 | b8 | 9a | 81 |
| 8b | eb | 51 | a4 |
| a1 | bf | bb | 73 |
| 1d | 5f | 49 | 7f |
| df | 77 | c2 | 3a |

| | | | |
|----|----|----|----|
| 7e | 9b | 55 | a3 |
| d8 | b3 | ae | 70 |
| fa | ae | 92 | 2e |
| 77 | fd | 0b | 0e |
| 45 | 04 | ef | 95 |
| 26 | 08 | 1e | ef |
| 4b | 0c | 5b | 09 |
| 9e | 6b | 5a | 1c |
| 31 | 57 | f8 | 6b |
| 53 | 08 | 6c | 94 |
| 1e | ab | 35 | 0f |
| 8b | 34 | 70 | 55 |
| 94 | 7c | de | 76 |
| 82 | d4 | 8d | 67 |
| 1c | 69 | b5 | 2b |
| 0d | 6c | 10 | 2c |
| a7 | 09 | f0 | 80 |
| 0d | 02 | 91 | 1f |
| c5 | f6 | 0e | 3e |
| dd | ed | 44 | 36 |
| 78 | 2c | a7 | b9 |
| 6a | 3f | 8a | 8d |
| 49 | 9f | 41 | 83 |
| 8e | 82 | d1 | 30 |
| af | be | 24 | bb |
| dc | 53 | 19 | e4 |
| fe | a1 | 22 | 91 |
| 50 | a0 | b7 | af |
| 8c | a7 | 43 | 89 |
| 5b | d0 | f4 | c7 |
| 62 | 44 | 10 | fb |
| 6f | a4 | 20 | 04 |
| 84 | a3 | b7 | 8a |
| 37 | 44 | 40 | fb |
| 8b | 8c | 71 | 99 |
| a4 | 6d | e8 | 18 |
| aa | 39 | 42 | 71 |
| a0 | c9 | a8 | c8 |
| 02 | 31 | 38 | 1d |
| 54 | 7c | 29 | 9e |
| 64 | ff | da | 9b |
| a4 | df | 64 | a4 |
| 6d | 4e | 38 | d0 |
| af | 04 | c6 | c1 |
| a0 | bf | 5f | 75 |
| e2 | 7f | 44 | c5 |
| af | fd | 3f | 5b |
| f8 | 37 | b7 | 44 |
| 8f | 58 | 28 | 17 |
| 62 | b8 | ab | a4 |
| 75 | 7d | fe | b9 |
| 10 | 16 | bc | 20 |
| ff | 9a | 79 | d5 |
| e7 | 2e | 5c | d4 |
| b9 | f5 | 09 | ee |
| ab | 0e | e2 | c4 |
| b0 | af | 83 | 7b |
| 73 | b0 | f2 | ca |
| 2b | 90 | 2d | 22 |

| | | | |
|----|----|----|----|
| 8c | e1 | bd | 3d |
| 12 | b6 | 00 | ee |
| 84 | 21 | 57 | c0 |
| ff | e4 | 47 | 68 |
| 8f | f1 | 00 | 1a |
| 7d | 69 | 8d | 9f |
| 91 | 13 | 1f | 14 |
| 0f | cc | 53 | 8e |
| 15 | 04 | 81 | 07 |
| dc | 12 | 38 | 7b |
| 02 | 9d | af | c1 |
| a9 | d6 | 66 | 24 |
| 5a | bd | 3c | eb |
| 02 | eb | ff | b8 |
| d3 | 2e | 0d | 48 |
| 7b | e2 | 39 | a5 |
| d0 | a8 | 89 | 0e |
| 80 | 23 | a6 | b4 |
| c2 | 47 | a1 | 97 |
| 52 | 15 | 75 | 33 |
| aa | 7d | e7 | 31 |
| d9 | 41 | b6 | ca |
| b6 | ab | d9 | d1 |
| a3 | af | 7a | 8e |
| 6b | 97 | 21 | b4 |
| 7f | 36 | e9 | 4d |
| 3f | 42 | bb | 09 |
| 94 | b3 | 58 | 19 |
| 7e | d6 | 89 | 39 |
| 45 | dc | 06 | 9c |
| 22 | f7 | 07 | ee |
| 6b | 8a | a6 | f9 |
| 93 | 8b | fa | b4 |
| c7 | 51 | 85 | 55 |
| b8 | 15 | 7c | d2 |
| 2c | 93 | 3b | 61 |
| 61 | 74 | 37 | 2c |
| 33 | d0 | 0f | d0 |
| c9 | 83 | c3 | ac |
| b0 | cb | f2 | 4a |
| 2a | 4f | 42 | d3 |
| af | 3e | 7f | 99 |
| f4 | 89 | 62 | c9 |
| 44 | e4 | 85 | 07 |
| 98 | aa | af | ec |
| ea | cc | d4 | b8 |
| cd | d8 | ff | 4f |
| c7 | cf | 1c | 99 |
| d9 | 07 | b1 | 18 |
| 92 | 2b | da | b7 |
| ef | fd | 2d | 54 |
| 7e | 7c | 1e | ee |
| c6 | 96 | 0c | a8 |
| 4f | cf | 97 | ef |
| 15 | 23 | 16 | 66 |
| 64 | 9f | 9f | 01 |
| 97 | 73 | 09 | 79 |
| 2a | e8 | 68 | 66 |
| f0 | 69 | 77 | cf |

B Summary of Transport Layer Packets

Table B-1 summarizes the values of the header fields for each Transport Layer Packet type. It is included for informative purposes only. Normative requirements regarding specific packet types can be found in the referenced subsections.

Table B-1. Transport Layer Packet Summary

| Usage | HopID | PDF | SuppID | Length | Payload |
|---------------------------------------|-------------------------|--------------------------------------|--------|--------|-----------------------------|
| Protocol Adapter Layer Packets | | | | | |
| Tunneled Packet | target HopID | determined by Protocol Adapter Layer | 0b | Varies | Protocol Adapter Layer data |
| Control Packets | | | | | |
| Read Request | 000h | 1h | 0b | 16 | See Section 6.4.2.3 |
| Read Response | 000h | 1h | 0b | Varies | See Section 6.4.2.4 |
| Write Request | 000h | 2h | 0b | Varies | See Section 6.4.2.5 |
| Write Response | 000h | 2h | 0b | 16 | See Section 6.4.2.6 |
| Notification Packet | 000h | 3h | 0b | 16 | See Section 6.4.2.7 |
| Notification Acknowledgment Packet | 000h | 4h | 0b | 12 | See Section 6.4.2.8 |
| Hot Plug Event Packet | 000h | 5h | 0b | 16 | See Section 6.4.2.9 |
| Inter-Domain Request | 000h | 6h | 0b | Varies | See Section 6.4.2.10 |
| Inter-Domain Response | 000h | 7h | 0b | Varies | See Section 6.4.2.11 |
| Link Management Packets | | | | | |
| Idle Packet | 001h | 0h | 1b | 0 | See Section 5.1.3.3.1 |
| Credit Grant Packet | 001h | 1h | 0b | Varies | See Section 5.1.3.3.2 |
| Path Credit Sync Packet | Target HopID (Not 001h) | 0h | 1b | 4 | See Section 5.1.3.3.3 |
| Shared Buffers Credit Sync Packet | 001h | 2h | 0b | 4 | See Section 5.1.3.3.4 |
| Time Sync Packets | | | | | |
| Follow-Up Packet | 003h | 1h | 0b | 60 | See Section 7.3.3.2 |
| Inter-Domain Time Stamp Packet | 003h | 2h | 0b | 28 | See Section 7.3.3.3 |

C Examples of Link Power Management Flows

C.1 Entry to Low Power States


This appendix contains examples of entry to CL2, CL1, or CL0s states.

Note: Re-timers take a passive role in entry flows – they track the flow but do not affect it. They are therefore omitted from the examples in this section.

C.1.1 Successful Entry to CL2 State

Figure C-1 shows an example of how the Adapters on a Link transition into the CL2 state. In the example, Adapter A initiates the entry to CL2 state. Adapter B does not have any objections asserted.

Figure C-1. Successful Entry to CL2 State




The following steps take place in Figure C-1:

1. Adapter A starts sending CL2_REQ Ordered Sets to Adapter B to initiate a transition to CL2 state.
2. Adapter B receives a CL2_REQ Ordered Set, then starts sending CL2_ACK Ordered Sets to Adapter A to indicate its ability to enter CL2 state.
3. Adapter A receives a CL2_ACK Ordered Set. It then shuts down its receiver and starts sending CL_OFF Ordered Sets to Adapter B.
4. Adapter B receives a CL_OFF Ordered Set. It then shuts down its receiver. After transmitting 375 CL2_ACK Ordered Sets, it shuts down its transmitter.
5. After Adapter A sends 375 CL_OFF Ordered Sets, it shuts down its transmitter within t_{TxOff} time.

C.1.2 Successful Entry to CL0s State

Figure C-2 shows an example of how the Adapters on a Link transition to the CL0s state. In the example, Adapter A initiates the entry to CL0s state. Adapter B rejects entry to CL2 and CL1 states because it has CL2 and CL1 objections asserted. CL0s is enabled in Adapter B (*CL0s Enable* bit is 1b).

Figure C-2. Successful Entry to CL0s State




The following steps take place in Figure C-2:

1. Adapter A starts sending `CL2_REQ` Ordered Sets to Adapter B to initiate a transition to CL2 state.
2. Adapter B receives a `CL2_REQ` Ordered Set, then starts sending `CL0s_ACK` Ordered Sets to Adapter A to indicate its ability to enter CL0s state.
3. Adapter A receives a `CL0s_ACK` Ordered Set. It starts sending `CL_OFF` Ordered Sets to Adapter B.
4. Adapter B receives a `CL_OFF` Ordered Set. It then shuts down its receiver.
5. After Adapter A sends 375 `CL_OFF` Ordered Sets, it shuts down its transmitter within t_{TxOff} time.

C.1.3 Rejection to Enter CL2 State

Figure C-3 shows an example where the Adapters fail to enter to a Low Power state. In the example, Adapter A initiates entry to CL2 state. Adapter B rejects entry to CL2 and CL1 states because it has CL2 and CL1 objections asserted. CL0s is not enabled in Adapter B (*CL0s Enable* bit is 0b).

Figure C-3. Failure to Enter CL2 State



The following steps take place in Figure C-3:


1. Adapter A starts sending `CL2_REQ` Ordered Sets to Adapter B to initiate a transition to CL2 state.

2. Adapter B receives a CL2_REQ Ordered Set, then starts sending CL_NACK Ordered Sets to Adapter A to indicate its objection to enter Low Power state.
3. Adapter A receives a CL_NACK Ordered Set. It then stops sending CL2_REQ Ordered Sets and resumes CL0 operation.

C.1.4 Concurrent Requests to Enter Low Power State

Figure C-4 shows an example where both Adapters on a Link attempt to transition to CL2 state at the same time. In the example, Adapter A has the *PM Secondary* bit set to 0b and Adapter B has the *PM Secondary* bit set to 1b.

Figure C-4. Concurrent Requests to Enter CL2 State




The following steps take place in Figure C-4:

1. Adapter A starts sending CL2_REQ Ordered Sets to Adapter B to initiate a transition to CL2 state.
2. Prior to receiving a CL2_REQ Ordered Set, Adapter B starts sending CL2_REQ Ordered Sets to Adapter A.
3. Adapter A receives a CL2_REQ Ordered Set. It stops sending CL2_REQ Ordered Sets and starts sending CL_NACK Ordered Sets.
4. Adapter B receives a CL_NACK Ordered Set. It then stops sending CL2_REQ Ordered Sets and resumes CL0 operation.
5. After Adapter A stops receiving CL2_REQ Ordered Sets, it stops sending CL_NACK Ordered Sets. It then starts sending CL2_REQ Ordered Sets again and the flow continues as in Section C.1.1.

C.1.5 CL2_REQ Ordered Sets are Not Received

Figure C-5 shows an example where Adapter A attempts to transition to CL2 state, but Adapter B detects a Link error instead of the CL2_REQ Ordered Set.

Figure C-5. Error in CL2_REQ Ordered Sets




The following steps take place in Figure C-5:

1. Adapter A starts sending CL2_REQ Ordered Sets to Adapter B to initiate a transition to CL2 state.
2. Adapter B detects a Link error. It then enters Training state and starts transmitting SLOS.
3. Adapter A either detects the Link error or receives an SLOS. It stops sending CL2_REQ Ordered Sets and enters Training state.

C.1.6 CL2_REQ Ordered Sets are Partially Received

Figure C-6 shows an example where Adapter A attempts to transition to CL2 state, but a Link error occurs before Adapter B receives all CL2_REQ Ordered Sets required for the transition.

Figure C-6. CL2_REQ Ordered Sets are Partially Received




The following steps take place in Figure C-6:

1. Adapter A starts sending CL2_REQ Ordered Sets to Adapter B to initiate a transition to CL2 state.
2. Adapter B receives a CL2_REQ Ordered Set. It starts sending CL2_ACK Ordered Sets to Adapter A, indicating its ability to enter CL2 state.
3. Adapter B detects a Link error on the Lane before 375 copies of the CL2_REQ Ordered Set are received. Adapter B disables RS-FEC on the receive direction to allow detection of SLOS.

4. Adapter A receives a CL2_ACK Ordered Set. It then shuts down its receiver. Adapter A starts sending CL_OFF Ordered Sets to Adapter B.
5. After Adapter A sends 375 CL_OFF Ordered Sets. It shuts down its transmitter within tTxOff time.
6. After Adapter B sends 375 CL2_ACK Ordered Sets, it shuts down its transmitter and receiver.

Figure C-7 shows an example where Adapter A attempts to transition to CL2 state, but a Link error occurs after Adapter B responds with CL_NACK Ordered Sets.

Figure C-7. Errors in CL2_REQ Reception and CL_NACK Response




The following steps take place in Figure C-7:

1. Adapter A starts sending CL2_REQ Ordered Sets to Adapter B to initiate a transition to CL2 state.
2. Adapter B receives a CL2_REQ Ordered Set, then starts sending CL_NACK Ordered Sets to Adapter A to indicate its objection to enter Low Power state.
3. Adapter B detects a Link error on the Lane before 375 copies of the CL_NACK Ordered Set have been transmitted. Adapter B transitions the Lane to Training state, then starts sending SLOS.
4. Adapter A receives a CL_NACK Ordered Set. It then stops sending CL2_REQ Ordered Sets and resumes CL0 operation.
5. Adapter A either detects the Link error or receives an SLOS, then enters Training state.

C.1.7 Error in CL2_ACK Ordered Sets

Figure C-8 shows an example where Adapter A attempts to transition to CL2 state, but detects a Link error before receiving a CL2_ACK Ordered Set.

Figure C-8. Error in CL2_ACK Ordered Sets




The following steps take place in Figure C-8:

1. Adapter A starts sending CL2_REQ Ordered Sets to Adapter B to initiate a transition to CL2 state.
2. Adapter B receives a CL2_REQ Ordered Set, then starts sending CL2_ACK Ordered Sets to Adapter A to indicate its ability to enter CL2 state.
3. Adapter A detects a Link error on the Lane before receiving a response to the CL2_REQ Ordered Sets. It transitions the Lane to Training state, then starts sending SLOS.
4. If RS-FEC is on, Adapter B detects link errors, disables RS-FEC, and then detects SLOS. If RS-FEC is off, Adapter B detects SLOS. After Adapter B detects SLOS, it enters the Training state.

C.1.8 Error in CL_OFF Ordered Sets

Figure C-9 shows an example where Adapter A attempts to transition CL2 state, but Adapter B detects a link error before receiving the CL_OFF Ordered Sets.

Figure C-9. Error in CL_OFF Ordered Sets



The following steps take place in Figure C-9:

1. Adapter A starts sending CL2_REQ Ordered Sets to Adapter B to initiate a transition to CL2 state.
2. Adapter B receives a CL2_REQ Ordered Set. It starts sending CL2_ACK Ordered Sets to Adapter A, indicating its ability to enter CL2 state.
3. Adapter A receives a CL2_ACK Ordered Set. It then shuts down its receiver. Adapter A starts sending CL_OFF Ordered Sets to Adapter B.
4. Adapter B detects a Link error on the Lane before 375 copies of the CL2_ACK Ordered Set are transmitted. Adapter B disables RS-FEC on its receive direction to allow detection of SLOS.
5. After Adapter A sends 375 CL_OFF Ordered Sets, it shuts down its transmitter.
6. After Adapter B sends 375 CL2_ACK Ordered Sets, it shuts down its transmitter and receiver.

C.2 Exit from Low Power States


This section contains examples of Adapter behavior during exit from CL2, CL1, or CL0s states.

Note: Re-timers take an active role in exit flows and are therefore shown in the examples in this section.

C.2.1 Example: Exit from CL0s State

Figure C-10 shows an example of how the Adapters on a Link exit from CL0s state. In the example, Adapter B initiates the CL0s exit. There are four Re-timers between Adapter A and Adapter B (Re-timer 1 to Re-timer 4).

Figure C-10. CLOs Exit



Note: The bar indicates Re-timer behavior in the direction of Adapter B.


The following steps take place in Figure C-10:

1. USB4™ Port B initiates exit from CLOs by sending a Low Frequency Periodic Signaling (LFPS) burst on all Lanes of the USB4 Port for the duration of at least 16 LFPS cycles. It then returns to Electrical Idle for tPreData, after which it starts transmitting SLOS1.
2. Upon detecting an LFPS burst on a Lane, each Re-timer sends a Low Frequency Periodic Signaling (LFPS) burst on the Lane for the duration of at least 16 LFPS cycles. It then returns to Electrical Idle for tPreData. The Re-timer then enables the receiver to start calibration, not earlier than tCLxIdleRx after the last LFPS cycle received. The Re-timer also starts transmitting locally-generated CL_WAKE1.X Ordered Set Symbols in the direction of the Lane exiting from CLOs, where X is the Re-timer Index in that direction.
3. On reception of 3 back-to-back CL_WAKE1.X Ordered Set Symbols, Adapter A transmits at least 8 CL_WAKE2.X Ordered Set Symbols before returning to service other traffic.
4. When a Re-timer with index X detects 3 back-to-back CL_WAKE2.X Ordered Set Symbols, it transitions the Lane in the direction exiting from CLOs to operate on the clock retrieved from the received Symbols rather than on its local clock. From this point on, the Re-timer forwards in the direction exiting from CLOs the Symbols it receives from the Lane and stops transmitting its locally-generated CL_WAKE1.X Ordered Set Symbols.
5. On detection of 3 back-to-back SLOS Symbols, Adapter A transmits 16 TS2 Ordered Sets.

C.2.2 Example: Exit from CL2 (or CL1) State

Figure C-11 shows an example of CL2 (or CL1) exit that involves 4 Re-timers between the ends of the Link. Adapter A initiates the exit from Low Power state. Each Re-timer has 2 indexes, one in each direction. For example, Re-timer 1-4 uses an index value = 1 when communicating with Adapter A and an index value = 4 when communicating with Re-timer 2-3.

Figure C-11. CL2 (or CL1) Exit



Notes:

- The bar to the left of a Re-timer indicates its behavior in the direction of Adapter B. The bar to the Right of a Re-timer indicates its behavior in the direction of Adapter A.
- The black arrows show the exit flow of Adapter A. The blue arrows show the exit flow of Adapter B.

The following steps take place in Figure C-11:

- Port A initiates exit from CL2 (or CL1) state by sending a Low Frequency Periodic Signaling (LFPS) burst on each Lane until the receiver detects an LFPS burst. It then

returns to Electrical Idle for tPreData, after which it starts transmitting SLOS1 Ordered Sets.

2. Upon detecting an LFPS burst on a Lane, each Re-timer sends a Low Frequency Periodic Signaling (LFPS) burst in both directions for a duration of at least 3 LFPS cycles. It then returns to Electrical Idle for tPreData. The Re-timer then enables the receiver to start calibration, not earlier than tCLxIdleRx after the last LFPS cycle is received. The Re-timer also starts transmitting CL_WAKE1.X Symbols on the Lane, where X is the index of the Re-timer in the respective direction.
3. Upon detecting an LFPS burst on a Lane, Adapter B sends a Low Frequency Periodic Signaling (LFPS) burst for a duration of at least 3 LFPS cycles. It then returns to Electrical Idle for tPreData. The Router then enables the receiver to start calibration, not earlier than tCLxIdleRx after the last LFPS cycle received. The Router also starts transmitting SLOS1 Ordered Sets on the Lane.
4. Upon reception of 3 back-to-back CL_WAKE1.X Symbols with identical Re-timer Index X, Adapter A and Adapter B each start transmitting CL_WAKE2.X Symbols on the Lane.
 - The Adapters ignore any received CL_WAKE2.X Symbols interleaved with CL_WAKE1.X Symbols when it determines the reception of back-to-back CL_WAKE1.X Symbols.
5. When a Re-timer with the receive side at index X detects 3 back-to-back CL_WAKE2.X Symbols, it transitions the Lane in the opposite direction to operate on the clock retrieved from the received Symbols rather than on its local clock. From this point on, the Re-timer forwards the Symbols it receives from the Lane and stops transmitting its locally-generated CL_WAKE1.X Symbols.
 - The Re-timer ignores any received CL_WAKE1.X Symbols interleaved with CL_WAKE2.X Symbols when it determines the reception of back-to-back CL_WAKE2.X Symbols.
6. When a Re-timer detects 3 back-to-back CL_WAKE2.X Symbols, and if the Re-timer is still using its local clock to transmit in this direction, the Re-timer starts toggling in this direction between sending the last 2 CL_WAKE2 Symbols received and 2 locally-generated Symbols.
7. When an Adapter detects 7 back-to-back CL_WAKE2.X Symbols, it transitions the Lane to Training state and start transmitting SLOS1.

D Serial Time Link Protocol (STLP)

The Serial Time Link Protocol transmits Host Router Time on a designated single-wire link (called TMU_CLK_OUT). TMU_CLK_OUT can be used for two purposes:

1. To test the Time Sync Protocol as part of silicon compliance testing.
2. To provide a high-precision time source to applications external to the Router. (e.g. for compliance testing or synchronizing multiple audio sampling devices).

The Serial Time Link Protocol also optionally includes a designated input single-wire link (called TMU_CLK_IN) that can receive a stream of Serial Time Link Packets and synchronize the Host Router Clock to the input in a similar manner as over an Inter-Domain Link (see Section 7.3.2).


A Router may optionally support the Serial Time Link Protocol. If a Router supports the Serial Time Link Protocol, it shall do so as defined in this section.

D.1 Time Synchronization

A Router shall periodically transmit a Serial Time Link Packet on TMU_CLK_OUT. The period between transmissions is implementation specific, but it is recommended that a Router transmit a Serial Time Link Packet every 16 μ s.

Figure D-1 shows the parameters for the Pulse Width Modulation that shall be used when transmitting on TMU_CLK_OUT.

Figure D-1. Pulse Width Modulation



If TMU_CLK_IN is used for time synchronization, a Router shall use the Inter-Domain equations in Section 7.4.2 to calculate time sync parameters. A Host Router shall use its TMU_CLK_IN input for time synchronization if the *Inter-Domain Enable* bit is 1b, the *IDTR* bit is 0b, and the *IDTI* bit is 0b. Otherwise, a Router shall not use TMU_CLK_IN for time synchronization.



CONNECTION MANAGER NOTE


A Connection Manager shall not enable TMU_CLK_IN and Inter-Domain Time Sync at the same time. A Router may want to take precautionary measures to prevent entering a bad state if both are enabled.

D.2 Serial Time Link Packet Format

A Serial Time Link Packet shall begin with preamble that is at least 8 cycles of "NULL" bit and shall end with a postamble that is at least 8 cycles of "NULL" bit.

The Preamble shall be followed by a Frame Marker equal to 0b00001111. This pattern marks the beginning of the packet.

Figure D-2. Serial Time Link Packet Structure



The end of first bit of Frame Marker is the Time Stamp Event later used by the UFP




Figure D-3 and Table D-1 define the contents of a Serial Time Link Packet.

A Router shall transmit bit 0 of a Serial Time Link Packet first, followed by subsequent bits in increasing order ending with bit 127.

Figure D-3. Serial Time Link Packet Format




Table D-1. Serial Time Link Packet Fields

| DW | Bits | Field | Description |
|----|-------|------------|---|
| 0 | 11:0 | Rsvd | Reserved |
| | 13:12 | Lock State | <p>This field indicates the State of the Router with respect to Host Router Time.</p> <p>Locking means that the Host Router Time is not yet consistent to the accuracy required by this specification and time stamps are not valid.</p> <p>Locked means that Host Router Time is consistent and time stamps are valid.</p> <p>00b – The source is Host Router with no Inter-Domain time sync (<i>Time Stamp</i> field is valid)</p> <p>01b – Locking (<i>Time Stamp</i> field is invalid)</p> <p>10b – Reserved</p> <p>11b – Locked (<i>Time Stamp</i> field is valid)</p> <p><i>Note: When in Locking state, TimeOffsetFromHR changes inconsistently and there is no clear trend for convergence. To calculate consistency, a derivative of the TimeOffsetFromHR shall be computed. TimeOffsetFromHR is expected to grow linearly, so if the growth is not linear, then LOCK is not achieved.</i></p> |

| DW | Bits | Field | Description |
|----|-------|-------------------------|---|
| 0 | 14 | Disruption State (DS) | This field indicates whether there is a Disruption. A disruption occurs when there is a sudden glitch in time continuity. 0b - Currently there is no disruption 1b - There was a disruption since the last serial packet transmission |
| | 15 | Rsvd | Reserved |
| | 31:16 | Time Stamp[79:64] | Bits [79:64] of the Time Stamp. The Time Stamp contains the Host Router Time value at the point when the first bit of the frame marker for the packet ends. |
| 1 | 31:0 | Time Stamp[63:32] | Bits [63:32] of the Time Stamp. The Time Stamp contains the Host Router Time value at the point when the first bit of the frame marker for the packet ends. |
| 2 | 31:0 | Time Stamp[31:0] | Bits [31:0] of the Time Stamp. The Time Stamp contains the Host Router Time value at the point when the first bit of the frame marker for the packet ends. |
| 3 | 31:0 | Cyclic Redundancy Check | CRC32 remainder value computed over the entire packet payload. The CRC32 computation in an Inter-Domain Time Stamp Packet shall be based on the following CRC: <ul style="list-style-type: none"> • Width: 32 • Poly: 1EDC 6F41h • Init: FFFF FFFFh • RefIn: True • RefOut: True • XorOut: FFFF FFFFh |

Figure D-4. TMU_CLK_OUT and TMU_CLK_IN Parameters





Figure D-5. Definition of TCO_{JTR} 

The minimum SSC as specified in Section 3.4.1.3 shall be present.

D.3 TMU_CLK_OUT and TMU_CLK_IN

Table D-2 contains the electrical specification that a Router shall use for the TMU_CLK_OUT and TMU_CLK_IN signals.

Table D-2. TMU_CLK_OUT and TMU_CLK_IN Specifications

| Symbol | Parameter | Min | Max | Units | Conditions |
|---|-----------------------------------|------|------|---------|--|
| TCO_{VOH} | TMU_CLK_OUT High Voltage | 2.2 | 3.47 | Volts | $TCO_{IOH} = -600\mu A$ (set by $3.4V/0.5M\Omega$). See Note 1. |
| TCO_{VOL} | TMU_CLK_OUT Low Voltage | -0.2 | 0.4 | Volts | $TCO_{IOL} = 600\mu A$ (set by $3.4V/10K\Omega$). See Note 1. |
| TCO_{TR} | TMU_CLK_OUT Rise Time | | 2.0 | ns | 10% to 90% Rise time. |
| TCO_{TF} | TMU_CLK_OUT Fall Time | | 2.0 | ns | 90% to 10% Fall time. |
| TCI_{VIH} | TMU_CLK_IN High Voltage Detection | 2.0 | 3.47 | Volts | See Note 2, Note 3. |
| TCI_{VIL} | TMU_CLK_IN Low Voltage Detection | -0.2 | 0.5 | Volts | See Note 4. |
| TCI_{IH} | TMU_CLK_IN High input current | | 25 | μA | $V_{in} = V_{DD}$. |
| TCI_{IL} | TMU_CLK_IN Low input current | | 0.4 | μA | $V_{in} = 0V$. |
| TCO_{JTR} | TMU_CLK_OUT Jitter | 0 | 2.5 | ns | Around ideal signal. |
| Notes: | | | | | |
| <ol style="list-style-type: none"> 1. This parameter shall be verified in both transaction and steady state. The steady state condition shall be measured with a continuous high or low level. The transaction state condition shall be measured when sending TMU_CLK_OUT data. Over/undershoot shall be ignored. 2. A buffer may be used between the connector and the Router to meet these logic levels. When present, this buffer shall meet $SBRX_{VIH}$ and $SBRX_{VIL}$ as defined in Table 3-1. 3. Logical high maps to V_{IH}. 4. Logical low maps to V_{IL}. | | | | | |

E Ingress Buffer Space

This Appendix contains examples of how a Router implementer can calculate the preferred buffer size in an Ingress Adapter for each Tunneled Protocol. Section E.1 explains the calculation for Protocols that target a specific bandwidth. Section E.2 explains the calculation for DP Main-Link Path.

E.1 Target Bandwidth Buffer Calculation

The formula below can be used to calculate the number of Ingress buffers that are needed to reach a target bandwidth for a specific Tunneled Protocol over a specific USB4 Link speed. The calculation should be repeated for every USB4 Link speed that the Router supports. The Router should implement enough buffer space to accommodate the maximum result.

$$\text{Number of Ingress Buffers} = (\text{Round Trip delay} / \text{Packet delay}) * (\text{Target Protocol BW} / \text{USB4 Link BW})$$

Packet delay is the transmit time of average sized Packet over the USB4 Link

Round Trip delay is the sum of the following:

1. The transmit time of average sized Packet over the USB4 Link.
2. The transmit time of a Credit Grant Packet over the USB4 Link.
3. The maximum latency of the Re-timers in both directions.
4. Time from Packet reception to Credit Grant Packet transmit.
5. Time from Credit Grant reception to Packet transmit.
6. The time a Packet travel through the cable in both directions.
7. The maximum skew between the 2 Lanes in both directions (if applicable).

The *Round Trip delay* can be calculated as follows:

$$\text{Round Trip Delay} = (\text{AVG_PS} / \text{ULS}) + (\text{CG size} / \text{ULS}) + (2 * \text{NoR} * \text{Max Re-timer Latency}) + (\text{PtoCG}) + (\text{CGtoP}) + (2 * \text{CD}) + (2 * 64)$$

where:

ULS = USB4 Link Speed

CD = Cable Delay

AVG_PS = Average Packet Size

NoR = Number of Re-timers

PtoCG = Time from Packet reception to Credit Grant Packet transmit

CGtoP = Time from Credit Grant Packet reception to Packet transmit

E.1.1 Example for USB3 Tunneling Ingress Buffer Calculation

The following parameters are assumed for the calculations:

- Target Protocol BW = 8Gbps for Gen2x1 ; 17Gbps for all other USB4 Links
- CD = 2 Meter Copper * 8 ns/meter
- AVG_PS = 183 Bytes
- NoR = 6

- PtoCG = 400 ns
- CGtoP = 300 ns

Round Trip delay (Gen2x1) = $(183 \cdot 8 / 10) + (8 \cdot 8 / 10) + (2 \cdot 6 \cdot 50) + (400) + (300) + (2 \cdot 2 \cdot 8) + (0)$
= ~1485ns

Round Trip delay (Gen2x2) = $(183 \cdot 8 / 20) + (8 \cdot 8 / 20) + (2 \cdot 6 \cdot 50) + (400) + (300) + (2 \cdot 2 \cdot 8) + (2 \cdot 6 \cdot 4) = \sim 1537\text{ns}$

Round Trip delay (Gen3x1) = $(183 \cdot 8 / 20) + (8 \cdot 8 / 20) + (2 \cdot 6 \cdot 30) + (400) + (300) + (2 \cdot 2 \cdot 8) + (0)$
= ~1169ns

Round Trip delay (40Gbps) = $(183 \cdot 8 / 40) + (8 \cdot 8 / 40) + (2 \cdot 6 \cdot 30) + (400) + (300) + (2 \cdot 2 \cdot 8) + (2 \cdot 6 \cdot 4) = \sim 1259\text{ns}$

The following number of Ingress buffers are needed for each USB4 Link speed, assuming 10% is taken off the USB4 Link bandwidth:

- Number of Buffers (Gen2x1) = $1485 / (183 \cdot 8 / 10) \cdot 8 / 9 = \sim 9$
- Number of Buffers (Gen2x2) = $1537 / (183 \cdot 8 / 20) \cdot 17 / 18 = \sim 19.8$
- Number of Buffers (Gen3x1) = $1169 / (183 \cdot 8 / 20) \cdot 17 / 18 = \sim 15.1$
- Number of Buffers (Gen3x2) = $1259 / (183 \cdot 8 / 40) \cdot 17 / 36 = \sim 16.2$

Therefore, for a Single-Lane Link, the Router needs 16 Ingress buffers to cover the maximum case (15.1), and for Dual-Lane Link it needs 20 Ingress buffers to cover the maximum case (19.8).

E.2 Ingress Buffers Calculation for DP Main Path

A DP Main-Link Path operates with the Flow Control Disabled scheme. The amount of buffers at the Ingress Lane Adapter is calculated to accommodate a temporary state where the DP Main-Link Path Packets are not drained immediately. The possible causes for delaying a DP Main-Link Packet are:

- Lower priority Tunneled Packet which started before the DP Main packet has arrived
- Control Packet
- A DP Main-Link Path Packet that arrives from another Adapter
- Credit Grant Packet
- Credit Sync Packet
- Time Sync Packet
- 16 Symbols for Enter or Exit CL0s

In Addition to the above buffers, a Router implementer needs to consider any internal delay from DP Main-Link Path Packet reception to DP Main-Link Path Packet dequeue from the Ingress Lane Adapter.