



GEN-Z CONSORTIUM™

Physical Layer Specification

Version 1.1

DISCLAIMER

This document is provided 'as is' with no warranties whatsoever, including any warranty of merchantability, noninfringement, fitness for any particular purpose, or any warranty otherwise arising out of any proposal, specification, or sample. Gen-Z Consortium™ disclaims all liability for infringement of proprietary rights, relating to use of information in this document. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Gen-Z is a trademark or registered trademark of the Gen-Z Consortium.

All other product names are trademarks, registered trademarks, or servicemarks of their respective owners.

DOWNLOAD DISCLAIMER

The Gen-Z Consortium (the 'Provider') is granting You a nonexclusive, worldwide, royalty-free, copyright license to this Document for evaluation purposes only, provided You abide by the terms of this Agreement. You understand that no assurances are provided that the Document does not infringe the intellectual property rights of any other entity. Neither the Provider nor any member thereof grants any other license grant, including a patent license, of any kind, whether expressed or implied or by estoppel. As a condition of exercising the rights and licenses granted under this Agreement, You assume sole responsibility to obtain any other intellectual property rights needed.

If You provide any feedback, comments, or inputs related to the Document, you agree to release, waive, discharge, defend, indemnify, and hold the Provider, its affiliates, officers, employees, agents, representatives, members and other third parties harmless from all liabilities, losses, damages, costs, and expenses (including attorneys' fees) on account of any claim, suit, action, demand, or proceeding made or brought against any such party, or on account of any feedback, comments, or inputs you provide to the Provider.

Change Summary

Revision	Date	Description
1.0	2019-1-18	1.0 Initial Release. Including the following clauses: <ul style="list-style-type: none">• Gen-Z Common• Gen-Z-E-NRZ-25G-Fabric• Gen-Z-E-NRZ-25G-Local• Gen-Z-E-NRZ-PCIe Gen 1-4
1.1	2019-9-17	Clauses/Features Added: <ul style="list-style-type: none">• Gen-Z-E-NRZ-PCIe enhanced to support PCIe Gen 5• Gen-Z-E-PAM4-50G-Fabric• Gen-Z-E-PAM4-50G-Local

Contents

CHANGE SUMMARY	3
CONTENTS	4
FIGURES	8
TABLES	10
1. PHYSICAL LAYER INTRODUCTION	13
1.1. Overview	13
1.2. Reference Documents.....	14
1.3. Documentation Conventions	15
1.3.1. <i>CLAUSE CONVENTIONS</i>	18
2. GLOSSARY	20
3. PHYSICAL LAYER ARCHITECTURE	27
3.1. Physical Layer Interface	28
4. GEN-Z COMMON SPECIFICATION	30
4.1. Introduction	30
4.2. Logical Sublayer.....	31
4.2.1. <i>LINK SERIALIZATION</i>	34
4.2.2. <i>SCRAMBLING / DESCRAMBLING</i>	35
4.2.3. <i>DATA STRIPING</i>	35
4.2.4. <i>PLA DATA WIDTH</i>	37
4.2.5. <i>PHYSICAL LAYER INITIALIZATION AND TRAINING</i>	38
4.2.6. <i>LANE REVERSAL</i>	45
4.2.7. <i>LANE POLARITY</i>	47
4.2.8. <i>LINK WIDTH</i>	49
4.2.9. <i>LINK BIST</i>	52
4.2.10. <i>PHYSICAL LAYER RETRAINING</i>	52
4.2.11. <i>PLTSM</i>	53
4.2.12. <i>CLOCK COMPENSATION</i>	60
4.2.13. <i>LOOPBACK</i>	60
4.2.14. <i>PHYSICAL LAYER AGGREGATION</i>	62
4.2.15. <i>RE-TIMERS</i>	63
4.3. Power Management	63
4.3.1. <i>OVERVIEW</i>	63
4.3.2. <i>LOW POWER OPERATION</i>	63
4.3.3. <i>PHYSICAL LAYER POWER MANAGEMENT STATES</i>	70
4.4. Electrical Sublayer	70
4.5. Management Control and Status.....	70
4.5.1. <i>INTERFACE PHYSICAL LAYER STRUCTURE</i>	70
4.5.2. <i>PHYSICAL LAYER MANAGEMENT</i>	71
5. GEN-Z-E-NRZ-25G-FABRIC SPECIFICATION	75
5.1. Introduction	75

5.2.	Logical Sublayer.....	75
5.2.1.	<i>LINK SERIALIZATION</i>	75
5.2.2.	<i>SCRAMBLING / DESCRAMBLING</i>	75
5.2.3.	<i>DATA STRIPING</i>	75
5.2.4.	<i>PLA DATA WIDTH</i>	76
5.2.5.	<i>PHYSICAL LAYER INITIALIZATION AND TRAINING</i>	76
5.2.6.	<i>LANE REVERSAL</i>	77
5.2.7.	<i>LANE POLARITY</i>	77
5.2.8.	<i>LINK WIDTH</i>	77
5.2.9.	<i>LINK BIST</i>	78
5.2.10.	<i>PHYSICAL LAYER RETRAINING</i>	78
5.2.11.	<i>PLTSM</i>	79
5.2.12.	<i>CLOCK COMPENSATION</i>	85
5.2.13.	<i>LOOPBACK</i>	86
5.2.14.	<i>PHYSICAL LAYER AGGREGATION</i>	86
5.2.15.	<i>RE-TIMERS</i>	86
5.3.	Power Management	86
5.3.1.	<i>OVERVIEW</i>	86
5.3.2.	<i>LOW POWER OPERATION</i>	86
5.3.3.	<i>PHYSICAL LAYER POWER MANAGEMENT STATES</i>	86
5.4.	Electrical Sublayer	88
5.4.1.	<i>ELECTRICAL INTERFACE</i>	88
5.4.2.	<i>HIGH SPEED SIGNALING SPECIFICATION</i>	89
5.4.3.	<i>CHANNEL CHARACTERISTICS</i>	89
5.4.4.	<i>MISCELLANEOUS SPECIFICATION</i>	90
5.5.	Management Control and Status.....	90
5.5.1.	<i>INTERFACE PHYSICAL LAYER STRUCTURE</i>	90
5.5.2.	<i>PHYSICAL LAYER MANAGEMENT</i>	130
6.	GEN-Z-E-NRZ-25G-LOCAL SPECIFICATION	131
6.1.	Introduction	131
6.2.	Logical Sublayer.....	131
6.2.1.	<i>LINK SERIALIZATION</i>	131
6.2.2.	<i>SCRAMBLING / DESCRAMBLING</i>	131
6.2.3.	<i>DATA STRIPING</i>	131
6.2.4.	<i>PLA DATA WIDTH</i>	131
6.2.5.	<i>PHYSICAL LAYER INITIALIZATION AND TRAINING</i>	131
6.2.6.	<i>LANE REVERSAL</i>	131
6.2.7.	<i>LANE POLARITY</i>	131
6.2.8.	<i>LINK WIDTH</i>	132
6.2.9.	<i>LINK BIST</i>	132
6.2.10.	<i>PHYSICAL LAYER RETRAINING</i>	132
6.2.11.	<i>PLTSM</i>	132
6.2.12.	<i>CLOCK COMPENSATION</i>	136
6.2.13.	<i>LOOPBACK</i>	136
6.2.14.	<i>PHYSICAL LAYER AGGREGATION</i>	136
6.2.15.	<i>RE-TIMERS</i>	136
6.3.	Power Management	136

6.4.	Electrical Sublayer	136
6.4.1.	<i>ELECTRICAL INTERFACE</i>	137
6.4.2.	<i>HIGH SPEED SIGNALING SPECIFICATION</i>	137
6.4.3.	<i>CHANNEL CHARACTERISTICS</i>	137
6.4.4.	<i>MISCELLANEOUS SPECIFICATION</i>	139
6.5.	Management Control and Status.....	139
6.5.1.	<i>INTERFACE PHYSICAL LAYER STRUCTURE</i>	139
6.5.2.	<i>PHYSICAL LAYER MANAGEMENT</i>	143
7.	GEN-Z-E-NRZ-PCIE SPECIFICATION	144
7.1.	Introduction	144
7.2.	Logical Sublayer.....	145
7.2.1.	<i>LINK SERIALIZATION</i>	145
7.2.2.	<i>SCRAMBLING / DESCRAMBLING</i>	150
7.2.3.	<i>DATA STRIPING</i>	150
7.2.4.	<i>PLA DATA WIDTH</i>	151
7.2.5.	<i>PHYSICAL LAYER INITIALIZATION AND TRAINING</i>	152
7.2.6.	<i>LANE REVERSAL</i>	154
7.2.7.	<i>LANE POLARITY</i>	155
7.2.8.	<i>LINK WIDTH</i>	155
7.2.9.	<i>LINK BIST</i>	155
7.2.10.	<i>PHYSICAL LAYER RETRAINING</i>	155
7.2.11.	<i>PLTSM</i>	155
7.2.12.	<i>CLOCK COMPENSATION</i>	155
7.2.13.	<i>LOOPBACK</i>	155
7.2.14.	<i>PHYSICAL LAYER AGGREGATION</i>	155
7.2.15.	<i>RE-TIMERS</i>	156
7.3.	Power Management	156
7.4.	Electrical Sublayer	156
7.5.	Management Configuration & Status	156
7.5.1.	<i>INTERFACE PHYSICAL LAYER STRUCTURE</i>	156
7.5.2.	<i>PHYSICAL LAYER MANAGEMENT</i>	169
8.	GEN-Z-E-PAM4-50G-FABRIC SPECIFICATION	170
8.1.	Introduction	170
8.2.	Logical Sublayer.....	170
8.2.1.	<i>LINK SERIALIZATION</i>	173
8.2.2.	<i>SCRAMBLING / DESCRAMBLING</i>	184
8.2.3.	<i>DATA STRIPING</i>	184
8.2.4.	<i>PLA DATA WIDTH</i>	184
8.2.5.	<i>PHYSICAL LAYER INITIALIZATION AND TRAINING</i>	184
8.2.6.	<i>LANE REVERSAL</i>	191
8.2.7.	<i>LANE POLARITY</i>	191
8.2.8.	<i>LINK WIDTH</i>	191
8.2.9.	<i>LINK BIST</i>	191
8.2.10.	<i>PHYSICAL LAYER RETRAINING</i>	191
8.2.11.	<i>PLTSM</i>	191
8.2.12.	<i>CLOCK COMPENSATION</i>	192

8.2.13.	LOOPBACK.....	193
8.2.14.	PHYSICAL LAYER AGGREGATION	193
8.2.15.	RE-TIMERS.....	193
8.3.	Power Management	193
8.4.	Electrical Sublayer.....	193
8.4.1.	ELECTRICAL INTERFACE	194
8.4.2.	HIGH SPEED SIGNALING SPECIFICATION	195
8.4.3.	CHANNEL CHARACTERISTICS	195
8.4.4.	MISCELLANEOUS SPECIFICATION.....	196
8.5.	Management Control and Status.....	196
8.5.1.	INTERFACE PHYSICAL LAYER STRUCTURE	196
8.5.2.	PHYSICAL LAYER MANAGEMENT.....	212
9.	GEN-Z-E-PAM4-50G-LOCAL SPECIFICATION	213
9.1.	Introduction	213
9.2.	Logical Sublayer.....	213
9.2.1.	LINK SERIALIZATION	213
9.2.2.	SCRAMBLING / DESCRAMBLING.....	213
9.2.3.	DATA STRIPING	213
9.2.4.	PLA DATA WIDTH	213
9.2.5.	PHYSICAL LAYER INITIALIZATION AND TRAINING	213
9.2.6.	LANE REVERSAL.....	213
9.2.7.	LANE POLARITY	213
9.2.8.	LINK WIDTH	214
9.2.9.	LINK BIST.....	214
9.2.10.	PHYSICAL LAYER RETRAINING.....	214
9.2.11.	PLTSM.....	214
9.2.12.	CLOCK COMPENSATION	214
9.2.13.	LOOPBACK.....	215
9.2.14.	PHYSICAL LAYER AGGREGATION	215
9.2.15.	RE-TIMERS.....	215
9.3.	Power Management	215
9.4.	Physical Layer Electrical Sublayer	215
9.4.1.	ELECTRICAL INTERFACE	216
9.4.2.	HIGH SPEED SIGNALING SPECIFICATION	217
9.4.3.	CHANNEL CHARACTERISTICS	217
9.4.4.	MISCELLANEOUS SPECIFICATION.....	218
9.5.	Management Control and Status.....	218
9.5.1.	INTERFACE PHYSICAL LAYER STRUCTURE	218
9.5.2.	PHYSICAL LAYER MANAGEMENT.....	221
	APPENDIX A GEN-Z AIC SPECIFICATION.....	222
A.1.	AIC Introduction	222
A.2.	AIC Loss Budget	222
A.2.1	DIE PAD INTERFACE LOSS BUDGETS	222
A.2.2	PACKAGE TO BOARD INTERFACE LOSS BUDGETS.....	223

Figures

FIGURE 1-1: GEN-Z PHYSICAL LAYER CLAUSES	13
FIGURE 3-1: PHYSICAL LAYER ARCHITECTURE.....	28
FIGURE 3-2: PLA INTERFACE SECTIONS	29
FIGURE 4-1: PHY COMMON SPECIFICATION CLAUSE REFERENCE.....	30
FIGURE 4-2: TRANSMIT (Tx) PHYSICAL LAYER LOGICAL SUBLAYER BLOCK (64B/66B).....	32
FIGURE 4-3: RECEIVE (Rx) PHYSICAL LAYER LOGICAL SUBLAYER BLOCK (64B/66B).....	33
FIGURE 4-4: SERIALIZER AND DESERIALIZER FUNCTIONS.....	34
FIGURE 4-5: DATA STRIPING WITH A 128-BIT PLA INTERFACE AND A 2-LANE LINK.....	36
FIGURE 4-6: DATA STRIPING WITH A 64-BIT PLA INTERFACE AND A 4-LANE LINK.....	37
FIGURE 4-7: DESKEW USING ALIGN HEADERS AND SETS	45
FIGURE 4-8: EQUATION FOR MAXIMUM TOLERATED SKEW BETWEEN GEN-Z LANES	45
FIGURE 4-9: STANDARD VERSUS LANE-REVERSED CONNECTION OF GEN-Z COMPONENTS	46
FIGURE 4-10: LANE-REVERSED UNIFORMITY EXAMPLES	46
FIGURE 4-11: LANE DIFFERENTIAL POLARITY REVERSAL	48
FIGURE 4-12: GEN-Z LANE POLARITY INVERSION FOR EFFICIENT ROUTING.....	49
FIGURE 4-13: ASYMMETRIC PHY EXAMPLES #1	51
FIGURE 4-14: ASYMMETRIC PHY EXAMPLES #2	51
FIGURE 4-15: PLTSM	54
FIGURE 4-16: LINK PROGRESSION THROUGH ALIGN PHYSICAL LAYER TRAINING STATE.....	57
FIGURE 4-17: MASTER LOOPBACK.....	61
FIGURE 4-18: SLAVE LOOPBACK.....	62
FIGURE 4-19: PHY-LP ENTRANCE.....	65
FIGURE 4-20: PHY-LP ENTRANCE SUPPORTING PHY-UP SEQUENCES.....	66
FIGURE 4-21: PHY-LP EXIT	67
FIGURE 4-22: PHY-UP LP ENTRANCE.....	68
FIGURE 4-23: PHY-UP LP ENTRANCE SUPPORTING PHY-UP SEQUENCES.....	68
FIGURE 4-24: PHY-UP LP EXIT	69
FIGURE 4-25: PHY-UP LP EXIT SUPPORTING PHY-UP SEQUENCES	69
FIGURE 4-26: INTERFACE PHY STRUCTURE COMMON, LANE, AND LOW POWER FIELDS	71
FIGURE 4-27: INTERFACE PHY STRUCTURE EXTENDED FEATURE FIELDS.....	71
FIGURE 4-28: PHYSICAL LAYER TRAINING FLOW.....	72
FIGURE 4-29: PHYSICAL LAYER STATUS SUB-PROCESS	73
FIGURE 4-30: PHYSICAL LAYER ADVANCED MANAGEMENT SUB-PROCESS.....	74
FIGURE 4-31: PHYSICAL LAYER RUNTIME RECONFIGURATION SUB-PROCESS.....	74
FIGURE 5-1: RETRAINING STATE MACHINE ARCS.....	78
FIGURE 5-2: GEN-Z-E-NRZ-25G-FABRIC PLTSM	79
FIGURE 5-3: GEN-Z SIGNAL DETECT SUB-STATE DIAGRAM.....	80
FIGURE 5-4: GEN-Z-E-NRZ-25G SYMMETRIC X4 DETECT	82
FIGURE 5-5: GEN-Z SIGNAL DETECT EXAMPLE BETWEEN TWO ENDPOINTS.....	84
FIGURE 5-6: GEN-Z-E-NRZ-25G-FABRIC PHY SPECIFIC FIELDS.....	95
FIGURE 5-7: GEN-Z-E-NRZ-25G-FABRIC PHY SPECIFIC FIELD LANE INDEXING	96
FIGURE 5-8: GEN-Z-E-NRZ-25G-FABRIC PHY SPECIFIC FIELD STATE INDEXING	97
FIGURE 6-1: GEN-Z-E-NRZ-25G-LOCAL PLTSM	133
FIGURE 6-2: SYMMETRIC X4 TO ASYMMETRIC Tx3:Rx1 AND Tx1:Rx3.....	134
FIGURE 6-3: ASYMMETRIC X2:X6 AND X1:X7	134
FIGURE 6-4: GEN-Z-E-NRZ-25G-LOCAL CHANNEL INSERTION LOSS.....	138
FIGURE 7-1: GEN-Z PCIe PHYSICAL LAYER	144
FIGURE 7-2: GEN-Z PACKETS ON X4 LINK WITH 8B/10B PCIe ENCODING	147
FIGURE 7-3: GEN-Z PACKETS ON X1 LINK WITH 128B/130B PCIe ENCODING	148

FIGURE 7-4: GEN-Z PACKETS ON X4 LINK WITH 128B/130B PCIe ENCODING	149
FIGURE 7-5: PHIT CRC ON AN 8-LANE PCIe PHYSICAL LAYER LINK	150
FIGURE 7-6: PCIe TO GEN-Z ALTERNATE PROTOCOL NEGOTIATION EXAMPLE.....	154
FIGURE 7-7: GEN-Z-E-NRZ-PCIe PHY SPECIFIC FIELDS	162
FIGURE 7-8: GEN-Z-E-NRZ-PCIe PHY SPECIFIC WINDOWED FIELDS.....	163
FIGURE 8-1: GEN-Z-E-PAM4-50G TRANSMIT (TX) PHYSICAL LAYER LOGICAL SUBLAYER BLOCK	171
FIGURE 8-2: GEN-Z-E-PAM4-50G RECEIVE (RX) PHYSICAL LAYER LOGICAL SUBLAYER BLOCK	172
FIGURE 8-3: CONSTRUCTION OF PHIT FEC 288	174
FIGURE 8-4: PHIT FEC 288 ON 4-LANE LINK	175
FIGURE 8-5: CONSTRUCTION OF PHIT FEC 320	178
FIGURE 8-6: PHIT FEC 320 ON 4-LANE LINK	180
FIGURE 8-7: GEN-Z-E-PAM4-50G-FABRIC ELECTRICAL SUBLAYER	194
FIGURE 9-1: GEN-Z-E-PAM4-50G-LOCAL ELECTRICAL SUBLAYER	216
FIGURE 9-2: GEN-Z-E-PAM4-50G-LOCAL CHANNEL INSERTION LOSS.....	218
FIGURE A-0-1: GEN-Z AIC REFERENCE	222

Tables

TABLE 1-1: DOCUMENT CLAUSE TERMINOLOGY	18
TABLE 4-1: LFSR SEED BY GEN-Z LANE NUMBER	35
TABLE 4-2: GEN-Z CORE STRIPING SPECIFICATION VERSUS LINK LANE COUNT	36
TABLE 4-3: PLA DATA WIDTH MULTIPLE (IN BYTES) VS NUMBER OF LANES USED	37
TABLE 4-4: PLA DATA WIDTH (IN BYTES) VS NUMBER OF LANES USED AND DESIRED MULTIPLIER	38
TABLE 4-5: PHY SIGNAL DETECT TRAINING SEQUENCES	39
TABLE 4-6: PHY ALIGN TRAINING SEQUENCES.....	40
TABLE 4-7: PHY IBIST TRAINING SEQUENCES	40
TABLE 4-8: PHY CONFIG TRAINING SEQUENCE	41
TABLE 4-9: LWR MODE DEFINITION	41
TABLE 4-10: 4-LANE LWR MODE	41
TABLE 4-11: PHY LOCK STATE TRAINING SEQUENCE.....	42
TABLE 4-12: DATA SEQUENCES.....	42
TABLE 4-13: PHY-UP SEQUENCES	42
TABLE 4-14: GEN-Z SUPPORTED ASYMMETRIC LINKS.....	50
TABLE 4-15: GEN-Z SUPPORTED ASYMMETRIC LINK LWR MODE	51
TABLE 4-16: 4-LANE LWR MODE EXAMPLES BASED ON LANE TRAINING.....	56
TABLE 4-17: LOW POWER PHY-UP PACKET MAPPING.....	64
TABLE 5-1: GEN-Z-E-NRZ-25G-FABRIC PHY ENCODING.....	75
TABLE 5-2: PLA LINK STATE TO PHY STATE MAPPING	76
TABLE 5-3: GEN-Z-E-NRZ-25G-FABRIC LANE-TO-LANE SKEW SPECIFICATION	77
TABLE 5-4: RETRAINING ARCS	78
TABLE 5-5: GEN-Z-E-NRZ-25G-FABRIC LOW POWER STATES	87
TABLE 5-6: GEN-Z-E-NRZ-25G-FABRIC INTERFACE PHY STRUCTURE SUPPORT EXCLUSIONS AND RESTRICTIONS	90
TABLE 5-7: GEN-Z-E-NRZ-25G-FABRIC EXTENDED FEATURE FIELDS.....	92
TABLE 5-8: GEN-Z-E-NRZ-25G-FABRIC EXTENDED FEATURE STATUS	93
TABLE 5-9: GEN-Z-E-NRZ-25G-FABRIC EXTENDED FEATURE CONTROL.....	93
TABLE 5-10: GEN-Z-E-NRZ-25G-FABRIC EXTENDED FEATURE CAPABILITY	94
TABLE 5-11: GEN-Z-E-NRZ-25G-FABRIC EXTENDED FEATURE REMOTE CAPABILITY	94
TABLE 5-12: GEN-Z-E-NRZ-25G-FABRIC PHY SPECIFIC FIELDS.....	97
TABLE 5-13: PHY SPECIFIC STATUS FIELDS.....	101
TABLE 5-14: PHY SUB-STATE MAPPING.....	102
TABLE 5-15: PHY SPECIFIC CONTROL FIELDS.....	103
TABLE 5-16: PHY SPECIFIC CAPABILITY FIELDS.....	103
TABLE 5-17: REMOTE PHY SPECIFIC CAPABILITY FIELDS	103
TABLE 5-18: PHY TX LANE STATUS FIELDS	103
TABLE 5-19: PHY TX LANE CONTROL FIELDS	106
TABLE 5-20: PHY RX LANE MARGIN FIELDS.....	107
TABLE 5-21: PHY RX LANE STATUS FIELDS	109
TABLE 5-22: PHY RX LANE CONTROL FIELDS.....	111
TABLE 5-23: PHY IDLE STATE STATUS FIELDS.....	112
TABLE 5-24: PHY IDLE STATE CONTROL FIELDS	112
TABLE 5-25: PHY IDLE STATE CAPABILITY FIELDS	113
TABLE 5-26: REMOTE PHY IDLE STATE CAPABILITY FIELDS.....	113
TABLE 5-27: PHY SIGNAL DETECT STATE STATUS FIELDS.....	114
TABLE 5-28: PHY SIGNAL DETECT STATE CONTROL FIELDS	114
TABLE 5-29: PHY SIGNAL DETECT STATE CAPABILITY FIELDS.....	116
TABLE 5-30: REMOTE PHY SIGNAL DETECT STATE CAPABILITY FIELDS.....	116
TABLE 5-31: PHY CHANNEL OPTIMIZE STATE STATUS FIELDS.....	117

TABLE 5-32: PHY CHANNEL OPTIMIZE STATE CONTROL FIELDS	117
TABLE 5-33: PHY CHANNEL OPTIMIZE STATE CAPABILITY FIELDS	118
TABLE 5-34: REMOTE PHY CHANNEL OPTIMIZE STATE CAPABILITY FIELDS.....	118
TABLE 5-35: PHY ALIGN STATE STATUS FIELDS	118
TABLE 5-36: PHY ALIGN STATE CONTROL FIELDS.....	119
TABLE 5-37: PHY ALIGN STATE CAPABILITY FIELDS	120
TABLE 5-38: REMOTE PHY ALIGN STATE CAPABILITY FIELDS	120
TABLE 5-39: PHY IBIST STATE STATUS FIELDS.....	120
TABLE 5-40: PHY IBIST STATE CONTROL FIELDS	121
TABLE 5-41: PHY IBIST STATE CAPABILITY FIELDS.....	122
TABLE 5-42: REMOTE PHY IBIST STATE CAPABILITY FIELDS.....	123
TABLE 5-43: PHY CONFIG STATE STATUS FIELDS	123
TABLE 5-44: PHY CONFIG STATE CONTROL FIELDS.....	124
TABLE 5-45: PHY CONFIG STATE CAPABILITY FIELDS	125
TABLE 5-46: REMOTE PHY CONFIG STATE CAPABILITY FIELDS	125
TABLE 5-47: PHY PHY-UP STATE STATUS FIELDS	125
TABLE 5-48: PHY PHY-UP STATE CONTROL FIELDS	126
TABLE 5-49: PHY PHY-UP STATE CAPABILITY FIELDS	126
TABLE 5-50: REMOTE PHY PHY-UP STATE CAPABILITY FIELDS.....	126
TABLE 5-51: PHY LP 1 STATE STATUS FIELDS	126
TABLE 5-52: PHY LP 1 STATE CONTROL FIELDS.....	127
TABLE 5-53: PHY LP 1 STATE CAPABILITY FIELDS	127
TABLE 5-54: REMOTE PHY LP 1 STATE CAPABILITY FIELDS	127
TABLE 5-55: PHY LP 2 STATE STATUS FIELDS	128
TABLE 5-56: PHY LP 2 STATE CONTROL FIELDS.....	128
TABLE 5-57: PHY LP 2 STATE CAPABILITY FIELDS	128
TABLE 5-58: REMOTE PHY LP 2 STATE CAPABILITY FIELDS	129
TABLE 5-59: PHY-UP LP 1 STATE STATUS FIELDS	129
TABLE 5-60: PHY-UP LP 1 STATE CONTROL FIELDS	129
TABLE 5-61: PHY-UP LP 1 STATE CAPABILITY FIELDS	130
TABLE 5-62: REMOTE PHY-UP LP 1 STATE CAPABILITY FIELDS.....	130
TABLE 6-1: GEN-Z-E-NRZ-25G-LOCAL INTERFACE PHY STRUCTURE SUPPORT EXCLUSIONS AND RESTRICTIONS.....	139
TABLE 6-2: GEN-Z-E-NRZ-25G-LOCAL EXTENDED FEATURE FIELDS.....	140
TABLE 6-3: GEN-Z-E-NRZ-25G-LOCAL EXTENDED FEATURE STATUS	141
TABLE 6-4: GEN-Z-E-NRZ-25G-LOCAL EXTENDED FEATURE CONTROL.....	142
TABLE 6-5: GEN-Z-E-NRZ-25G-LOCAL EXTENDED FEATURE CAPABILITY	142
TABLE 6-6: GEN-Z-E-NRZ-25G-LOCAL EXTENDED FEATURE REMOTE CAPABILITY	142
TABLE 6-7: GEN-Z-E-NRZ-25G-LOCAL UNIQUE INTERFACE PHY STRUCTURE REQUIREMENTS.....	143
TABLE 7-1: PLA DATA WIDTH MULTIPLE (IN BYTES) VS NUMBER OF PCIe LANES USED	151
TABLE 7-2: PLA DATA WIDTH (IN BYTES) VS NUMBER OF PCIe LANES USED AND DESIRED MULTIPLIER.....	151
TABLE 7-3: GEN-Z-E-NRZ-PCIe PORT ORIENTATION SUPPORT	152
TABLE 7-4: PLA LINK STATE TO LTSSM STATE MAPPING	153
TABLE 7-5: GEN-Z-E-NRZ-PCIe INTERFACE PHY STRUCTURE SUPPORT EXCLUSIONS AND RESTRICTIONS	157
TABLE 7-6: GEN-Z-E-NRZ-PCIe EXTENDED FEATURE FIELDS	159
TABLE 7-7: GEN-Z-E-NRZ-PCIe EXTENDED FEATURE STATUS	159
TABLE 7-8: GEN-Z-E-NRZ-PCIe EXTENDED FEATURE CONTROL	160
TABLE 7-9: GEN-Z-E-NRZ-PCIe EXTENDED FEATURE CAPABILITY	161
TABLE 7-10: GEN-Z-E-NRZ-PCIe EXTENDED FEATURE REMOTE CAPABILITY.....	161
TABLE 7-11: GEN-Z-E-NRZ-PCIe PHY-SPECIFIC FIELDS	163
TABLE 7-12: PCIe PHY STATUS FIELDS	168
TABLE 7-13: PCIe PHY CONTROL FIELDS	168
TABLE 7-14: PCIe PHY CAPABILITY FIELDS	168
TABLE 7-15: REMOTE PCIe PHY CAPABILITY FIELDS.....	169

TABLE 8-1: PHIT FEC 288 PROBABILITY OF DETECTION	177
TABLE 8-2: PHIT FEC 320 PROBABILITY OF DETECTION	183
TABLE 8-3: GEN-Z-E-PAM4-50G PHY DETECT TRAINING SEQUENCES	185
TABLE 8-4: GEN-Z-E-PAM4-50G PHY ALIGN TRAINING SEQUENCES	185
TABLE 8-5: GEN-Z-E-PAM4-50G PHY IBIST TRAINING SEQUENCES.....	186
TABLE 8-6: GEN-Z-E-PAM4-50G PHY CONFIGURATION STATE TRAINING SEQUENCE	186
TABLE 8-7: GEN-Z-E-PAM4-50G PHY LOCK STATE TRAINING SEQUENCE.....	187
TABLE 8-8: DATA SEQUENCES.....	187
TABLE 8-9: PHY-UP SEQUENCES	188
TABLE 8-8: CLOCK COMPENSATION SEQUENCES.....	189
TABLE 8-8: GEN-Z-E-PAM4-50G-FABRIC LANE-TO-LANE SKEW	190
TABLE 8-9: GEN-Z-E-PAM4-50G-FABRIC INTERFACE PHY STRUCTURE SUPPORT EXCLUSIONS AND RESTRICTIONS.....	196
TABLE 8-10: GEN-Z-E-PAM4-50G-FABRIC EXTENDED FEATURE FIELDS.....	198
TABLE 8-11: GEN-Z-E-PAM4-50G-FABRIC EXTENDED FEATURE STATUS	198
TABLE 8-12: GEN-Z-E-PAM4-50G-FABRIC EXTENDED FEATURE CONTROL.....	199
TABLE 8-13: GEN-Z-E-PAM4-50G-FABRIC EXTENDED FEATURE CAPABILITY	200
TABLE 8-14: GEN-Z-E-PAM4-50G-FABRIC EXTENDED FEATURE REMOTE CAPABILITY	201
TABLE 8-15: GEN-Z-E-PAM4-50G-FABRIC UNIQUE INTERFACE PHY STRUCTURE REQUIREMENTS.....	202
TABLE 8-16: GEN-Z-E-PAM4-50G PHY SPECIFIC STATUS FIELDS.....	202
TABLE 8-17: GEN-Z-E-PAM4-50G PHY SPECIFIC CONTROL FIELDS	203
TABLE 8-18: GEN-Z-E-PAM4-50G PHY SPECIFIC CAPABILITY FIELDS.....	205
TABLE 8-19: GEN-Z-E-PAM4-50G REMOTE PHY SPECIFIC CAPABILITY FIELDS.....	205
TABLE 8-20: GEN-Z-E-PAM4-50G PHY TX LANE STATUS FIELDS	205
TABLE 8-21: GEN-Z-E-PAM4-50G PHY TX LANE CONTROL FIELDS.....	208
TABLE 8-22: GEN-Z-E-PAM4-50G PHY RX LANE STATUS FIELDS.....	210
TABLE 9-1: GEN-Z-E-PAM4-50G-LOCAL INTERFACE PHY STRUCTURE SUPPORT EXCLUSIONS AND RESTRICTIONS.....	219
TABLE 9-2: GEN-Z-E-PAM4-50G-LOCAL UNIQUE INTERFACE PHY STRUCTURE REQUIREMENTS.....	220
TABLE A-1: PECFF AND CEM AIC DIE PAD INTERFACE DESIGN GUIDELINES	223
TABLE A-2: GEN-Z PECFF AND CEM AIC PACKAGE TO BOARD INTERFACE DESIGN GUIDELINES.....	224

1. Physical Layer Introduction

1.1. Overview

This document specifies logical, functional, electrical, and channel characteristics for a multitude of physical layers (PHY) that enable communication from the Gen-Z Core layers to other devices via a specific media and signaling interface. The Gen-Z Physical Layer Specification is formatted as a clause-based document and enables separate and subsequent development cycles for individual physical layer specification clauses required for different applications.

The Gen-Z Common Specification clause specifies logical and functional requirements for general application. If the Gen-Z Common Specification clause is not satisfactory for a given individual physical layer clause, then the individual physical layer clause shall take precedence.

Though each physical layer clause can communicate with a Gen-Z Core via the Physical Layer Abstraction (PLA) interface, they may interoperate only with another device supporting the same physical layer clause specified within this document. *Gen-Z Physical Layer Clauses* illustrates how a single Gen-Z Core can communicate with a multitude of devices over different Physical Layers using conceptually named Clauses X, Y, and Z. *Gen-Z Physical Layer Clauses* also illustrates how clauses shall be compatible in order to establish a communication channel, using Clause X and Clause Z as an example.

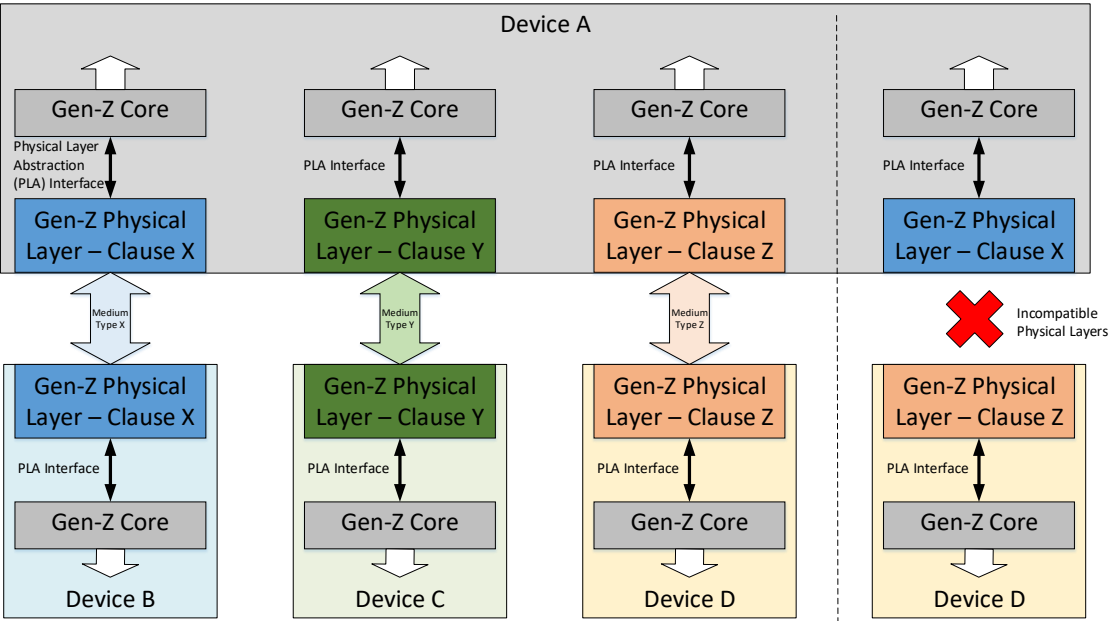


Figure 1-1: Gen-Z Physical Layer Clauses

This Gen-Z Physical Layer Specification document includes clauses for:

- A Gen-Z Common Specification specifies general functional and logical requirements for Gen-Z Physical Layers.
- Gen-Z-E-NRZ-25G-Fabric specifies a Gen-Z physical layer capable of a line-rate at 25.78125 GT/s per lane inclusive of 64b/66b encoding (raw data-rate of 25 Gbps), uses Non-Return to Zero (NRZ) signaling over Fabric media (Long Reach), and provides a BER of 10^{-12} or better. The

physical layer operates with an electrical interface compatible with the Channel Operating Margin (COM) that has a recommended Insertion Loss¹ (IL) no greater than 30 dB at 12.89 GHz.

- Gen-Z-E-NRZ-25G-Local specifies a Gen-Z physical layer capable of a line-rate at 25.78125 GT/s per lane inclusive of 64b/66b encoding (raw data-rate of 25 Gbps), uses Non-Return to Zero (NRZ) signaling over Local media (Very Short Reach), and provides a BER of 10⁻¹⁵ or better. The physical layer consists of an electrical interface compatible with a recommended Insertion Loss¹ (IL) no greater than 10 dB and no less than 4 dB at 12.89 GHz.
- Gen-Z-E-NRZ-PCIe[®] specifies a Gen-Z physical layer capable of operating at PCI Express[®] speeds up to 32 GT/s per lane. The physical layer supports the following chapters in the *PCI Express Base Specification*:
 - Physical Layer Logical Block Specification
 - Power Management
 - Electrical Sub-block
- Gen-Z-E-PAM4-50G-Fabric defines a Gen-Z physical layer capable of a line-rate at 53.125 GT/s per lane inclusive of the overhead using 4-level Pulse Amplitude Modulation (PAM4) signaling over Fabric media (Medium Reach). One of two different Phit Forward Error Correction (Phit FEC) encodings are used, Phit FEC 288 which uses a BCH(288, 260) codeword or Phit FEC 320 which uses a BCH(320, 260) codeword. Phit FEC 288 has a raw data-rate of 47.222 Gbps and supports a raw BER of 10⁻⁹ or better. Phit FEC 320 has a raw data-rate of 42.5 Gbps and supports a raw BER of 10⁻⁷ or better. Both types of Phit FEC have a corrected BER of 10⁻¹⁵ or better. A compliant physical layer consists of an electrical interface compatible with the Channel Operating Margin (COM) that has a recommended Insertion Loss¹ (IL) no greater than 20 dB at 13.2813 GHz.
- Gen-Z-E-PAM4-50G-Local defines a Gen-Z physical layer capable of a line-rate at 53.125 GT/s per lane inclusive of the overhead with Forward Error Correction (raw data-rate of 47.222 Gbps) using 4 level Pulse Amplitude Modulation (PAM4) signaling over Local mediums (Very Short Reach) providing a raw BER of 10⁻⁹ or better and a corrected BER of 10⁻¹⁵ or better. A compliant physical layer consists of an electrical interface compatible with the Channel Operating Margin (COM) that has a recommended Insertion Loss¹ (IL) no greater than 10 dB and no less than 4 dB at 13.2813 GHz.

1.2. Reference Documents

Gen-Z Core Specification, <http://genzconsortium.org>

PCIe Enclosure Compatible Form Factor (PECF) Specification, <http://genzconsortium.org>

Gen-Z Management Architecture Specification, <http://genzconsortium.org>

IEEE 802.3 Standard for Ethernet, <http://www.ieee802.org/3/>

Optical Internetworking Forum - Common Electrical I/O (OIF-CEI) Implementation Agreements, www.oiforum.com

PCI Express Base Specification, <http://www.pcisig.com>

PCI Express Card Electromechanical Specification, <http://www.pcisig.com>

¹ The Insertion Loss budget for this specification clause follows IEEE 802.3 standard conventions: test points are at the package-to-board interface.

1.3. Documentation Conventions

Shall, Should, May, and Can

This specification adheres to Section 13.1 of the IEEE Specifications Style Manual, which dictates use of the words 'shall', 'should', 'may', and 'can' in the development of documentation, as follows:

- 5 • The word *shall* is used to indicate mandatory requirements strictly to be followed in order to conform to the Specification and from which no deviation is permitted (*shall* equals *is required to*).
- The use of the word *must* is deprecated and shall not be used when stating mandatory requirements; *must* is used only to describe unavoidable situations.
- 10 • The use of the word *will* is deprecated and shall not be used when stating mandatory requirements; *will* is only used in statements of fact.
- The word *should* is used to indicate that among several possibilities one is recommended as particularly suitable, without mentioning or excluding others; or that a certain course of action is preferred but not necessarily required; or that (in the negative form) a certain course of action is deprecated but not prohibited (*should* equals *is recommended that*).
- 15 • The word *may* is used to indicate a course of action permissible within the limits of the Specification (*may* equals *is permitted*).
- The word *can* is used for statements of possibility and capability, whether material, physical, or causal (*can* equals *is able to*).

Capitalization

Some terms are capitalized to distinguish their definition in the context of this document from their common English meaning. Words not capitalized and not defined within this document's Glossary have their common English meaning.

Bit fields are capitalized to improve legibility.

The first letter of a term composed of multiple words concatenated as follows (A_B_C) is capitalized.

Operation Code and Operation Class names are capitalized.

Numbers and Number Bases

Unless explicitly stated otherwise by this specification, numerical values without qualifiers are decimal. This specification uses the following qualifiers:

- Hexadecimal numbers are written with a '0x' prefix followed by a mix of digits 0 through 9 and / or upper case English letters A through F, e.g., 0xFFFF or 0x80. Hexadecimal numbers larger than four digits are represented with a space dividing each group of four digits, as in 0x1E FFFF EFFF.
- Binary numbers are written with a lower case 'b' suffix, e.g., 1001b and 10b. Binary numbers larger than four digits are written with a space dividing each group of four digits, as in 1000 0101 0010b.

A dash between two numbers represents a range of values, e.g., 0-7 represents zero to seven inclusive.

A colon between two numbers represents a range of bits, e.g., Bits 7:0 represents a binary value of bits seven to zero inclusive.

Standard Units of Time

Unit	Description
ms	Millisecond— 10^{-3} seconds
μ s	Microsecond— 10^{-6} seconds
ns	Nanosecond— 10^{-9} seconds
ps	Picosecond— 10^{-12} seconds
fs	Femtosecond— 10^{-15} seconds

Reserved

The following applies to the term ‘Reserved’:

- The contents, state, or information are not specified at this time.
- Any field, feature, capability, etc. marked ‘Reserved’ is subject to change.
- Components shall not use any field, feature, capability, etc. that is marked ‘Reserved’.
- All RsvdP control structure fields are reserved for future read-write implementations. Field is read-only and shall return zero when read. Software shall preserve the value read for subsequent writes.
- All RsvdZ control structure fields are reserved for future RW1C implementations. Field is read-only and shall return zero when read. Software shall use 0x0 for writes.
- All Reserved control structure fields shall return 0x0 when read. Software written to this specific specification shall not write to Reserved control fields. Hardware shall ignore all writes to Reserved control structure fields. These fields are reserved for future implementations.
- Reserved encodings of any control structure, packet fields, etc. shall not be used.
- Any implementation dependence on a ‘Reserved’ field value, or encoding will result in a non-compliant implementation. The functionality of such an implementation cannot be guaranteed in this or any future revision of this specification.
- A Reserved protocol field shall be transmitted as zero and ignored upon receipt.

Read-only Control Structure Fields

Hardware shall ignore writes to read-only (RO) control structure fields.

Write-only Control Structure Fields

If a write-only (WO) Control Space field is read, then the read shall return zero and shall not be treated as an error.

Developer Notes

Developer Notes are informative. They are included for clarification and illustration only. These notes are delineated by: **Developer Note:** *Text*

State Machine Conventions

State machines describe the behavior of a component, an interface, a link, or the physical layer. A state machine does not imply the internal design or implementation. State machines use the following conventions:

- Each state machine is represented by a rectangular box.
 - The top section of the box contains the state name.
 - The bottom section of the box contains the actions which occur in the state.

5

- Transition arrows indicate state transitions which are made when the expression next to the arrow is satisfied.
- A transition arrow which does not originate in a state indicates a global transition. Such a transition will occur regardless of the current state.
- If no exit condition for state is met, then the state machine remains in the current state.
- A logical “OR” is represented by a “+”.
- A logical “AND” is represented by a “*”.
- A logical “INVERT” is represent by a “~”.
- If a conflict occurs, state machine figures take precedence over descriptive text.

Flow Chart Conventions

Flow charts illustrate packet processing, validation, etc. Flow charts use the following conventions:

- A ‘Go to’ represents a jump to another flow chart to complete processing.
- A ‘Perform XXX ()’ represents a branch to a flow chart that performs additional processing before returning to continue processing in the current flow chart. Conceptually, this is similar to invoking a sub-function.

Table Entry Conventions

Unless explicitly stated within this specification, common table entries are defined as follows.

Table Entry	Description
M	Mandatory—the field or associated semantic is treated as a ‘shall’.
O	Optional—the field or associated semantic is treated as a ‘shall if implemented’.
MC	Mandatory Conditional—If the described condition is true, then the field or associated semantic is treated as a ‘shall’.

Endian

This specification uses little endian bit and byte order.

Label Conventions

- If a capability is unsupported (e.g., a read-only capability field or bit indicates unsupported), then the corresponding / dependent control fields or bits shall be Reserved.
- Unless explicitly stated otherwise by this specification, if software sets a bit or a field to an unsupported value or enables an unsupported capability or feature, then the component behavior may be non-deterministic.
- Control structures and volatile media shall be updated as specified in *Component Reset* and *Interface Reset* of the *Gen-Z Core Specification* when a reset or power-on event occurs.
- Unless explicitly stated otherwise by this specification, if the label associated with a bit or field contains the word “Support”, then zero indicates the corresponding functionality is unsupported and a non-zero value indicates the corresponding functionality is supported. If functionality is unsupported, then the bit or field shall be hardwired to zero.
- Unless explicitly stated otherwise by this specification, if the label associated with a bit or field contains the word “Enable”, then zero indicates the corresponding functionality is disabled and a non-zero value indicates the corresponding functionality is enabled. If a field is associated with an unsupported functionality, then the enable bit or field shall be hardwired to zero.

- Unless explicitly stated otherwise by this specification, if the label associated with a bit or field contains the word “Valid”, then zero indicates the corresponding field is invalid and a non-zero value indicates the corresponding field is valid. If a field is associated with an unsupported functionality, then the valid bit or field shall be hardwired to zero.
- Unless explicitly stated otherwise by this specification, if the label associated with a bit or field contains the words “In Progress”, then zero indicates the corresponding action is not currently occurring and a non-zero indicates the corresponding action is currently occurring. These bits or fields are automatically updated by the component.

Writing Style

The writing style in Gen-Z specifications requires writers to:

- Use the active voice
- Use articles wherever possible
- Use simple verb tenses
- Use language consistently
- Avoid long compound words
- Use relatively short sentences (maximum of 20-25 words wherever possible)
- Use relatively short paragraphs (6 relatively short sentences)
- Use single and multi-level bullets to simplify and delineate information, e.g., requirements

1.3.1. Clause Conventions

Document Clause Terminology is used to construct and interpret clause names for each individual physical layer clause. The generic name shall be as follows:

Gen-Z-<Medium>-<Signaling-Type>-<Bit-Rate>-<Topology>

For example, Gen-Z-E-NRZ-8G/16G-Local is interpreted as follows: Specifies a physical layer clause using an electrical interface with a Non-Return to Zero signaling type at both 8 Gbps and 16 Gbps for a Gen-Z Local topology.

Table 1-1. Document Clause Terminology

Medium	
E	Electrical Interconnect
O	Optical Interconnect
Signaling Type	
NRZ	Non-Return to Zero
PAM4	Pulse Amplitude Modulation—4 Level
Bit-Rate	
#[M,G]	The designated # for each bit-rate may not explicitly detail the exact data-rate (minus encoding) or baud-rate (plus encoding) per channel. The # value is used in simplified and general terms to provide ease of identifying a clause specification.

Or	<ul style="list-style-type: none"> • M = Megabit • G = Gigabit
<Protocol Speed, Revision, or Generation>	<p>Multiple bit-rates can be documented within a single clause specification as such '8G/16G' when applicable.</p> <p>Additionally, a clause can label its bit-rate according to the speed, revision, or generation of a specific physical layer protocol.</p>
Or <blank>	<p>This field can be intentionally blank for specific fully leveraged physical layer protocols to indicate support for current and future generations.</p>
Topology	
Fabric	<p>Clause specification for a Fabric-based topology, typically demanding complex requirements to satisfy complicated channels or media (i.e. longer reach, higher Insertion Loss)</p>
Local	<p>Clause specification for a Local-based topology, typically demanding simpler requirements to satisfy less complicated channels or media (i.e. shorter reach, lower Insertion Loss)</p>
<Protocol>	<p>Clause specification for a specific fully leveraged physical layer protocol.</p> <p>For example, PCI Express would set <protocol> to "PCIe".</p>

2. Glossary

Term	Definition
IEEE 802.3	A working group within the Institute of Electrical and Electronic Engineers (IEEE) that produces standards for wired Ethernet.
Active State Power Management (ASPM)	A power management protocol used by PCI Express devices.
add-in card (AIC)	A card inserted into a connector and mounted in an enclosure slot.
aggregation	Aggregation refers to grouping multiple interfaces into a single active interface (this is the antithesis of interface bifurcation where a single link is split into multiple distinct interfaces).
baud rate	The number of signal changes (pulse or symbol) per second. Synonymous with line-rate. Represented as Baud per second (Bd/s or Bdps).
BIST	Built In Self-Test
Bit Error Ratio (BER)	The number of bit errors divided by the total number of transferred bits over a specific observed time interval
Bit lock	A period in physical layer training when the inbound receiving data and clock are locked to a reliable level of integrity and fidelity
bit rate	The number of bits transferred per unit of time
CDR	Clock and Data Recovery
CEM	PCI Express Card Electromechanical Form Factor
Continuous Timer Linear Equalizer (CTLE)	An equalization technique, often used by a receiver (Rx), employing a linear filter to compensate for frequency responses across a channel or medium.

Term	Definition
Cyclic Redundancy Check (CRC)	A code added to data that is used to detect errors that occur during transmission, storage, or retrieval
Decision Feedback Equalizer (DFE)	An equalization technique used by a receiver (Rx) employing an adaptive filter and feedback of detected symbols to compensate for degradation across a channel or medium.
equalization	A technique used to improve the performance of a communications signal to compensate for degradation (ISI) across a channel or medium. Equalization can be performed by a transmitter, receiver, or both with methods to configure statically during initialization, dynamically during physical layer training, or adaptively anytime the link is transmitting information.
Fast Training Sequence (FTS)	PHY sequences formatted to allow quick advancement through physical layer training.
Feed Forward Equalizer (FFE)	An equalization technique, often used by a transmitter (Tx), employing a finite impulse response (FIR) to precondition a signal to compensate for degradation across a channel or medium.
flit	Flow control unit for data transferred to/from the Gen-Z Core
Forward Error Correction (FEC)	A technique used for controlling errors in data transmission over unreliable or noisy communication channels.
Interconnect Built-In Self-Test (IBIST)	A test technique using known data patterns to evaluate the quality of a physical layer's link or interconnect.
Inter-Symbol Interference (ISI)	The distortion of a signal across a channel or medium in which one symbol interferes with subsequent symbols in either time, phase, or amplitude making communication less reliable.
Link-level Reliability (LLR)	Data validity checking at an individual link level
lane	A point-to-point signal channel that connects one transmitter and one receiver. Transmitters send electrical or photonic signals to receivers.

Term	Definition
line-rate	The actual number of signal changes (pulse or symbol) per second. Synonymous with baud rate.
Linear Feedback Shift Register (LFSR)	A shift register whose input bit is a linear function of its previous state
link	A set of one or more lanes used to exchange packets between two interfaces. A link is defined as the connection between two ports.
LTSSM	PCI Express Link Training and Status State Machine
local	Refers to the physical layer port on the near end.
LSB	Least Significant Bit is the bit position in a binary number having the lowest value. This is also known as the low-order bit.
LWR	Link Width Reduction
media	Plural form of medium (mediums) used with communication channels.
medium	A physical communication channel used to convey an information signal, such as cables, wire line, or fiber-optics.
multi-layered board (MLB)	A printed circuit board (PCB) consisting of multiple layers and electrical connections (vias).
MSB	Most Significant Bit is the bit position in a binary number having the greatest value. This is also known as the high-order bit.
NAI	Non-operational, aggregated interface—an interface that has been combined with one or more interfaces that have been configured to operate as a single interface. Though its resources (e.g., transmit and receives lanes, flow-control buffers, etc.) can be used by the SAI, the NAI cannot be used for any purpose.
NRZ	Non-Return to Zero (NRZ). Synonymous with PAM2

Term	Definition
PAM#	Pulse Amplitude Modulation where # defines the numbers of levels.
PCI Express (PCIe)	PCI Express is an I/O interconnect and protocol specified by the PCI-SIG.
PCI-SIG	Peripheral Component Interconnect Special Interest Group
PECFF	PCIe Enclosure Compatible Form Factor
PHY	Physical Layer
Physical Medium Dependent (PMD) Sublayer	The PMD defines the details for transmission and reception of signals across a physical medium. In IEEE 802.3, the PMD additionally defines the logical methods and electrical parameters for equalization optimization.
PHY Sequences	A physical layer packet exchanged between neighboring devices during physical layer training and initialization. Also known as Ordered-Sets.
Phit	Physical digit. A Phit may require one or more physical layer clock cycles to transfer across a link.
PLA	Physical Layer Abstraction interface defined by the <i>Gen-Z Core Specification</i>
PLTSM	Physical Layer Training State Machine
port	A set of one or more lanes on a single interface.
ppm	Parts Per Million
Ordered-Set	A physical layer packet exchanged between neighboring devices during physical layer training and initialization.
remote	Refers to the physical layer port on the far end.

Term	Definition
Requester	<p>A component that generates request packets, other than a <i>CTL-UE Packet</i> or an <i>Unsolicited Event (UE) Packet</i>, and receives and executes any corresponding response packets.</p> <p>Throughout this document, wherever 'Requester' is used, the specification equally applies to a Requester-Responder when acting as a Requester.</p>
Responder	<p>A component that receives and executes request packets and generates any corresponding response packets.</p> <p>Throughout this document, wherever 'Responder' is used, the specification equally applies to a Requester-Responder when acting as a Responder.</p>
RO (Read-only)	<p>An attribute where a field is read-only, i.e., it cannot be modified.</p> <p>If the feature associated with the field is unsupported, then the field shall be hardwired to zero.</p>
ROS (Read-only Sticky)	<p>An attribute where a field is RO with sticky semantics.</p> <p>If the feature associated with the field is unsupported, then the field shall be hardwired to zero.</p>
Reed-Solomon Forward Error Correction (RS-FEC)	<p>An error correction code used in various physical layer standards.</p>
RW (Read-Write)	<p>An attribute where a field is read-write, i.e., the field can be modified.</p> <p>If the feature associated with the field is unsupported, then the field shall be hardwired to zero.</p>
RW1C (Read-Write 1 to Clear)	<p>A bit field attribute where a read returns the bit's value, and writing a value of 1b to the bit clears its value.</p> <p>Unless explicitly stated otherwise by this specification, writing a value of 0b shall have no effect on the field.</p> <p>If the feature associated with the field is unsupported, then the field shall be hardwired to zero.</p>
RW1CS (Read-Write 1 to Clear, Sticky)	<p>A bit field attribute of RW1C with sticky semantics.</p> <p>If the feature associated with the field is unsupported, then the field shall be hardwired to zero.</p>
RWS (Read-Write Sticky)	<p>A bit field attribute of RW with Sticky semantics.</p>

Term	Definition
	If the feature associated with the field is unsupported, then the field shall be hardwired to zero.
Receiver (Rx)	A lane's receiving destination of an electrical or photonic signal in a point-to-point connection.
SAI	Single, aggregated interface—the sole operational interface among the set of interfaces that have configured to operate as a single interface.
sticky	Bit or field attribute where the value is preserved across a Warm Component Reset or a Warm Interface Reset.
switch	A component or component-integrated functionality that performs packet relay between component interfaces.
symbol	Unit of information transferred in the channel/medium on a single transition. For NRZ signaling, a symbol is a single bit (logical 0 or 1). For multi-level signaling, a symbol can consist of more than one bit.
transfer rate	The amount of information transferred per second, accounting for signal modulation. Represented as Transfers per second (T/s or Tps) or bits per second (b/s or bps).
Transmitter (Tx)	A lane's driving source of an electrical or photonic signal in a point-to-point connection.
Unit Interval (UI)	UI is the time interval taken to transmit one symbol on a single lane of a link. UI duration is a function of the transfer rate, not the width of a link.
via	An electrical connection between physical layers in a PCB.
WO (Write-only)	An attribute where a field is write-only. Reads of a write-only field shall return zero. If the feature associated with the field is unsupported, then the field shall be hardwired to zero.
x# link	Defines a link with # numbers of lanes. For example, x4 defines a link with four lanes.

3. Physical Layer Architecture

Gen-Z Physical layer clauses consists of four major architectural sublayer clauses as illustrated in *Physical Layer Architecture*:

- 5 • Physical Layer Logical Sublayer—Specifies the logical functions of the Gen-Z Physical Layer including details for:
 - Link / physical layer initialization
 - Link / physical layer training and retraining flows
 - Lane aggregation within a link
 - Lane-to-lane deskew
 - 10 ○ Lane polarity
 - Link reversal
 - Clock and Data Recovery (CDR)
 - Bit lock and alignment
 - Data serialization and deserialization
 - 15 ○ Data scrambling and descrambling
 - Data encoding and decoding
 - Lane-to-lane data striping
 - Methods for test
- 20 • Physical Layer Electrical Sublayer—Specifies the electrical interface requirements of the Gen-Z Physical Layer as it communicates across a physical medium or channel and includes details for:
 - Link / physical layer signaling type
 - Lane signaling speed
 - Methods for channel optimization
 - Transmitter output signaling
 - 25 ○ Receiver input tolerances
 - Allowable channel characteristics
- 30 • Physical Layer Power Management—Specifies the functional requirements for power management of the Gen-Z Physical Layer including details for:
 - Functional states and flows
 - Electrical characteristics for each power management state
- Physical Layer Management Control and Status—Specifies the functional interface structures for software to configure and report status of the physical layer. The interface to and from the Physical Management Control and Status is not defined and implementation specific.

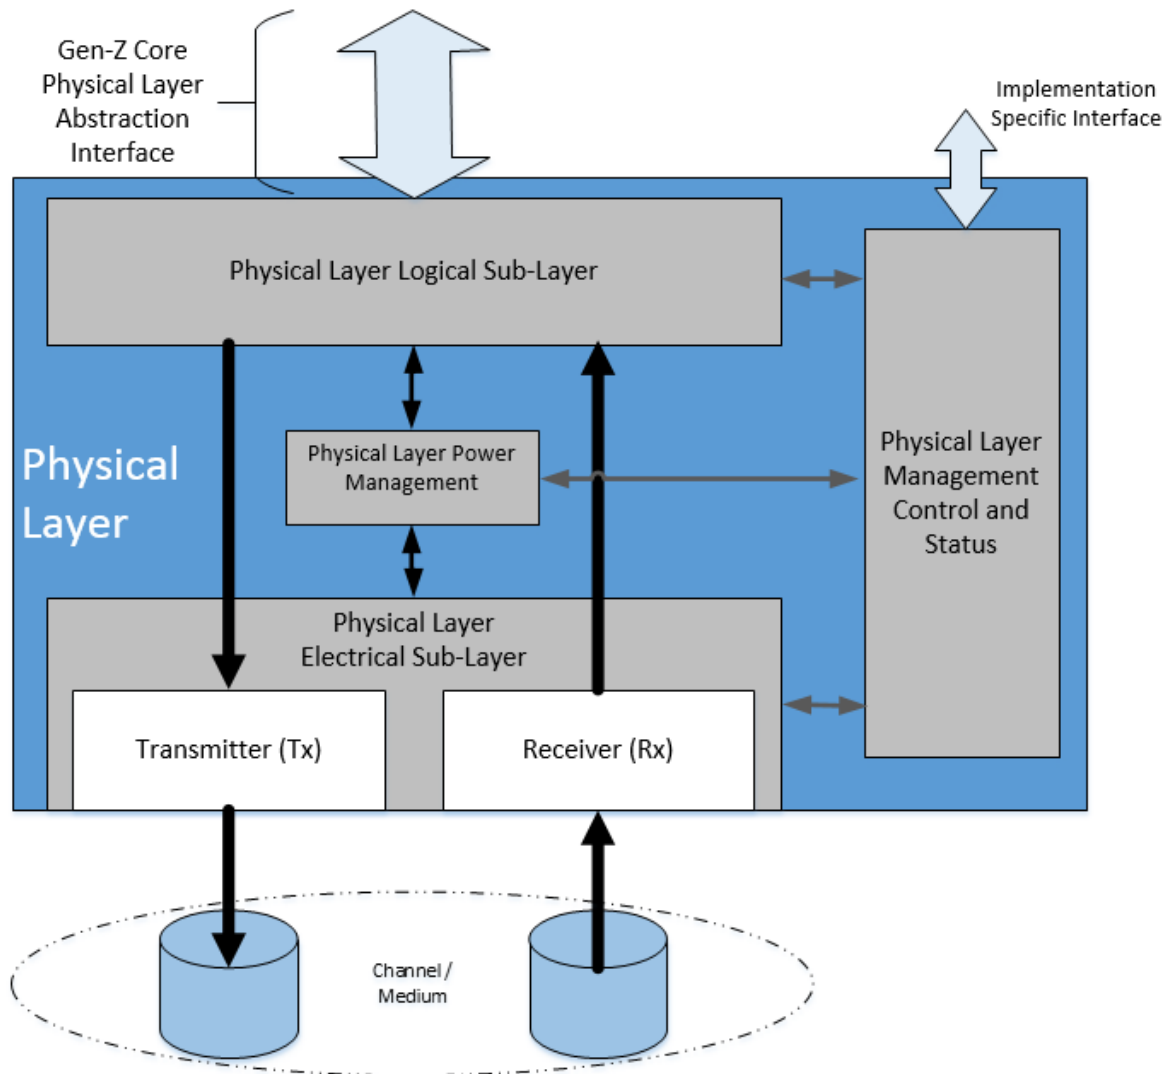


Figure 3-1: Physical Layer Architecture

3.1. Physical Layer Interface

Though each Gen-Z Physical Layer clause is uniquely specified, each shall communicate to the Gen-Z Core across the *Gen-Z Core Specification* PLA interface. The PLA Interface is divided into three sections as illustrated in *PLA Interface Sections*:

- Control and Status Interface: Directs and manages the Gen-Z Physical Layer
- Transmit Interface: Manages the bandwidth of the outbound (transmit) PLA data
- Receive Interface: Receives and forwards the inbound (receive) PLA data

See the *Gen-Z Core Specification* for additional PLA interface details and timing diagrams.

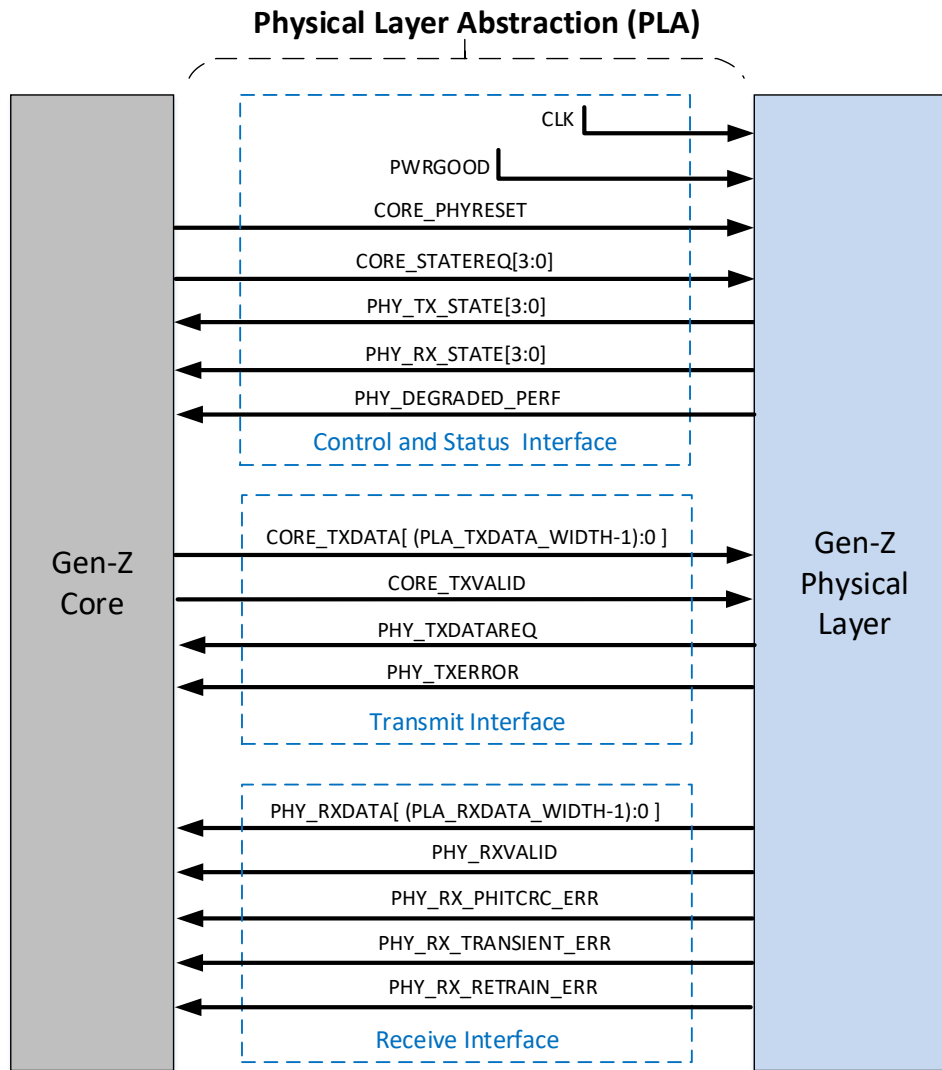


Figure 3-2: PLA Interface Sections

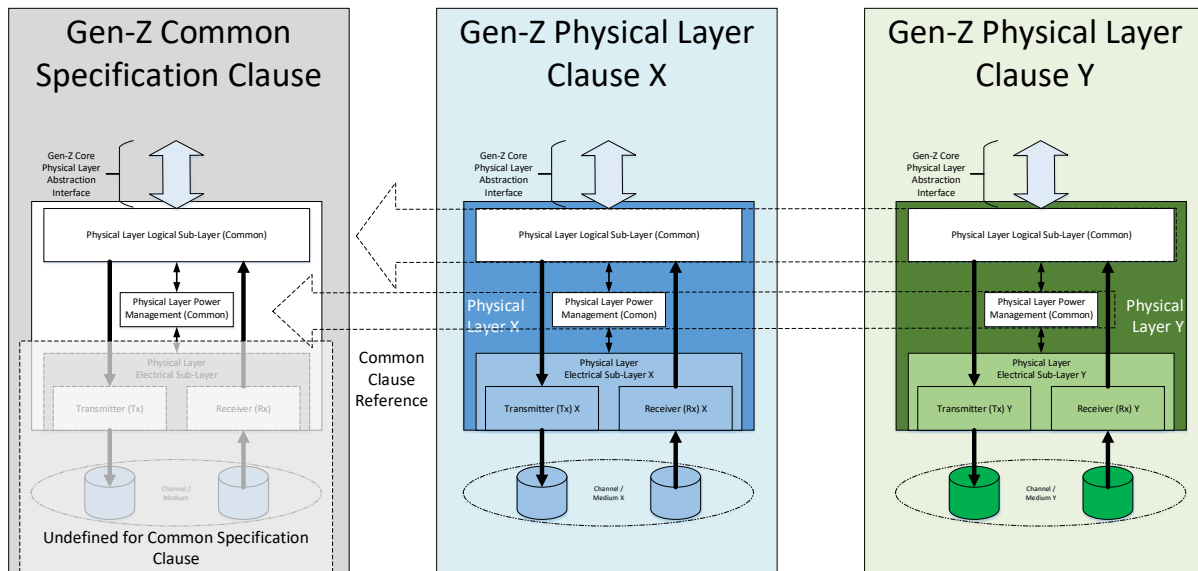
4. Gen-Z Common Specification

4.1. Introduction

5 The Gen-Z Common Specification clause specifies link and lane-level protocols that are independent of transmission media and signaling rates. The Gen-Z Common Specification clause specifies features such as initialization and training state machines, and power management that may be common across multiple physical layers. A physical layer clause references the Gen-Z Common Specification clause in cases where Logical Sublayer and Power Management features can be reused. The major differentiating characteristic for Gen-Z Physical Layer clauses is the Electrical Sublayer; hence, the Gen-Z Common Specification clause does not include any such content. *PHY Common Specification Clause Reference* illustrates how the Gen-Z Common Specification clause can be referenced from individual physical layer clauses.

10

Where appropriate, the Gen-Z Common Specification clause documents and clarifies the purpose for logical functions in their respective section headings.



15

Figure 4-1: PHY Common Specification Clause Reference

4.2. Logical Sublayer

The logical sublayer specifies the logical functions between the PLA and the electrical sublayer. The logical sublayer is responsible for link and lane transmit and receive operations for striping, scrambling, encoding, and serialization.

- 5 The *Transmit (Tx) Physical Layer Logical Sublayer Block* and *Receive (Rx) Physical Layer Logical Sublayer Block* figures illustrate an implementation that includes a 128-bit PLA interface, 64b/66b line code and a two-lane link for illustration purposes only. Each of the logical sublayer functions represented in these figures is specified in the subsequent sections.

10 The *Transmit (Tx) Physical Layer Logical Sublayer Block* figure illustrates the functions and data flow through the logical sublayer for the outbound (transmit) side.

The steps to transmit a data packet are as follows:

1. Transmit data is sent from the Gen-Z Core to the Tx Data register over the PLA Interface.
2. The data is striped as specified in *Data Striping*.
3. The data is scrambled and encoded.
- 15 4. The data is transferred to the serializer.
5. The data is transmitted serially by the lane drivers over the electrical or optical physical media.

The initialization, training, and control flow through the logical sublayer requirements for transmitting *Physical Layer (PHY) Sequences* are as follows:

1. PHY Sequences shall not bypass the Data Striping function.
- 20 2. PHY Sequences may be scrambled and may be controlled by the Enable signal to the LFSR scrambler. When disabled, the scrambler shall be bypassed and shall not advance.
3. The PHY Sequences may then be encoded, serialized and transmitted in the same manner as the data sets.

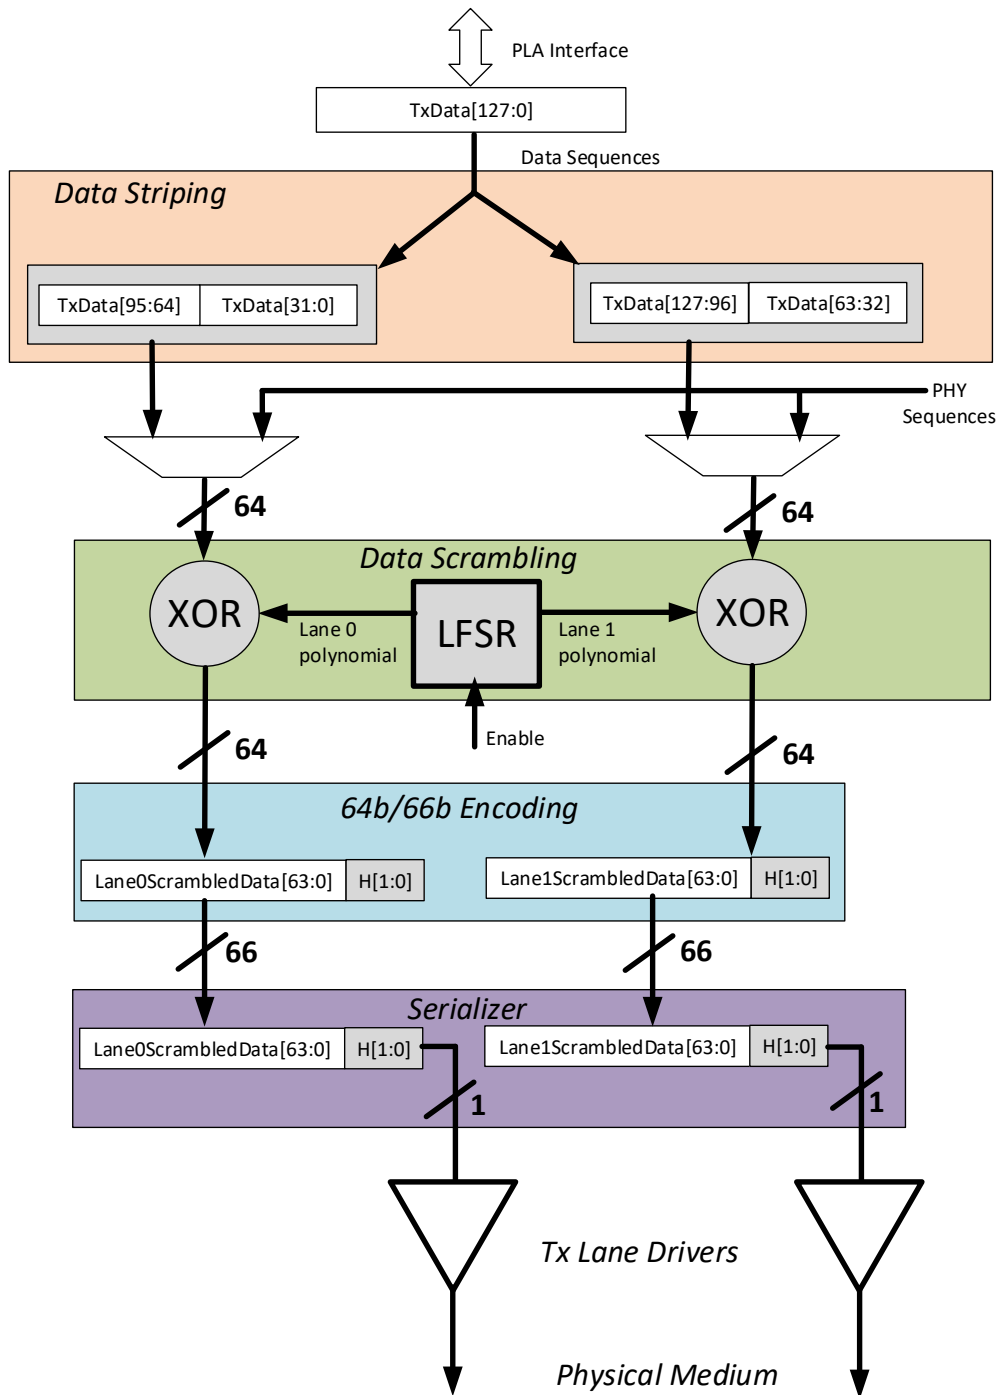


Figure 4-2: Transmit (Tx) Physical Layer Logical Sublayer Block (64b/66b)

The Receive (Rx) Physical Layer Logical Sublayer Block figure illustrates the functions and data flow through the logical sublayer for the inbound (receive) side, which is in the opposite order of the

5

outbound (transmit) path. The output is a 128-bit data word which is sent to the Gen-Z Core via the PLA interface.

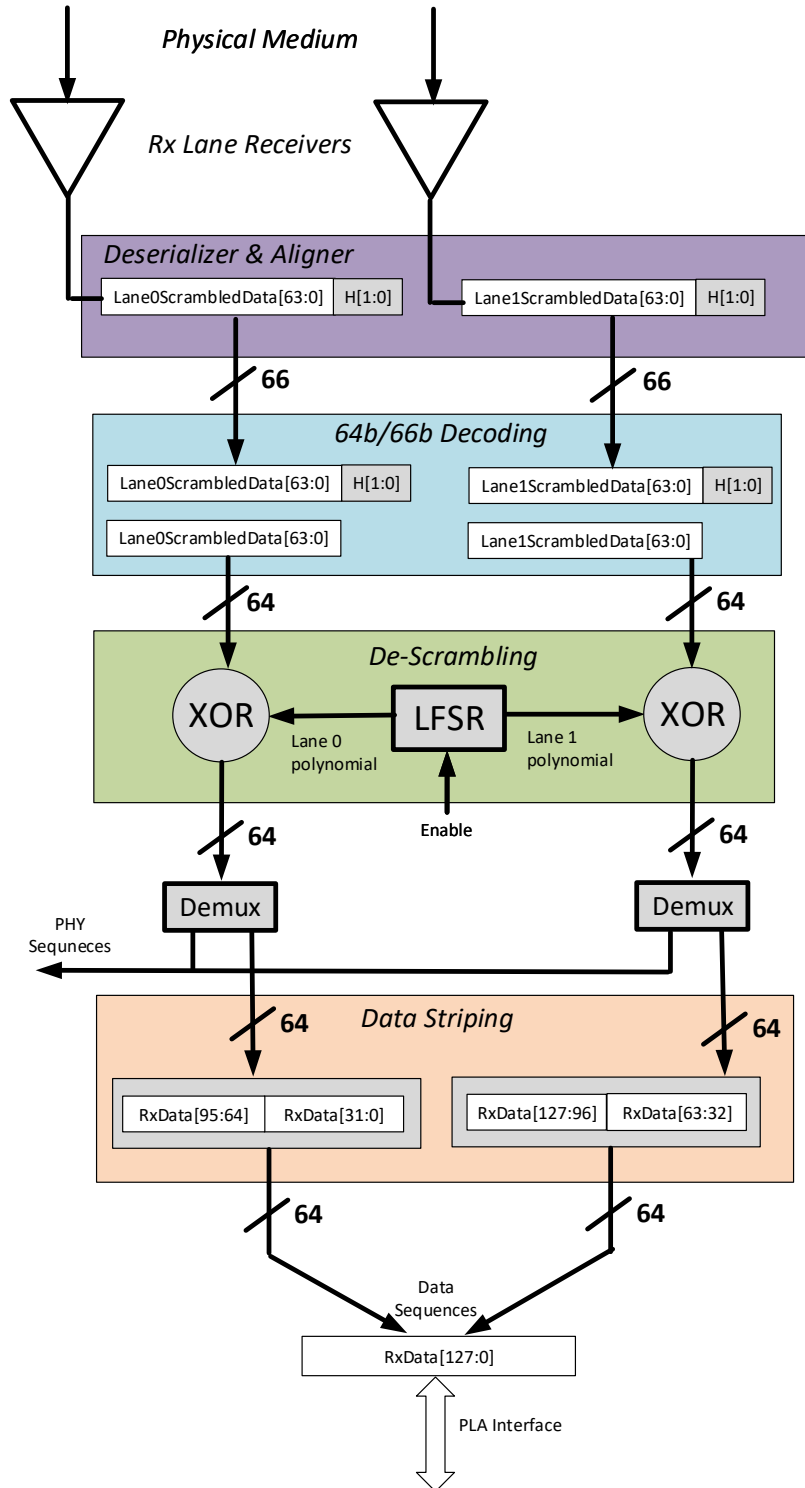


Figure 4-3: Receive (Rx) Physical Layer Logical Sublayer Block (64b/66b)

4.2.1. Link Serialization

The logical sublayer transmitter uses a serializer to convert a parallel data word into a serial bit stream, and, on the receive side, uses a deserializer to convert the serial bit stream into a parallel data word.

Serializer and Deserializer Functions illustrates a conceptual representation of the serializer and deserializer functions. On the transmit side of the link, the serializer starts with N-bit parallel data and serializes the data by shifting out one bit per UI for transmission over the link (a right shift operation). Unless explicitly stated otherwise by an individual physical layer clause, physical layer clauses referencing the Gen-Z Common Specification clause shall serialize the transmitted data with the least significant bit (LSB), bit 0, of the data first and the most significant bit (MSB), bit N-1, last. *Serializer and Deserializer Functions* illustrate this operation showing bit 0 being transmitted at UI time 1 and bit N-1 at UI time N.

On the receiver side, the deserializer samples one bit per UI time and shifts the received bits into an N-bit parallel data structure. The deserializer in *Serializer and Deserializer Functions* is also a right-shift operation. Unless explicitly stated otherwise by an independent clause, physical layer clauses referencing the Gen-Z Common Specification clause shall deserialize received data with the least significant bit first and shall fill its register until data bit N - 1. Upon receiving the N-bit data, the parallel data is sent to the decoder, descrambler, and striped before sent to the Gen-Z Core via the PLA interface.

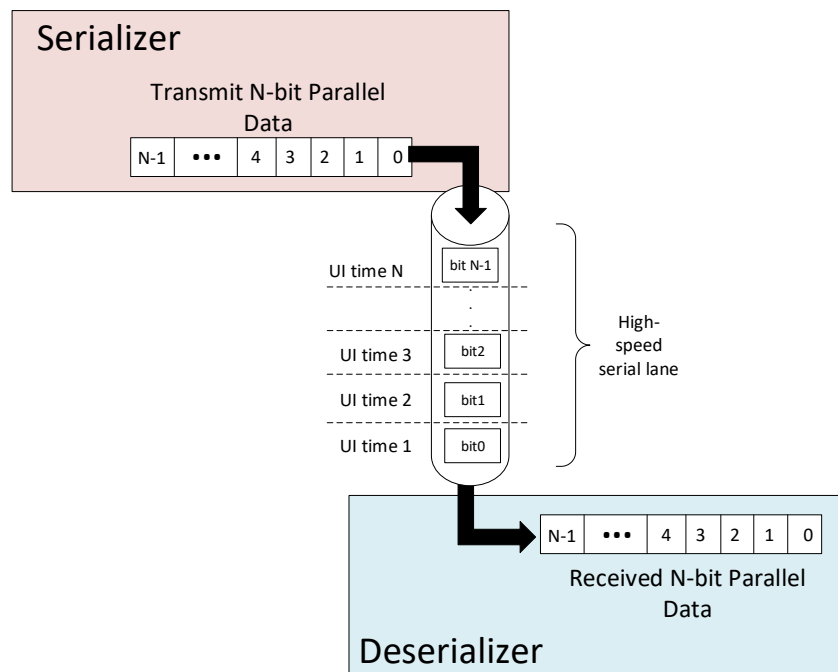


Figure 4-4: Serializer and Deserializer Functions

4.2.1.1. Encoding / Decoding

Support and details for line coding are specified in each individual physical layer clause. Line codes can assist a physical layer in providing DC balanced data patterns, reducing the largest length of contiguous data patterns, providing error detection/correction, or differentiating between physical layer

information and core data. Gen-Z Physical layers may implement either or a combination of an 8b/10b, 128b/130b, or 64b/66b encoding line code or Forward Error Correction (FEC) code.

4.2.1.2. Framing

Support and details for framing is specified in each individual physical layer clause.

5 4.2.2. Scrambling / Descrambling

Physical layer clauses referencing the Gen-Z Common Specification clause shall employ data scrambling to maintain data transition density for timing recovery and for energy dispersal. The scrambler shall utilize a linear feedback shift register (LFSR) implementing the PCIe polynomial $G(x) = X^{23} + X^{21} + X^{16} + X^8 + X^5 + X^2 + 1$ as documented in the *PCI Express Base Specification*.

10 Energy is dispersed across the link by assigning a unique initial LFSR seed to each lane. *LFSR Seed by Gen-Z Lane Number* specifies the LFSR seed by lane number modulo 8 ($N\%8$). Each transmitter shall use the LFSR seed corresponding to each lane number, e.g., transmitter lane 27 uses $N = 3$ ($27 \text{ MOD } 8 = 3$) with LFSR initialized to 0x18C0DB.

15 The same LFSR seed shall be used for scrambling and de-scrambling data. During the Signal Detect state of physical layer training, the receiver shall determine the lane number of the remote transmitter by detecting if the lanes of the link are connected in a reversed order (see *Lane Reversal*). The remote transmitter lane number determines the starting seed of the descrambler LFSR using *LFSR Seed by Gen-Z Lane Number*.

Table 4-1: LFSR Seed by Gen-Z Lane Number

Tx Lane Number Mod 8 ($N\%8$)	LFSR Seed
0	0x1DBFBC
1	0x0607BB
2	0x1EC760
3	0x18C0DB
4	0x010F12
5	0x19CFC9
6	0x0277CE
7	0x1BB807

20 4.2.3. Data Striping

Data striping specifies the how data from the Gen-Z Core (via the PLA) is divided into smaller byte-level words for transmission over the implemented lanes. Data striping improves error detection when used in conjunction with end-to-end cyclical redundancy checks (ECRC) and prelude cyclical redundancy checks (PCRC) to detect burst and lane failures (see ECRC and PCRC in the *Gen-Z Core Specification*).

Physical layer clauses referencing the Gen-Z Common Specification clause shall stripe data from the PLA. Physical layer-specific control and training packets (PHY Sequences) shall not be striped and shall be transmitted in parallel over all lanes in the link (see *Transmit (Tx) Physical Layer Logical Sublayer Block* which illustrates that training and control data packets bypass the striping block in the transmitter via a bypass multiplexer).

The Gen-Z Core specifies the striping of data packets as a function of the number of implemented lanes as specified in *Gen-Z Core Striping Specification versus Link Lane Count*.

Table 4-2: Gen-Z Core Striping Specification versus Link Lane Count

Number of Lanes in the Gen-Z Link	1	2	4	8	16	32	64	128	256
Data Striping Requirement [bytes]	4	4	1	1	1	1	1	1	1

The PLA data word in one and two-lane implementations is parceled into 4-byte phits to be transmitted by each lane. In a two-lane implementation phits are transmitted alternately between lane 0 and lane 1 as illustrated in *Data Striping with a 128-bit PLA interface and a 2-lane link*.

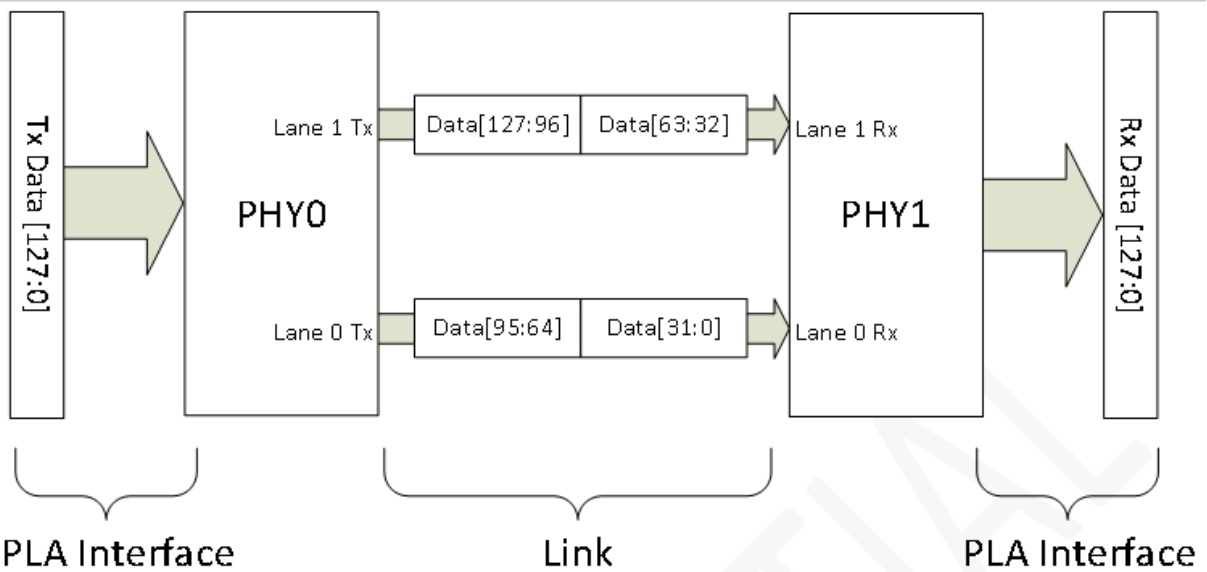


Figure 4-5: Data Striping with a 128-bit PLA interface and a 2-lane link

Physical layer implementations with more than two lanes shall use single byte striping as illustrated in *Data striping with a 64-bit PLA interface and a 4-lane Link*.

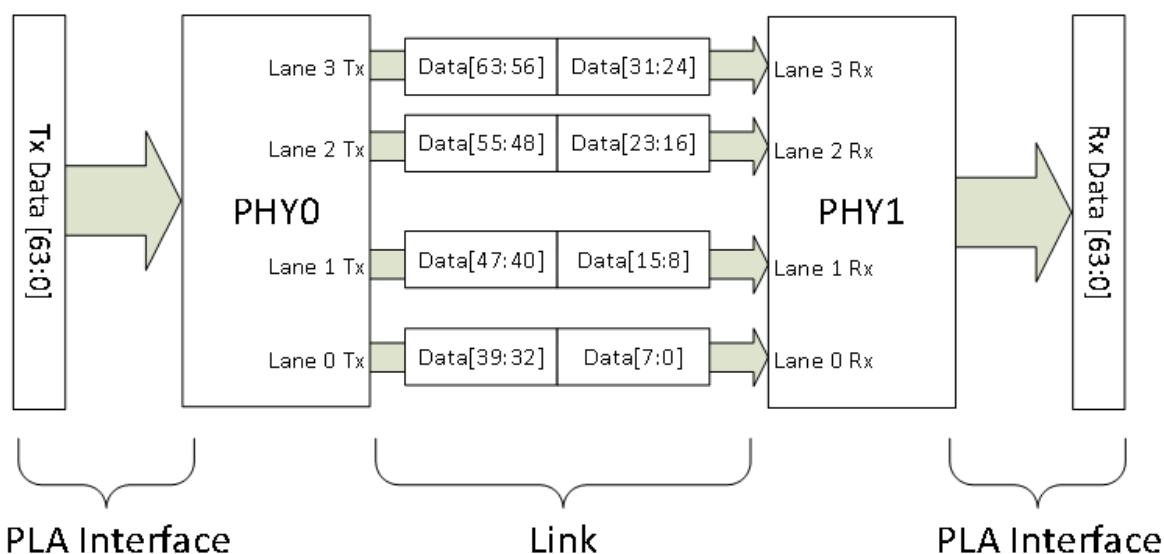


Figure 4-6: Data striping with a 64-bit PLA interface and a 4-lane Link

4.2.4. PLA Data Width

The PLA data width requirement is determined using a three-step process.

1. The data striping requirement is determined as specified in *Data Striping*.
2. The PLA data multiple is determined as a function of data striping and the number of lanes and may be as specified in *PLA Data Width Multiple (in bytes) vs Number of Lanes Used*. Note this table is not inclusive of all possible PLA data widths and is provided only as an example for how to apply the *Gen-Z Core Specification*.
3. The physical layer implementer chooses a PLA data width that may be a multiple of the number of bytes specified in *PLA Data Width (in bytes) vs Number of Lanes Used and Desired Multiplier*. Note this table is not inclusive of all possible PLA data widths and is provided only as an example for how to apply the *Gen-Z Core Specification*.

Table 4-3: PLA Data Width Multiple (in bytes) vs Number of Lanes Used

Number of Lanes in the Gen-Z Link	1	2	4	8	16	32	64	128	256
PLA Data width multiple [bytes]	4	8	4	8	16	32	64	128	256

- 15 *PLA Data Width Multiple (in bytes) vs Number of Lanes Used* provides examples of Gen-Z compliant PLA data widths. Any PLA Data multiplier specified by the *Gen-Z Core Specification* may be used as long as the final width is a multiple of the number of bytes specified in *PLA Data Width Multiple (in bytes) vs Number of Lanes Used*.

Table 4-4: PLA Data Width (in bytes) vs Number of Lanes Used and Desired Multiplier

Number of Lanes in the Gen-Z Link	1	2	4	8	16	32	64	128	256
PLA Data Multiplier = 1	4	8	4	8	16	32	64	128	256
PLA Data Multiplier = 2	8	16	8	16	32	64	128	256	512
PLA Data Multiplier = 3	12	24	12	24	48	96	192	384	768
PLA Data Multiplier = 4	16	32	16	32	64	128	256	512	1024
PLA Data Multiplier = 8	32	64	32	64	128	256	512	1024	2048

Each row of *PLA Data Width (in bytes) vs Number of Lanes Used and Desired Multiplier* is derived by multiplying the PLA Data Multiplier for each row by the byte multiple given in *PLA Data Width Multiple (in bytes) vs Number of Lanes Used*.

5 4.2.5. Physical Layer Initialization and Training

Physical layer initialization and training defines the process by which a link or the connection of two physical layers transitions from initial power-on or a system-directed reset to transmitting and receiving valid data at the specified data-rate and BER. Physical layer training enables a local port or physical layer to discover connectivity to a remote port, to establish a communication link, and to optimize channel configuration to provide robust data integrity.

Gen-Z physical layer implementations may support low power state transitions to optimize power consumption for changes in periodic data throughput requirements. Transitions between low power and normal operational states may require physical layer retraining as specified in this section.

4.2.5.1. PLA State Map

15 Support for PLA states and their mapping is specified in each individual physical layer clause.

4.2.5.2. Physical Layer Link versus Lane Training

Some stages of Gen-Z physical layer training require link-level knowledge of the data across all lanes and these are classified as link-level training stages, whereas those which can complete with only lane-level data are classified as lane-level training tasks. Some lane and link training functions are applicable only to a specific physical medium. Therefore, Gen-Z physical layer implementations shall adhere to the training requirements of their specific clause. A list of common link and lane-level tasks are enumerated below:

- Link-level training:
 - Lane reversal

- Lane polarity (for differential links only)
 - Link width and asymmetry
 - Link bit rate
 - Lane-to-lane de-skew
- 5 ▪ Lane-level training:
- Lane bit lock
 - Lane symbol lock
 - Lane channel optimization (equalization for electrical links only)

4.2.5.3. **Physical Layer (PHY) Sequences**

10 Physical Layer (PHY) Sequences are defined as data used specifically for communicating information between both physical layers across a link. PHY Sequences shall not be forwarded across the PLA to the Gen-Z Core.

15 Line coding provides a prescribed method to differentiate between Gen-Z Core data and Gen-Z physical layer data (or information). During physical layer training, only PHY Sequences shall be exchanged independent of line code presence. When line coding is present, physical layers may stall Gen-Z Core data streams to transmit and receive PHY Sequences as necessary, e.g., PHY-Up Sequences that enable a physical layer to compensate for clock tolerance differences or to drive autonomous low-power state transitions. For physical layers without line coding or any other overhead, only Gen-Z Core data sets shall be exchanged once a physical layer has trained to PHY-Up.

20 The following PHY Sequences enable a Gen-Z physical layer receiver to acquire symbol lock, detect lane polarity (in differential implementations), determine link width, discover lane reversal, and transfer physical layer control messages. The “Scrambled” column in Training Sequence tables indicates whether or not the Training Sequence shall be scrambled; the “Word<N>” columns are hexadecimal 32-bit values.

25 4.2.5.3.1. **Signal Detect Training Sequences**

The physical layer (PHY) Detect Header and Detect Sets shall be 128 bits composed of 32-bit words as specified as in *PHY Signal Detect Training Sequences (SDTS)*.

Table 4-5: PHY Signal Detect Training Sequences

Type	Scrambled	Word 0	Word 1	Word 2	Word 3
Header	No	0x0000FFFF	0x0000FFFF	0x0000FFFF	0x0000FFFF
Set	Yes	0x3C3C3C3C	0x0	0x0	0x0

30 Detect Headers and Sets shall be exchanged during the Signal Detect physical layer training state. To enable the remote physical layer port to detect the low frequency content being transmitted, Detect Headers shall not be scrambled. Unless explicitly stated otherwise by this specification, upon reception of a Detect Header each lane shall reset its scrambler. The number of Detect Headers transmitted and the frequency between Detect Sets shall be as specified in each individual physical layer clause.

Developer Note: Typically, a single Detect Header is transmitted followed by a series of Detect Sets. To provide flexibility in Detect Header identification in the Signal Detect state, the number of repeated Detect Headers transmitted can be configured through the Interface PHY Structure. Increasing the Detect Header repeat count can provide faster, more robust detection at the expense of decreased high-frequency transitions which benefits CDR.

4.2.5.3.2. Align Training Sequences

The physical layer (PHY) Align Header and Align Set shall be 128 bits composed of 32-bit words as specified in *PHY Align Training Sequences*.

Table 4-6: PHY Align Training Sequences

Type	Scrambled	Word 0	Word 1	Word 2	Word 3
Header	No	0x0000FFFF	0x0000FFFF	0x0000FFFF	0x0000FFFF
Set	Yes	0x1E1E1E1E	0x0	<Ack>	0x0

Align Headers and Sets shall be exchanged only in the Align physical layer training state. To enable the remote physical layer port to detect lane polarity and perform symbol alignment on a known pattern, Align Headers shall not be scrambled. Upon reception of every Align Header, each lane shall reset its scrambler. The Align Header frequency between Align Sets shall be as specified in each individual physical layer clause.

Align Sets of the Align Training Sequence shall be scrambled. Initially, Word 2 of the Align Set is set to not acknowledge, <Ack> = 0x00000000, by the transmitter. Once a receiver has successfully received 8 consecutive Align Sets, its associated local transmitter shall begin transmitting the acknowledge signal, <Ack> = 0xFFFFFFFF, to the remote receiver in Align Set Word 2.

4.2.5.3.3. IBIST Training Sequences

The physical layer (PHY) IBIST Header and IBIST Pattern shall be 128 bits composed of 32-bit words as specified in *PHY IBIST Training Sequence*.

Table 4-7: PHY IBIST Training Sequences

Type	Scrambled	Word 0	Word 1	Word 2	Word 3
Header	Yes	0xE5E5E5E5	0x0	0x0	0x0
Pattern	Yes	0x0	0x0	0x0	0x0

IBIST Training Sequences shall be exchanged in the IBIST (Interconnect BIST) physical layer training state. The IBIST Header shall be followed by a programmable number of IBIST Patterns, which are 128 bits of scrambled 0s. The receiver uses the scrambled 0s to observe recovered signal health using an implementation-specific algorithm.

4.2.5.3.4. Config Training Sequence

The physical layer (PHY) Config Training Sequence shall be 128 bits composed of 32-bit words as specified in *PHY Config Training Sequence*.

Table 4-8: PHY Config Training Sequence

Type	Scrambled	Word 0	Word 1	Word 2	Word 3
Config	Yes	0x2D2D2D2D	0x0	<width>	0x0

5 Config Training Sequences shall be exchanged in the Config state. The physical layer receive side is responsible for determining link configuration and for negotiating with its link partner which specific lanes are valid. The receiver shall account for link width and lane reversal to populate the <width> field. The <width> field shall be as follows:

- <width> = 8{~lwr_mode[1:0], lwr_mode[1:0]}
- 10 • lwr_mode[1:0] shall be configured as specified in *LWR Mode Definition*.

Table 4-9: LWR Mode Definition

Lanes Supported	lwr_mode[1:0]
All Lanes (full implementation of the link)	0x3
Upper half of the Transmitter Lanes	0x2
Lower half of the Transmitter Lanes	0x1
None of the link (NACK)	0x0

4-lane LWR Mode illustrates an example of lwr_mode[1:0] assignment based upon which lanes trained successfully in the physical layer Align state and if lane reversal was detected.

Table 4-10: 4-lane LWR Mode

Rx Lane 0	Rx Lane 1	Rx Lane 2	Rx Lane 3	Reversal	lwr_mode
Trained	Trained	Trained	Trained	Trained	0x3
FailedToTrain	Trained	Trained	FailedToTrain	FailedToTrain	0x0
Trained	Trained	FailedToTrain	Trained	Trained	0x1
Trained	FailedToTrain	Trained	Trained	Trained	0x2

15 4.2.5.3.5. Lock Training Sequence

The physical layer (PHY) Lock Training Sequence shall be 128 bits composed of 32-bit words as specified in *PHY Lock State Training Sequence*. When the local logical sublayer is ready to transition to the PHY-

Up state, it shall set the <width> field (Word 2) in the Config Training Sequence and shall then transmit a single Lock Training Sequence indicating the beginning of PHY-Up traffic.

Table 4-11: PHY Lock State Training Sequence

Type	Scramble	Word 0	Word 1	Word 2	Word 3
Lock	Yes	0xA4A4A4A4	0x0	0x0	0x0

4.2.5.3.6. PHY-Up and Data Sequences

- 5 When physical line-coding is enabled, *PHY-Up Sequences* and *Data Sequences* provide a method to differentiate between Gen-Z Core data and physical layer-specific information for control and configuration when in a PHY-Up state. Each individual physical clause should refer to their *Encoding / Decoding* section for further information.

Table 4-12: Data Sequences

Type	Encoding Header	Scramble	Striped	Word 0	Word 1	Word 2 (128b/130b)	Word 3 (128b/130b)
Data	10b	Yes	Yes	Data[31:0]	Data[63:32]	Data[95:64]	Data[127:96]

- 10 The Word 2 and Word 3 columns shall be used only in 128b/130b encoding. When using 128b/130b encoding Data Sequences shall contain Word 0-Word 3. When using 64b/66b encoding, Data Sequences shall contain only Word 0 and Word 1.

The following are the features and requirements associated with the transmission of Data Sequences:

- 15
- Data Sequences shall consist of Gen-Z Core data from the PLA, shall be striped, according to the data striping rules, and shall be scrambled.
 - Data Sequences shall use an Encoding Header as specified in *Data Sequences*.

The following are the features and requirements associated with the reception of Data Sequences:

- 20
- Data Sequences shall match the Encoding Header as specified in *Data Sequences*. Data Sequences shall be de-scrambled, shall be re-striped, according to the data striping rules, and shall be transmitted across the PLA as Gen-Z Core data.

Table 4-13: PHY-Up Sequences

Type	Encoding Header	Scramble	Striped	Word 0	Word 1	Word 2 (128b/130b)	Word 3 (128b/130b)
Idle	01b	Yes	No	0xAAAAAAAA A	0xAAAAAAAA A	0xAAAAAAAA A	0xAAAAAAAA A
LP	01b	Yes	No	0xBBBBBBBB	<LPID>	0xBBBBBBBB	<LPID>
PHY Down	01b	Yes	No	0xCCCCCCCC	0xCCCCCCCC	0xCCCCCCCC	0xCCCCCCCC

Type	Encoding Header	Scramble	Striped	Word 0	Word 1	Word 2 (128b/130b)	Word 3 (128b/130b)
Retrain	01b	Yes	No	0x99999999	<RetrainID>	0x99999999	<RetrainID>
Clk Comp	01b	No	No	0x66666666	<SetsRemain>	0x66666666	<SetsRemain>

The Word 2 and Word 3 columns shall be used only in 128b/130b encoding. When using 128b/130b encoding PHY-Up Sequences shall contain Word 0-Word 3. When using 64b/66b encoding, PHY-Up Sequences shall contain only Word 0 and Word 1. In 128b/130b encoded PHY-Up Sequences Word 0 and Word 2 shall contain the same control information and Word 1 and Word 3 shall contain the same control information.

The following are the features and requirements associated with the transmission of PHY-Up Sequences:

- PHY-Up Sequences shall consist of Gen-Z physical layer information, shall not be striped, and shall be scrambled as specified in *PHY-Up Sequences*.
- PHY-Up Sequences shall use an Encoding Header as specified in *PHY-Up Sequences*.
- PHY-Up Sequence of the same type shall be transmitted on all lanes within a links, and where applicable, any encoded fields shall match such as <LPID>, <RetrainID>, and <SetsRemain>.
- For additional resiliency with narrower link widths:
 - One-lane (x1) links shall transmit four consecutive PHY-Up Sequences.
 - Two-lane (x2) links shall transmit two consecutive PHY-Up Sequences.

The following are the features and requirements associated with the reception of PHY-Up Sequences:

- PHY-Up Sequences shall match the Encoding Header as specified in *PHY-Up Sequences*.
- PHY-Up Sequences shall be de-scrambled as specified in *PHY-Up Sequences*, shall be decoded as Gen-Z physical layer information, and shall not be transmitted across the PLA to the Gen-Z Core.
- PHY-Up Sequences shall be considered valid if all their types match, and where applicable, any encoded fields match such as <LPID>, <RetrainID>, and <SetsRemain> on more than half of the active lanes within a link.
- For additional resiliency with narrower link widths:
 - One-lane (x1) links shall validate that four consecutive PHY-Up Sequences match, inclusive of encoded fields.
 - Two-lane (x2) links shall validate that two consecutive PHY-Up Sequences match, inclusive of encoded fields, on both lanes.
- If an invalid PHY-Up Sequence is detected, then:
 - If Phit-CRC is enabled, then shall assert PHY_RX_RETRAIN_ERR and stop processing incoming traffic until physical layer retraining completes.
 - If Phit-CRC is not enabled, then shall assert PHY_RX_TRANSIENT_ERR and drop the PHY-Up Sequence or PHY-Up Sequences for one-lane (x1) and two-lane (x2) links.

Local and remote physical layers may request conflicting PHY-Up Sequence types that can affect the state of the physical layer link. Reconciliation of these conflicts shall have the following priority:

1. PHY Down – highest priority
2. Retrain
3. LP (Low Power) – lowest priority

The details for field encodings for specific PHY-Up Sequences can be found in the following sections:

- LP (Low Power) <LPID> = *Low Power Operation*
- Retrain <RetrainID> = *Physical Layer Retraining*
- Clk Comp <SetsRemain> = *Clock Compensation*

Developer Note: Each individual physical layer clause specifies if the LP PHY-Up Sequence infers retraining or explicitly requires an additional Retrain PHY-Up Sequence.

4.2.5.4. **Error Recovery**

Error recovery features enable Gen-Z physical layers to continue transmitting data in the presence of bit errors. Error recovery mechanisms may be disabled to assist in debug and analysis. Gen-Z physical layer error recovery features may include:

- Disable Physical Layer Auto-retraining—If enabled (0x0), then the PLTSM shall autonomously attempt to retrain. If disabled (0x1), then the PLTSM shall not autonomously retrain and shall halt in the physical layer training state that failed to meet the requirements of its exit condition. The Disable Physical Layer Auto-retraining sub-field exists in the PHY Control field of the Interface PHY Structure as specified in the *Gen-Z Core Specification*.
- Re-transmission of failed packets—The Gen-Z Core may detect and re-transmit failed data packets.
- Physical layer directed retrain—If the accumulated errors exceeds a programmable threshold in the Gen-Z Core, then the Gen-Z Core may force a physical layer retrain.

4.2.5.5. **Lane-to-Lane Deskew**

A physical layer specification enables different skew (or delay) between physical layer lanes to ease the burden of layout on package and board designers. During physical layer training, skew between lanes is discovered and compensation is made to ensure error-free data transfers. Specialized training data packets called Align Headers and Align Sets are utilized for this purpose. While in the Align state each lane shall continuously transmit *Align Training Sequences* consisting of an Align Header followed by a configurable number of Align Sets to enable the receiver to detect its current location between two Align Headers and de-skew to align to the nearest header.

Deskew Using Align Headers and Sets illustrates the concept of measuring the distance between two headers and performing deskew. Lane 0 and Lane 1 transmitters simultaneously and continuously transmit *Align Training Sequences*. The receiver measures its skew and performs its alignment by moving its sample point to the nearest Align Header. The sample point for Lane 0 receiver shall be before Align Set N/2 towards the first Align Header. Similarly, Lane 1 shall detect that it is sampling after Align Set number N/2 and de-skew by moving its sample point later in time towards the later Align Header.

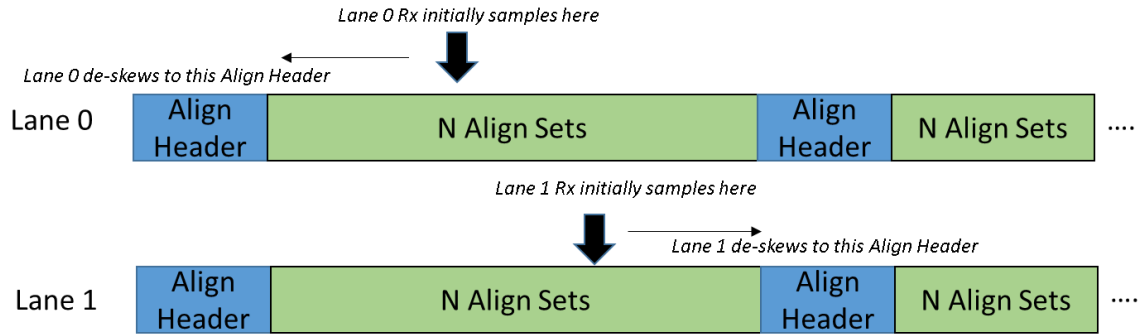


Figure 4-7: Deskew Using Align Headers and Sets

With this methodology, Gen-Z physical layers detect and de-skew lanes with an amount of skew proportional to the number of Align Sets transmitted in this training stage. *Equation for Maximum Tolerated Skew between Gen-Z Lanes* provides the calculation of the amount of lane skew which can be corrected to the data-rate and the number of Align Sets transmitted between Align Headers during training.

$$\text{Maximum Lane Skew [seconds]} < \text{Number of Align Sets} \times \frac{128 \text{ UI}}{\text{Align Set}} \times \frac{T \text{ seconds}}{\text{UI}} \times 0.5$$

Figure 4-8: Equation for Maximum Tolerated Skew between Gen-Z Lanes

In the equation above, the factor 0.5 indicates that the amount of lane skew shall be less than half of the deskew buffer depth so that there is no aliasing between adjacent Align Sets, and that each lane deskew engine can deskew towards the nearest Align Header as illustrated in *Deskew Using Align Headers and Sets*.

Physical layers referencing the Gen-Z Common Specification clause may support 8, 16, 32, or 64 Align Sets to be transmitted between Align Headers. This provides ample flexibility for skew budgeting and potential data-rate changes possible with individual physical layer clauses. There is a tradeoff between the max lane-to-lane skew and on-chip de-skew buffer size, and as such, each individual physical layer clause specifies the amount of skew between lanes and the number of Align Sets between Align Headers.

4.2.6. Lane Reversal

In a standard lane connection between two physical layers, Lane 0 of physical layer A is connected to Lane 0 of physical layer B. Similarly, Lane N of physical layer A is connected to Lane N of physical layer B, where N is the highest numbered lane in a link.

There are specific board placement circumstances where requiring the standard connection between physical layers can consume more board routing resource and increase cost.

Lane reversal provides increased board design flexibility to efficiently route between physical layers. In a lane-reversed link the lane number ordering is reversed between two connected interfaces. This enables a board design to use direct point-to-point routes for the lanes even if the physical layer package pinouts are reversed. *Standard versus Lane-Reversed Connection of Gen-Z Components (a)* illustrates a standard connection, and *Standard versus Lane-Reversed Connection of Gen-Z Components (b)* illustrates a reversed connection between physical layers with transmitter Lane 0 connected to Lane N of the receiver and Lane 1 of the transmitter connected to Lane N-1 of the receiver.

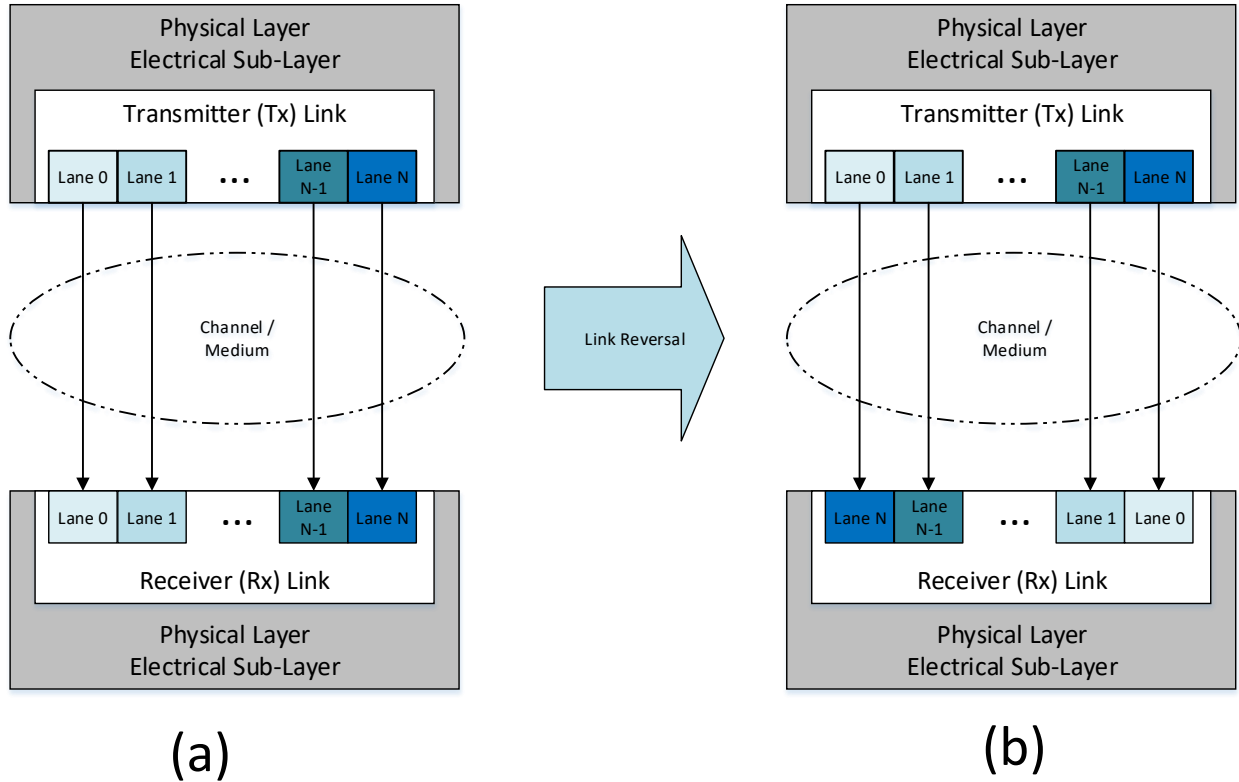


Figure 4-9: Standard versus Lane-Reversed Connection of Gen-Z Components

Lane reversal may occur at multiple points along a connection between physical layers, such as between package pins and C4 bumps on the silicon die or in the printed circuit board route between components. Lane reversal may be uniform (a) or non-uniform (b), independent per direction, as highlighted in *Lane-Reversed Uniformity Examples*. Gen-Z physical layer training enables link partners to discover lane reversal and adapt automatically to ensure correct data transfer.

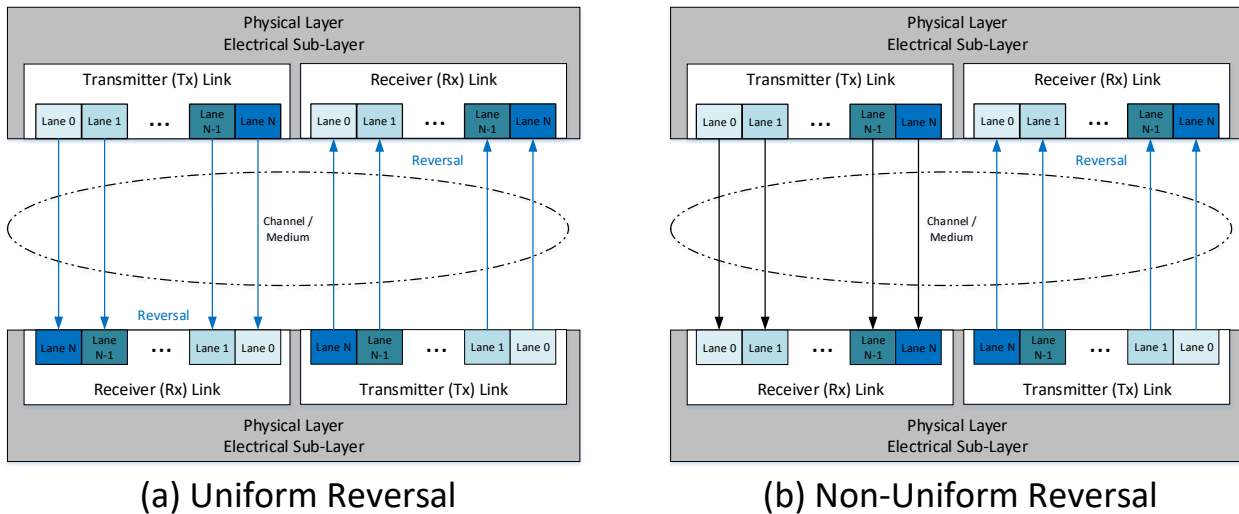


Figure 4-10: Lane-Reversed Uniformity Examples

All lane reversed connections within a link shall adhere strictly to sequential ordering, such as DeviceA_TX_Lane[X:0] → DeviceB_RX_Lane[0:X]. Lane remapping shall be contiguous.

Support for lane reversal including rules and policies for lane asymmetry and independent (non-uniform) or uniform directionality is specified in each individual physical layer clause.

5 **4.2.7. Lane Polarity**

Support for lane polarity including rules and policies is specified in each individual physical layer clause. Gen-Z physical layers may not require differential connectivity per lane. However, Gen-Z physical layers referencing the Gen-Z Common Specification should provide a mechanism to handle lane polarity inversion.

10 There may be cases where it is inefficient or costly to connect two devices' physical layers in a manner that maintains differential lane polarity across the link as illustrated in *Lane Differential Polarity Reversal*. *Lane Differential Polarity Reversal (a)* illustrates the standard case where the polarity of two Gen-Z physical layers are identical and a direct point-to-point board route can be used. In this scenario, connecting the positive pin of the transmitter lane to the positive pin of the receiver and connecting the
15 negative pin of the transmitter lane to the negative pin of the receiver is simple and efficient.

Lane Differential Polarity Reversal (b) illustrates the case where the lane polarity of two different physical layers is opposite. This case requires the board design to swizzle the lane route in order to maintain the differential polarity between the two physical layers. Maintaining the lane polarity in (b) can require more complicated board design and can consume more routing resource.

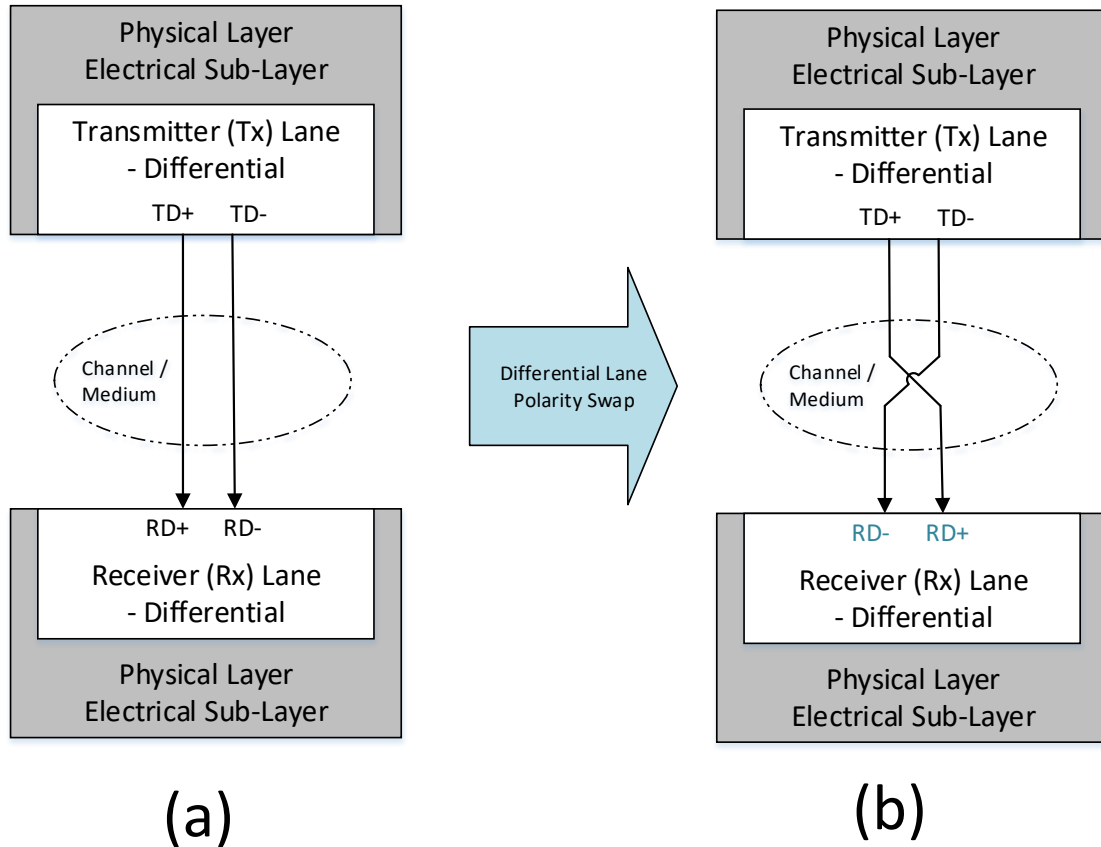


Figure 4-11: Lane Differential Polarity Reversal

5 To enable efficient resource use, physical layers supporting the Gen-Z Common Specification should provide support for lanes to connect with reversed polarity as illustrated in *Gen-Z Lane Polarity Inversion for Efficient Routing*. Notice that the receiver lane polarity is opposite of the transmitter's and straight-line routes are used to connect the Gen-Z components.

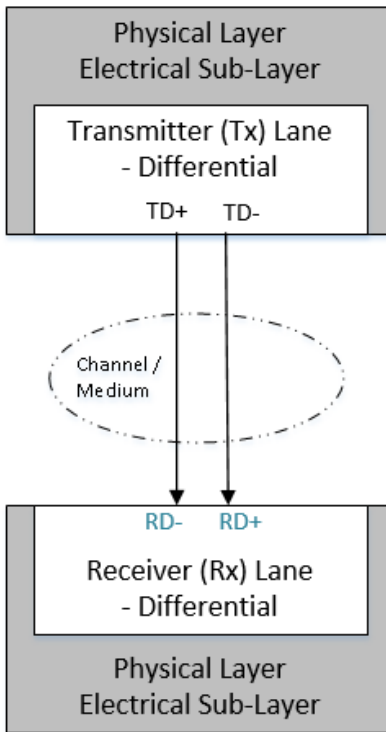


Figure 4-12: Gen-Z Lane Polarity Inversion for Efficient Routing

The Gen-Z PLTSM can detect lane polarity inversion as illustrated in *Gen-Z Lane Polarity Inversion for Efficient Routing* by observing the inbound data in the Align physical layer training state. Once polarity is detected, the physical layer should correct for lane differential inversion by inverting the received data in its receiver.

4.2.8. Link Width

The link width of a physical layer is defined as the number of physical lanes in its implementation. Physical layers referencing the Gen-Z Common Specification may support link widths of 1, 2, 4, 8, 16, 32, 64, 128 and 256 lanes.

Two independent Gen-Z devices may implement the same physical layer with different link widths. The PLTSM specified in the Gen-Z Common Specification specifies a mechanism to negotiate to a common link width.

Developer Note: Physical layers that support PHY-Up low power states can reduce link width to provide power savings at the expense of performance.

4.2.8.1. Asymmetric Links

Asymmetric link support including rules and policies is specified in each individual physical layer clause.

An asymmetric link is a physical layer with a different number of transmit versus receive lanes. Asymmetric links enable physical layers to optimize performance, power, and system resources based on the required data bandwidth per direction across a link.

5 Gen-Z physical layers referencing the Gen-Z Common Specification may support asymmetric transmit and receive widths and ratios as specified in *Gen-Z Supported Asymmetric Links*. The nomenclature of the table entries is: TxM:RxN, where M is the Tx lane count and N is the Rx lane count for a given link width and asymmetry ratio. An asymmetric Tx3:Rx1 to Tx1:Rx3 example is illustrated in (a) *Asymmetric PHY Examples #1*. Depending upon mechanical form factors and system implementations, a non-fully utilized symmetric physical layer can be connected to an asymmetric physical layer as illustrated in (b) *Asymmetric PHY Examples #1*.
 10 Alternative examples for (a) Tx7:Rx1 to Tx1:Rx7 and (b) Tx6:Rx2 to Tx2:Rx6 are illustrated in *Asymmetric PHY Examples #2*.

Table 4-14: Gen-Z Supported Asymmetric Links

Link Width	Tx:Rx Ratio 1:3, 3:1	Tx:Rx Ratio 1:7, 7:1
2	N/A	N/A
4	Tx3:Rx1 or Tx1:Rx3	N/A
8	Tx2:Rx6 or Tx6:Rx2	Tx1:Rx7 or Tx7:Rx1
16	Tx4:Rx12 or Tx12:Rx4	Tx2:Rx14 or Tx14:Rx2
32	Tx8:Rx24 or Tx24:Rx8	Tx4:Rx28 or Tx28:Rx4
64	Tx16:Rx48 or Tx48:Rx16	Tx8:Rx56 or Tx56:Rx8
128	Tx32:Rx96 or Tx96:Rx32	Tx16:Rx112 or Tx112:Rx16
256	Tx64:Rx192 or Tx192:Rx64	Tx32:Rx224 or Tx224:Rx32

15 **Developer Note:** With asymmetric links, the Channel Optimize physical layer training state can take longer depending upon the asymmetry ratio. Due to the difference in Tx and Rx lane counts, independent lanes might be unable to exchange Tx or Rx direction-specific settings when a physical layer uses a symmetric-based Channel Optimize function. Other aspects of the Gen-Z physical layer specifications, such as striping and CRC can be more complex.

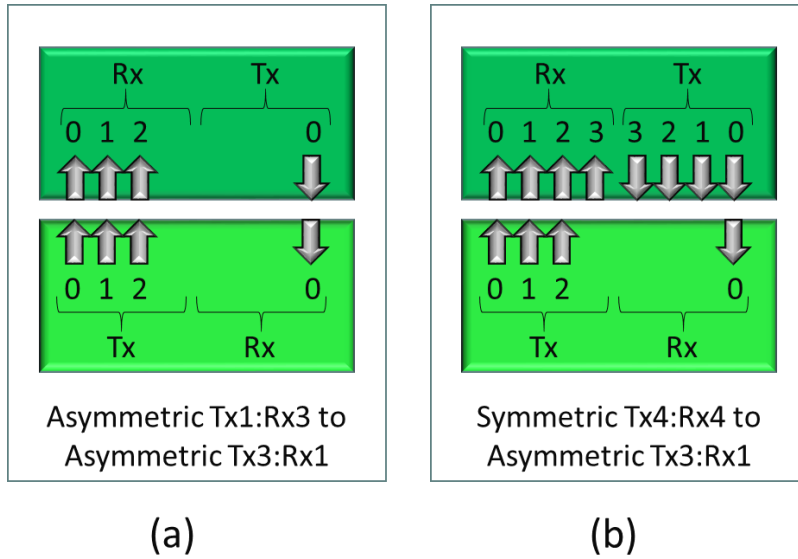


Figure 4-13: Asymmetric PHY Examples #1

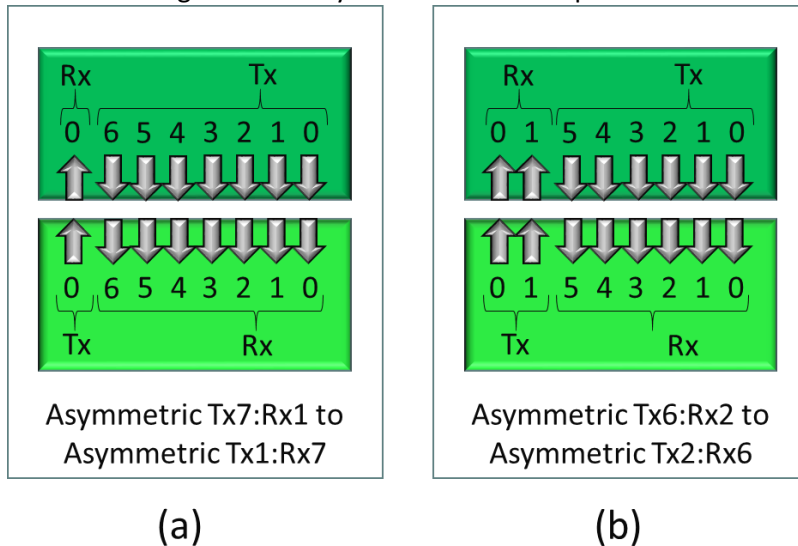


Figure 4-14: Asymmetric PHY Examples #2

Power savings via Link Width Reduction (LWR) using asymmetric links may be supported as specified in *Gen-Z Supported Asymmetric Link LWR Mode*. In addition to these configurations, asymmetric links of any ratio may reduce to a single lane link (Tx1:Rx1).

Table 4-15: Gen-Z Supported Asymmetric Link LWR Mode

Link Width	Tx:Rx Ratio 1:3, 3:1 LWR	Tx:Rx Ratio 1:7, 7:1 LWR
2	N/A	N/A
4	Tx1:Rx1	N/A
8	Tx1:Rx3 or Tx3:Rx1	Tx1:Rx3 or Tx3:Rx1

Link Width	Tx:Rx Ratio 1:3, 3:1 LWR	Tx:Rx Ratio 1:7, 7:1 LWR
16	Tx2:Rx6 or Tx6:Rx2	Tx1:Rx7 or Tx7:Rx1
32	Tx4:Rx12 or Tx12:Rx4	Tx2:Rx14 or Tx14:Rx2
64	Tx8:Rx24 or Tx24:Rx8	Tx4:Rx28 or Tx28:Rx4
128	Tx16:Rx48 or Tx48:Rx16	Tx8:Rx56 or Tx56:Rx8
256	Tx32:Rx96 or Tx96:Rx32	Tx16:Rx112 or Tx112:Rx16

4.2.9. Link BIST

Link built-in self-test (BIST) enables evaluation of the signal quality of a physical layer's lane. Link BIST may provide programmable patterns and overrides which can aid in debugging and lane integrity margining. The *Interconnect BIST (Optional)* state from the PLTSM provides a mechanism for running Link BIST.

Link BIST support including configuration for patterns and run time shall be as specified in each individual physical layer clause.

4.2.10. Physical Layer Retraining

The Gen-Z physical layer logical sublayer shall support retraining. Retraining may be directed by software, firmware or any physical layer policies necessary to improve the signal quality or enable physical layer reconfiguration. Retraining provides the link an autonomous mechanism to recover from data integrity problems. Retraining arcs enable a physical layer to re-optimize its configuration for a specific channel, such as receiver re-equalization for electrical links or quick reconfiguration of a physical layer for a low power transition.

Four possible arc types of physical layer retraining may be supported by a physical layer as specified the *Common, Lane, and Low Power Interface PHY Structure*. Physical layer retraining may be initiated by:

- A retrain-request from the Gen-Z Core across the PLA via a state transition to PHY-Down-Retrain.
- Reception of a Retrain PHY-Up Packet from the remote physical layer.
- Internal physical layer control to enable advanced features, such as PHY-Up LP (low power) states.
- An explicit control request for physical layer retraining from the *Common, Lane, and Low Power Interface PHY Structure*.

The following are the features and requirements associated with retraining for all physical layers:

- All physical layers shall support Retrain Arc 1. Retrain Arc 1 is the transition from the *PHY-Up Process* sub-state to the *Align entry* sub-state of the PLTSM.
- All Gen-Z Core retrain-requests from the PLA via a state transition to PHY-Down-Retrain shall result in a Retrain Arc 1.
- Retrain Arc 2, Arc 3, and Arc 4 may be supported by physical layers that require line-coding (enables PHY-Up Sequences). Support and details for Arc 2, Arc 3, and Arc 4 shall be as specified in each individual physical layer clause.

- Retrain arc numbers shall be prioritized with Retrain Arc 1 the highest and Retrain Arc 4 the lowest.
 - Retrain arc numbers may not imply any functional ordering in the PLTSM.

The following are the features and requirements associated with retraining for physical layers that support line-coding and PHY-Up Sequences:

- Retrain packets and their <RetrainID> as described in *PHY-Up Sequences* shall be as specified:
 - ArcID[3:0] = (Retrain Arc # - 1)
 - n[15:0] = {ArcID[3:0], ArcID[3:0], ArcID[3:0], ArcID[3:0]}
 - <RetrainID> = {~n[15], n[15], ~n[14], n[14], ... ~n[0], n[0]}
 - ~ = inverted bit value
- When a physical layer is directed to retrain, the logical sublayer shall transmit retraining PHY-Up Sequences with the ArcID[3:0] set to the appropriate request.
- A physical layer shall continue to transmit retraining PHY-Up Sequences until logical sublayer has transmitted a minimum of 4 and received a minimum of 2 with matching ArcID[3:0].
 - If only one physical layer retraining arc is supported, then retraining arc reception and acknowledgement via a retraining PHY-Up Sequence shall be optional.
- If more than one physical layer retraining arc is supported and there is a conflict in arc number requests, then the highest priority ArcID[3:0] (lowest number) request shall take precedence.

The following are the features and requirements associated with retraining for physical layers without encoding to support PHY-Up Sequences:

- The Gen-Z Core shall initiate and coordinate all retraining events via a Link CTL packet.
- Retrain Arc 1 shall be the only transition supported.

A physical layer that is link-width reduced (LWR) and not in a low-power mode shall enable all of its transmitters during physical layer retraining. Receivers may attempt to retrain to full width or stay in a reduced-width configuration.

Developer Note: For physical layers where retraining events do not correct transient errors, a transition back to the Channel Optimize state may be required. Entry back to the Channel Optimize state can be achieved by disabling and re-enabling the physical layer with one of the following two methods:

- A directed transition from the Gen-Z PLA to PHY Idle and back to PHY-Up.
- Through direct control of the physical layer via the Interface PHY Structure.

4.2.11. PLTSM

The following physical layer training states shall be required by all physical layers referencing the Gen-Z Common Specification, specified by solid lines in *PLTSM*.

- PHY Idle
- Signal Detect
- Align
- Config (Configuration)
- PHY-Up

Physical layers referencing the Gen-Z Common Specification may require other physical layer training states and in such a case, are specified in their individual physical layer clause. For example, the Channel Optimize state, where transmit equalization can be performed on electrical links, may be required in its

individual physical layer clause. The physical layer training states include Channel Optimize, IBIST, and PHY Low Power as illustrated with a blue dashed line in *PLTSM*.

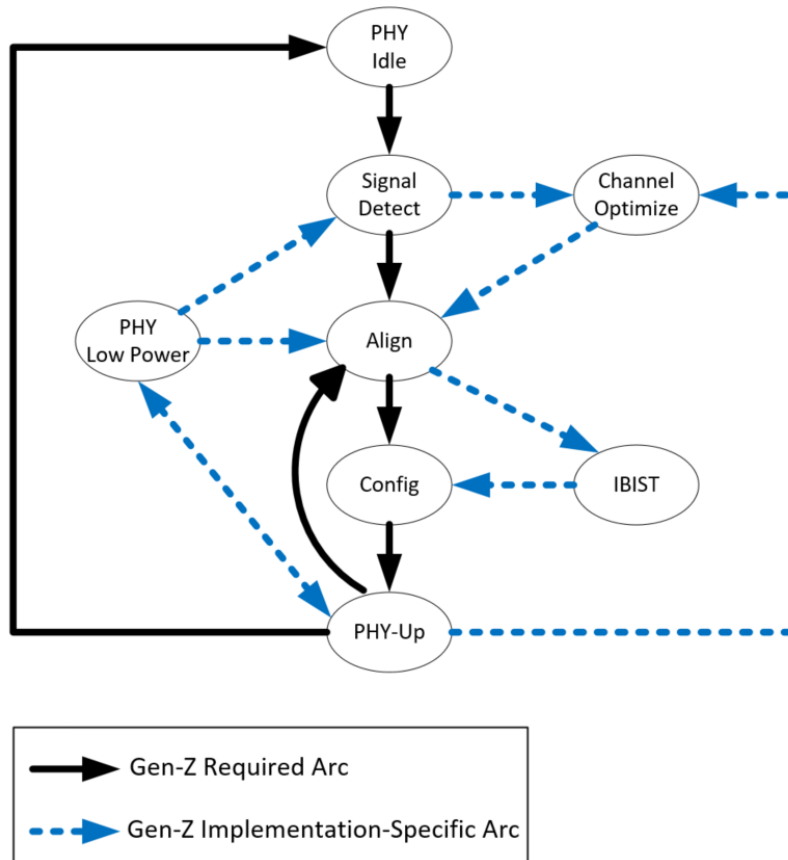


Figure 4-15: PLTSM

A detailed description of each PLTSM state and the conditions of the associated arcs is provided below.

4.2.11.1. **PHY Idle (Required)**

A physical layer shall remain in the PHY Idle state until the Gen-Z Core requests the physical layer to change its state across the PLA.

- 10 The physical layer shall enter the PHY Idle state on initial power on or when the Gen-Z Core directs the PLTSM across the PLA.

Developer Note: The PHY Idle physical layer training state is intended to be the lowest power physical layer state when the majority of the physical layer circuits are disabled.

4.2.11.2. **Signal Detect (Required)**

- 15 Signal detection refers to the mechanism by which a physical layer indicates it is available to establish a connection to a remote physical layer. Signal detection occurs in the Signal Detect physical layer

training state. Individual physical layer clauses can support electrical or optical physical layers, which may have different mechanisms for detecting the remote physical layer. As such, the features and requirements associated with the Signal Detect state shall be as specified in each individual physical layer clause.

5 Upon successful completion of Signal Detect, the PLTSM may progress to the Align state.

4.2.11.3. **Channel Optimize**

The features and requirements associated with the Channel Optimize state are specified in each individual physical layer clause.

4.2.11.4. **Align (Required)**

10 4.2.11.4.1. **Align entry**

The Align state shall be reached in one of two arcs:

- From the Signal Detect state in physical layer training.
- From the PHY-Up state due to a retrain.

15 Retraining from PHY-Up may use Fast Training Sequences (FTS) to achieve quick realignment for directed low power transitions or non-severe error recovery.

4.2.11.4.2. **Align process**

20 Upon entry to the Align state, each physical layer's lane transmitter shall continuously transmit *Align Training Sequences*. To compensate for lane skew from route and repeater differences, the physical layer shall transmit 8 to 64, inclusive, Align Sets after the Align Header. The amount of skew tolerated between lanes is a function of the number of Align Sets transmitted and is specified in each individual physical layer clause (see *Equation for Maximum Tolerated Skew between Gen-Z Lanes*).

System software shall configure both link partners to utilize the same Align Set count (4x to 64x) prior to physical layer training. To enable each individual lane to align to the Align Header and acquire symbol lock while in the Align state, Align Training Sequences shall not be striped across lanes.

25 A physical layer's transmitter shall initially set the Align Set Word 2:<Ack> field to 0x00000000. After the logical sublayer has detected/received a minimum of eight consecutive Align Headers from Align Training Sequence, the transmitter shall begin transmitting 0xFFFFFFFF in the Word 2:<Ack> field. The transmitter shall not scramble Align Headers and shall scramble Align Sets.

30 Upon entry into the Align state, the physical layer's receiver shall enable all implementation circuits and compensation mechanisms necessary to acquire bit lock. If a previous PLTSM state has not detected polarity inversion already, then the receiver PHY logical sublayer shall monitor the incoming data stream to look for the Align Header and the 1's complement (inversion) of the Align Header. If the inverted Align Header is found, then the receiver shall detect polarity inversion of the differential lane and shall then invert all subsequent received data. Once the Align Header is detected and polarity inversion is accounted for, the receiver shall begin descrambling the incoming Align Sets using the LFSR Seed table (see *LFSR Seed by Gen-Z Lane Number*) and its lane number based on the lane reversal status.

35 After receiving two consecutive Align Headers the receiver logical layer may begin deskewing the active lanes, detecting lane reversal (if applicable), and determining link width.

Each individual physical layer clause may choose to identify lane reversal using alternative methods prior to the Align state. If such methods are not defined, then lane reversal can be determined in the Align process state, in which case the requirements shall be specified in that individual physical layer clause.

Developer Note: For physical layer clauses that expect lane reversal to be determined in the Align process sub-state, each lane can first try its non-reversed LFSR Seed by lane number. If after an evaluation period, consecutive Align Sets are not appropriately detected, each lane can then try its reversed LFSR Seed by lane number. Reconciliation of reversal, such as the consecutive number detected or lane reversal majority detection, is outside the scope of this specification.

The number of lanes supported by a link, i.e., link width, is determined by the receive side of the physical layer. Unless explicitly stated otherwise by this specification, physical layers supporting the Gen-Z Common Specification clause shall run at full width (all lanes) or half width. Half-width links shall include only all contiguous lanes in the lower half or only all contiguous lanes in the upper half of the link.

Physical layer receivers shall monitor all lanes to detect which lanes train successfully. *4-Lane LWR Mode Examples Based on Lane Training* provides four examples of link width determination of a 4-lane link. In each example, each individual Rx lane has successfully Trained or FailedToTrain. In each example, the receiver observes the training results for the four lanes and determines link width reduction (LWR), if necessary, given in the 'Link Width' column. In Example 1, all four lanes have successfully trained and a full link can be utilized. In Example 2, one lane in each half of the link has trained successfully, while the other lane failed. Since a half-width link includes only all upper or only all lower lanes in a link, there is no supported Link Width, and the link fails to train. In Example 3, Rx Lane 0 and Rx Lane 1 both successfully trained, whereas Rx Lane 2 failed, and the result is a half-width link composed of the lower lanes. Example 4 illustrates a similar situation and results in a half-width link composed of the upper lanes.

Table 4-16: 4-Lane LWR Mode Examples Based on Lane Training

Example	Rx Lane 0	Rx Lane 1	Rx Lane 2	Rx Lane 3	Link Width
1	Trained	Trained	Trained	Trained	Full
2	FailedToTrain	Trained	Trained	FailedToTrain	Failed
3	Trained	Trained	FailedToTrain	Trained	Half-width: Lower
4	Trained	FailedToTrain	Trained	Trained	Half-width: Upper

The goal of the Align state is to achieve bit lock, symbol lock, lane deskew, and lane polarity detection. *Link Progression through Align Physical Layer Training State* illustrates a flow by which implementations should sequence through the Align state; a physical layer clause may implement lane reversal in the Signal Detect state.




Figure 4-16: Link Progression through Align Physical Layer Training State

4.2.11.4.3. Align Exit

The PLTSM exits the Align state as follows:

- 5 • The logical sublayer attempts to align all lanes (full width) to the Align Header until the Align timer expires. Receiver Lanes which are unaligned at the timer expiration shall be disabled by the logical sublayer so they may be mapped out in the Config state when determining the link width.
- 10 • Lanes which have received minimum of 4 Align Sets with the <Ack> field = 0xFFFFFFFF and transmitted a minimum 8 Align Sets with <Ack> = 0xFFFFFFFF when the Align timer expires are considered aligned and shall be included as logical sublayer attempts to create a full or half-width link.
- The logical sublayer shall then determine the largest supported inbound link width, full-width, half-width with lower lanes, or half-width with upper lanes.
- 15 • If the Align timer expires twice and there are no aligned lanes, then the PLTSM shall set the Physical Layer Training Status sub-field to 0x2 (Training has failed) in the PHY Status field of the Interface PHY Structure as specified in the *Gen-Z Core Specification*.
 - If physical layer training failed and the Disable Physical Layer Auto-retraining sub-field is set to 0x1 in the PHY Control field of the Interface PHY Structure as specified in the *Gen-Z Core Specification*, then the PLTSM shall remain in the Align State to aid in debug.
 - If physical layer training failed and the Disable Physical Layer Auto-retraining sub-field is set to 0x0 in the PHY Control field of the Interface PHY Structure as specified in the *Gen-*
- 20

Z Core Specification, then the link shall return to the IDLE state and shall autonomously attempt to retrain until training is successful or power is removed

After successful Align completion, the PLTSM shall advance to the Configuration (Config) State unless IBIST state transition is enabled. Details for enabling the IBIST state shall be as specified in each individual physical layer clause.

4.2.11.5. **Interconnect BIST (Optional)**

Interconnect Built-In Self-Test (IBIST) provides a method to quickly evaluate a lane's quality during physical layer training and to aid in debugging and margining.

4.2.11.5.1. **IBIST Entry**

When enabled by both the local and remote physical layers, the IBIST state shall be entered from the Align state.

4.2.11.5.2. **IBIST Process**

Upon entry of the IBIST state, each individual lane transmitter shall transmit a single, scrambled 128 UI IBIST Header from an IBIST Training Sequence, as specified in *PHY IBIST Training Sequences* followed by a series of scrambled IBIST Patterns. Upon reception of the IBIST Header, an aligned receiving lane shall evaluate the scrambled IBIST Patterns for errors. The number of patterns evaluated can be configured as specified in each individual physical layer clause.

Entry into IBIST shall be governed by explicit matching on the IBIST Header. If an IBIST Header error occurs that prevents IBIST state entry, then the following applies:

- If all receiving lanes encounter errors, then physical layer training shall timeout.
- If a physical layer has multiple receiving lanes and at least one matches the IBIST Header, then the physical layer may take one of the following actions:
 - If link width reduction is enabled, then a logical sublayer may map out the lanes that fail to enter the IBIST state.
 - Restrict all lanes from entering IBIST which shall result in a physical layer timeout.
 - Ignore lanes that do not enter the IBIST state, report status in the Rx Lane Status field of the PHY Specific Interface PHY Structure, and exit the IBIST state when all active IBIST lanes complete their evaluation period.

4.2.11.5.3. **IBIST Exit**

The PLTSM shall exit the IBIST state after the configured number of IBIST Patterns have been received from an IBIST Training Sequence as specified in the IBIST Control Fields of the *Unique Interface Physical Layer Structure*.

Upon exit from the IBIST state, the PLTSM shall transition to the Config State. If errors or insufficient margin is determined in the IBIST state, then a physical layer shall disable lanes which can be mapped out in the Config state.

Developer Note: For test, debug, or error analysis, a field in the Interface PHY Structure for IBIST configuration is provided to stay perpetually in the IBIST state during physical layer training.

4.2.11.6. **Config (Required)**

The Gen-Z Common Specification clause specifies the minimum requirements for the transmission and reception of Config Training Sequences. Individual physical layer clauses may require additional Configuration Training Sequences.

5 **4.2.11.6.1. Config Entry**

The Config state shall be reached by one of two arcs:

- From the Align state.
- From the IBIST state, if enabled.

10 The operating Tx and Rx width are determined in the Config state. The logical sublayer should use the receiver (inbound) results from the Align state to inform the remote transmitter of the lanes which to support in the link. The physical layer's transmitter shall disable lanes that are not used.

4.2.11.6.2. Config Process

15 Upon entry into the Config state, each individual lane transmitter shall continuously transmit 128 UI Config Training Sequences as specified in *PHY Config Training Sequence* in parallel. Striping across lanes shall not be used with Config Training Sequences. Initially, the transmitter shall set Word 2:<width> to 0x00000000 in the Config Training Sequence. After a minimum of eight consecutive Config Training Sequences have been received, the local logical sublayer sets the <width> field in the Config Training Sequence. After the local logical sublayer has received a minimum of four consecutive Config Training Sequences with valid non-zero <width> fields and the transmitter has transmitted a minimum of eight
20 Config Training Sequences with non-zero <width> field, the transmitter shall transmit a single Lock Training Sequence.

4.2.11.6.3. Config Exit

25 After both transmitting and receiving Lock Training Sequences, the PLTSM shall advance to the PHY-Up state, and the link is fully operational. The logical sublayer should inform its local transmitter of which lanes to disable with the <width> field upon exit of the Config state.

4.2.11.7. **PHY-Up (Required)**

The PHY-Up state indicates that physical layer training has completed and that the link is operational, i.e., can exchange Gen-Z Core data and messages.

30 The Gen-Z Common Specification clause specifies a minimum set of physical layer sequences when line-coding is present. Individual physical layer clauses may define additional physical sequences.

For individual physical layer clauses that do not require line-coding, no physical layer-specific sequences exist in the PHY-Up state, and Gen-Z Link CTL packets are used to direct the physical layer.

4.2.11.7.1. PHY-Up Entry

The Physical Layer State Machine shall enter the PHY-Up state in one of two ways:

- 35 • From the Config state.
- If supported, then from one of the PHY Low Power states

4.2.11.7.2. PHY-Up Process

If in the PHY-Up state and line-coding is enabled, then the logical sublayer may transmit *PHY-Up Sequences*.

4.2.11.7.3. PHY-Up Exit

- 5 The PLTSM shall exit the PHY-Up state under any of the following conditions:
- On a Physical layer reset or system power down, the Physical Layer State Machine shall return to the PHY Idle state.
 - The Gen-Z Core directs the Physical Layer State Machine to a supported PHY LP state or to retrain via the PLA.
 - 10 • Bit errors are detected and the physical layer autonomously retrains by returning to the Align state.

4.2.12. Clock Compensation

15 In order to compensate for variation in reference clock sources for non-synchronous systems, clock compensations PHY-Up Sequences shall be periodically transmitted during physical layer training and in the *PHY-Up (Required)* state.

For non-synchronous systems, a physical layer shall transmit a series of clock compensation PHY-Up Sequences as illustrated in *PHY-Up Sequences*. Clock compensation PHY-Up Sequences shall not be scrambled; this enables a scrambler/descrambler pair at opposition physical layer ends to remain synchronized to each other when clock compensation PHY-Up Sequences are added or deleted. To maintain good DC bias and bit transition density without scrambling, the <SetsRemain> word(s) shall be encoded as follows:

- `cc_rem_value[3:0]` = Remaining Clock Compensation PHY-Up Sequences
- `n[15:0]` = {`cc_rem_value[3:0]`,`cc_rem_value[3:0]`,`cc_rem_value[3:0]`,`cc_rem_value[3:0]`}
- <SetsRemain> = {`~n[15]`, `n[15]`, `~n[14]`, `n[14]`, ... `~n[0]`, `n[0]`}
- `~` = inverted bit value

25 Upon reception of a clock compensation physical packet, a physical layer may add or remove clock compensation PHY-Up Sequences. If PHY-Up Sequences are added, then the new clock compensation PHY-Up Sequence shall increment the value of field `cc_rem_value[3:0]` by the number of clock compensation sets added. If PHY-Up Sequences are removed, then the value of field `cc_rem_value[3:0]` shall be decremented by the number of clock compensation sets removed.

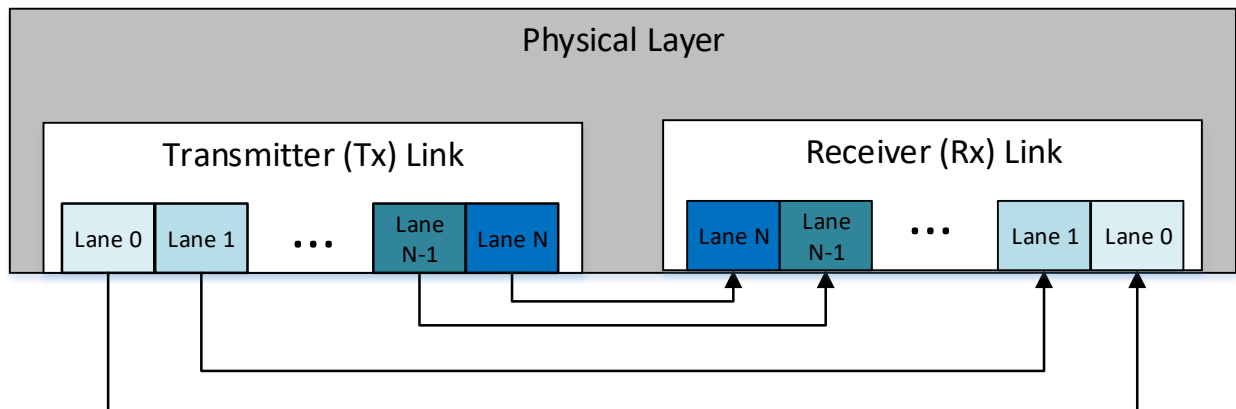
30 As the reference clock requirements are unique to each individual physical layer clause, the amount transmitted, the frequency of transmission, the rules for addition and removal, and the states required for clock compensation physical packets shall be as specified within their individual physical layer clause.

4.2.13. Loopback

35 Physical layer loopback is a functional feature in which a port can perform physical layer training or other special test modes by coordinating and exchanging information with either itself (master loopback) or the remote physical layer (slave loopback). Support for loopback configurations shall be as specified in each individual physical layer clause.

4.2.13.1. **Master Loopback (TX → RX)**

Master Loopback is a configuration in which a physical layer's transmit data-path is locally connected to its receive data-path. This may occur at different or multiple levels within a design, such as an internal parallel interface or externally at a serial interface.



5

Figure 4-17: Master Loopback

4.2.13.2. **Slave Loopback (Rx → TX)**

Slave Loopback is defined as a configuration in which a physical layer's receive data-path is connected to its transmit data-path. The primary use case for this topology is to enable connections to external test equipment or other components to evaluate and margin a remote receiver. The configuration for slave loopback shall be enabled statically prior to physical layer training and may require specific reference clock topologies (i.e. shared reference), special state transition control, and/or limits on lane width.

10

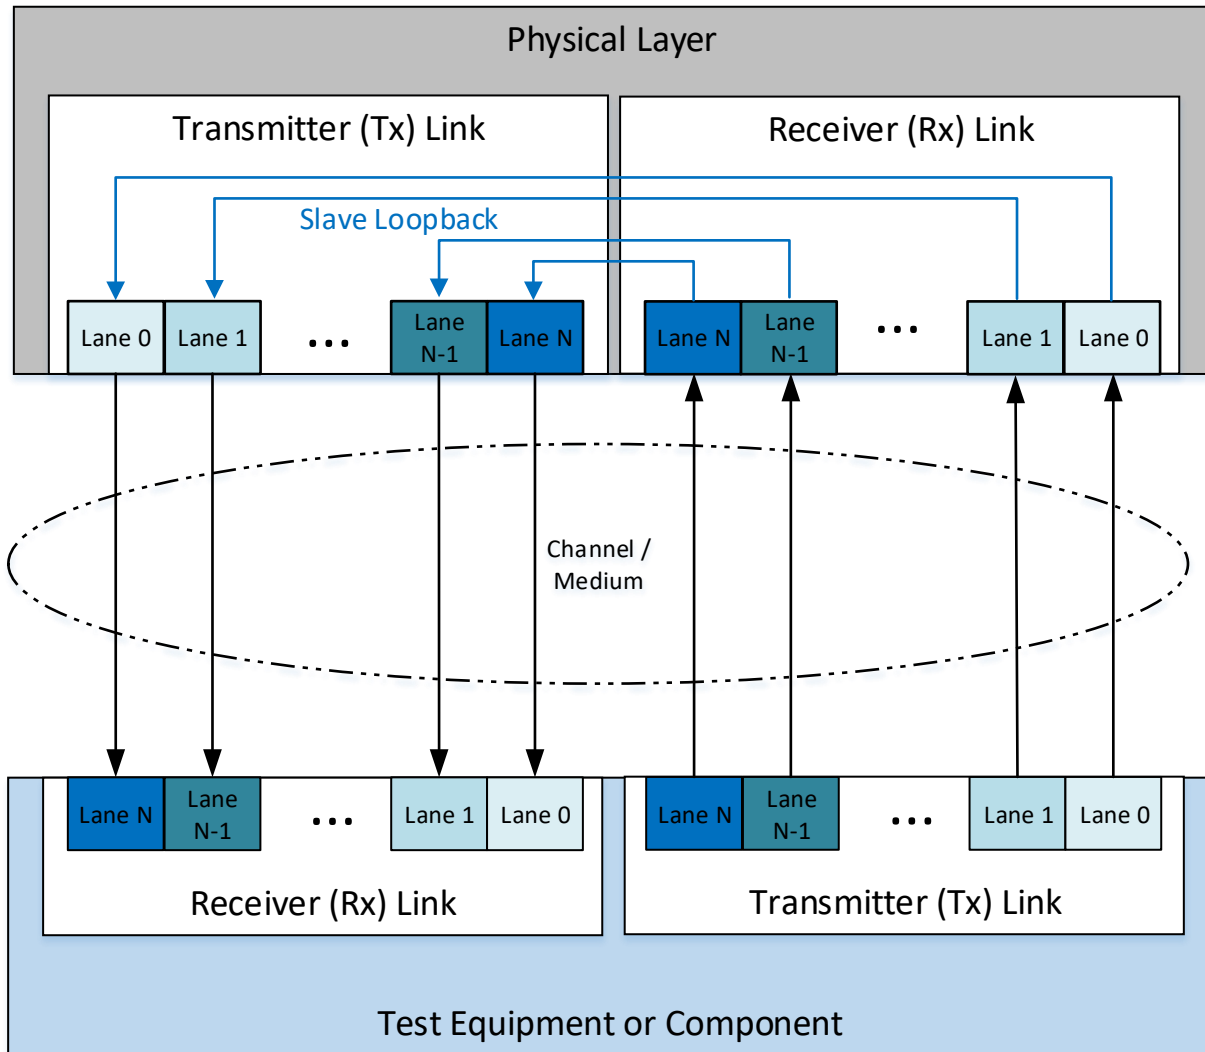


Figure 4-18: Slave Loopback

4.2.14. Physical Layer Aggregation

The *Gen-Z Core Specification* specifies a mechanism for grouping (or aggregating) smaller, narrower Interface Structures, including the Interface PHY Structures, into larger, wider single aggregated interfaces (SAI).

5

Support for static aggregation or disaggregation of a physical layer and its interfaces shall be as specified in each individual physical Layer clause.

A physical layer shall include its own PLA interface and Interface PHY Structure for each operational Gen-Z Core Interface Structure.

10

Developer Note: *Bifurcation* is a common term with physical layers and is considered the opposite of aggregation. In this context, bifurcation is equivalent to disaggregation.

4.2.15. Re-timers

Re-timers may be used to extend the reach of physical layers. The support for re-timers shall be as specified in each individual physical layer clause.

4.3. Power Management

4.3.1. Overview

The Gen-Z PLA and Interface PHY Structure enable support for eight low power states—four PHY Low Power Idle states and four PHY Low Power Up states that enable data transmission and/or reception at a reduced bandwidth. The PLA specifies the state mapping while the Interface PHY Structure specifies the Control, Status, and Capability fields for each low power state.

4.3.2. Low Power Operation

Support for the four PHY Low Power Idle (PHY LP#) and four PHY Low Power Up (PHY-Up LP#) may be specified in each individual physical layer clause's Power Management section and in their respective *Common, Lane, and Low Power Interface PHY Structure*.

The following are the features and requirements associated with low power operation:

- PHY LP and PHY-Up LP entrance state transitions may be initiated by:
 - A Gen-Z Core physical layer low power state request across the PLA.
 - If encoding supports PHY-Up Sequences, then:
 - Autonomously by hardware by transmitting PHY-Up Sequences of type PHY LP with a LPID = PHY LP# or PHY-Up LP#.
 - On reception of PHY-Up Sequences of type LP with a LPID = PHY LP# or PHY-Up LP#.
- PHY LP and PHY-Up LP exit state transition may be initiated by:
 - A Gen-Z Core physical layer low-power state request across the PLA.
 - An explicit control request for exiting a low-power state from the *Common, Lane, and Low Power Interface PHY Structure*.
 - On reception of a Signal Detect trigger when in a PHY-LP state.
 - If encoding supports PHY-Up Sequences, then:
 - Autonomously in hardware by transmitting PHY-Up Sequences of type LP with a LPID = PHY-Up.
 - On reception of a PHY-Up Sequences with of type LP with a LPID = PHY-Up when in a PHY-Up LP state.
- All supported low-power states may be enabled or disabled via the *Common, Lane, and Low Power Interface PHY Structure*.
- Gen-Z Core PHY-Up-LP# entrance and exit requests from the PLA may require retraining; see each individual physical layer clause.
- Exiting low-power states shall have a higher priority than all low-power up or idle states.
- PHY-Up LP# shall have priority over PHY LP#.
 - Low-power state numbers shall be prioritized with #1 the highest and #4 the lowest.

- When exiting a PHY-LP# state as directed by the Gen-Z Core from the PLA, retraining may be required to enter the PHY-Up state. If so, the details for retraining shall be specified in each individual physical layer clause.
- When entering or exiting a PHY-Up LP# state as directed by the Gen-Z Core from the PLA, retraining may be required from and to the PHY-Up state. If so, the details for retraining shall be specified in each individual physical layer clause.

The following are the features and requirements associated with low power for physical layers with encoding or framing to enable PHY-Up Sequences:

- Low power packets and their <LPID> are detailed in the *PHY-Up Sequences*.
 - LP_State[3:0] = see *Low Power PHY-Up Packet Mapping*
 - n[15:0] = {LP_State[3:0], LP_State[3:0], LP_State[3:0], LP_State[3:0]}
 - <LPID> = {~n[15], n[15], ~n[14], n[14], ... ~n[0], n[0]} (~ = inverted bit value)
- A physical layer shall continue to transmit PHY-Up Sequences of type LP until the logical sublayer has transmitted a minimum of 4 and received a minimum of 2 with matching LP_State[3:0].
- If more than one physical layer low-power state is supported and there is a conflict in arc number requests, then the highest priority LP_State[3:0] (lowest number) request shall take precedence.

The following are the features and requirements associated with low power for physical layers without encoding to enable PHY-Up Sequences:

- The Gen-Z Core shall initiate and coordinate all PHY-LP# and PHY-Up LP# entrance events via a Link CTL packet.
- The Gen-Z Core shall initiate and coordinate all PHY-Up LP# exit events via a Link CTL packet.

Table 4-17: Low Power PHY-Up Packet Mapping

Low Power State	LP_State[3:0] Mapping	Notes
PHY-Up	0x0	Exit from any PHY LP or PHY-Up LP state
PHY-Up LP1	0x1	Entrance to PHY-Up LP1 state
PHY-Up LP2	0x2	Entrance to PHY-Up LP2 state
PHY-Up LP3	0x3	Entrance to PHY-Up LP3 state
PHY-Up LP4	0x4	Entrance to PHY-Up LP4 state
PHY LP1	0x9	Entrance to PHY LP1 state
PHY LP2	0xA	Entrance to PHY LP2 state
PHY LP3	0xB	Entrance to PHY LP3 state
PHY LP4	0xC	Entrance to PHY LP4 state

4.3.2.1. Physical Layer Low Power States

4.3.2.1.1. PHY-LP Entrance

PHY-LP entrance shall be directed by the Gen-Z Core across the PLA interface following the transmission and acknowledgment of an Enter-Link-LP Link CTL packet as illustrated in *PHY-LP Entrance*.

- 5 The following sequence of steps shall be taken to enter PHY-LP:
1. Once the physical layer has received a request across the PLA to enter a PHY-LP# state:
 - a. If physical line-coding is enabled, then logical sublayer shall transmit *PHY-Up Sequences* of type LP with the appropriate LPID for its PHY LP# for 8 μ s.
 - b. If physical line-coding is not enabled, then the logical sublayer shall transmit data from the PLA for 8 μ s.
 - 10 2. The physical layer shall enter the PHY-LP# state and shall report its status across the PLA to the Gen-Z Core.
 3. After waiting 1 ms, the physical layer shall enable its receivers for Signal Detect as specified in its individual physical layer clause.

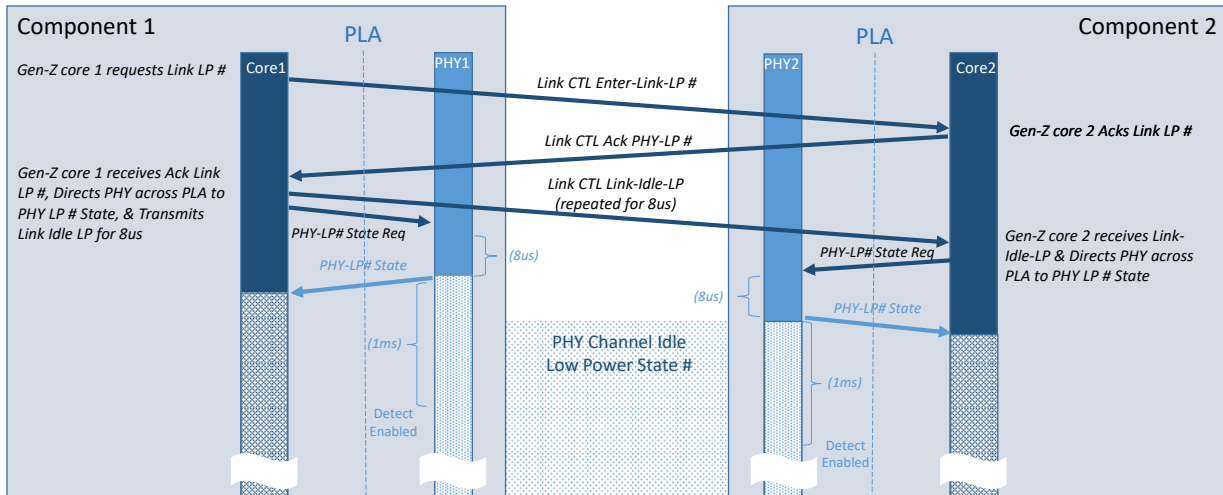
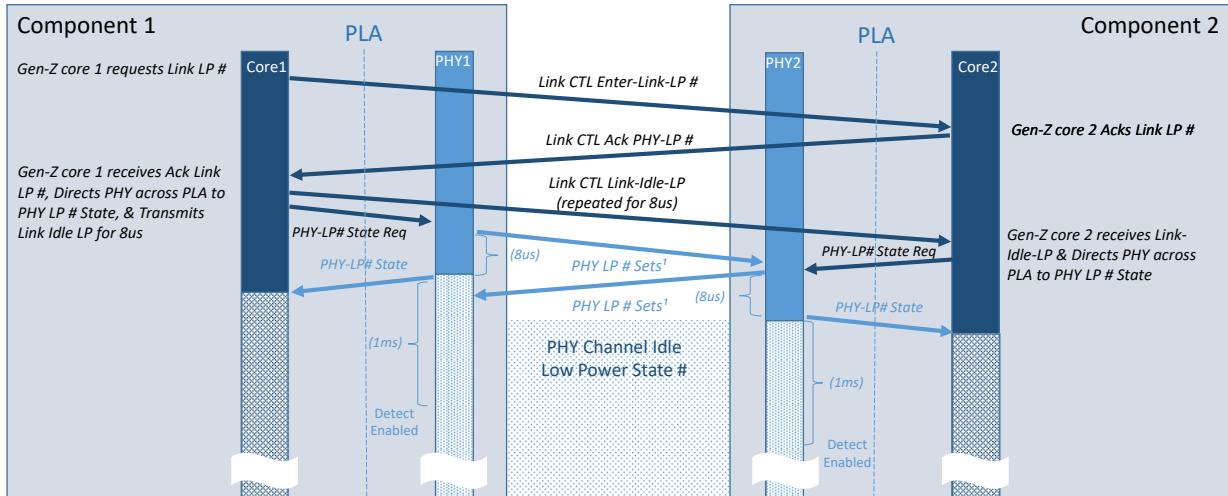


Figure 4-19: PHY-LP Entrance

Developer Note: For physical layers without line-coding and if the Gen-Z Core is not transmitting Link-Idle-LP Link CTL packets, it is recommended that the Gen-Z Core transmit Idles until the Gen-Z Core receives a PHY-LP# status response from the physical layer across the PLA.

It is recommended that the interface's Gen-Z Core initiating the low-power transition continuously transmit Link-Idle-LP Link CTL packets. When physical line-coding is enabled, it is recommended that the physical layer ignore Gen-Z Core data and transmit PHY-Up Sequences of type LP.



¹Physical Framing/Encoding Enabled

Figure 4-20: PHY-LP Entrance Supporting PHY-Up Sequences

Developer Note: The total latency to enter PHY-LP states is a function of the following:

- Time from Core 1's reception of Link ACK Link CTL packet to transmitting Link-Idle-LP Link CTL packet and directing PHY1 to enter a PHY-LP# state across the PLA.
- And the greater of:
 - PHY1's entrance:
 - Minimum time for Core 1 to transmit Link-Idle-LP Link CTL packet or with line-coding, for PHY1 to transmit PHY-Up Sequences of type LP (LPID = PHY LP#).
 - PHY1's PHY-LP# entrance latency as specified in its Interface PHY Structure.
 - Or PHY2's entrance:
 - Time for the Link-Idle-LP Link CTL packet to cross the channel or media.
 - Time from Core 2's reception of the Link-Idle-LP Link CTL packet to directing PHY2 to enter a PHY-LP# state across the PLA.
 - Minimum time for PHY2 to forward IDLEs or with line-coding, or to transmit PHY-Up Sequences of type LP (LPID = PHY LP#).
 - PHY2's PHY-LP# entrance latency as specified in its Interface PHY Structure.

4.3.2.1.2. PHY-LP Exit

PHY-LP state exit may be directed by the Gen-Z Core as a state request change across the PLA as illustrated in *PHY-LP Exit*. The physical layer that does not initiate the PHY-LP exit shall be capable of detecting an inbound PHY Signal Detect Training Sequence and reinitiating physical layer training to the PHY-Up state without direction from the Gen-Z Core across the PLA.

The following sequence of steps shall be taken to exit PHY-LP:

1. Once the physical layer has received a request across the PLA to enter a PHY-LP state, the logical sublayer shall start transmitting PHY Signal Detect Training Sequences.
2. The remote component's receiver shall detect the presence of a signal and shall begin transmitting PHY Signal Detect Training Sequences.
3. After the completion of physical layer training the PHY shall enter the PHY-Up state and shall report its status across the PLA to the Gen-Z Core.
4. The remote component's Gen-Z Core PHY state request shall transition to PHY-Up across the PLA.

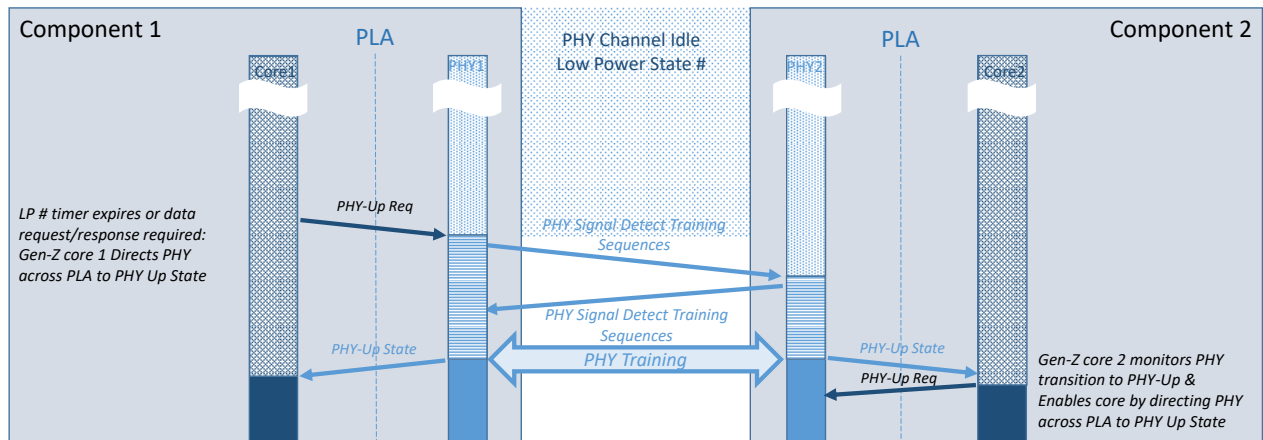


Figure 4-21: PHY-LP Exit

Developer Note: The total latency to exit PHY-LP states is a function of the following:

- Time for Core 1 to direct PHY1 to enter a PHY-Up# state across the PLA.
- Time for PHY1's PHY Signal Detect Training Sequences to cross the channel or media.
- Time for PHY2 to receive and begin transmission of the PHY Signal Detect Sequence.
- Time for PHY2's PHY Signal Detect Training Sequences to cross the channel or media.
- Time for physical layer training to complete which may be the greater of:
 - PHY1's PHY-LP# exit latency as specified in its Interface PHY Structure.
 - PHY2's PHY-LP# exit latency as specified in its Interface PHY Structure.

4.3.2.2. Physical Layer Low Power Up States

4.3.2.2.1. PHY-Up LP Entrance

PHY-Up LP state entrance shall be directed by the Gen-Z Core across the PLA interface following the transmission and acknowledgment of an Enter-Link-UP-LP Link CTL packet as illustrated in *PHY-Up LP Entrance*.

The following sequence of steps shall be taken to enter PHY-Up LP:

1. Once the physical layer has received a request across the PLA to enter a PHY-Up LP# state:
 - a. If physical line-coding is enabled, then the logical sublayer shall transmit PHY-Up LP# packets as specified in *PHY-Up Sequences* and shall follow the rules specified in *Low Power Operation* as illustrated in *PHY-Up LP Entrance Supporting PHY-Up Sequences*.
 - b. If physical line-coding is not enabled, then the logical sublayer shall transmit data from the PLA for 1 μ s as illustrated in *PHY-Up LP Entrance*.
2. The physical layer shall begin retraining.
3. After the physical layer state has trained to a degraded performance, low-power state, the PLTSM shall enter the PHY-Up LP# state and shall report its status across the PLA to the Gen-Z Core.

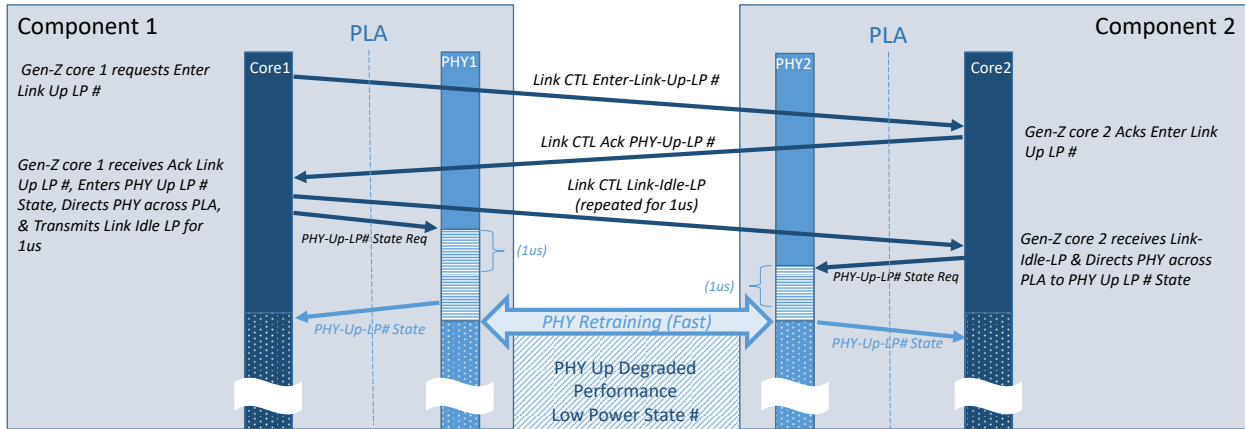
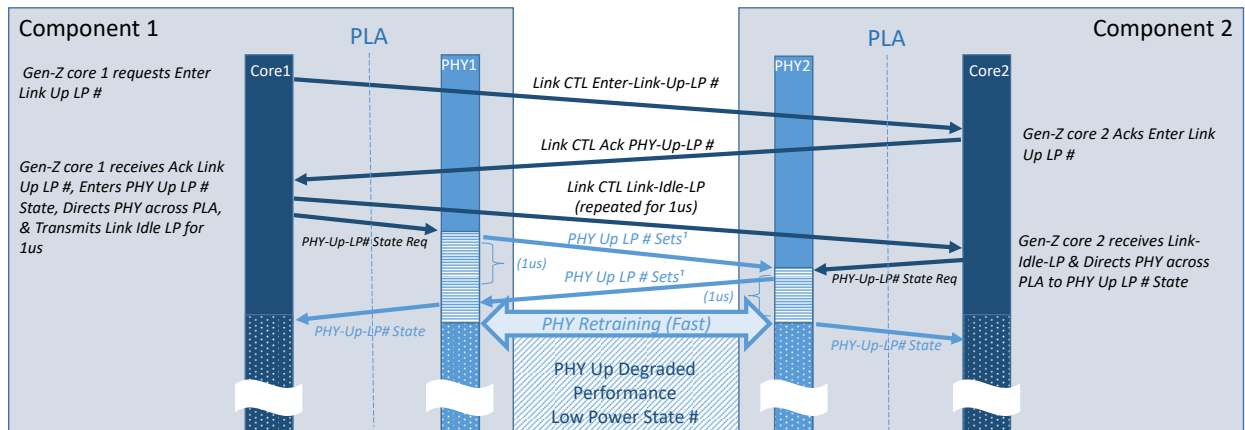


Figure 4-22: PHY-Up LP Entrance



¹Physical Framing/Encoding Enabled

Figure 4-23: PHY-Up LP Entrance Supporting PHY-Up Sequences

Developer Note: The total latency to enter (or exit) PHY-Up LP states is a function of the following:

- Time from Core 1's reception of a Link ACK Link CTL packet to transmitting Link-Idle-LP Link CTL packets and directing PHY1 to enter (or exit) a PHY-Up LP# state across the PLA.
- And the greater of:
 - PHY1's entrance (or exit):
 - Minimum time for Core1 to transmit Link-Idle-LP Link CTL packets or with line-coding, for PHY1 to transmit PHY-Up Sequences of type LP (LPID = PHY-Up LP# for entrance or PHY-Up for exit).
 - Or PHY2's entrance (or exit):
 - Time for the Link-Idle-LP Link CTL packet to cross the channel or media.
 - Time from Core2's reception of the Link-Idle-LP Link CTL packet to directing PHY2 to enter (or exit) Low Power Up state across the PLA.
 - Minimum time for PHY2 to forward IDLEs or with line-coding, to transmit PHY-Up Sequences of type LP (LPID = PHY-Up LP# for entrance or PHY-Up for exit).
- Time for physical layer training to complete which may be the greater of:
 - PHY1's PHY-Up LP# entrance (or exit) latency as specified in its Interface PHY Structure.
 - PHY2's PHY-Up LP# entrance (or exit) latency as specified in its Interface PHY Structure.

4.3.2.2.2. PHY-Up LP Exit

PHY-UP LP state exits shall be directed by the Gen-Z Core across the PLA interface following the transmission and acknowledgment of an Exit-Link-UP-LP Link CTL packet as illustrated in *PHY-Up LP Exit*.

The following sequence of steps shall be taken to exit PHY-UP LP:

- 5 1. Once the physical layer has received a request across the PLA to enter a PHY-Up state:
 - a. If physical line-coding is enabled, then the logical sublayer shall transmit PHY-Up Sequences as specified in *PHY-Up Sequences* and follow the rules specified in *Low Power Operation* as illustrated in *PHY-Up LP Exit Supporting PHY-Up Sequences*.
 - 10 b. If physical line-coding is not enabled, then the logical sublayer shall transmit data from the PLA for 1 μ s as illustrated in *PHY-Up LP Exit*.
2. The physical layer shall begin retraining.
3. After the physical layer state has transitioned to the PHY-Up state, the PLTSM shall enter the PHY-Up state and shall report its status across the PLA to the Gen-Z Core.

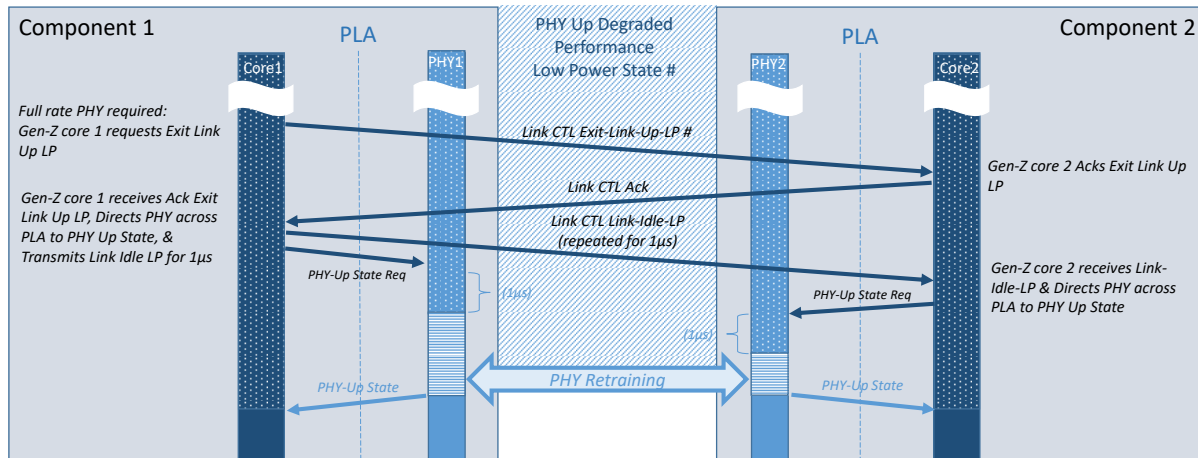


Figure 4-24: PHY-Up LP Exit

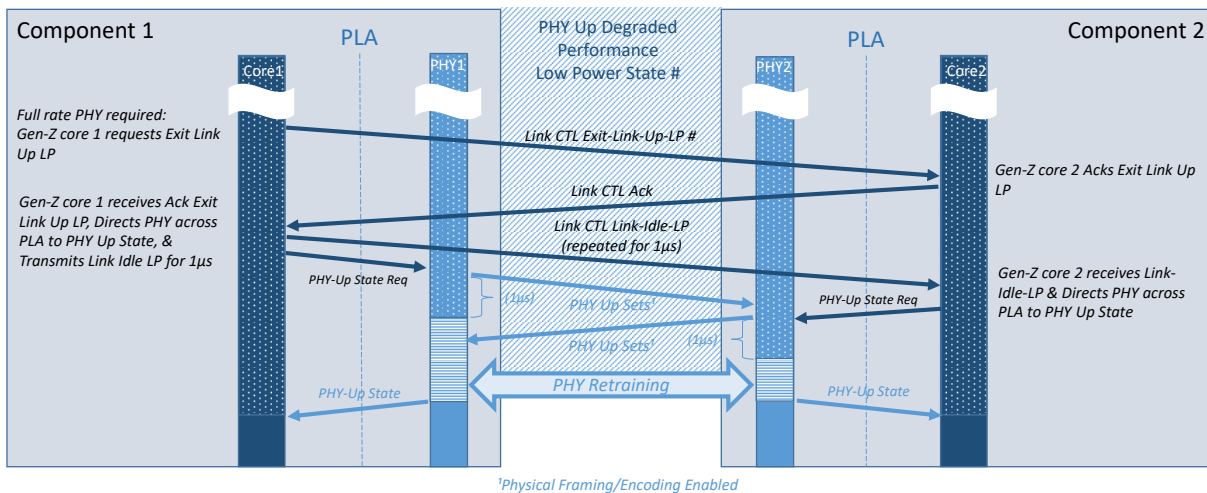


Figure 4-25: PHY-Up LP Exit Supporting PHY-Up Sequences

Developer Note: Calculating the exit latency of PHY-Up LP states is similar to entrance and is enumerated in the Developer Note section of PHY-Up LP Entrance.

4.3.3. Physical Layer Power Management States

Low-power modes depend upon their specific physical layer and shall be as specified within each individual physical layer clause.

4.3.3.1. Physical Layer Low Power Idle States

5 Support for PHY Low Power Idle states shall be as specified in each individual physical layer clause.

4.3.3.2. Physical Layer Low Power Up States

Support for PHY Low Power Up states shall be as specified in each individual physical layer clause.

4.4. Electrical Sublayer

10 Since the Electrical Sublayer is a primary differentiating characteristic of a physical layer, the Gen-Z Common Specification clause shall not include any content in this section.

4.5. Management Control and Status

4.5.1. Interface Physical Layer Structure

4.5.1.1. Interface PHY Structure

15 The Interface PHY is common to all physical layers and is used by software for various management activities. The Interface PHY structure shall be as specified in the *Gen-Z Core Specification*.

4.5.1.1.1. Common, Lane, and Low Power Interface PHY Structure

20 Common, Lane, and Low Power fields shall be allocated for each Interface PHY Structure as specified in the *Gen-Z Core Specification*. However, not all fields may be required for each individual physical layer clause. Thus each individual physical layer clause shall provide details on specific values, exclusions, or modified behavior for the *Interface PHY Structure Common, Lane, and Low Power Fields*.

Interface PHY Structure

+7		+6		+5		+4		+3		+2		+1		+0				
7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0			
R0				TxDLY	PHY Type				Size				Vers	Type				< Byte 0x0
VD PTR								Next Interface PHY PTR								< Byte 0x8		
PHY Control								PHY Status								< Byte 0x10		
PHY CAP 1 Control								PHY CAP 1								< Byte 0x18		
R1								PHY Events								< Byte 0x20		
R2																< Byte 0x28		

Interface PHY Structure Common Fields (Required)

PHY Lane Control								PHY Lane Status								< Byte 0x30
PHY Remote Lane CAP								PHY Lane CAP								< Byte 0x38

Interface PHY Structure Lane Fields (Required)

PHY LP Timing CAP								PHY LP CAP								< Byte 0x40
-------------------	--	--	--	--	--	--	--	------------	--	--	--	--	--	--	--	-------------

Interface PHY Structure Low Power Fields (Required)

PHY UP LP Timing CAP								PHY UP LP CAP								< Byte 0x48
----------------------	--	--	--	--	--	--	--	---------------	--	--	--	--	--	--	--	-------------

Interface PHY Structure PHY Up Low Power Fields (Required)

Figure 4-26: Interface PHY Structure Common, Lane, and Low Power Fields

4.5.1.1.2. Extended Feature Interface PHY Structure

Extended feature fields shall be allocated for each Interface PHY Structure and are unique to an individual physical layer. The details for the *Interface PHY Structure Extended Feature Fields* shall be as specified in each individual physical layer clause.

PHY Extended Control								PHY Extended Status								< Byte 0x50
PHY Remote Extended CAP								PHY Extended CAP								< Byte 0x58

Interface PHY Structure Extended Feature Fields (Required)

Figure 4-27: Interface PHY Structure Extended Feature Fields

4.5.1.2. Unique Interface Physical Layer Structure

The Gen-Z Common Specification clause shall not specify unique physical layer structures as those are only relevant to individual physical layer clauses.

4.5.2. Physical Layer Management

The management of a physical layer may be controlled by the out of band management interface specified in the *Gen-Z Core Specification*. Management is responsible for understanding the supported local and remote physical layer clauses to configure and initialize the physical layers to train to the PHY-Up state. The details are illustrated in the *Physical Layer Training Flow*.

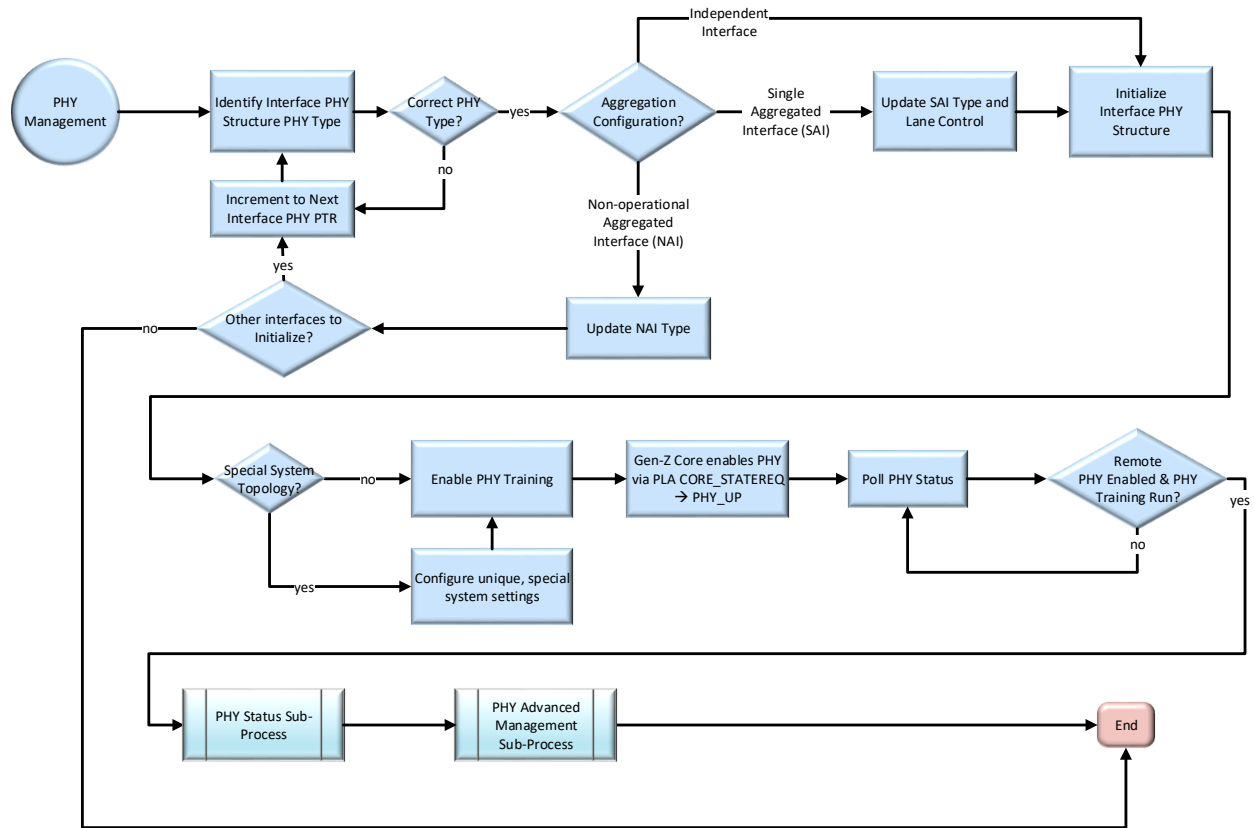


Figure 4-28: Physical Layer Training Flow

4.5.2.1. *Physical Layer Training Status*

During physical layer training, status may be monitored to provide additional information for error or degraded performance error conditions as illustrated in *Physical Layer Status Sub-Process*.

- 5 The algorithms or flows for error and degraded performance error can be implementation-specific unless specified in an individual physical layer clause or alternative Gen-Z document.

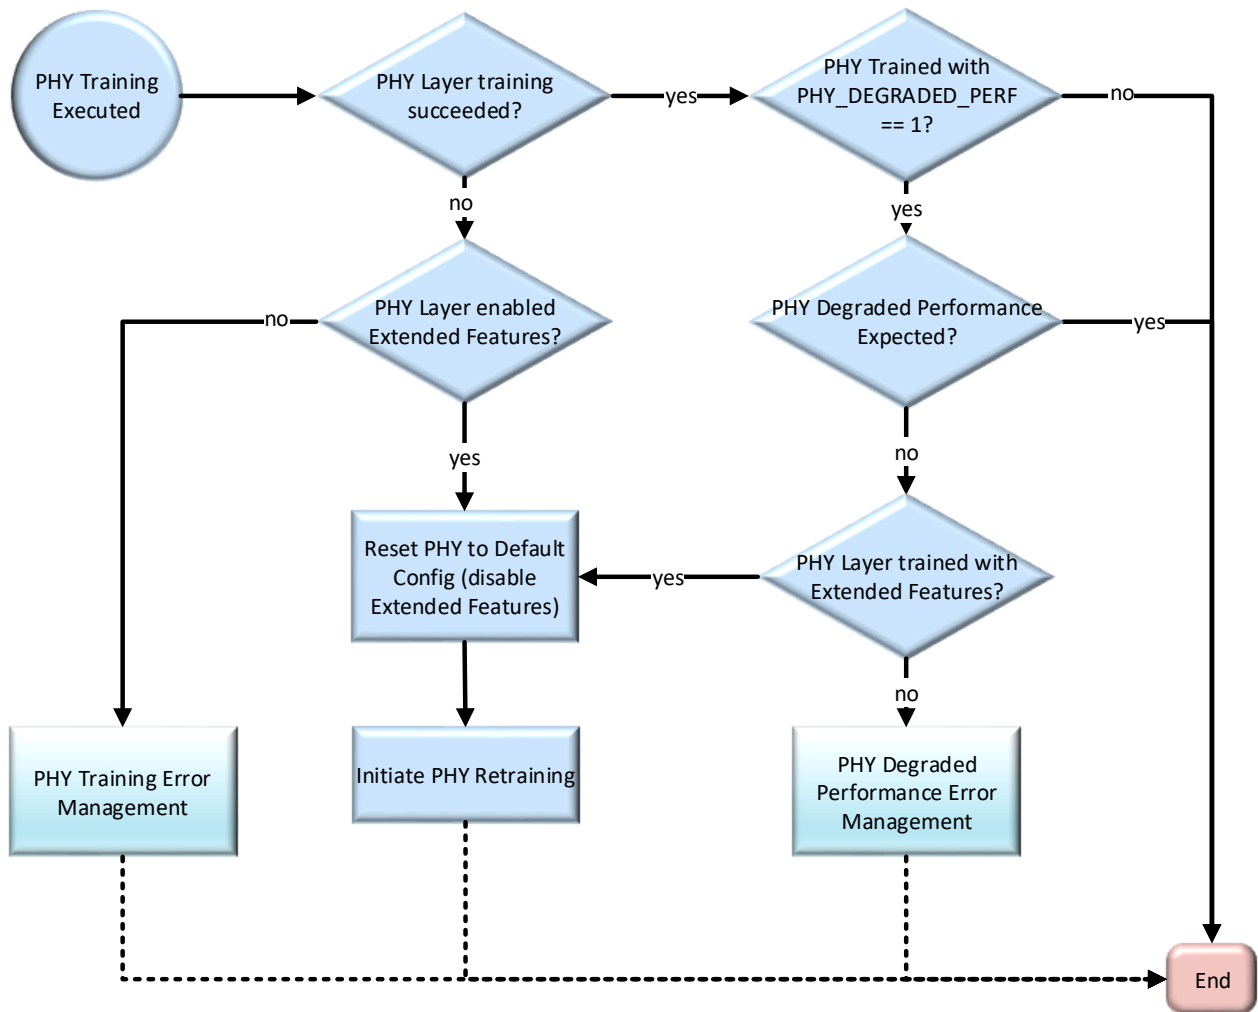


Figure 4-29: Physical Layer Status Sub-Process

4.5.2.2. **Physical Layer Advanced Management**

Gen-Z physical layer clauses may provide options for advanced management to support reconfiguration based on mutually supported baseline or extended feature capabilities in the *Interface Physical Layer Structure*. *Physical Layer Advanced Management Sub-Process* illustrates a flow for supporting advanced management.

The negotiation mechanism for identifying the primary manager for physical layers with independent controlling entities shall be as specified in the *Gen-Z Management Architecture Specification*.

The algorithms or flows for advanced feature management can be implementation-specific unless specified in an individual physical layer clause or alternative Gen-Z document.

PHY Advanced Management Sub-Process

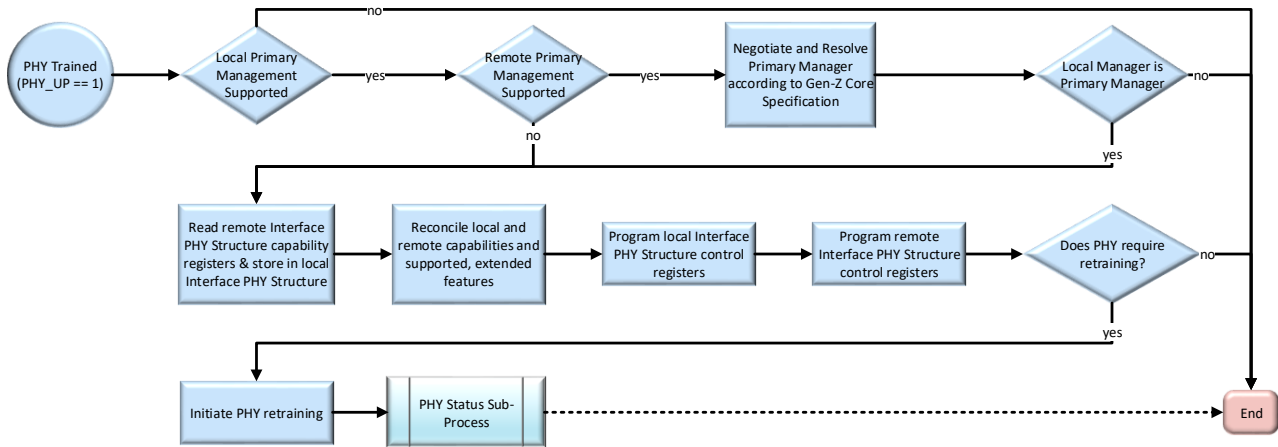


Figure 4-30: Physical Layer Advanced Management Sub-Process

4.5.2.3. *Physical Layer Runtime Reconfiguration*

Gen-Z physical layer clauses may provide options for advanced management to support runtime reconfiguration based on temporal solution requirements, such as reducing symmetric link width to improve power savings or enabling asymmetry for specific workloads.

Either the remote or local controlling entity for the physical layer can initiate a runtime reconfiguration. To prevent possible error recovery for degraded performance reconfigurations on the remote physical layer, the local controlling entity should configure the remote Interface PHY structure as illustrated in *Physical Layer Runtime Reconfiguration Sub-Process*.

Unless specified in an individual physical layer clause or alternative Gen-Z document, the algorithms or flows for runtime reconfiguration can be implementation-specific.

PHY Runtime Reconfiguration Sub-Process




Figure 4-31: Physical Layer Runtime Reconfiguration Sub-Process

5. Gen-Z-E-NRZ-25G-Fabric Specification

5.1. Introduction

5 Gen-Z-E-NRZ-25G-Fabric specifies a Gen-Z physical layer capable of a line-rate at 25.78125 GT/s per lane inclusive of 64b/66b encoding (raw data-rate of 25 Gbps) and using Non-Return to Zero (NRZ) signaling over Fabric media (Long Reach) providing a BER of 10^{-12} or better. A compliant physical layer consists of an electrical interface compatible with the Channel Operating Margin (COM) that has a recommended Insertion Loss¹ (IL) no greater than 30 dB at 12.89 GHz.

5.2. Logical Sublayer

5.2.1. Link Serialization

10 Refer to the Gen-Z Common Specification clause section on *Link Serialization*.

5.2.1.1. Encoding / Decoding

64b/66b encoding shall be used in this physical layer clause. 64b/66b encoding enables the use of PHY Sequences via a two-bit sync-header as specified in *Gen-Z-E-NRZ-25G-Fabric PHY Encoding*.

Table 5-1: Gen-Z-E-NRZ-25G-Fabric PHY Encoding

Sync Header[1:0]	Information Type
10b	Data
01b	PHY Sequences
00b or 11b	Invalid, treated as a receiver error

Refer to the associated *Extended Feature Interface PHY Structure* for additional options.

15 5.2.1.2. Framing

Explicit framing beyond encoding shall not be used with this physical layer clause.

5.2.2. Scrambling / Descrambling

Refer to the Gen-Z Common Specification clause section on *Scrambling / Descrambling*.

5.2.3. Data Striping

20 Refer to the Gen-Z Common Specification clause section on *Data Striping*.

5.2.4. PLA Data Width

Refer to the Gen-Z Common Specification clause section on *PLA Data Width*.

5.2.5. Physical Layer Initialization and Training

Refer to the Gen-Z Common Specification clause section on *Physical Layer Initialization and Training*.

5.2.5.1. PLA State Map

The PLA interface is used to control and report the status of physical layer states. The supported PLA states and their mapping shall be as specified in *PLA Link State to PHY State Mapping*.

Table 5-2: PLA Link State to PHY State Mapping

PLA State	LTSSM State
PHY-Down	Disabled
PHY-Up	P0
PHY-LP1	P1 (Standby)
PHY-LP2	P2 (Sleep)
PHY-Up-LP1	Link Width Reduction (LWR)
PHY-Down-Retrain	Retrain
All Others	Unused

5.2.5.2. Physical Layer Link versus Lane Training

Refer to the Gen-Z Common Specification clause section on *Physical Layer Link versus Lane Training*.

5.2.5.3. Physical Layer (PHY) Sequences

Refer to the Gen-Z Common Specification clause section on *Physical Layer (PHY) Sequences*. All PHY Sequences shall follow the rules specified in *Encoding / Decoding*.

5.2.5.4. Error Recovery

Refer to the Gen-Z Common Specification clause section on *Error Recovery*.

5.2.5.5. Lane-to-Lane Deskew

Refer to the Gen-Z Common Specification clause section for details on *Lane-to-Lane Deskew*.

The lane-to-lane skew shall meet the requirements specified in *Gen-Z-E-NRZ-25G-Fabric Lane-to-Lane Skew Specification*.

Table 5-3: Gen-Z-E-NRZ-25G-Fabric Lane-to-Lane Skew Specification

Skew Type	Notes	Value	Units
TX lane to lane output skew	Between any two lanes on a single PHY port	1.25 (max)	ns
Rx lane to lane skew	Across all lanes on a single PHY port including lane-to-lane variations due to channel and repeater delays	5.0 (max)	ns

5.2.6. Lane Reversal

Lane reversal, both uniform and non-uniform, shall be supported for this physical layer clause.

5 Detection of lane reversal shall occur in the *Signal Detect (Required)* state of the PLTSM. Reversal detection shall occur before *Channel Optimize (Required)* so any independent lane adjustments occur on the appropriate lane.

Developer Note: *In a closed system, lane reversal may be explicitly programmed via the Unique Interface Physical Layer Structure of the Interface PHY Structure.*

5.2.7. Lane Polarity

Lane polarity shall be supported for this physical layer clause. Detection of lane polarity shall occur in the *Channel Optimize (Required)* state of the PLTSM.

10 **Developer Note:** *In a closed system where transmitter and receiver equalization settings are explicitly programmed and the Channel Optimize state is bypassed, then lane polarity can be determined autonomously in the Align (Required) of the PLTSM or can be explicitly programmed via the Unique Interface Physical Layer Structure of the Interface PHY Structure.*

5.2.8. Link Width

Symmetric link widths of 1, 2, 4, 8, 16, or 32 transmit and receive lane may be supported by this physical layer clause.

20 **Developer Note:** *In a closed system where transmitter and receiver equalization settings are explicitly programmed and the Channel Optimize state is bypassed, then asymmetric operation can be configured via the Unique Interface Physical Layer Structure of the Interface PHY Structure.*

25 *Alternatively, after symmetric physical layer training, asymmetric operation can be explicitly programmed by reducing lanes and disabling them via the Common, Lane, and Low Power Interface PHY Structure.*

5.2.9. Link BIST

Refer to the Gen-Z Common Specification clause section on *Link BIST*.

5.2.10. Physical Layer Retraining

5 This specification specifies support for 2 retraining arcs from the PHY-Up state. Retrain Arc 1 shall be supported, and Arc 2 should be supported. Each retraining arc provides tradeoffs in physical layer reconfiguration versus exit latencies, which are specified in *Retraining Arcs* and illustrated in *Retraining State Machine Arcs*.

10 The Gen-Z Common Specification clause *Physical Layer Retraining* section provides details on the retraining packet definition and *Link Progression through Align Physical Layer Training State* provides details on the intermediate arcs in the Align State.

Table 5-4: Retraining Arcs

Retraining Arc	State Transition	Expected Latency	Reconfiguration
Retrain Arc 1 (Default)	PHY-Up → Align	microseconds	Full Receiver Reset
Retrain Arc 2	PHY-Up → Align Symbol Lock	< 1 μs	Receiver Realignment

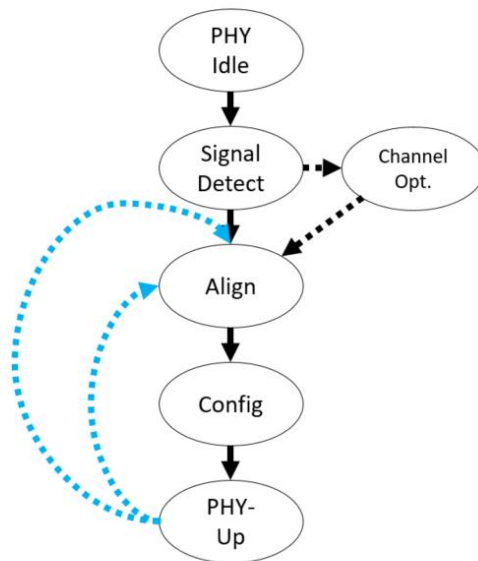


Figure 5-1: Retraining State Machine Arcs

Developer Note: Entrance and exits from for Low Power Up states that encompass link width reduction can require the use of retraining arcs.

5.2.11. PLTSM

The following training states shall be required by PLTSMs supporting this physical layer clause and as illustrated in the *Gen-Z-E-NRZ-25G-Fabric* :

- PHY Idle
- Signal Detect
- Channel Optimize
- Align
- Configure (Config)
- PHY-Up

10 The physical layer states IBIST and PHY Low Power may be supported. The arcs to optional training states are illustrated with a blue dashed line in *Gen-Z-E-NRZ-25G-Fabric* .

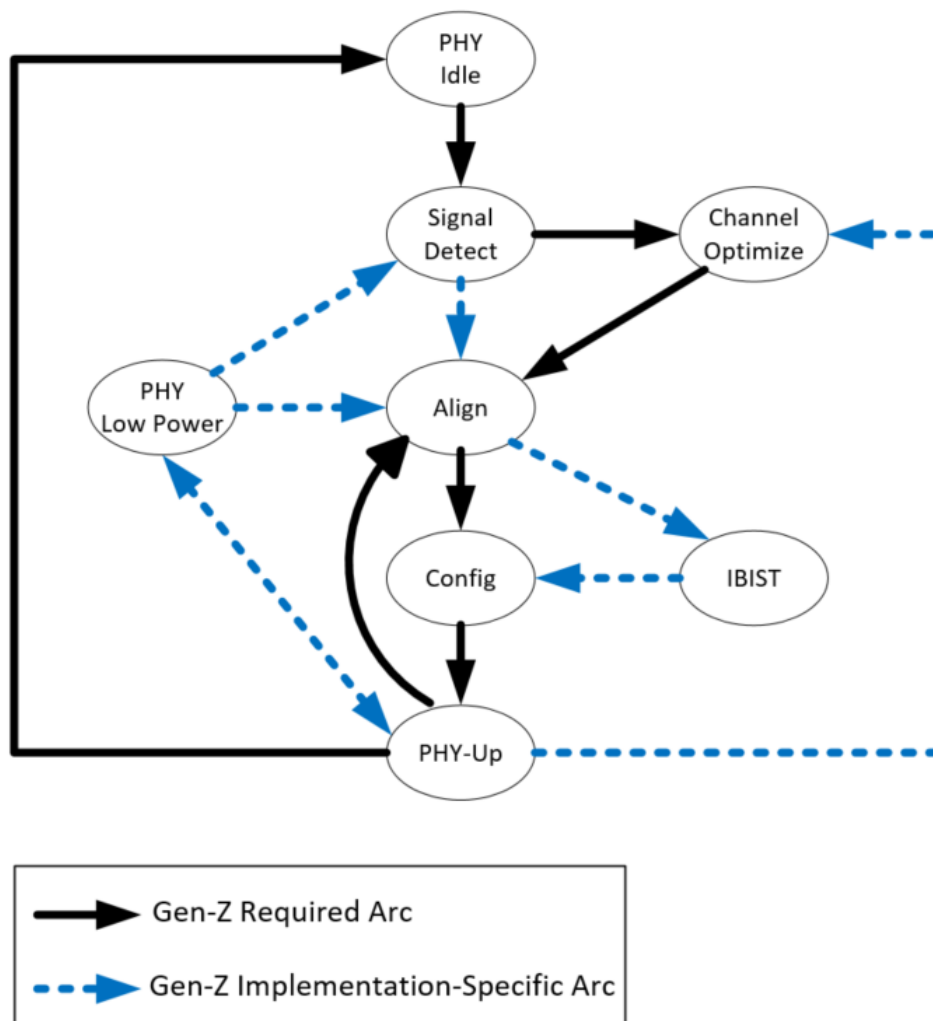


Figure 5-2: Gen-Z-E-NRZ-25G-Fabric PLTSM

A detailed description of each PLTSM state and the conditions of the associated arcs is provided below.

5.2.11.1. PHY Idle (Required)

Refer to the Gen-Z Common Specification clause section on *PHY Idle (Required)*.

5.2.11.2. Signal Detect (Required)

5.2.11.2.1. Signal Detect Entry

- 5 The Signal Detect state may be entered from PHY Idle or when the Gen-Z Core requests the physical layer to transition to a state other than PHY Idle on the CORE_STATEREQ signal of the PLA interface.

5.2.11.2.2. Signal Detect Process

- 10 The Gen-Z Detect state diagram is illustrated in *Gen-Z Signal Detect Sub-State Diagram*. The encapsulating state, Signal Detect, is comprised of the following sub-states: Detect Send, Detect Idle, Detect Wait Reversal, and Detect Wait. The Signal Detect Sub-State Machine may be implementation specific, however, the transmission of the Signal Detect Training Sequences and their associated timing shall be followed.

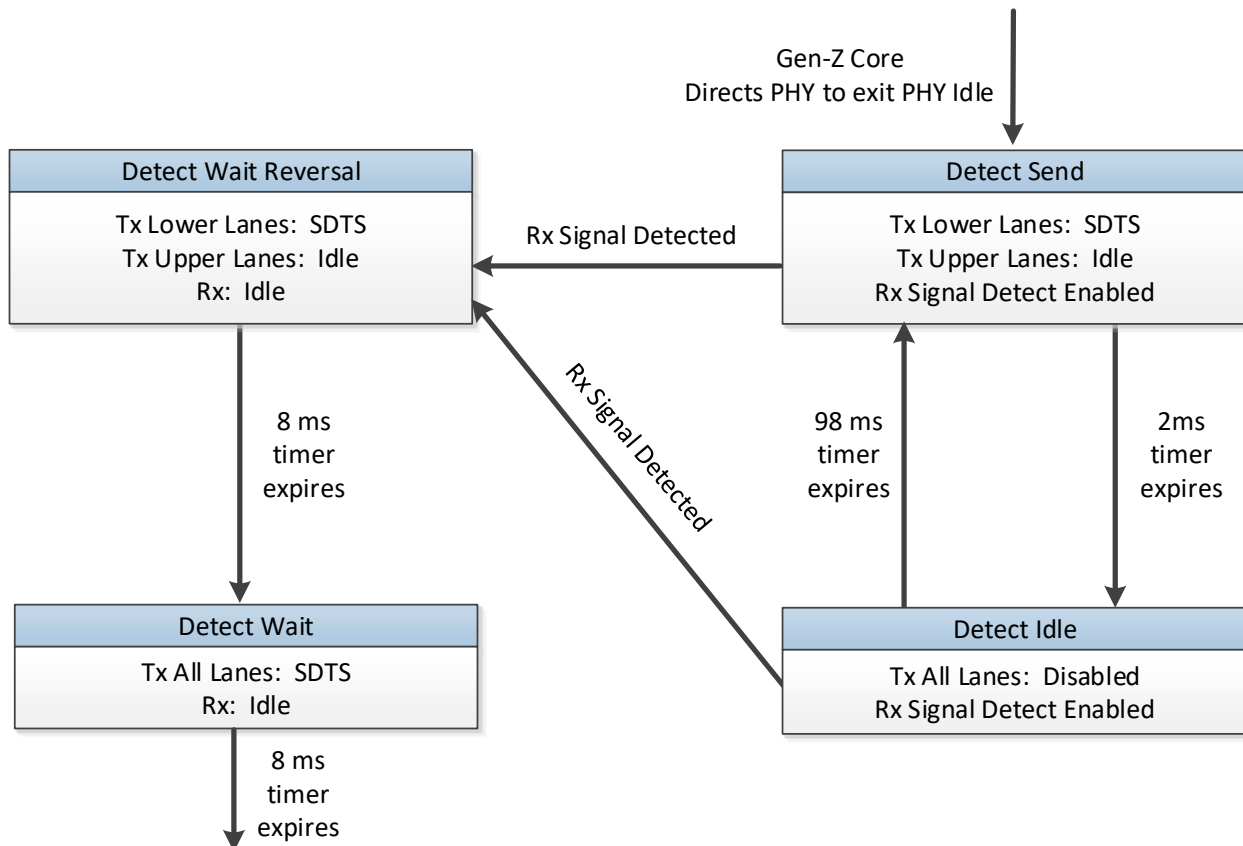


Figure 5-3: Gen-Z Signal Detect Sub-State Diagram

- 15 The definition and function of the lowest-level Signal Detect states is as follows:

Detect Send

While in the Detect Send sub-state,

5

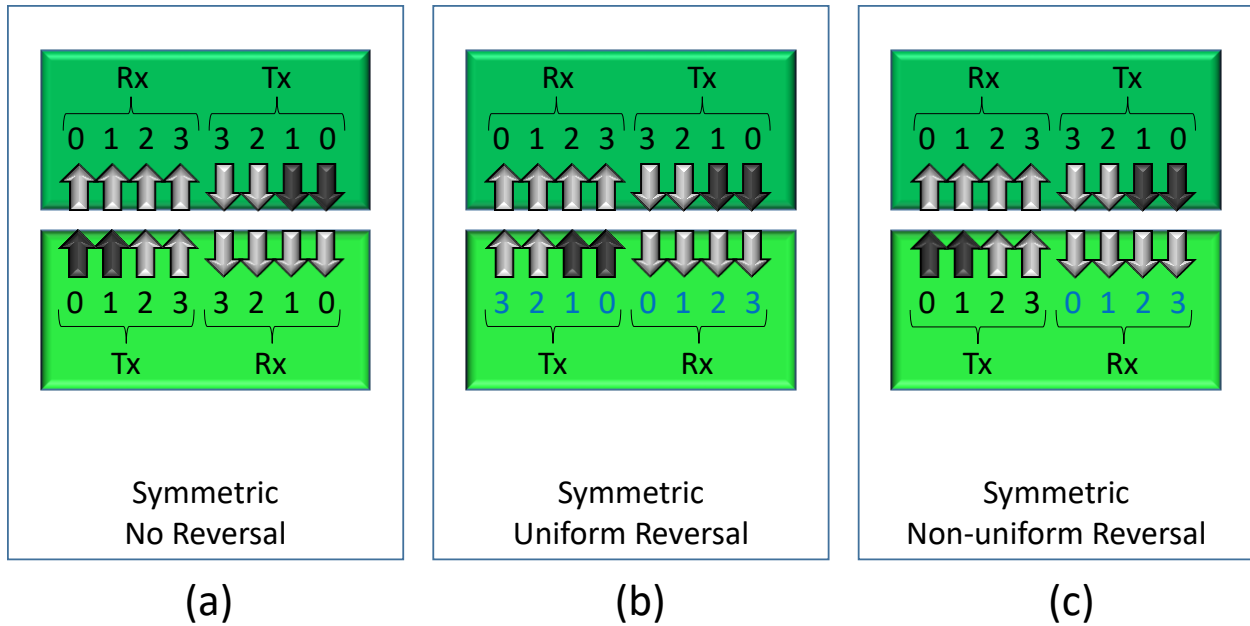
- The physical layer shall enable its transmitter (Tx) on the lower lanes[N/2-1:0] only, where N is the total number of transmitting lanes, and shall enable the receiver (Rx) detection circuit on all lanes as illustrated in *Gen-Z-E-NRZ-25G Symmetric x4 Detect* with the Detect Tx lanes darkened. Lanes[N/2-1:0] transmitters shall continuously transmit Signal Detect Training Sequences (SDTS) as specified as in *PHY Signal Detect Training Sequences* for 2 ms to signal to the remote physical layer.
- The physical layer shall put the upper lanes[N:N/2] of its transmitter in an electrical idle condition, where N is the total number of transmitting lanes.
- The receiver detection circuits shall monitor the local receive pins for the SDTS transmitted by the remote physical layer.
- The Rx lanes shall detect endpoint connectivity and shall use this information to determine if lane reversal is present in the connection between the physical layers.
- If any lower Rx lane[N/2-1:0] detects a signal, then the link in that direction shall not be reversed as illustrated in both upper and lower physical layers of (a) and the upper physical layer of (c) from *Gen-Z-E-NRZ-25G Symmetric x4 Detect Examples*. However, if an upper Rx lane[N-1:N/2] detects a signal, then the link in that direction shall be reversed as illustrated in both upper and lower physical layers of (b) and the lower physical layer of (c) from *Gen-Z-E-NRZ-25G Symmetric x4 Detect Examples*.
- If the local Rx detection circuits do not detect the SDTS within 2 ms, then the PHY shall transition to the Detect Idle sub-state. If the detection circuits identifies the SDTS before the 2 ms timer expires, then the PHY shall transition to the Detect Wait Reversal sub-state.

10

15

20

Developer Note: For physical layer link widths greater than 3 lanes, the transmission of the Signal Detect Training Sequences on all lanes[N/2-1:0] provides a redundant path to link width reduction in the case lane0 has faults.



= DETECT lane(s) during PHY training

Figure 5-4: Gen-Z-E-NRZ-25G Symmetric x4 Detect

Detect Idle

The Detect Idle sub-state can reduce the physical layer power consumption while the logical sublayer awaits an inbound signal from the remote physical layer. In the Detect Idle sub-state, the transmitter circuits shall be powered down, and the receiver detect circuits shall be enabled. The physical layer shall stay in this sub-state until the receiver detects SDTS or a 98 ms timeout expires. If SDTS are detected, then the Detect Wait Reversal sub-state shall be entered. If the 98ms timer expires, then the Detect Send sub-state shall be entered.

Detect Wait Reversal

The physical layer shall enter the Detect Wait Reversal sub-state after the local receiver has identified SDTS from the remote transmitter. The purpose of this sub-state is to allow the remote receiver time to detect a signal and determine the reversal state of the link. The transmitters shall continue to transmit SDTS for 8 ms on the lower half of lanes of the link, and then the physical layer shall transition to the Detect Wait sub-state.

Developer Note: A receiver's detection circuit can be implementation specific. Given all transmitters may not be sending information until the Detect Wait sub-state, receivers should be idle during the Detect Wait Reversal and Detect Wait sub-states.

Detect Wait

The physical layer shall enable all transmitters on the link and send SDTS upon transitioning into the Detect Wait sub-state. The purpose of this sub-state is to ensure all transmitters are sending information before enabling the receivers in the next PLTSM state. The transmitters shall continue to

transmit SDTS for 8 ms on all lanes, and then the physical layer shall transition to the next state of the PLTSM.

Developer Note: *Gen-Z Signal Detect Example between Two Endpoints illustrates an example handshake between two physical layers as they transition through the Detect State (see Gen-Z Signal Detect Sub-State Diagram). In Gen-Z Signal Detect Example between Two Endpoints, time advances from top to bottom. The communication between Endpoint 1 (EP1) and Endpoint 2 (EP2) is represented by arrows originating with the endpoint that transmits a Signal Detect Training Sequence and terminating at the reception endpoint. Gen-Z Signal Detect Sub-State Diagram also illustrates the physical layer training sub-state of each endpoint as it transitions from PHY Idle through Detect and on to the Channel Optimize state.*

On initial power up (or explicit direction by the Gen-Z Core) the two Gen-Z physical layer endpoints begin in the PHY Idle state. In Gen-Z Signal Detect Example between Two Endpoints, the Gen-Z Core directs Endpoint 1 to exit the PHY Idle state, and the logical sublayer enters the Detect Send sub-state. Upon entering the Detect Send sub-state, EP1 repeatedly transmits SDTS on the lower lanes to EP2 for 2 ms. The Detect Header is formatted for the purpose of detect, containing a low frequency signal which can trigger a slow speed data pattern detection circuit in the remote receiver.

In the Gen-Z Signal Detect Example between Two Endpoints, note that EP2 is still in PHY Idle when EP1 transmits the first SDTS. Therefore, EP2 does not respond, and EP1 enters Detect Idle sub-state. After waiting in Detect Idle for 98 ms, EP1 enters the Detect Send sub-state again and continuously transmits SDTS for 2 ms. EP2 is still in PHY Idle and does not respond.

When the Gen-Z Core requests PHY-Up in EP2, Endpoint 2 transitions to the Detect Send sub-state and begins transmitting SDTS to EP1 on the lower lanes. EP1 is in the Detect Idle sub-state with its signal detection circuit is enabled. EP1 detects the low-frequency signal in the SDTS from EP2 and EP1 and advances to the Detect Wait Reversal sub-state without waiting the full 98 ms of Detect Idle time. EP2 is still in Detect Send when the SDTS from EP1 trigger the slow-speed data detection circuit. EP2 advances to the Detect Wait Reversal sub-state.

In the Detect Wait Reversal sub-state, both endpoints have detected the presence of the remote physical layer, and after transmitting SDTS for 8 ms on the lower lanes, advance to the Detect Wait sub-state.

In the Detect Wait sub-state, both endpoints transmit SDTS on all lanes for 8 ms and advance to the Next Physical Layer Training State.

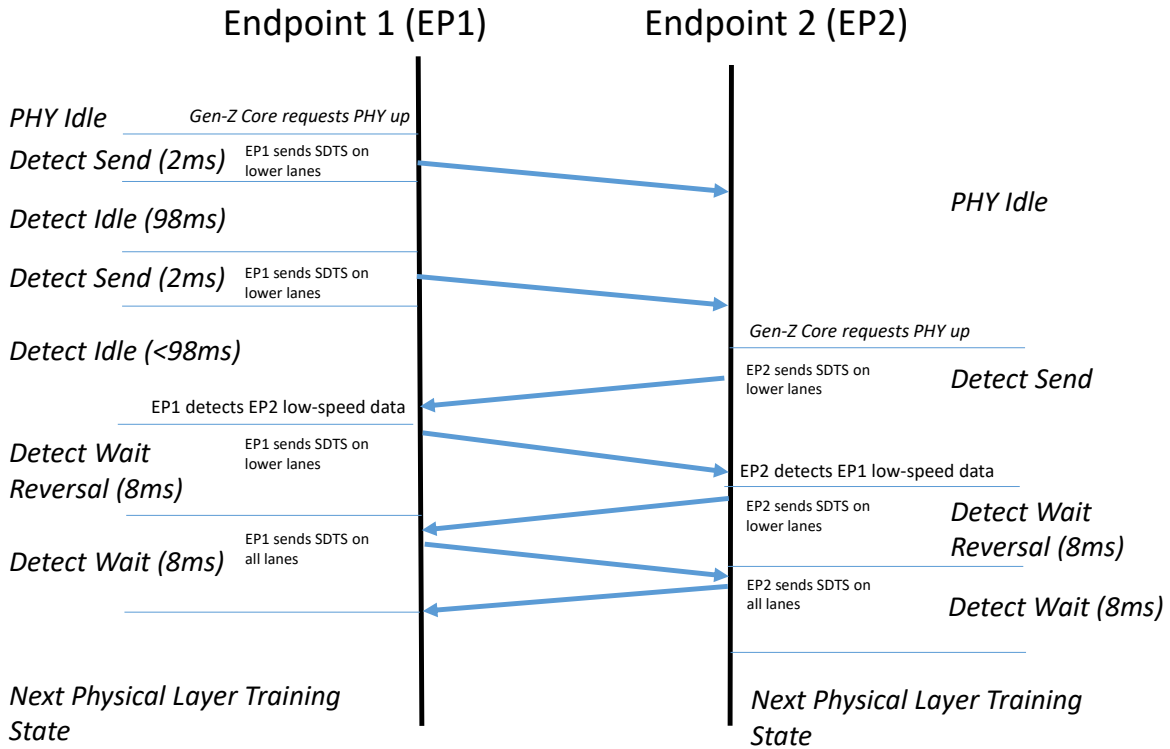


Figure 5-5: Gen-Z Signal Detect Example between Two Endpoints

5.2.11.2.3. Signal Detect Exit

The PLTSM shall exit the Signal Detect state after the following conditions have been met:

- Signal Detect Training Sequences from remote physical layer been detected by the local physical layer.
- After transmitting SDTS on all lanes for 8 ms in the Detect Wait sub-state.

Upon exiting the Signal Detect state, the PLTSM shall transition to the Channel Optimize state.

5.2.11.3. Channel Optimize (Required)

Channel Optimize shall be a required state for this individual physical layer clause. Channel Optimize specifies an in-band parameter exchange between physical layers to enable transmitter and receiver equalization adaptation.

Developer Note: In closed systems where fixed equalization settings can be statically programmed prior to physical layer training, the Interface PHY Structure provides an option to bypass the Channel Optimize state.

5.2.11.3.1. Channel Optimize Entry

On entry both physical layer transmitters and receivers shall be enabled. This state is reached from the Signal Detect state.

5.2.11.3.2. Channel Optimize Process

Refer to the following topics and clauses from the *IEEE Standard for Ethernet 802.3™-2015* Specification for the Channel Optimize process:

- KR PMD (in-band exchange protocol): clause 72.6
- Tx Equalization: clause 93.8

5.2.11.3.3. Channel Optimize Exit

Upon completion of the Channel Optimize state, the independent lane equalization settings shall be used by the transmitters and receivers and the PLTSM shall transition to the Align state.

5.2.11.4. *Align (Required)*

Refer to the Gen-Z Common Specification clause section on *Align (Required)*.

5.2.11.5. *Interconnect BIST (Optional)*

Refer to the Gen-Z Common Specification clause section on *Interconnect BIST (Optional)*.

5.2.11.6. *Config (Required)*

Refer to the Gen-Z Common Specification clause section on *Config (Required)*.

5.2.11.7. *PHY-Up (Required)*

Refer to the Gen-Z Common Specification clause section on *PHY-Up (Required)*.

5.2.12. Clock Compensation

To compensate for a maximum clock tolerance offset of 200 ppm (± 100 ppm) with this individual physical layer clause, the number of consecutive clock compensation PHY-Up Sequences and their frequency shall be as specified in the *PHY Specific Control Fields* of the *Unique Interface Physical Layer Structure*.

Refer to the Gen-Z Common Specification clause section on *Clock Compensation* for rules and policies on transmitting and receiving, including addition or removal, of clock compensation PHY-Up Sequences.

Developer Note: For closed systems with known clocking topologies, options exist in the *PHY Specific Control Fields* of the *Unique Interface Physical Layer Structure* to disable, increase, or decrease the use of clock compensation PHY-Up Sequences. For example in a system where both the local and remote physical layer share a common reference source, the use of clock compensation PHY-Up Sequences can be disabled to reduce physical layer overhead and increase Gen-Z Core data performance.

5.2.13. Loopback

5.2.13.1. Master Loopback (TX → RX)

An implementation may support master loopback but this individual physical layer clause shall not specify a configuration setting or a physical layer state, including entry and exit to that state.

5 **Developer Note:** *The PLTSM does not prevent a master loopback connection at the external pins. It is recommended that implementations support and test this topology.*

5.2.13.2. Slave Loopback (Rx → TX)

Though an implementation may support slave loopback, this physical layer clause does not specify a configuration setting or a physical layer state, including entry and exit to that state.

10 5.2.14. Physical Layer Aggregation

An implementation may support static aggregation of multiple physical layers into a single aggregated interface (SAI).

5.2.15. Re-timers

15 Re-timers capable of implementing the PLTSM specified in this individual physical layer clause are considered protocol-aware. Protocol-aware re-timers can be used in topologies where the optimum equalization coefficients are not known prior to initialization and extended channel reach is required.

Developer Note: *Generic re-timers, which are not Gen-Z physical layer protocol-aware, can be used in topologies where equalization coefficients can be statically set by manageability or software and the Channel Optimize state can be bypassed.*

20 5.3. Power Management

5.3.1. Overview

Physical layers implementing this individual physical layer clause may support the PHY-LP states P1 (standby) and P2 (sleep), and the PHY-Up LP (LWR) state, dynamic link width reduction. Form factors supporting this individual physical layer clause may require low power states.

25 5.3.2. Low Power Operation

Refer to the Gen-Z Common Specification clause section on *Low Power Operation*.

5.3.3. Physical Layer Power Management States

Each low power state provides tradeoffs in power savings versus entrance and exit latencies as specified in *Gen-Z-E-NRZ-25G-Fabric Low Power States*.

Table 5-5: Gen-Z-E-NRZ-25G-Fabric Low Power States

PHY Power State	PLL Status	TX/RX Status	Data Transmission	Exit Latency Target	Power Savings Target
PHY LP1: P1	On	Off	No	5-50 μ s	Medium
PHY LP2: P2	Off	Off	No	<500 μ s	High
PHY-Up LP1: LWR	On	On + Idle	Yes	1–5 μ s	Low to Medium

5.3.3.1. PHY-LP Sub-States

5.3.3.1.1. P1 (Standby)

Physical layers implementing this individual physical layer clause may support the PHY-LP P1 state. If supported, the P1 state shall map to the PLA LP 1 state and shall be governed by the rules specified in the Gen-Z Common Specification clause section on *Low Power Operation*.

5.3.3.1.2. P2 (Sleep)

Physical layers implementing this individual physical layer clause may support the PHY-LP P2 state. If supported, the P2 state shall map to the PLA LP 2 state and shall be governed by the rules specified in the Gen-Z Common Specification clause section on *Low Power Operation*.

5.3.3.2. PHY-Up LP Sub-States

5.3.3.2.1. Dynamic Link Width Reduction (LWR)

Physical layers implementing this individual physical layer clause may support the dynamic link width reduction PHY-Up LP state. If supported, then the dynamic LWR state shall map to the PLA PHY-Up LP 1 state and shall be governed by the rules specified in the Gen-Z Common Specification clause section on *Low Power Operation*.

The following are the features and requirements associated with the dynamic LWR PHY-Up LP state entrance:

- Entrance to the dynamic low power state shall require retraining. Retrain Arc 2 as specified in *Retraining Arcs* should be implemented for entrance into this low-power state.
- After transmission of 4 PHY-Up LP sequences with (LP_State=PHY-Up LP1) and reception of 2 PHY-Up LP sequences with (LP_State=PHY-Up LP1) as specified in *PHY-Up Sequences*, the physical layer shall transmit PHY-Up LP sequences with (LP_State=PHY-Up LP1) for 100 ns on all lanes and shall then begin retraining with the associated reduced lanes disabled.
- The PHY-Up LP sequence with (LP_State = PHY-Up LP1) indicates retraining so PHY-Up Retraining sequences shall not be transmitted.
- Transmit equalization settings for disabled lanes shall be preserved for use on exit.

The following are the features and requirements associated with the dynamic LWR PHY-Up LP state exit:

- Exit from the dynamic LWR state shall require retraining. Retrain Arc 1 as specified in *Retraining Arcs* shall be required for exit from this low power state.
- After transmission of 4 PHY-Up sequences with (LP_State=PHY-Up) and reception of 2 PHY-Up LP sequences with (LP_State=PHY-Up) as specified in *PHY-Up Sequences*, the physical layer shall enable disabled lanes and shall transmit PHY-Up sequences with (LP_State=PHY-Up) for 500 ns on all lanes.
- Re-enabled lanes shall use their previously preserved transmit equalization settings.
- The PHY-Up LP sequence with (LP_State=PHY-Up) indicates retraining, hence PHY-Up Retraining sequences shall not be transmitted.

5.4. Electrical Sublayer

This Gen-Z physical layer electrical sublayer shall meet all single lane/channel requirements specified in *IEEE Standard for Ethernet 802.3by™-2016* for the following clause:

- Clause 111. Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-KR and 25GBASE-KR-S

Many requirements of *IEEE Standard for Ethernet 802.3by™-2016* Clause 111 reference *IEEE Standard for Ethernet 802.3™-2015* for the following clauses when related to the operation of a single lane/channel:

- Clause 93. Physical Medium Dependent (PMD) sublayer and baseband medium, type 100GBASE-KR4
- Annex 93A (normative) Specification methods for electrical channels
- Annex 93B (informative) Electrical backplane reference model
- Annex 93C (normative) Receiver interference tolerance

5.4.1. Electrical Interface

5.4.1.1. Overview

The electrical interface for this physical layer clause shall be based on high-speed, low-voltage differential signaling where each connection is point-to-point and signaling is unidirectional.

5.4.1.1.1. Forwarded Clock

This physical layer clause shall not require a forwarded clock.

5.4.1.2. Lane Signaling Rate

The signaling rate for this physical layer clause shall be as specified in *IEEE Standard for Ethernet 802.3™-2015*, clause 93.8.1.2 at 25.78125 GBd/s \pm 100 Parts Per Million (ppm) per lane.

5.4.1.3. Lane Bit Error Rate

This physical layer clause shall operate in the no-FEC mode as specified by guideline 'c' in *IEEE Standard for Ethernet 802.3by™-2016*, clause 111.1 with an overall bit error ratio (BER) requirement of 10^{-12} or better.

5.4.1.4. **Lane Equalization**

Tx Equalization coefficient range and step size requirements shall be as specified in *IEEE Standard for Ethernet 802.3™-2015* clause 93.8.

5.4.2. **High Speed Signaling Specification**

5 5.4.2.1. **Transmitter Specification**

The transmitter for this physical layer clause shall be as specified for a single lane in *IEEE Standard for Ethernet 802.3™-2015*, clause 93.8.1 with the same exception noted in *IEEE Standard for Ethernet 802.3by™-2016*, clause 111.8.2.

5.4.2.2. **Receiver Specification**

10 The receiver for this physical layer clause shall be as specified for a single lane in *IEEE Standard for Ethernet 802.3™-2015*, clause 93.8.2 with additional requirements specified in *IEEE Standard for Ethernet 802.3by™-2016*, clause 111.8.3 with an exception that no-FEC mode shall be supported for test requirements as specified in Table 111-6.

5.4.3. **Channel Characteristics**

15 5.4.3.1. **Channel Operating Margin**

The Channel Operating Margin (COM) for this physical layer clause shall be as specified in *IEEE Standard for Ethernet 802.3by™-2016*, clause 111.9.2 for no-FEC mode. The channels for this physical layer clause shall meet this COM requirement with the Target Detector Error Ratio (DER_0) parameter set to 10^{-12} .

5.4.3.2. **Insertion Loss**

20 The Insertion Loss for this physical layer clause shall be as specified in *IEEE Standard for Ethernet 802.3™-2015*, clause 93.9.2.

5.4.3.3. **Return Loss**

The Return Loss for this physical layer clause shall be as specified in *IEEE Standard for Ethernet 802.3™-2015*, clause 93.9.3.

25 5.4.3.4. **Coupling**

The transmitter for this physical layer clause shall be AC-coupled to the receiver as specified in *IEEE Standard for Ethernet 802.3™-2015*, clause 93.9.4.

5.4.4. Miscellaneous Specification

5.4.4.1. Reference Clock

The clock frequency source of the transmitter for this physical layer clause shall achieve a baud rate with an offset of ± 100 Parts Per Million (ppm) relative to the baud rate of the receiver.

- 5 **Developer Note:** Though the reference clock frequency is not explicitly defined, common reference clock frequencies for physical layers of this type are 312.5 MHz and 156.25 MHz.

5.5. Management Control and Status

5.5.1. Interface Physical Layer Structure

- 10 The Interface PHY Structure is specified generically in the *Gen-Z Core Specification* to cover multiple physical layers. The mapping for the Interface PHY Structure associated with this individual physical layer clause is specified below.

5.5.1.1. Common Interface Physical Layer Structure

5.5.1.1.1. Common, Lane, and Low Power Interface PHY Structure

- 15 The following specifies the features and requirements associated with the Gen-Z-E-NRZ-25G-Fabric specification. The *Gen-Z-E-NRZ-25G-Fabric Interface PHY Structure Support Exclusions and Restrictions* specifies restricted or excluded fields or sub-fields from the *Interface PHY Structure Common, Lane, and Low Power Fields*.

Table 5-6: Gen-Z-E-NRZ-25G-Fabric Interface PHY Structure Support Exclusions and Restrictions

Support	Sub-field / Description	Interface PHY Structure Field	Supported Value(s)
Restricted	PHY Clause 5—25G Fabric	PHY Type	0x0
Restricted	Physical Layer Operational Status 0x0—PHY-Down, uninitialized state 0x1—PHY-Up 0x2—PHY-Down-Retrain 0x3–0x6—Reserved 0x7—PHY LP 1 → Mapped to P1 0x8—PHY LP 2 → Mapped to P2 0x9–0xA Reserved 0xB—PHY-Up LP 1 → Mapped to Dynamic Link Width Reduction 0xC–0xF—Reserved	PHY Status	See Description
Restricted	Previous Physical Layer Operational Status 0x0—PHY-Down, uninitialized state	PHY Status	See Description

Support	Sub-field / Description	Interface PHY Structure Field	Supported Value(s)
	0x1—PHY-Up 0x2—PHY-Down-Retrain 0x3-0x6—Reserved 0x7—PHY LP 1 → Mapped to P1 0x8—PHY LP 2 → Mapped to P2 0x9-0xA—Reserved 0xB—PHY-Up LP 1 → Mapped to Dynamic Link Width Reduction 0xC-0xF—Reserved		
Restricted	Physical Layer Retraining Arc Retraining Arc 1 shall be supported, and Retraining Arc 2 should be supported for this physical layer clause.	PHY Control	0x0-0x1
Excluded	Retrain Arc 3-4 Support	PHY Cap 1	0b
Restricted	Enable Phit CRC Encoding Phit CRC encoding shall be enabled by default for this physical layer clause	PHY Cap 1 Control	1b (Default)
Restricted	Asymmetric Lane Status Asymmetry shall not be supported for this physical layer clause	PHY Lane Status	0x0
Restricted	Enable Lane Asymmetry	PHY Lane Control	0x0
Restricted	Asymmetric Lane Support	PHY Lane Cap	0x0
Restricted	Reversal Support Only uniform reversal shall be supported for this physical layer clause	PHY Lane Cap	0x0-0x2
Restricted	Asymmetric Lane with Reversal Support	PHY Lane Cap	0b
Restricted	Remote Asymmetric Lane Support	PHY Remote Lane Cap	0x0
Restricted	Remote Reversal Support Only uniform reversal shall be supported for this physical layer clause	PHY Remote Lane Cap	0x0-0x2

Support	Sub-field / Description	Interface PHY Structure Field	Supported Value(s)
Restricted	Remote Asymmetric Lane with Reversal Support	PHY Remote Lane Cap	0b
Excluded	Entry, Exit Latency PHY-LP 3-4	PHY Low Power Timing Capability	0x0
Excluded	PHY-LP 3-4 Support	PHY Low Power CAP	0b
Excluded	Entry, Exit Latency PHY-Up-LP 2-4	PHY-Up Low Power Timing Capability	0x0
Excluded	PHY-Up-LP 2-4 Support	PHY-Up Low Power CAP	0b

5.5.1.1.2. Extended Feature Interface PHY Structure

The following specifies the features and requirements associated with the Interface PHY Structure for Extended Feature fields for the Gen-Z-E-NRZ-25G-Fabric specification.

Table 5-7: Gen-Z-E-NRZ-25G-Fabric Extended Feature Fields

Field Name	Size (bits)	Value / Bit Location	M / O	Access	Description
PHY Extended Status	32	-	M	RO	See <i>Extended Feature Status</i>
PHY Extended Control	32	-	M	RW	See <i>Extended Feature Control</i>
PHY Extended Cap	32	-	M	RO	See <i>Extended Feature Capability</i>
PHY Remote Extended Cap	32	-	M	RW	See <i>Extended Feature Remote Capability</i>

Table 5-8: Gen-Z-E-NRZ-25G-Fabric Extended Feature Status

Bit Location	Access	Description
0	RO	Physical Layer Extended Feature Operational Status 0b—PHY Extended Features disabled 1b—PHY Extended Features enabled
2:1	RO	Physical Layer Extended Feature Training Status 0x0—Training with Extended Features has not occurred 0x1—Training with Extended Features succeeded 0x2—Training with Extended Features has failed 0x3—Reserved
5:3	RO	Current Forward Error Correction (FEC) Status 0x0—No-FEC 0x1—BASE-R FEC (802.3 Clause 74) 0x2—RS-FEC (802.3 Clause 91, 108) 0x3-0x7—Reserved
8:6	RO	Current Transfer Rate Status 0x0—25.78125 GT/s 0x1—20.0 GT/s 0x2—28.125 GT/s 0x3-0x7—Reserved
31:9	-	RsvdZ

Table 5-9: Gen-Z-E-NRZ-25G-Fabric Extended Feature Control

Bit Location	Access	Description
2:0	RW	Forward Error Correction (FEC) Mode—Determines the active FEC type when Extended Features are enabled. 0x0—No-FEC (Default) 0x1—BASE-R FEC (802.3 Clause 74) 0x2—RS-FEC (802.3 Clause 91, 108) 0x3-0x7—Reserved
5:3	RW	Transfer Rate Mode—Determines the active transfer rate when Extended Features are enabled. 0x0—25.78125 GT/s (Default) 0x1—20.0 GT/s 0x2—28.125 GT/s

Bit Location	Access	Description
31:6		0x3-0x7—Reserved
	-	RsvdP

Table 5-10: Gen-Z-E-NRZ-25G-Fabric Extended Feature Capability

Bit Location	Access	Description
0	RO	BASE-R FEC (802.3 Clause 74) Support
1	RO	RS-FEC (802.3 Clause 91, 108) Support
3:2	-	RsvdZ
4	RO	20.0 GT/s Transfer Rate Support
5	RO	28.125 GT/s Transfer Rate Support
31:6	-	RsvdZ

Table 5-11: Gen-Z-E-NRZ-25G-Fabric Extended Feature Remote Capability

Bit Location	Access	Description
0	RO	BASE-R FEC (802.3 Clause 74) Support
1	RO	RS-FEC (802.3 Clause 91, 108) Support
3:2	-	RsvdZ
4	RO	20.0 GT/s Transfer Rate Support
5	RO	28.125 GT/s Transfer Rate Support
31:6	-	RsvdZ

5.5.1.2. ***Unique Interface Physical Layer Structure***

The following specifies the features and requirements associated with the Interface PHY Structure for PHY-specific fields for the Gen-Z-E-NRZ-25G-Fabric specification.

Interface PHY Structure: PHY-Specific Configuration Space

+7		+6		+5		+4		+3		+2		+1		+0																																									
7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
PHY Specific Control								PHY Specific Status								< Byte 0x60																																							
PHY Specific Remote Capability								PHY Specific Capability								< Byte 0x68																																							
R6								R5				PHY TX Lane Index				< Byte 0x70																																							
PHY TX Lane Control								PHY TX Lane Status								< Byte 0x78																																							
PHY RX Lane Margin								R7				PHY RX Lane Index				< Byte 0x80																																							
PHY RX Lane Control								PHY RX Lane Status								< Byte 0x88																																							
R9								R8				PHY State Index				< Byte 0x90																																							
Remote PHY State Capability				PHY State Capability				PHY State Control				PHY State Status				< Byte 0x98																																							

Figure 5-6: Gen-Z-E-NRZ-25G-Fabric PHY Specific Fields

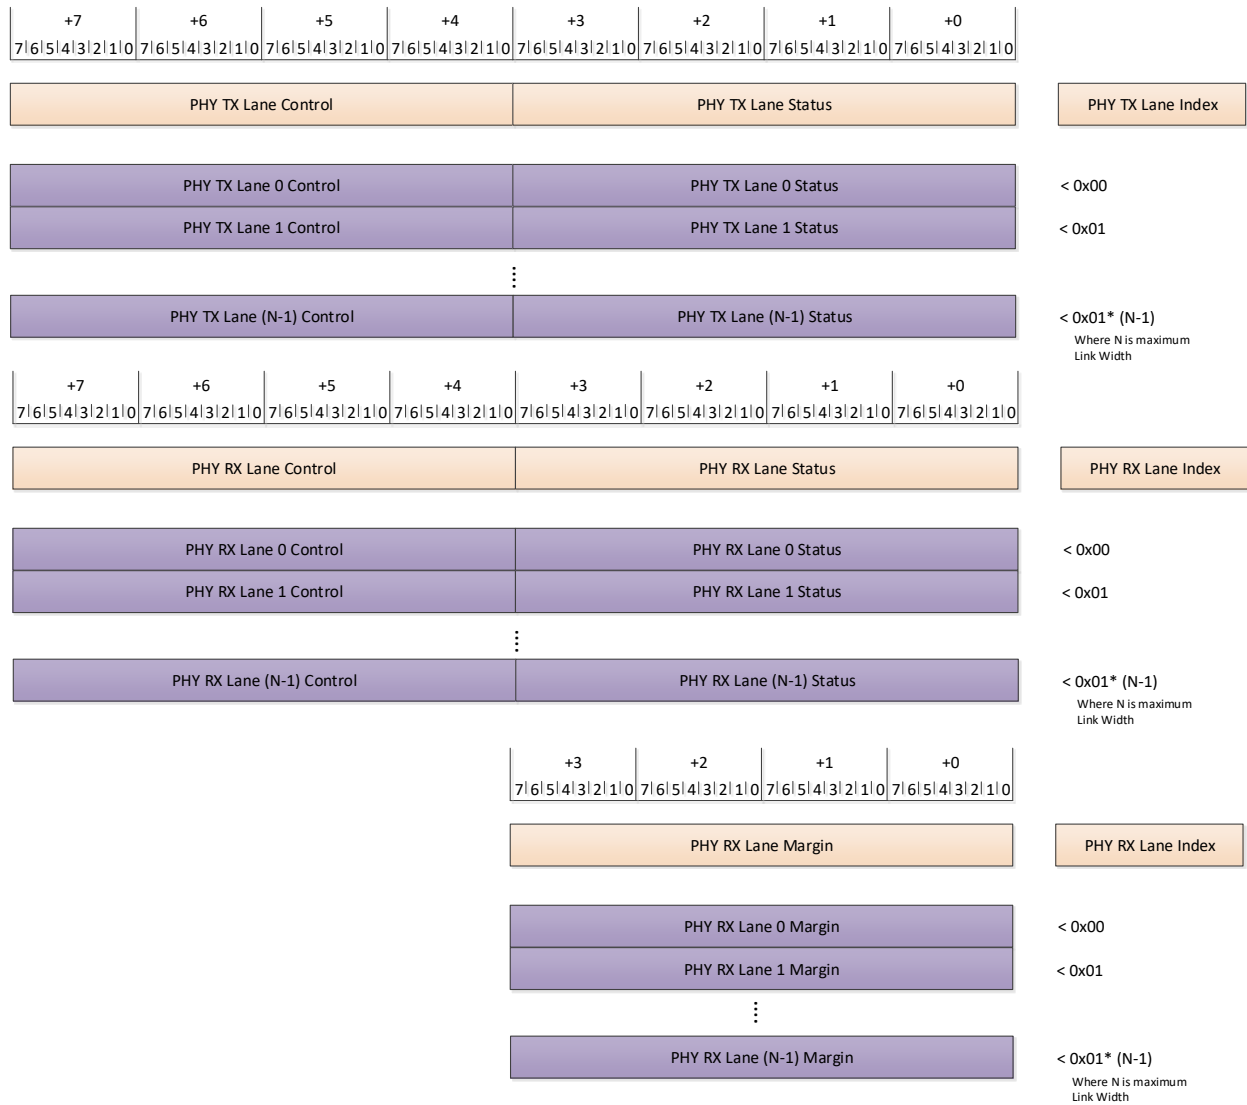


Figure 5-7: Gen-Z-E-NRZ-25G-Fabric PHY Specific Field Lane Indexing

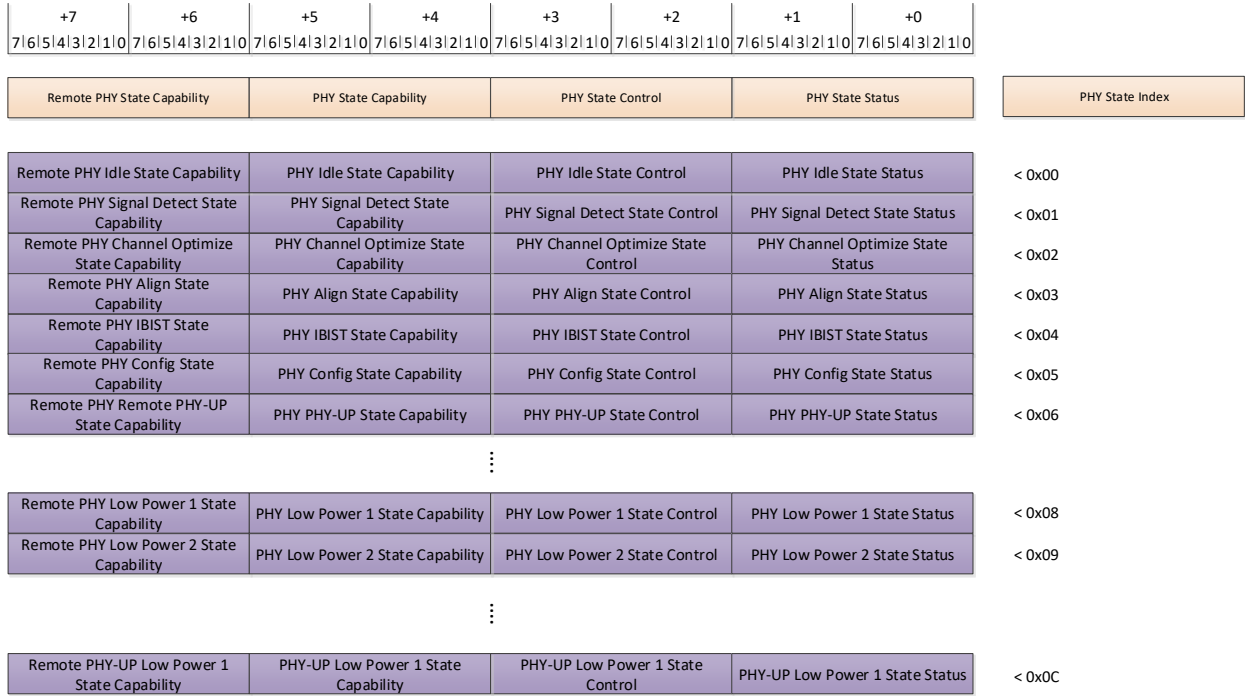


Figure 5-8: Gen-Z-E-NRZ-25G-Fabric PHY Specific Field State Indexing

Table 5-12: Gen-Z-E-NRZ-25G-Fabric PHY Specific Fields

Field Name	Size (bits)	Value / Bit Location	M / O	Access	Description
PHY Specific Status	32	-	M	RO	See <i>PHY Specific Status Fields</i>
PHY Specific Control	32	-	M	RW	See <i>PHY Specific Control Fields</i>
PHY Specific Capability	32	-	M	RO	See <i>PHY Specific Capability Fields</i>
PHY Remote Specific Capability	32	-	M	RW	See <i>Remote PHY Specific Capability Fields</i>
PHY TX Lane Index	12	-	M	RW	This field index selects the lane number for the PHY TX Lane Status and PHY TX Lane Control interface fields as illustrated in <i>Gen-Z-E-NRZ-25G-Fabric PHY Specific Field Lane Indexing</i> .

Field Name	Size (bits)	Value / Bit Location	M / O	Access	Description
PHY TX Lane Status	32	-	M	RO	See <i>PHY TX Lane Status Fields</i> Developer Note: Recommend that the <i>Physical Layer State</i> field be read to validate the current PLTSM state. The current state determines if the fields in this structure are valid.
PHY TX Lane Control	32	-	M	RW	See <i>PHY TX Lane Control Fields</i>
PHY Rx Lane Index	12	-	M	RW	This field index selects the lane number for the PHY Rx Lane Margin, PHY Rx Lane Status, and PHY Rx Lane Control interface fields as illustrated in <i>Gen-Z-E-NRZ-25G-Fabric PHY Specific Field Lane Indexing</i> .
PHY Rx Lane Margin	32	-	M	RW	See <i>PHY Rx Lane Margin Fields</i> Developer Note: Recommend that the <i>Physical Layer State</i> field be read to validate the current PLTSM state. The current state determines if the fields in this structure are valid.
PHY Rx Lane Status	32	-	M	RO	See <i>PHY Rx Lane Status Fields</i> Developer Note: Recommend that the <i>Physical Layer State</i> field be read to validate the current PLTSM state. The current state determines if the fields in this structure are valid.
PHY Rx Lane Control	32	-	M	RW	See <i>PHY Rx Lane Control Fields</i>
PHY State Index	16	-	M	RW	This field index selects the lane number for the PHY State Status, PHY State Control, PHY State Capability, and Remote PHY State Capability interface fields as illustrated in <i>Gen-Z-E-NRZ-25G-Fabric PHY Specific Field State Indexing</i> . PHY State Index mapping: 0x0—Idle State

Field Name	Size (bits)	Value / Bit Location	M / O	Access	Description
					0x1—Signal Detect State 0x2—Channel Optimize State 0x3—Align State 0x4—IBIST State 0x5—Config State 0x6—PHY-Up State 0x7—Reserved 0x8—LP 1 State 0x9—LP 2 State 0xA-0xB—Reserved 0xC—PHY-Up LP 1 State 0xD-0xFFFF—Reserved
PHY State Status	16	-	M	RO	Indicates status for the following states based on the value of the PHY State Index field—For individual details on each state, see: <i>PHY Idle State Status Fields</i> <i>PHY Signal Detect State Status Fields</i> <i>PHY Channel Optimize State Status Fields</i> <i>PHY Align State Status Fields</i> <i>PHY IBIST State Status Fields</i> <i>PHY Config State Status Fields</i> <i>PHY PHY-Up State Status Fields</i> <i>PHY LP 1 State Status Fields</i> <i>PHY LP 2 State Status Fields</i> <i>PHY-Up LP 1 State Status Fields</i> Developer Note: Recommend that the Physical Layer State field be read to validate the PLTSM has progressed to or beyond the observed state to determine if the fields in this structure are valid.
PHY State Control	16	-	M	RW	Determines control for the following states based on the value of the PHY State Index field—For individual details on each state, see: <i>PHY Idle State Control Fields</i> <i>PHY Signal Detect State Control Fields</i> <i>PHY Channel Optimize State Control Fields</i> <i>PHY Align State Control Fields</i> <i>PHY IBIST State Control Fields</i> <i>PHY Config State Control Fields</i> <i>PHY PHY-Up State Control Fields</i> <i>PHY LP 1 State Control Fields</i>

Field Name	Size (bits)	Value / Bit Location	M / O	Access	Description
PHY State Capability					<i>PHY LP 2 State Control Fields</i> <i>PHY-Up LP 1 State Control Fields</i>
	16	-	M	RO	Indicates capability features for the following states based on the value of the PHY State Index field—For individual details on each state, see: <i>PHY Idle State Capability Fields</i> <i>PHY Signal Detect State Capability Fields</i> <i>PHY Channel Optimize State Capability Fields</i> <i>PHY Align State Capability Fields</i> <i>PHY IBIST State Capability Fields</i> <i>PHY Config State Capability Fields</i> <i>PHY PHY-Up State Capability Fields</i> <i>PHY LP 1 State Capability Fields</i> <i>PHY LP 2 State Capability Fields</i> <i>PHY-Up LP 1 State Capability Fields</i>
Remote PHY State Capability	16	-	M	RW	Indicates remote PHY capability features for the following states based on the value of the PHY State Index field—For individual details on each state, see: <i>Remote PHY Idle State Capability Fields</i> <i>Remote PHY Signal Detect State Capability Fields</i> <i>Remote PHY Channel Optimize State Capability Fields</i> <i>Remote PHY Align State Capability Fields</i> <i>Remote PHY IBIST State Capability Fields</i> <i>Remote PHY Config State Capability Fields</i> <i>Remote PHY PHY-Up State Capability Fields</i> <i>Remote PHY LP 1 State Capability Fields</i> <i>Remote PHY LP 2 State Capability Fields</i> <i>Remote PHY-Up LP 1 State Capability Fields</i>
R5	20	-	-	-	RsvdP
R6	32	-	-	-	RsvdP
R7	20	-	-	-	RsvdP
R8	16	-	-	-	RsvdP

Field Name	Size (bits)	Value / Bit Location	M / O	Access	Description
R9	32	-	-	-	RsvdP

Table 5-13: PHY Specific Status Fields

Bit Location	Access	M/O	Description
3:0	RO	M	Physical Layer State—Indicates current state of the physical layer 0x0—PHY Idle 0x1—Signal Detect 0x2—Channel Optimize 0x3—Align 0x4—Config 0x5—IBIST (optional) 0x6—PHY-Up 0x7—Reserved 0x8—LP 1 0x9—LP 2 0xA-0xB—Reserved 0xC—PHY-Up LP 1 0xD-0x1F—Reserved
7:4	RO	O	Physical Layer Sub-state—See <i>PHY Sub-state Mapping</i>
11:8	RO	M	Previous Physical Layer State—If automatic re-initialization control is enabled, then these bits shall indicate the previous final state of the physical layer prior to reattempting physical layer training. The values are equivalent to those of the Physical Layer State field.
15:12	RO	O	Previous Physical Layer Sub-state—If automatic re-initialization control is enabled, then this field indicates the previous final sub-state of the physical layer prior to reattempting physical layer training. The values are equivalent to those of the Physical Layer Sub-state field.
31:16	RO	-	RsvdZ

Table 5-14: PHY Sub-state Mapping

PHY State	PHY State Mapping	PHY Sub-state	PHY Sub-state Mapping	Notes
PHY Idle	0x0	n/a	-	
Channel Optimize	0x1	n/a	-	
Signal Detect	0x2	Detect Send	0x0	Transmitting Signal Detect Training Sequences
		Detect Idle	0x1	Off
		Detect	0x2	Exiting Signal Detect
Align	0x3	Bit Lock	0x0	Locking to Data
		Equalization	0x1	Rx functions for equalization
		Symbol Lock	0x2	Locking to parallel symbols
		Deskew	0x3	Lane-to-lane deskew
IBIST	0x4	n/a	-	
Config	0x5	n/a	-	
PHY-Up	0x6	n/a	-	
Low Power 1	0x8	PHY LP1 Entry	0x0	Transitioning to PHY LP1
		PHY LP1 Exit	0x1	Transitioning from PHY LP1
		PHY LP1	0x2	In PHY LP1
Low Power 2	0x9	PHY LP2 Entry	0x0	Transitioning to PHY LP2
		PHY LP2 Exit	0x1	Transitioning from PHY LP2
		PHY LP2	0x2	In PHY LP2
PHY-Up Low Power 1	0xC	PHY-Up LP1 Entry	0x0	Transitioning to PHY-Up LP1
		PHY-Up LP1 Exit	0x1	Transitioning from PHY-Up LP1
		PHY-Up LP1 Active	0x2	In PHY-Up LP1

Table 5-15: PHY Specific Control Fields

Bit Location	Access	M/O	Description
1:0	RW	M	PHY-Up Clk Comp Sequence Count—Determines the number of consecutive PHY-Up Clk Comp sequences to transmit to enable a physical layer to autonomously compensate for clock differences with the remote PHY. 0x0—1 0x1—2 0x2—3 (Default) 0x3—4
3:2	RW	M	PHY-Up Clk Compensation Sequence Frequency—Determines the minimum frequency for PHY-Up Clk Comp sequences to transmit to enable a physical layer to autonomously compensate for clock differences with the remote PHY. 0x0—1 per 3200 blocks 0x1—1 per 1600 blocks (Default) 0x2—1 per 800 blocks 0x3—1 per 400 blocks A block shall be 66 UI for 64b/66b encoding and 130 UI for 128b/130b encoding.
15:4	RW	-	RsvdP

Table 5-16: PHY Specific Capability Fields

Bit Location	Access	M/O	Description
31:0	RO	-	RsvdZ

Table 5-17: Remote PHY Specific Capability Fields

Bit Location	Access	M/O	Description
31:0	RW	-	RsvdP

Table 5-18: PHY TX Lane Status Fields

Bit Location	Access	M/O	Description
1:0	RO	M	Cursor Status Mode—Indicates the method by which the Tx equalization cursors were configured. 0x0—Automatic PMD in-band exchange 0x1—Direct Interface PHY Structure control

Bit Location	Access	M/O	Description
			0x2-0x3—Reserved
5:2	RO	-	RsvdZ
9:6	RO	M	<p>Current Pre-cursor Local Setting</p> <p>If Cursor Status Mode used Automatic PMD in-band exchange (0x0), then these bits shall indicate the adjustment count for the local pre-cursor:</p> <p>0x0—Reserved</p> <p>0x1-0x7—Final positive adjustments from current initialization mode (Initialize or Preset). Example, 0x3 → 3 positive adjustments.</p> <p>0x8—Reserved</p> <p>0x9-0xF—Final negative (minus 8) adjustments from current initialization mode (Initialize or Preset). Example, 0xA → 10-8 = 2 negative adjustments.</p> <p>If Cursor Status Mode used Direct Interface PHY Structure control (0x1), then these bits shall indicate the following for the local pre-cursor:</p> <p>0x0—c(-1) ratio 0</p> <p>0x1—c(-1) ratio -0.05</p> <p>0x2—c(-1) ratio -0.1</p> <p>0x3—c(-1) ratio -0.15</p> <p>0x4-0xF—Reserved</p>
13:10	RO	M	<p>Current Post-cursor Local Setting</p> <p>If Cursor Status Mode used Automatic PMD in-band exchange (0x0), then these bits shall indicate the adjustment count for the local post-cursor:</p> <p>0x0—Reserved</p> <p>0x1-0x7—Final positive adjustments from current initialization mode (Initialize or Preset). Example, 0x3 → 3 positive adjustments.</p> <p>0x8—Reserved</p> <p>0x9-0xF—Final negative (minus 8) adjustments from current initialization mode (Initialize or Preset). Example, 0xA → 10-8 = 2 negative adjustments.</p> <p>If Cursor Status Mode used Direct Interface PHY Structure control (0x1), then these bits shall indicate the following for the local post-cursor:</p> <p>0x0—c(1) ratio 0</p>

Bit Location	Access	M/O	Description
			0x1—c(1) ratio -0.05 0x2—c(1) ratio -0.1 0x3—c(1) ratio -0.15 0x4—c(1) ratio -0.2 0x5—c(1) ratio -0.25 0x6-0xF—Reserved
17:14	RO	M	Current Output Margin / Primary-cursor Local Setting If Cursor Status Mode used Automatic PMD in-band exchange g (0x0), then these bits shall indicate the adjustment count for the local primary-cursor: 0x0—Reserved 0x1-0x7—Final positive adjustments from current initialization mode (Initialize or Preset). Example, 0x3 → 3 positive adjustments. 0x8—Reserved 0x9-0xF—Final negative (minus 8) adjustments from current initialization mode (Initialize or Preset). Example, 0xA → 10-8 = 2 negative adjustments. If Cursor Status Mode used Direct Interface PHY Structure control (0x1), then these bits shall indicate the following for the local output margin: 0x0—Normal output swing 0x1—~87.5% output swing 0x2—~75% output swing 0x3—~62.5% output swing 0x4—~50% output swing 0x5—~37.5% output swing 0x6—~25% output swing 0x7-0xF—Reserved
19:18	RO	M	Current Initialization Status—Indicates the starting initialization PMD setting mode. 0x0—PMD Initialize settings 0x1—PMD Preset settings 0x2-0x3—Reserved
21:20	RO	M	Local Equalization Status 0x0—Local PMD equalization not run 0x1—Local PMD equalization in progress 0x2—Local PMD equalization completed with no errors 0x3—Local PMD equalization completed with errors

Bit Location	Access	M/O	Description
23:22	RO	M	Remote Equalization Status 0x0—Remote PMD equalization not run 0x1—Remote PMD equalization in progress 0x2—Remote PMD equalization completed with no errors 0x3—Remote PMD equalization completed with errors
31:24	RO	-	RsvdZ

Table 5-19: PHY TX Lane Control Fields

Bit Location	Access	M/O	Description
1:0	RW	M	Cursor Control Mode 0x0—Automatic PMD in-band exchange (Default) 0x1—Direct Interface PHY Structure control 0x2-0x3—Reserved
5:2	RW	-	RsvdP
9:6	RW	M	Pre-cursor Local Setting If Cursor Control Mode is configured for automatic PMD in-band exchange, then these bits shall be invalid and set to 0x0. If Cursor Control Mode is configured for Direct Interface PHY Structure control, then these bits shall determine the local pre-cursor setting as: 0x0—c(-1) ratio 0 0x1—c(-1) ratio -0.05 0x2—c(-1) ratio -0.1 0x3—c(-1) ratio -0.15 0x4-0xF—Reserved
13:10	RW	M	Post-cursor Local Setting If Cursor Control Mode is configured for automatic PMD in-band exchange, then these bits shall be invalid and set to 0x0. If Cursor Control Mode is configured for Direct Interface PHY Structure control, then these bits shall determine the local post-cursor setting as: 0x0—c(1) ratio 0 0x1—c(1) ratio -0.05 0x2—c(1) ratio -0.1 0x3—c(1) ratio -0.15

Bit Location	Access	M/O	Description
17:14			0x4—c(1) ratio -0.2 0x5—c(1) ratio -0.25 0x6-0xF—Reserved
	RW	M	Output Margin / Primary-cursor Local Setting If Cursor Control Mode is configured for automatic PMD in-band exchange, then these bits shall be invalid and set to 0x0. If Cursor Control Mode is configured for Direct Interface PHY Structure control, then these bits shall determine the local primary output margin as: 0x0—Normal output swing 0x1—~87.5% output swing 0x2—~75% output swing 0x3—~62.5% output swing 0x4—~50% output swing 0x5—~37.5% output swing 0x6—~25% output swing 0x7-0xF—Reserved
19:18	RW	M	Initialization Mode—Determines the starting initialization PMD setting. 0x0—PMD Initialize settings 0x1—PMD Preset settings 0x2-0x3—Reserved
31:20	RW	-	RsvdP

Table 5-20: PHY Rx Lane Margin Fields

Bit Location	Access	M/O	Description
1:0	RW	O	Rx Lane Margin Control—Determines the type of Rx Lane Margining being testing. 0x0—Disabled 0x1—Timing offset Rx margining enabled 0x2—Voltage offset Rx margining enabled 0x3—Synchronous timing and voltage offset Rx margining enabled
2	RW	O	Timing Offset Direction 0b—Data sampler shifted earlier in time (left) 1b—Data sampler shifted later in time (right)

Bit Location	Access	M/O	Description
8:3	RW	O	<p>Timing Offset Magnitude—Determines the ratio in timing offsets with a full UI = 64 total steps.</p> <p>0x0—Reserved 0x1—1/64 step 0x2—2/64 steps 0x3—3/64 steps ... 0x20—32/64 steps (1/2 UI) ... 0x28—40/64 steps 0x29-0x3F—Reserved</p> <p>Developer Note: Assuming data sampling is centered, 32 steps earlier or later would reach the transition/edge sampling. The range of timing offset support can be implementation-specific and can support steps beyond the transition/edge by up to 8 steps.</p>
9	RW	O	<p>Voltage Offset Direction</p> <p>0b—Data sampler is shifted higher in voltage offset (positive) 1b—Data sampler is shifted lower in voltage offset (negative)</p>
15:10	RW	O	<p>Voltage Offset Magnitude—Determines the ratio in voltage offsets with 64 total steps. The range of voltage offset support is implementation specific.</p> <p>0x0—Reserved 0x1—1/64 step 0x2—2/64 steps 0x3—3/64 steps ... 0x20—32/64 steps (1/2 UI) ... 0x28—40/64 steps 0x29-0x3F--Reserved</p>
16	RO	O	Timing Offset Rx margining Support
17	RO	O	Voltage Offset Rx margining Support
18	RO	O	Synchronous Timing and Voltage Offset Rx Margining Support
19	RO	O	Margin Offset Configuration Status—Indicates the validity of the Rx Lane Margining setting (timing, voltage, or both) programmed for an implementation.

Bit Location	Access	M/O	Description
23:20			If an implementation cannot support the programmed setting (timing, voltage, or both), then this field should report Invalid (1b). 0b—Valid timing, voltage, or both timing and voltage offset setting 1b—Invalid timing, voltage, or both timing and voltage offset setting
	RO	O	Rx Margining Status—Indicates the estimated Bit Error Ratio (BER) for the current Rx Lane Margining programmed settings. 0x0—Rx margining not run 0x1—Rx margining run; BER < $\sim 10^{-15}$ 0x2—Rx margining run; BER = $\sim 10^{-14}$ 0x3—Rx margining run; BER = $\sim 10^{-13}$ 0x4—Rx margining run; BER = $\sim 10^{-12}$ 0x5—Rx margining run; BER = $\sim 10^{-11}$ 0x6—Rx margining run; BER = $\sim 10^{-10}$ 0x7—Rx margining run; BER = $\sim 10^{-9}$ 0x8—Rx margining run; BER = $\sim 10^{-8}$ 0x9—Rx margining run; BER = $\sim 10^{-7}$ 0xA—Rx margining run; BER = $\sim 10^{-6}$ 0xB—Rx margining run; BER = $\sim 10^{-5}$ 0xB—Rx margining run; BER = $\sim 10^{-4}$ 0xD-0xF—Reserved
	RO	-	RsvdZ

Table 5-21: PHY Rx Lane Status Fields

Bit Location	Access	M/O	Description
0	RO	M	Signal Detected Status—Indicates the signal detection status of the indexed Rx lane from the Signal Detect state. This bit should set to 0b when the Physical Layer State field is in PHY Idle (0x0). 0b—Signal Not Detected 1b—Signal Detected
1	RO	M	Current Lane Polarity Status—Indicates the polarity status of the indexed Rx lane. This bit should set to 0b when the Physical Layer State field is in PHY Idle (0x0). 0b—Polarity Normal 1b—Polarity Reversed
2	RO	O	CDR Status—Indicates the Rx CDR status of the indexed Rx lane. This bit should set to 0b when the Physical Layer State field is in PHY Idle (0x0).

Bit Location	Access	M/O	Description
3			0b—CDR Not Locked 1b—CDR Locked
	RO	O	Rx Equalization Status 0b—Equalization Not Run 1b—Equalization Completed
4	RO	M	Rx Symbol Alignment Status—Indicates the Rx symbol alignment status of the indexed Rx lane. This bit should set to 0b when the Physical Layer State field is in PHY Idle (0x0). 0b—Symbol Alignment Not Achieved 1b—Symbol Alignment Achieved
5	RO	M	Lane Error Status—Indicates the Rx error status of the indexed Rx lane during physical layer training. This bit should set to 0b when the Physical Layer State field is in PHY Idle (0x0). 0b—No Error Detected 1b—Errors Detected
7:6	RO	-	RsvdZ
11:8	RO	O	Rx IBIST status—If IBIST is supported and enabled, then this field indicates the estimated Bit Error Ratio (BER) from the IBIST state for the indexed Rx lane. 0x0—IBIST not run 0x1—IBIST run; BER < $\sim 10^{-15}$ 0x2—IBIST run; BER = $\sim 10^{-14}$ 0x3—IBIST run; BER = $\sim 10^{-13}$ 0x4—IBIST run; BER = $\sim 10^{-12}$ 0x5—IBIST run; BER = $\sim 10^{-11}$ 0x6—IBIST run; BER = $\sim 10^{-10}$ 0x7—IBIST run; BER = $\sim 10^{-9}$ 0x8—IBIST run; BER = $\sim 10^{-8}$ 0x9-0xF—Reserved
15:12	RO	-	RsvdZ
19:16	RO	O	Requested Pre-cursor Remote Setting—Indicates the locally requested pre-cursor setting for the remote transmitter for the indexed Rx lane. 0x0—c(-1) ratio 0 0x1—c(-1) ratio -0.05 0x2—c(-1) ratio -0.1

Bit Location	Access	M/O	Description
23:20			0x3—c(-1) ratio -0.15 0x4-0xF—Reserved
	RO	O	Requested Post-cursor Remote Setting—Indicates the locally requested post-cursor setting for the remote transmitter for the indexed Rx lane. 0x0—c(1) ratio 0 0x1—c(1) ratio -0.05 0x2—c(1) ratio -0.1 0x3—c(1) ratio -0.15 0x4—c(1) ratio -0.2 0x5—c(1) ratio -0.25 0x6-0xF—Reserved
27:24	RO	O	Requested Output Margin / Primary-cursor Remote Setting—Indicates the locally requested primary output margin setting for the remote transmitter for the indexed Rx lane. 0x0—Normal output swing 0x1—~87.5% output swing 0x2—~75% output swing 0x3—~62.5% output swing 0x4—~50% output swing 0x5—~37.5% output swing 0x6—~25% output swing 0x7-0xF—Reserved
28	RO	O	Change in Remote Equalization Requested—Indicates status for equalization change requests from the local physical layer to the remote transmitter for the indexed Rx lane. 0b—No Request 1b—Request for Remote Equalization Setting Change
31:29	RO	-	RsvdZ

Table 5-22: PHY Rx Lane Control Fields

Bit Location	Access	M/O	Description
1:0	RW	O	Signal Detect Control—Forces or restricts the indexed Rx lane’s detection status when in the Signal Detect state. 0x0—Normal; Automatic Signal Detect (Default) 0x1—Reserved 0x2—Force No Signal Detected

Bit Location	Access	M/O	Description
3:2			0x3—Force Signal Detected
	RW	O	Polarity Lane Control—Determines the configuration of the indexed Rx lane’s polarity, automatically or statically (Fixed). 0x0—Fixed, Normal Polarity 0x1—Fixed, Inverted Polarity 0x2—Automatically Tuned (Default) 0x3—Reserved
8:4	RW	O	Recommended CTLE Peaking—Determines the indexed Rx lane’s CTLE peaking setting. 0x0—Reserved 0x1—1 dB 0x2—2 dB ... 0x1E—30 dB 0x1F—31 dB
31:9	RW	-	RsvdP

Table 5-23: PHY Idle State Status Fields

Bit Location	Access	M/O	Description
0	RO	O	PHY Full Power Down Status—Indicates the status of the power down mode of the physical layer when in a PHY Idle status. 0b—PHY Idle Power 1b—PHY Idle Power Down Mode
15:1	RO	-	RsvdZ

Table 5-24: PHY Idle State Control Fields

Bit Location	Access	M/O	Description
0	RW	O	PHY Full Power Down Control—Determines the power down mode for the physical layer when in the PHY Idle state. Full power down mode can put the PHY in an ultra-low power state. This mode shall only be controlled by local sideband manageability and shall not be written in-band to prevent disabling the remote physical layer, possibly indefinitely or until device/component reset. Prior to a physical layer being enabled, this bit shall be set to 0b. 0b—PHY Idle Power

Bit Location	Access	M/O	Description
			1b—PHY Idle Power Down Mode Developer Note: For physical layers that have no remote connection, an implementation can use this field to achieve an extremely low-power state by disabling features such as internal/external supplies, bias currents, or PLLs. Latency expectations for exiting the PHY Idle Power Down Mode can consume a significant period of time.
15:1	RW	-	RsvdP

Table 5-25: PHY Idle State Capability Fields

Bit Location	Access	M/O	Description
0	RO	O	PHY Full Power Down Support
2:1	RO	O	PHY Full Power Down Exit Latency—Indicates the exit latency for the local PHY when in a PHY Idle Full Power Down Mode. 0x0—250 ms 0x1—500 ms 0x2—1s 0x3—2s
15:3	RO	-	RsvdZ

Table 5-26: Remote PHY Idle State Capability Fields

Bit Location	Access	M/O	Description
0	RW	O	Remote PHY Full Power Down Support
2:1	RW	O	PHY Full Power Down Exit Latency—Determines the exit latency for the remote PHY when in a PHY Idle Full Power Down Mode 0x0—250 ms 0x1—500 ms 0x2—1s 0x3—2s
15:3	RW	-	RsvdP

Table 5-27: PHY Signal Detect State Status Fields

Bit Location	Access	M/O	Description
0	RO	M	<p>Signal Detect State Status—Indicates the combined (OR) Rx lane detection status from the Signal Detect state and shall be invalid when the Physical Layer State field is in PHY Idle (0x0).</p> <p>0b—Signal Not Detected 1b—Signal Detected</p> <p>Developer Note: Recommend that the Physical Layer State field be read to validate the PLTSM has progressed to or beyond this state to determine if the fields in this structure are valid.</p>
1	RO	M	<p>Signal Detect State Error Status—Indicates error status from the Signal Detect state. Signal Detect state errors can occur if signals are detected but a valid full width or half width link cannot be resolved.</p> <p>If the PLTSM has not progressed to or beyond the Signal Detect state, then this bit shall be invalid.</p> <p>0b—Signal Detect State Exited Without Errors 1b—Signal Detect state Exited With Errors</p>
2	RO	M	<p>Signal Detect reversal status</p> <p>If the PLTSM has not progressed to or beyond the Signal Detect state, then this bit shall be invalid.</p> <p>0b—Reversal Not Detected 1b—Reversal Detected</p>
15:3	RO	-	RsvdZ

Table 5-28: PHY Signal Detect State Control Fields

Bit Location	Access	M/O	Description
0	WO	O	<p>Force Signal Detect—Forces detection of a signal to the PLTSM so it may progress past the Signal Detect state.</p> <p>1b—Force signal detection</p>
2:1	RW	O	<p>Detect Send Sub-state Timer Reconfiguration—Determines the active, transmit Signal Detect sub-state time. This field can provide flexibility in power savings vs. latency to exit the Signal Detect state.</p> <p>0x0—500 μs 0x1—1 ms 0x2—2 ms (Default)</p>

Bit Location	Access	M/O	Description
4:3			0x3—4 ms
	RW	O	Detect Idle Sub-state Timer Reconfiguration—Determines the inactive, low power Signal Detect sub-state time. This field can provide flexibility in power savings vs. latency to exit the Signal Detect state. 0x0—24.5 ms 0x1—49 ms 0x2—98 ms (Default) 0x3—196 ms
6:5	RW	O	Detect Sub-state Timer Reconfiguration—Determines the active Signal Detect sub-state transition time. This field can provide flexibility in powering on physical layer circuits vs. latency to exit the Signal Detect state. 0x0—2 ms 0x1—4 ms 0x2—8 ms (Default) 0x3—16 ms
8:7	RW	O	Detect Header Repeat Count This field can provide flexibility in signal detection time vs. high speed data transitions in the Signal Detect state. 0x0—Single Detect Header 0x1—2 back-to-back Detect Headers (Default) 0x2—4 back-to-back Detect Headers 0x3—Reserved
10:9	RW	O	Signal Detect Header Frequency This field can provide flexibility in signal detection time versus high speed data transitions in the Signal Detect state. 0x0—Detect Header(s) repeat after 8 Detect Sets (Default) 0x1—Detect Header(s) repeat after 16 Detect Sets 0x2—Detect Header(s) repeat after 32 Detect Sets 0x3—Detect Header(s) repeat after 64 Detect Sets
11	RW	O	Detect Header Scrambler Reset Control—Determines when the scrambler resets during the Signal Detect state. 0b—Detect Headers do not reset the scrambler 1b—Detect Headers reset the scrambler (Default)

Bit Location	Access	M/O	Description
			Developer Note: If an implementation relies on the Detect Set data, then disabling the scrambler reset can enable more variation in the data pattern being transmitted during the Signal Detect state.
15:12	RW	-	RsvdP

Table 5-29: PHY Signal Detect State Capability Fields

Bit Location	Access	M/O	Description
0	RO	O	Force Signal Detect Control Support
1	RO	O	Detect Send Sub-state Timer Reconfiguration Support
2	RO	O	Detect Idle Sub-state Timer Reconfiguration Support
3	RO	O	Detect Sub-state Timer Reconfiguration Support
4	RO	O	Detect Header Repeat Count Reconfiguration Support
5	RO	O	Detect Header Frequency Reconfiguration Support
6	RO	O	Detect Header Scrambler Reset Support 0b—Unsupported: Detect Headers reset scrambler for Detect Sets 1b—Support: Free running scrambler on Detect Sets
15:7	RO	-	RsvdZ

Table 5-30: Remote PHY Signal Detect State Capability Fields

Bit Location	Access	M/O	Description
0	RW	O	Remote PHY Force Signal Detect Control Support
1	RW	O	Remote PHY Detect Send Sub-state Timer Reconfiguration Support
2	RW	O	Remote PHY Detect Idle Sub-state Timer Reconfiguration Support
3	RW	O	Remote PHY Detect sub-state timer reconfiguration support
4	RW	O	Remote PHY Detect Header Repeat Count Reconfiguration Support
5	RW	O	Remote PHY Detect Header Frequency Reconfiguration Support
6	RO	O	Remote PHY Detect Header Scrambler Reset Support 0b—No support; Detect Headers reset scrambler for Detect Sets 1b—Support for free running scrambler on Detect Sets

Bit Location	Access	M/O	Description
15:7	RW	-	RsvdP

Table 5-31: PHY Channel Optimize State Status Fields

Bit Location	Access	M/O	Description
1:0	RO	M	<p>Channel Optimize State Status</p> <p>0x0—Channel Optimize state has not been entered or is in progress. 0x1—Channel Optimize state has run and exited 0x2—Channel Optimize state was bypassed 0x3—Reserved</p> <p>Developer Note: Recommend that the Physical Layer State field be read to validate the PLTSM has progressed to or beyond this state to determine if the fields in this structure are valid.</p>
3:2	RO	M	<p>Channel Optimize State Error Status</p> <p>If the Channel Optimize state was bypassed or the PLTSM has not progressed to or beyond the Channel Optimize state, then these bits shall be invalid.</p> <p>0x0—Channel Optimize state exited without errors 0x1—Channel Optimize state exited with errors 0x2—Channel Optimize state timed out 0x3—Reserved</p>
15:3	RO	-	RsvdZ

Table 5-32: PHY Channel Optimize State Control Fields

Bit Location	Access	M/O	Description
0	RW	O	<p>Channel Optimize State Bypass Control</p> <p>0b—Channel Optimize state bypassed 1b—Channel Optimize state enabled (Default)</p> <p>Developer Note: A physical layer can use this field to bypass the Channel Optimize state in closed systems when Tx equalization cursors can be explicitly configured.</p>
3:1	RW	O	<p>Channel Optimize State Timeout Reconfiguration</p> <p>This field can provide flexibility in tuning accuracy vs. latency to exit the Channel Optimize state.</p> <p>0x0—3.125 ms</p>

Bit Location	Access	M/O	Description
			0x1—6.25 ms 0x2—12.5 ms 0x3—25 ms 0x4—50 ms (Default) 0x5—100 ms 0x6—200 ms 0x7—400 ms
15:4	RW	-	RsvdP

Table 5-33: PHY Channel Optimize State Capability Fields

Bit Location	Access	M/O	Description
0	RO	O	Channel Optimize State Bypass Control Support
1	RO	O	Channel Optimize State Timeout Reconfiguration Support
15:2	RO	-	RsvdZ

Table 5-34: Remote PHY Channel Optimize State Capability Fields

Bit Location	Access	M/O	Description
0	RW	O	Remote PHY Channel Optimize State Bypass Control Support
1	RW	O	Remote PHY Channel Optimize State Timeout Reconfiguration Support
15:2	RW	-	RsvdP

Table 5-35: PHY Align State Status Fields

Bit Location	Access	M/O	Description
0	RO	M	Align State Status 0b—Align state has not been entered or is in progress. 1b—Align state has run and exited Developer Note: Recommend that the Physical Layer State field be read to validate the PLTSM has progressed to or beyond this state to determine if the fields in this structure are valid.
2:1	RO	M	Align State Error Status If the PLTSM has not progressed to or beyond Align state, then these bits shall be invalid.

Bit Location	Access	M/O	Description
3			0x0—Align state exited without errors 0x1—Align state exited with errors 0x2—Align state timed out 0x3—Reserved
	RO	M	Align Acknowledgements Received If the PLTSM has not progressed to or beyond the Align state, then this bit shall be invalid. 0b—Not Received 1b—Received
4	RO	M	Align Acknowledgements Sent If the PLTSM has not progressed to or beyond the Align state, then this bit shall be invalid. 0b—Not Sent 1b—Sent
15:5	RO	-	RsvdZ

Table 5-36: PHY Align State Control Fields

Bit Location	Access	M/O	Description
1:0	RW	O	Align Header frequency This field can provide flexibility in alignment time versus lane-to-lane skew depth in the Align state. 0x0—Align Header repeat after 8 Align Sets 0x1—Align Header repeat after 16 Align Sets (Default) 0x2—Align Header repeat after 32 Align Sets 0x3—Align Header repeat after 64 Align Sets
4:2	RW	O	Align state minimum timeout This field can provide flexibility in lane alignment time vs. latency to exit the Align state. The minimum timeout shall specify the time after which the Align state may: <ul style="list-style-type: none"> • Timeout, if no lanes have reached alignment • Execute LWR with lanes that have reached alignment and transition to the Config state If all lanes reach alignment within a physical layer, then may transition from the Align state to the Config state. 0x0—500 μ s

Bit Location	Access	M/O	Description
5			0x1—1 ms 0x2—2 ms 0x3—4 ms 0x4—8 ms (Default) 0x5—16 ms 0x6—32 ms 0x7—64 ms
	RW	O	Align State Timeout Control—Disables the Align state timeout. This field can provide flexibility for test and debug purposes. 0b—Timeout disabled 1b—Timeout enabled (Default)
	RW	-	RsvdP

Table 5-37: PHY Align State Capability Fields

Bit Location	Access	M/O	Description
0	RW	O	Align Header Frequency Reconfiguration Support
1	RW	O	Align State Minimum Timeout Reconfiguration Support
2	RW	O	Align State Timeout Control Support
15:2	RW	-	RsvdP

Table 5-38: Remote PHY Align State Capability Fields

Bit Location	Access	M/O	Description
0	RW	O	Remote PHY Align Header Frequency Reconfiguration Support
1	RW	O	Remote PHY Align State Minimum Timeout Reconfiguration Support
2	RW	O	Remote PHY Align State Timeout control Support
15:2	RW	-	RsvdP

Table 5-39: PHY IBIST State Status Fields

Bit Location	Access	M/O	Description
1:0	RO	O	IBIST State Status 0x0—IBIST state has not been entered.

Bit Location	Access	M/O	Description
			0x1—IBIST state is in progress 0x2—IBIST state has run and exited 0x3—IBIST state was bypassed Developer Note: Recommend that the Physical Layer State field be read to validate the PLTSM has progressed to or beyond this state to determine if the fields in this structure are valid.
2	RO	O	IBIST Error Status—Indicates if an IBIST error was detected on any Rx lane. IBIST errors are reported on a per lane basis in <i>PHY Rx Lane Status Fields</i> . If the PLTSM has not progressed to or beyond the IBIST state, then this bit shall be invalid. 0b—IBIST state exited without errors 1b—IBIST state exited with errors
3	RO	O	IBIST Lane De-configure Status—Indicates if an implementation de-configured any Rx lane due to IBIST errors. If the PLTSM has not progressed to or beyond the IBIST state, then this bit shall be invalid. 0b—IBIST did not de-configure a Rx lane 1b—IBIST de-configured a Rx lane
4	RO	O	IBIST Header Status—Indicates detection status of an IBIST Header. If the Physical Layer State field reports IBIST, then this bit shall be validated to ensure an error was not present in the IBIST Header preventing entry into checking of the IBIST Pattern. If the PLTSM has not progressed to or beyond the IBIST state, then this bit shall be invalid. 0b—IBIST Header not identified 1b—IBIST Header received
15:5	RO	-	RsvdZ

Table 5-40: PHY IBIST State Control Fields

Bit Location	Access	M/O	Description
0	RW	O	IBIST State Control Enable—Determines if the PLTSM can execute the IBIST state. If the IBIST state is supported by the local and remote physical layers, then this bit shall contain the same value in both the local and remote Interface PHY Structures.

Bit Location	Access	M/O	Description
1			0b—IBIST state bypassed (Default) 1b—IBIST state enabled
	RW	O	IBIST Lane De-configure Control—Determines if Rx lanes are de-configured when errors are detected in the IBIST state. 0b—Do not de-configure a Rx lane on IBIST errors 1b—De-configure a Rx lane on IBIST errors
5:2	RW	O	IBIST Pattern PRBS—Determines the PRBS settings for scrambling the IBIST Pattern. If the IBIST state is supported by the local and remote physical layers, then these bits shall contain the same value in both the local and remote Interface PHY Structures. 0x0—PRBS7 0x1—PRBS11 0x2—PRBS15 0x3—PRBS 23 (Default) 0x4—PRBS31 0x5-0xF--Reserved
8:6	RW	O	IBIST Pattern Test Period—Determines the number of IBIST test patterns during the IBIST state. If the IBIST state is supported by the local and remote physical layers, then these bits shall contain the same value in both the local and remote Interface PHY Structures. 0x0—2 ¹⁰ IBIST Patterns 0x1—2 ¹¹ IBIST Patterns 0x2—2 ¹² IBIST Patterns (Default) 0x3—2 ¹⁴ IBIST Patterns 0x4—2 ¹⁶ IBIST Patterns 0x5—2 ¹⁸ IBIST Patterns 0x6—2 ²⁰ IBIST Patterns 0x7—Continuous
15:9	RW	-	RsvdP

Table 5-41: PHY IBIST State Capability Fields

Bit Location	Access	M/O	Description
0	RO	O	IBIST State Control Support

Bit Location	Access	M/O	Description
1	RO	O	IBIST Lane De-configure Support
2	RO	O	IBIST Pattern PRBS7 Support
3	RO	O	IBIST Pattern PRBS11 Support
4	RO	O	IBIST Pattern PRBS15 Support
5	RO	O	IBIST Pattern PRBS31 Support
15:6	RO	-	RsvdZ

Table 5-42: Remote PHY IBIST State Capability Fields

Bit Location	Access	M/O	Description
0	RO	O	Remote PHY IBIST State Control Support
1	RO	O	Remote PHY IBIST Lane De-configure Support
2	RO	O	Remote IBIST Pattern PRBS7 Support
3	RO	O	Remote IBIST Pattern PRBS11 Support
4	RO	O	Remote IBIST Pattern PRBS15 Support
5	RO	O	Remote IBIST Pattern PRBS31 Support
15:6	RO	-	RsvdZ

Table 5-43: PHY Config State Status Fields

Bit Location	Access	M/O	Description
0	RO	M	Config State Status 0b—Config state has not been entered or is in progress 1b—Config state has run and exited
1	RO	M	Config State Error Status If the PLTSM has not progressed to or beyond the Config state, then this bit shall be invalid. 0b—Config state exited without errors 1b—Config state exited with errors

Bit Location	Access	M/O	Description
2	RO	M	<p>Config Width Acknowledgements Received</p> <p>If the PLTSM has not progressed to or beyond the Config state, then this bit shall be invalid.</p> <p>0b—Not received 1b—Received</p>
3	RO	M	<p>Config Width Acknowledgements Sent</p> <p>If the PLTSM has not progressed to or beyond the Config state, then this bit shall be invalid.</p> <p>0b—Not sent 1b—Sent</p>
5:4	RO	M	<p>Tx Link Width Reduction Status</p> <p>If the PLTSM has not progressed to or beyond the Config state, then these bits shall be invalid.</p> <p>0x0—All lanes disabled 0x1—Upper lanes disabled 0x2—Lower lanes disabled 0x3—All lanes enabled</p>
7:6	RO	M	<p>Rx Link Width Reduction Status</p> <p>If the PLTSM has not progressed to or beyond the Config state, then these bits shall be invalid.</p> <p>0x0—All lanes disabled 0x1—Upper lanes disabled 0x2—Lower lanes disabled 0x3—All lanes enabled</p>
15:8	RO	-	RsvdZ

Table 5-44: PHY Config State Control Fields

Bit Location	Access	M/O	Description
0	RW	O	<p>Config State Timeout Control—Determines if the Config state uses a timeout.</p> <p>This field can provide flexibility for test and debug purposes.</p> <p>0b—Config state timeout disabled (Default) 1b—Config state timeout enabled</p>

Bit Location	Access	M/O	Description
2:1	RW	O	Config State Minimum Timeout This field can be used in conjunction with the Config State Timeout Control to specify a timeout after entering the Config State. 0x0—1 μ s 0x1—5 μ s 0x2—10 μ s 0x3—20 μ s
15:3	RW	-	RsvdP

Table 5-45: PHY Config State Capability Fields

Bit Location	Access	M/O	Description
0	RO	O	Config State Timeout Control Support
1	RO	O	Config State Minimum Timeout Support
15:2	RO	-	RsvdZ

Table 5-46: Remote PHY Config State Capability Fields

Bit Location	Access	M/O	Description
0	RW	O	Remote PHY Config State Timeout Control Support
1	RW	O	Remote PHY Config State Minimum Timeout Support
15:2	RW	-	RsvdZ

Table 5-47: PHY PHY-Up State Status Fields

Bit Location	Access	M/O	Description
7:0	RO	M	PHY-Up Errors—This counter shall be incremented by one (modulo 2^8) unless saturated whenever a physical layer error is encountered. Error counts are reset using the Clear PHY-Up Status field in the PHY PHY-Up State Control structure or when a physical layer retrains.
15:8	RO	O	Previous PHY-Up Errors—This counter shall report the previous state of the PHY-Up Errors fields when reset either via the Clear PHY-Up Status field in the PHY PHY-Up State Control structure or when a physical layer retrains.

Table 5-48: PHY PHY-Up State Control Fields

Bit Location	Access	M/O	Description
0	WO	M	Clear PHY-Up Status 1b—Reset fields
3:1	RW	O	Retraining Error Threshold—Determines the number of errors required before a physical layer may autonomously retrain. The Disable Physical Layer Auto-retraining field in the PHY Control structure may disable this feature so that physical layer retraining then becomes the responsibility of the Gen-Z Core. 0x0—Any error 0x1—BER > 10 ⁻¹² 0x2—BER > 10 ⁻¹¹ 0x3—BER > 10 ⁻¹⁰ 0x4-0x7—Reserved
15:4	RW	-	RsvdP

Table 5-49: PHY PHY-Up State Capability Fields

Bit Location	Access	M/O	Description
0	RO	O	Retraining Error Threshold Support
15:1	RO	-	RsvdZ

Table 5-50: Remote PHY PHY-Up State Capability Fields

Bit Location	Access	M/O	Description
0	RW	O	Remote PHY Retraining Error Threshold Support
15:1	RW	-	RsvdP

Table 5-51: PHY LP 1 State Status Fields

Bit Location	Access	M/O	Description
9:0	RO	O	PHY LP 1 Events—This counter shall be incremented by one (modulo 2 ¹⁰) unless saturated whenever a physical layer LP 1 state is entered. Event counts are reset using the Clear LP 1 Status field in the PHY LP 1 State Control structure.
15:10	RO	-	RsvdZ

Table 5-52: PHY LP 1 State Control Fields

Bit Location	Access	M/O	Description
0	WO	O	Clear LP 1 Status 1b—Reset fields
2:1	RW	O	LP 1 Entrance Time—Determines the time required to transmit PHY-Up Sequences of type LP with LPID = Low Power 1 prior to beginning the transition to the LP 1 state. 0x0—2 μ s 0x1—4 μ s 0x2—8 μ s (Default) 0x3—16 μ s
15:3	RW	-	RsvdP

Table 5-53: PHY LP 1 State Capability Fields

Bit Location	Access	M/O	Description
2:0	RO	O	Estimated Power Savings for LP 1 State 0x0—None 0x1—1 to 20% reduction 0x2—21 to 40% reduction 0x3—41 to 60% reduction 0x4—61 to 80% reduction 0x5—81 to 99% reduction 0x6-0x7—Reserved
15:3	RO	-	RsvdZ

Table 5-54: Remote PHY LP 1 State Capability Fields

Bit Location	Access	M/O	Description
2:0	RW	O	Remote PHY Estimated Power Savings for LP 1 State 0x0—None 0x1—1 to 20% reduction 0x2—21 to 40% reduction 0x3—41 to 60% reduction 0x4—61 to 80% reduction 0x5—81 to 99% reduction 0x6-0x7—Reserved
15:3	RW	-	RsvdP

Table 5-55: PHY LP 2 State Status Fields

Bit Location	Access	M/O	Description
9:0	RO	O	PHY LP 2 Events—This counter shall be incremented by one (modulo 2 ¹⁰) unless saturated whenever a physical layer LP 2 state is entered. Event counts are reset using the Clear LP 2 Status field in the PHY LP 2 State Control structure.
15:10	RO	-	RsvdZ

Table 5-56: PHY LP 2 State Control Fields

Bit Location	Access	M/O	Description
0	WO	O	Clear LP 2 Status 1b—Reset fields
2:1	RW	O	LP 2 Entrance Time—Determines the time required to transmit PHY-Up Sequences of type LP with LPID = Low Power 2 prior to beginning the transition to the LP 2 state. 0x0—2 μ s 0x1—4 μ s 0x2—8 μ s (Default) 0x3—16 μ s
15:3	RW	-	RsvdP

Table 5-57: PHY LP 2 State Capability Fields

Bit Location	Access	M/O	Description
2:0	RO	O	Estimated Power Savings for LP 2 State 0x0—None 0x1—1 to 20% reduction 0x2—21 to 40% reduction 0x3—41 to 60% reduction 0x4—61 to 80% reduction 0x5—81 to 99% reduction 0x6-0x7—Reserved
15:3	RO	-	RsvdZ

Table 5-58: Remote PHY LP 2 State Capability Fields

Bit Location	Access	M/O	Description
2:0	RW	O	Remote PHY Estimated Power Savings for LP 2 State 0x0—None 0x1—1 to 20% reduction 0x2—21 to 40% reduction 0x3—41 to 60% reduction 0x4—61 to 80% reduction 0x5—81 to 99% reduction 0x6-0x7—Reserved
15:3	RW	-	RsvdP

Table 5-59: PHY-Up LP 1 State Status Fields

Bit Location	Access	M/O	Description
9:0	RO	O	PHY-Up LP 1 Events—This counter shall be incremented by one (modulo 2^{10}) unless saturated whenever a physical layer PHY-Up LP 1 state is entered. Event counts are reset using the Clear LP 1 Status field in the PHY-Up LP 1 State Control structure.
15:10	RO	-	RsvdZ

Table 5-60: PHY-Up LP 1 State Control Fields

Bit Location	Access	M/O	Description
0	WO	O	Clear PHY-Up LP 1 Status 1b—Reset fields
2:1	RW	O	PHY-Up LP 1 Entrance Time—Determines the time required to transmit PHY-Up Sequences of type LP with LPID = PHY-Up Low Power 1 prior to initiating retraining to enter PHY-Up LP 1 state. 0x0—250 ns 0x1—500 ns 0x2—1 μ s (Default) 0x3—2 μ s
4:3	RW	O	PHY-Up LP 1 Exit Time—Determines the time required to transmit PHY-Up Sequences of type LP with LPID = PHY-Up prior to initiating retraining to return to the PHY-Up state. 0x0—250 ns 0x1—500 ns

Bit Location	Access	M/O	Description
15:5			0x2—1 μ s (Default) 0x3—2 μ s
	RW	-	RsvdP

Table 5-61: PHY-Up LP 1 State Capability Fields

Bit Location	Access	M/O	Description
2:0	RO	O	Estimated Power Savings for PHY-Up LP 1 State 0x0—None 0x1—1 to 20% reduction 0x2—21 to 40% reduction 0x3—41 to 60% reduction 0x4—61 to 80% reduction 0x5—81 to 99% reduction 0x6-0x7—Reserved
15:3	RO	-	RsvdZ

Table 5-62: Remote PHY-Up LP 1 State Capability Fields

Bit Location	Access	M/O	Description
2:0	RW	O	Remote PHY Estimated Power Savings for PHY-Up LP 1 State 0x0—None 0x1—1 to 20% reduction 0x2—21 to 40% reduction 0x3—41 to 60% reduction 0x4—61 to 80% reduction 0x5—81 to 99% reduction 0x6-0x7—Reserved
15:3	RW	-	RsvdP

5.5.2. Physical Layer Management

Refer to the Gen-Z Common Specification clause section on *Physical Layer Management*.

6. Gen-Z-E-NRZ-25G-Local Specification

6.1. Introduction

5 Gen-Z-E-NRZ-25G-Local specifies a physical layer capable of a line-rate at 25.78125 GT/s per lane inclusive of 64b/66b encoding (raw data-rate of 25 Gbps) and using Non-Return to Zero (NRZ) signaling over Local media (Very Short Reach) providing a BER of 10^{-15} or better. A compliant physical layer consists of an electrical interface compatible with a recommended Insertion Loss¹ (IL) no greater than 10 dB and no less than 4 dB at 12.89 GHz.

6.2. Logical Sublayer

6.2.1. Link Serialization

10 Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Link Serialization*.

6.2.2. Scrambling / Descrambling

Refer to the Gen-Z Common Specification clause section on *Scrambling / Descrambling*.

6.2.3. Data Striping

Refer to the Gen-Z Common Specification clause section on *Data Striping*.

15 6.2.4. PLA Data Width

Refer to the Gen-Z Common Specification clause section on *PLA Data Width*.

6.2.5. Physical Layer Initialization and Training

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Physical Layer Initialization and Training*.

20 6.2.6. Lane Reversal

Refer to the Gen-Z Common Specification clause section on *Lane Reversal*.

6.2.7. Lane Polarity

Lane polarity shall be supported for this physical layer clause. Detection of lane polarity shall occur in the *Align (Required)* state of the PLTSM.

25 **Developer Note:** *In a closed system, lane polarity can be explicitly programmed via the Unique Interface Physical Layer Structure of the Interface PHY Structure. In a system where lane polarity is not*

explicitly programmed and the Channel Optimize state is enabled for the negotiation of transmitter and receiver equalization settings, then lane polarity should be determined autonomously in the Channel Optimize (Optional) state of the PLTSM.

6.2.8. Link Width

5 Symmetric link widths of 1, 2, 4, 8, 16, or 32 transmit and receive lane may be supported by this physical layer clause. Asymmetry may supported up to an accumulated total link width of 64 lanes as specified in the Gen-Z Common Specification *Link Width* section.

10 **Developer Note:** *The Channel Optimize state can be enabled for this physical layer clause to determine transmitter and receiver equalization settings. If so, then recommend that symmetric operation be enabled initially to enable each lane to appropriately optimize its equalization setting. After symmetric physical layer training, asymmetric operation can be explicitly configured by reducing lanes and disabling them via the Common, Lane, and Low Power Interface PHY Structure.*

6.2.9. Link BIST

Refer to the *Gen-Z Common Specification* clause section on *Link BIST*.

15 6.2.10. Physical Layer Retraining

Refer to the *Gen-Z-E-NRZ-25G-Fabric Specification* clause section on *Physical Layer Retraining*.

6.2.11. PLTSM

The following training states shall be required this PLTSMs supporting this physical layer clause as illustrated in the *Gen-Z-E-NRZ-25G-Local* :

- 20
- PHY Idle
 - Signal Detect
 - Align
 - Configure
 - PHY-Up

25 The physical layer states Channel Optimize, IBIST, and PHY Low Power may be supported. The arcs to optional training states are illustrated with a blue dashed in *Gen-Z-E-NRZ-25G-Local* .

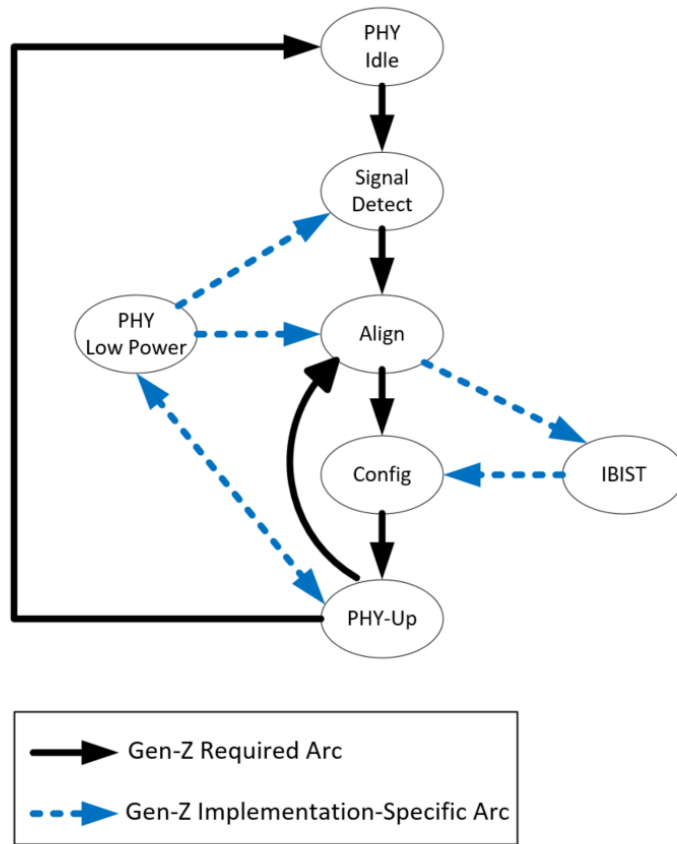


Figure 6-1: Gen-Z-E-NRZ-25G-Local PLTSM

A detailed description of each PLTSM state and the conditions of the associated arcs is provided below.

6.2.11.1. ***PHY Idle (Required)***

- 5 Refer to the Gen-Z Common Specification clause section on *PHY Idle (Required)*.

6.2.11.2. ***Signal Detect (Required)***

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Signal Detect (Required)* with the below noted inclusions for asymmetric links in *Detect Send* of *Signal Detect process*.

6.2.11.2.1. **Signal Detect entry**

- 10 Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Signal Detect Entry*.

6.2.11.2.2. **Signal Detect process**

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Signal Detect Process* with the following inclusions for asymmetric links in *Detect Send*.

Detect Send

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Detect Send* with the following inclusions for asymmetric links.

- 5 If the transmitter has an odd number of lanes where N is greater than or equal to 3 for asymmetry, then the central lane shall not enable its transmitter as illustrated in (a) and (b) of *Symmetric x4 to Asymmetric Tx3:Rx1 and Tx1:Rx3* and (c) and (d) *Asymmetric x2:x6 and x1:x7*. If the receiver has an odd number of lanes where N is greater than or equal to 3 for asymmetry, then the central lane shall ignore the output of its data detect circuit, such as lane1 of the lower, asymmetric Rx3 in (c) and (d) of *Symmetric x4 to Asymmetric Tx3:Rx1 and Tx1:Rx3*.

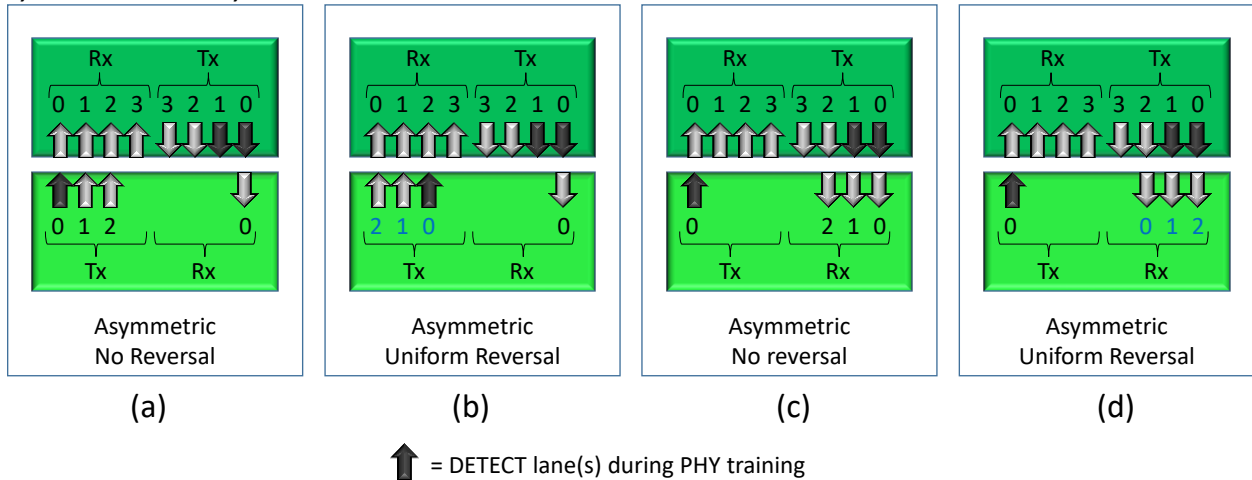


Figure 6-2: Symmetric x4 to Asymmetric Tx3:Rx1 and Tx1:Rx3

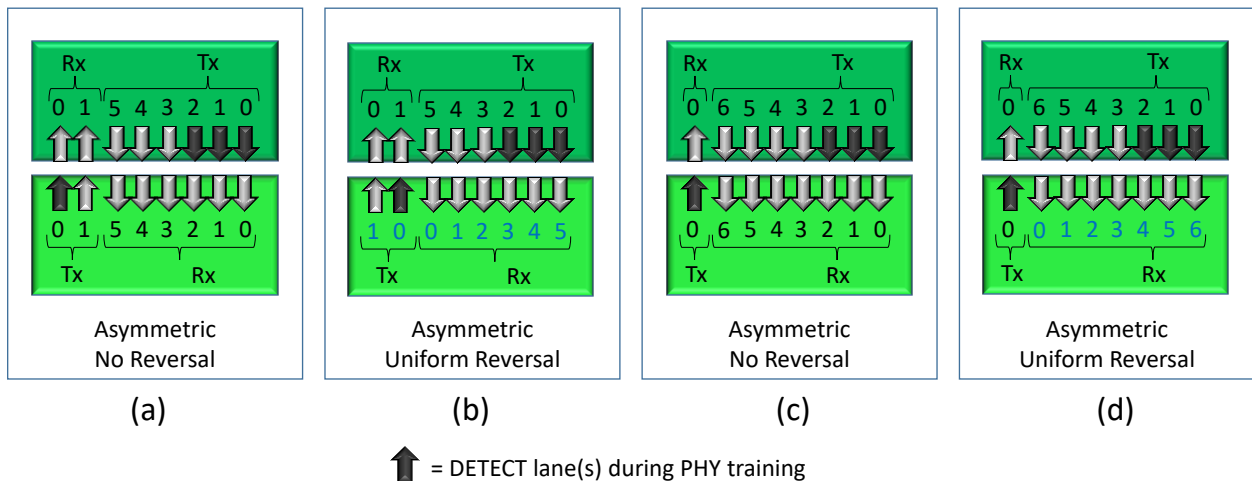


Figure 6-3: Asymmetric x2:x6 and x1:x7

Detect Idle

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Detect Idle*.

Detect Wait Reversal

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Detect Wait Reversal*.

Detect Wait

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Detect Wait*.

5 **6.2.11.2.3. Signal Detect exit**

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Signal Detect Exit*.

6.2.11.3. Channel Optimize (Optional)

Channel Optimize state may be implemented for this physical layer but the transition to this state shall be disabled by default.

10 **6.2.11.3.1. Channel Optimize entry**

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Channel Optimize Entry*. The Channel Optimize state may be entered from the Signal Detect state if enabled for both the local and remote physical layer via in the *Unique Interface Physical Layer Structure*.

6.2.11.3.2. Channel Optimize process

15 Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Channel Optimize Process*.

6.2.11.3.3. Channel Optimize exit

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Channel Optimize Exit*.

6.2.11.4. Align (Required)

Refer to the Gen-Z Common Specification clause section on *Align (Required)*.

20 **6.2.11.5. Interconnect BIST (Optional)**

Refer to the Gen-Z Common Specification clause section on *Interconnect BIST (Optional)*.

6.2.11.6. Config (Required)

Refer to the Gen-Z Common Specification clause section on *Config (Required)*.

6.2.11.7. PHY-Up (Required)

25 Refer to the Gen-Z Common Specification clause section on *PHY-Up (Required)*.

6.2.12. Clock Compensation

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Clock Compensation*.

6.2.13. Loopback

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Loopback*.

5 6.2.14. Physical Layer Aggregation

An implementation can support static aggregation of multiple physical layers into a single aggregated interface (SAI).

6.2.15. Re-timers

Re-timers for this specification clause shall not be supported.

10 **Developer Note:** *If a topology using this specification clause exceeds the channels requirements, then recommend that developers consider using a physical layer capable of supporting the Gen-Z-E-NRZ-25G-Fabric specification.*

6.3. Power Management

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Power Management*.

15 6.4. Electrical Sublayer

This Gen-Z physical layer electrical sublayer shall meet all single lane/channel requirements specified in the *IEEE Standard for Ethernet 802.3by™-2016* for the following clause:

- Annex 109A (normative) Chip-to-chip 25 Gigabit Attachment Unit Interface (25GAUI C2C)

The only exceptions are the following:

- 20
- Support for COM with a channel Insertion Loss¹ (IL) shall be no greater than 10 dB at 12.89 GHz.
 - COM requirements for the decision feedback equalizer (DFE) shall be modified.

As a service to the reader, many requirements of *IEEE Standard for Ethernet 802.3by™-2016* Annex 109A reference *IEEE Standard for Ethernet 802.3™-2015* for the following clauses when related to the operation of a single lane/channel:

- 25
- Clause 93. Physical Medium Dependent (PMD) sublayer and baseband medium, type 100GBASE-KR4
 - Annex 83D (normative) Chip-to-chip 100Gb/s four-lane Attachment Unit Interface (CAUI-4)
 - Annex 93A (normative) Specification methods for electrical channels
 - Annex 93B (informative) Electrical backplane reference model
 - Annex 93C (normative) Receiver interference tolerance
- 30

6.4.1. Electrical Interface

6.4.1.1. Overview

The electrical interface for this physical layer clause shall be based on high speed, low voltage differential signaling where each connection is point-to-point and signaling is unidirectional.

5 6.4.1.1.1. Forwarded Clock

This physical layer clause shall not require a forwarded clock.

6.4.1.2. Lane Signaling Rate

The signaling rate for this physical layer clause shall be as specified in Table 83D-1 of *IEEE Standard for Ethernet 802.3™-2015* clause 83E.3.1 at 25.78125 Gb/s \pm 100 Parts Per Million (ppm) per lane .

10 6.4.1.3. Lane Bit Error Rate

This physical layer clause shall operate in a no-FEC mode with an overall bit error ratio (BER) of 10^{-15} or better.

6.4.1.4. Lane Equalization

15 Tx Equalization coefficient range and step size requirements shall be as specified in *IEEE Standard for Ethernet 802.3™-2015* clause 83D.3.1.1. The interface physical layer structure should provide fields to configure the transmit equalization coefficients, which can be manipulated by manageability prior to the Align physical layer training state.

6.4.2. High Speed Signaling Specification

6.4.2.1. Transmitter Specification

20 The transmitter for this physical layer clause shall meet be as specified in *IEEE Standard for Ethernet 802.3™-2015*, clause 83D.3.1.

6.4.2.2. Receiver Specification

25 The receiver for this physical layer clause shall be as specified in Table 83D-4 of the *IEEE Standard for Ethernet 802.3™-2015*, clause 83D.3.3 with an exception that only Test 2 values in Table 83D-5 shall be required for receiver interference tolerance.

6.4.3. Channel Characteristics

6.4.3.1. Channel Operating Margin

The Channel Operating Margin (COM) for this physical layer clause shall be as specified for a single lane in *IEEE Standard for Ethernet 802.3™-2015*, clause 83D.4. The channels for this physical layer clause shall

meet this COM requirement with the Target Detector Error Ratio (DER_0) parameter value set to 10^{-15} and the decision feedback equalizer (DFE) length, N_b , set to 3 but with the first tap set to zero.

6.4.3.2. Insertion Loss

The channel insertion loss ($IL(f)$ in dB) for this physical layer clause should meet the equation illustrated in *Gen-Z-E-NRZ-25G-Local Channel Insertion Loss*.

$$IL(f) \leq \begin{cases} 1.5 + 1.7119\sqrt{f} + 0.1825f & 0.5 \leq f \leq f_b/2 \\ -37.695 + 3.7f & f_b/2 \leq f \leq f_b \end{cases} \text{ (dB)}$$

$$IL(f) \geq \begin{cases} 0 & 0.5 \leq f \leq 1 \\ 1/3(f - 1) & 1 \leq f \leq 17.5 \\ 5.5 & 17.5 < f \leq f_b \end{cases} \text{ (dB)}$$

Where,

f is the frequency in GHz

f_b is the signaling rate (25.78125) in Gb/s

$IL(f)$ is the insertion loss at frequency f




Figure 6-4: Gen-Z-E-NRZ-25G-Local Channel Insertion Loss

6.4.3.3. Return Loss

The Return Loss for this physical layer clause shall be as specified in *IEEE Standard for Ethernet 802.3™-2015*, clause 93.9.3.

6.4.3.4. Coupling

The transmitter for this physical layer clause shall be AC-coupled to the receiver as specified in *IEEE Standard for Ethernet 802.3™-2015*, clause 93.9.4.

6.4.4. Miscellaneous Specification

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Miscellaneous Specification*.

6.5. Management Control and Status

6.5.1. Interface Physical Layer Structure

6.5.1.1. Common Interface Physical Layer Structure

6.5.1.1.1. Common, Lane, and Low Power Interface PHY Structure

The following specifies the features and requirements associated with the Gen-Z-E-NRZ-25G-Local specification. The *Gen-Z-E-NRZ-25G-Local Interface PHY Structure Support Exclusions and Restrictions* indicates restricted or excluded fields or sub-fields from the *Interface PHY Structure Common, Lane, and Low Power Fields*.

Table 6-1: Gen-Z-E-NRZ-25G-Local Interface PHY Structure Support Exclusions and Restrictions

Support	Sub-field / Description	Interface PHY Structure Field	Supported Value(s)
Restricted	PHY Clause 6—25G Local	PHY Type	0x1
Restricted	Physical Layer Operational Status 0x0—PHY-Down, uninitialized state 0x1—PHY-Up 0x2—PHY-Down-Retrain 0x3-0x6—Reserved 0x7—PHY LP 1→ Mapped to P1 0x8—PHY LP 2 → Mapped to P2 0x9-0xA—Reserved 0xB—PHY-Up LP 1 → Mapped to Dynamic Link Width Reduction 0xC-0xF—Reserved	PHY Status	See Description
Restricted	Previous Physical Layer Operational Status 0x0—PHY-Down, uninitialized state 0x1—PHY-Up 0x2—PHY-Down-Retrain 0x3-0x6 Reserved 0x7—PHY LP 1→ Mapped to P1 0x8—PHY LP 2 → Mapped to P2 0x9-0xA—Reserved 0xB—PHY-Up LP 1 → Mapped to Dynamic Link Width Reduction 0xC-0xF—Reserved	PHY Status	See Description

Support	Sub-field / Description	Interface PHY Structure Field	Supported Value(s)
Restricted	Physical Layer Retraining Arc Only Retraining Arc 1 and Arc 2 are supported for this physical layer clause	PHY Control	0x0-0x1
Excluded	Retrain Arc 3-4 Support	PHY Cap 1	0b
Restricted	Enable Phit CRC Encoding Phit CRC encoding shall be disabled by default for this physical layer clause	PHY Cap 1 Control	0b (Default)
Restricted	Asymmetric Lane with Reversal Support	PHY Lane Cap	0b
Restricted	Remote Asymmetric Lane with Reversal Support	PHY Remote Lane Cap	0x0
Excluded	Entry, Exit Latency PHY-LP 3-4	PHY Low Power Timing Capability	0x0
Excluded	PHY-LP 3-4 Support	PHY Low Power CAP	0b
Excluded	Entry, Exit Latency PHY-Up-LP 2-4	PHY-Up Low Power Timing Capability	0x0
Excluded	PHY-Up-LP 2-4 Support	PHY-Up Low Power CAP	0b

6.5.1.1.2. Extended Feature Interface PHY Structure

The following specifies the features and requirements associated with the Interface PHY Structure for Extended Feature fields for the Gen-Z-E-NRZ-25G-Local specification.

Table 6-2: Gen-Z-E-NRZ-25G-Local Extended Feature Fields

Field Name	Size (bits)	Value / Bit Location	M / O	Access	Description
PHY Extended Status	32	-	M	RO	See <i>Extended Feature Status</i>

Field Name	Size (bits)	Value / Bit Location	M / O	Access	Description
PHY Extended Control	32	-	M	RW	See <i>Extended Feature Control</i>
PHY Extended Cap	32	-	M	RO	See <i>Extended Feature Capability</i>
PHY Remote Extended Cap	32	-	M	RW	See <i>Gen-Z-E-NRZ-25G-Local Extended Feature Remote Capability</i>

Table 6-3: Gen-Z-E-NRZ-25G-Local Extended Feature Status

Bit Location	Access	Description
0	RO	Physical Layer Extended Feature Operational Status 0b—PHY Extended Features disabled 1b—PHY Extended Features enabled
2:1	RO	Physical Layer Extended Feature Training Status 0x0—Training with Extended Features has not occurred 0x1—Training with Extended Features succeeded 0x2—Training with Extended Features has failed 0x3—Reserved
5:3	RO	Physical Data Encoding Mode 0x0—64b/66b 0x1—None 0x2—128b/130b 0x3-0x7—Reserved
8:6	RO	Current Transfer Rate Mode 0x0—25.78125 GT/s 0x1—20.0 GT/s 0x2—28.125 GT/s 0x3-0x7—Reserved
31:9	-	RsvdZ

Table 6-4: Gen-Z-E-NRZ-25G-Local Extended Feature Control

Bit Location	Access	Description
2:0	RW	Physical Data Encoding Mode—Determines the active line-coding mode when Extended Features are enabled. 0x0—64b/66b (Default) 0x1—None 0x2—128b/130b 0x3-0x7—Reserved
5:3	RW	Current Transfer Rate Mode—Determines the active transfer rate when Extended Features are enabled. 0x0—25.78125 GT/s (Default) 0x1—20.0 GT/s 0x2—28.125 GT/s 0x3-0x7—Reserved
31:6	-	RsvdP

Table 6-5: Gen-Z-E-NRZ-25G-Local Extended Feature Capability

Bit Location	Access	Description
0	RO	No Physical Data Encoding Support
1	RO	128b/130b Physical Data Encoding Support
2	-	Reserved
3	RO	20.0 GT/s Transfer Rate Support
4	RO	28.125 GT/s Transfer Rate Support
31:5	-	RsvdZ

Table 6-6: Gen-Z-E-NRZ-25G-Local Extended Feature Remote Capability

Bit Location	Access	Description
0	RO	No Physical Data Encoding Support
1	RO	128b/130b Physical Data Encoding Support
2	-	RsvdZ

Bit Location	Access	Description
3	RO	20.0 GT/s Transfer Rate Support
4	RO	28.125 GT/s Transfer Rate Support
31:5	-	RsvdZ

6.5.1.2. *Unique Interface Physical Layer Structure*

The Unique Interface Physical Layer structure for the Gen-Z-E-NRZ-25G-Local specification references the *Unique Interface Physical Layer Structure* from the Gen-Z-E-NRZ-25G-Fabric specification clause. The *Gen-Z-E-NRZ-25G-Local Unique Interface PHY Structure Requirements* shall specify any differences.

Table 6-7: Gen-Z-E-NRZ-25G-Local Unique Interface PHY Structure Requirements

Differences	Sub-field / Description	Interface PHY Structure Field
Default Value	Cursor Control Mode 0x0—Automatic PMD in-band exchange 0x1—Direct Interface PHY Structure control (Default) 0x2-0x3—Reserved	<i>PHY TX Lane Control Fields</i>
Default Value	Channel Optimize State Bypass Control The Channel Optimize state may be implemented for this physical layer clause. 0b—Channel Optimize state bypassed (Default) 1b—Channel Optimize state enabled	<i>PHY Channel Optimize State Control Fields</i>

5 6.5.2. **Physical Layer Management**

Refer to the Gen-Z Common Specification clause section on *Physical Layer Management*.

7. Gen-Z-E-NRZ-PCIe Specification

7.1. Introduction

The physical layer specified in this clause shall not require any features detailed in the Gen-Z Common Specification clause.

5 Gen-Z-E-NRZ-PCIe specifies a Gen-Z physical layer capable of operating at PCI Express speeds up to 32 GT/s using Non-Return to Zero signaling. A compliant physical layer shall meet all requirements specified in the *PCI Express Base Specification* for the following chapters:

- Physical Layer Logical Block Specification
- Power Management
- 10 • Electrical Sub-block

The only exception and addition are the following:

- Framing of Gen-Z packets shall be as specified in the *Framing* section of this physical layer clause.
- Crosslink support shall be required for Gen-Z physical layers on components that support Switch or both Gen-Z Requestor and Responder functionality as specified in *Physical Layer Initialization and Training*.
- 15

Gen-Z PCIe Physical Layer illustrates the Gen-Z Core and PLA connection in this clause. The PLA-to-PCIe interface block is implementation specific.

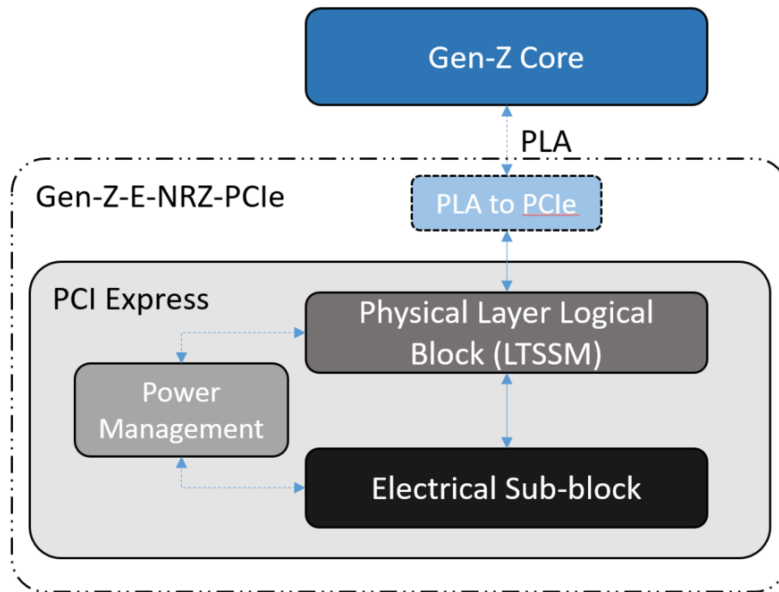


Figure 7-1: Gen-Z PCIe Physical Layer

20

7.2. Logical Sublayer

Logical Sublayer exceptions and additions from the *PCI Express Base Specification* shall be as specified in this section.

7.2.1. Link Serialization

5 Serialization and deserialization shall be as specified in *PCI Express Base Specification*.

7.2.1.1. Encoding / Decoding

Encoding and decoding shall be as specified in *PCI Express Base Specification*.

7.2.1.2. Framing

10 The Gen-Z PCIe Physical Layer shall support PCIe data framing using the 8b/10b and 128b/130b PCIe encoding rules. PCIe framing is used to differentiate Gen-Z packet traffic from PCIe Ordered-Sets. All 8b/10b and 128b/130b data framing rules specified in the *PCI Express Base Specification* for the PCIe Physical Layer Logical Block Specification shall apply with the following exceptions:

- SDP or EDB tokens shall not be generated.
- PMUX packets shall not be supported.
- 15 • Physical Layer TLPs shall contain Gen-Z packets rather than PCIe TLP traffic.
- DLLP packets shall not be generated.
- LCRC shall not be generated by the transmitter or checked at the receiver. Instead, Gen-Z packet data shall occupy the location where LCRC is normally located.
- 20 • TLP sequence numbers shall always be 0, and PCIe Packet Retry shall not be used. Instead, Gen-Z retry mechanisms shall be used for transient error recovery, and Gen-Z ECRC, PCRC and optional Phit CRC shall be used for Gen-Z packet protection and packet boundary determination.

Once the Gen-Z PCIe Physical Layer transitions to the PHY-Up state, the Gen-Z packets shall be transmitted using the following rules:

- 25 • TLPs shall start with a STP and be fully encapsulated with one or more consecutive Gen-Z packets.
- Gen-Z packets may span TLP boundaries.
- TLP length shall be within the minimum and maximum TLP length requirements as specified by the *PCI Express Base Specification*.
- 30 • PCIe Ordered-Sets and PCIe Framing Tokens shall be removed from received data before transmitting the received data to Gen-Z Core.
- SKP OS may be inserted in the data stream between TLPs. Gen-Z packet transmission shall be periodically halted to insert SKP Ordered-Sets as specified by the *PCI Express Base Specification*.
- Gen-Z packet transmission shall halt when the data stream needs end to transmit Ordered-Sets for link state transitions.
- 35 • If 8b/10b encoding is used, then:
 - The start of a TLP shall begin with a STP symbol followed by two bytes of Gen-Z packet stream (which replace the Sequence Number) followed by a Gen-Z packet stream while replaces the TLP.

5

- Gen-Z packets shall be scrambled and encoded as data symbols using 8b/10b scrambling and encoding rules for TLP packets as specified by the *PCI Express Base Specification*.
- At the end of the TLP, the four bytes of LCRC shall be replaced by four bytes of a Gen-Z packet stream and shall be followed by an END symbol.
- Received 8b/10b data symbols shall be decoded and descrambled as specified by the *PCI Express Base Specification* to identify Gen-Z packets.
- If receiver errors as specified by the *PCI Express Base Specification* are detected, e.g. 8b/10b decode errors, then:
 - The PLA signal PHY_RX_PHITCRC_ERR shall be asserted for one cycle if Phit CRC encoding is enabled.
 - The PLA signal PHY_RX_TRANSIENT_ERR shall be asserted for one cycle if Phit CRC encoding is disabled.

10

- If 128b/130b encoding is used, then:

15

- Start of Data Stream (SDS) Ordered-Set shall be transmitted to start the data stream following physical layer training.
- The data stream may contain TLPs, IDL tokens, or an EDS token.
- TLPs shall start with an STP token with the TLP length, FCRC, and parity and a TLP sequence number of 0. The STP shall be followed by scrambled and 128b/130b encoded data blocks containing Gen-Z packets from the Gen-Z Core.
- All framing rules specified in the *PCI Express Base Specification* regarding STP, IDL, EDS, and Ordered-Sets shall be followed.
- Received 128b/130b data symbols shall be decoded and descrambled from TLPs as specified by the *PCI Express Base Specification* to identify Gen-Z packets.
- If a framing error is detected as specified in the *PCI Express Base Specification*, then the receiver shall signal an error to the Gen-Z Core by asserting the PLA signal PHY_RX_RETRAIN_ERR and then follow the framing error recovery sequence specified in the *PCI Express Base Specification*.

20

25

30

Developer Note: The number of Gen-Z packets encapsulated within a TLP is a function of the TLP length and the Gen-Z packet length, which may be various sizes.

It is recommended that long TLPs, which generate the fewest IDL tokens between TLPs, be used to reduce PCIe Framing bandwidth overhead.

An example of link traffic on an x4 link with 8b/10b encoding is illustrated in *Gen-Z Packets on x4 link with 8b/10b PCIe Encoding*.

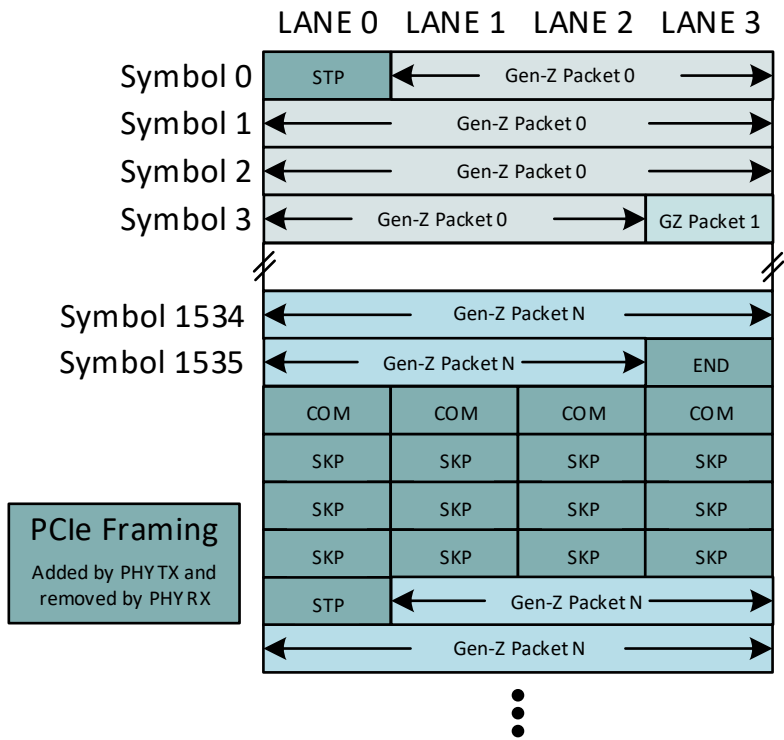


Figure 7-2: Gen-Z Packets on x4 link with 8b/10b PCIe Encoding

Gen-Z Packets on x1 link with 128b/130b PCIe Encoding illustrates an example of link traffic on an x1 link with 128b/130b encoding.

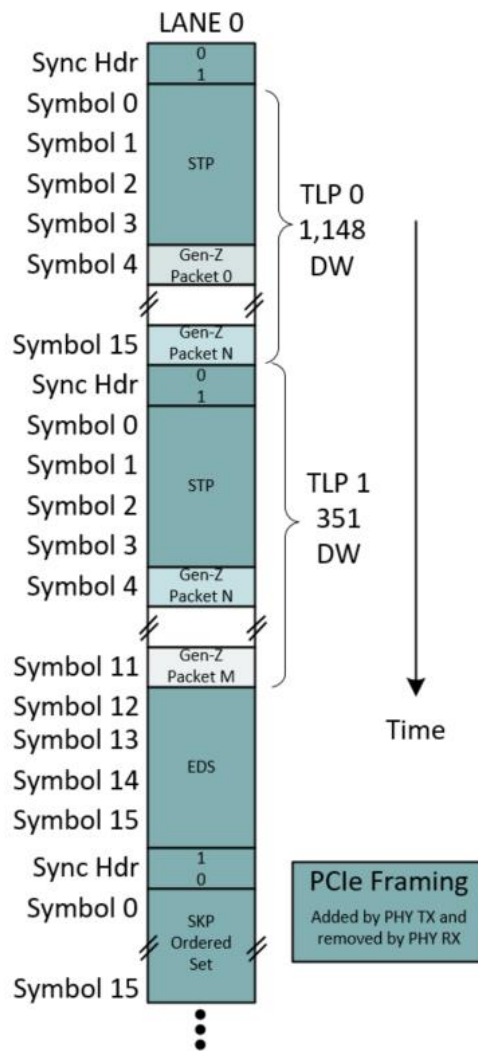


Figure 7-3: Gen-Z Packets on x1 link with 128b/130b PCIe Encoding

5 *Gen-Z Packets on x4 link with 128b/130b PCIe Encoding* illustrates an example of link traffic on an x4 link with 128b/130b encoding.

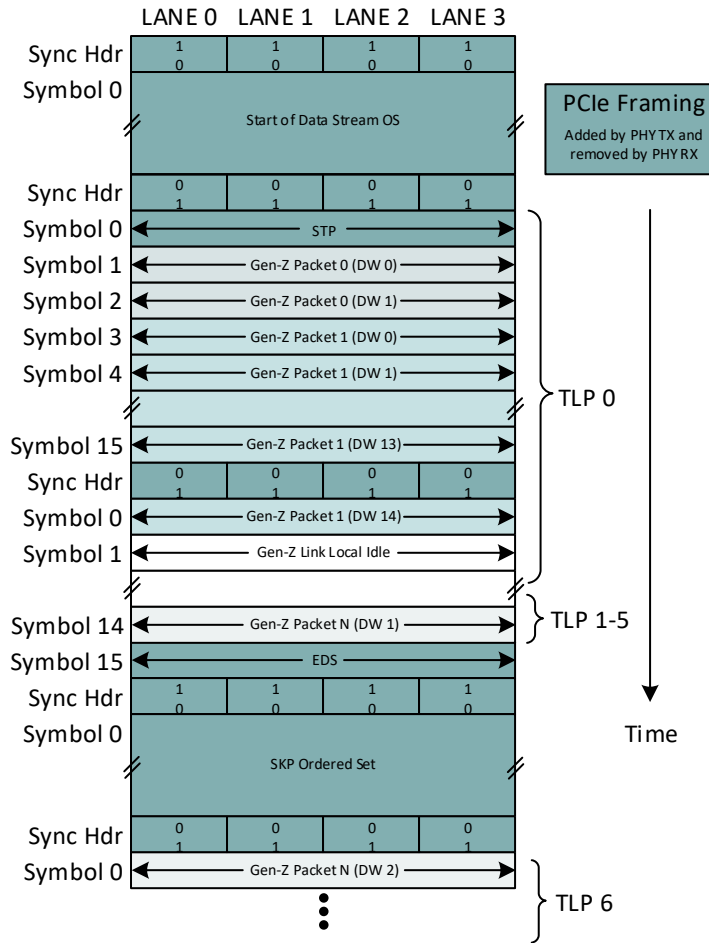


Figure 7-4: Gen-Z Packets on x4 link with 128b/130b PCIe Encoding

Phit CRC on an 8-lane PCIe Physical Layer link illustrates an example of how a phit may be transported on an 8-lane link with a PCIe Physical Layer. Only Gen-Z packets shall be included in the Phit. PCIe symbols, sync headers, Order Sets, and framing tokens shall not be included in the Phit.

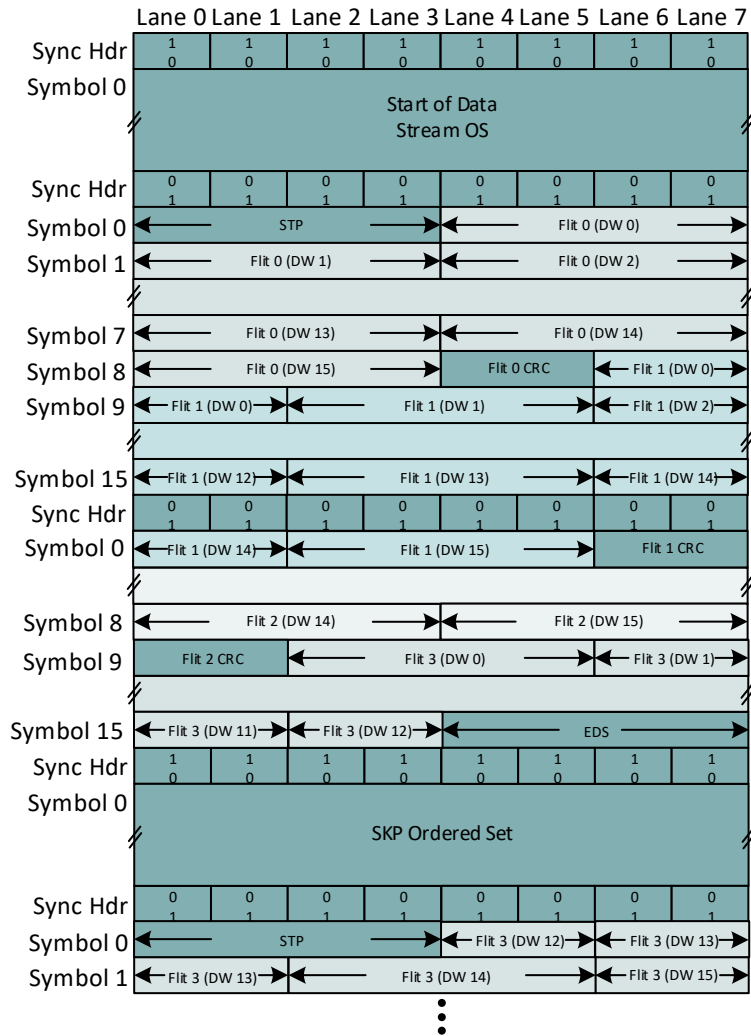


Figure 7-5: Phit CRC on an 8-lane PCIe Physical Layer link

7.2.1.3. Precoding

PCI Express Gen 5 introduced precoding to help reduce bit flips caused by DFE-induced burst errors. Transmitters are required to support precoding, and receivers are encouraged to support precoding on channels with high DFE first tap coefficients (e.g. $H1/H0 > 0.5$).

7.2.2. Scrambling / Descrambling

Scrambling and descrambling for this physical layer clause shall be as specified in *PCI Express Base Specification*.

7.2.3. Data Striping

Data striping for this physical layer clause shall be as specified in *PCI Express Base Specification*.

7.2.4. PLA Data Width

The PLA data width requirement is determined using a three-step process.

1. The data striping requirement for PCI Express physical layers is specified in *Data Striping*.
2. The PLA data multiple is determined as a function of data striping and the number of lanes supported by PCI Express physical layers and may be as specified in *PLA Data Width Multiple (in bytes) vs Number of PCIe Lanes Used*. Note this table is not inclusive of all possible PLA data widths and is provided only as an example for how to apply the *Gen-Z Core Specification*.
3. The physical layer implementer chooses a PLA data width that may be a multiple of the number of bytes specified in *PLA Data Width (in bytes) vs Number of PCIe Lanes Used and Desired Multiplier*. Note this table is not inclusive of all possible PLA data widths and is provided only as an example for how to apply the *Gen-Z Core Specification*.

Table 7-1: PLA Data Width Multiple (in bytes) vs Number of PCIe Lanes Used

Number of Lanes in the Gen-Z Link	1	2	4	8	12	16	32
PLA Data width multiple [bytes]	4	4	4	8	12	16	32

Table 7-2: PLA Data Width (in bytes) vs Number of PCIe Lanes Used and Desired Multiplier

Number of Lanes in the Gen-Z Link	1	2	4	8	12	16	32
PLA Data Multiplier = 1	4	4	4	8	12	16	32
PLA Data Multiplier = 2	8	8	8	16	24	32	64
PLA Data Multiplier = 3	12	12	12	24	36	48	96
PLA Data Multiplier = 4	16	16	16	32	48	64	128

PLA Data Multiplier = 8	32	32	32	64	96	128	256
-------------------------	----	----	----	----	----	-----	-----

7.2.5. Physical Layer Initialization and Training

Physical layer initialization and training for this physical layer clause shall be as specified in *PCI Express Base Specification*.

At Gen3 speeds and higher, the physical layer specified in this clause requires components to determine their port orientation, upstream facing or downstream facing, so the PCIe LTSSM Equalization state can transition through its sequential phases. Traditional system topologies using PCI Express are constructed to statically resolve upstream and downstream facing port orientation for each physical layer in a link. However, systems composed of Gen-Z components may use PCIe physical layers in a new-type of topology to create a link where port orientation can't easily resolve.

Crosslink as specified in the *PCI Express Base Specification* provides a method to reconcile port orientation during initial Gen1 physical layer training. As specified in *Gen-Z-E-NRZ-PCIe Port Orientation Support*, crosslink should be supported by Gen-Z switches and Gen-Z components containing both Requestor and Responder functionality. For Gen-Z components with either Requestor or Responder functionality, crosslink may be supported.

In new Gen-Z system topologies where port orientation cannot be resolved, management software shall be responsible for programming port orientation unless crosslink is supported by both the remote and local physical layers. For port orientation configuration of Gen-Z components with PCIe physical layers, refer to the *Gen-Z Management Architecture Specification*.

Table 7-3: Gen-Z-E-NRZ-PCIe Port Orientation Support

Gen-Z Functionality per PHY	Downstream (Facing) Port	Upstream (Facing) Port	Crosslink
Requestor Only	Required	Permitted	Permitted
Responder Only	Permitted	Required	Permitted
Requestor and Responder	Required	Required	Recommended
Switch	Required	Required	Recommended

Developer Note: For conceptual purposes only, a Gen-Z Requestor can equate to a PCIe Root Port, and a Gen-Z Responder can equate to a PCIe End Point. PCIe PHY port orientation when connecting a Gen-Z Requestor to a Gen-Z Responder would look similar to a Root Port to End Point connection. However, a Gen-Z device may have both Requestor and Responder functionality, and in Gen-Z systems, such dual-function devices may connect their physical layers together. Complexities arise when those devices use PCIe physical layers as port orientation now should be reconciled so only one is an upstream facing port while the other is a downstream facing port.

7.2.5.1. **PLA State Map**

The PCIe Physical Layer shall use the Link Training and Status State Machine (LTSSM) and physical layer training and state transition protocol as specified in the *PCIe Base Specification*. A Gen-Z Core shall control and report the status of LTSSM states across the PLA interface. *PLA Link State to LTSSM State Mapping* specifies the mapping of LTSSM states to PLA Link States.

Table 7-4: PLA Link State to LTSSM State Mapping

PLA State	LTSSM State
PHY-Down	Disabled
PHY-Up	L0
PHY-LP1	L1
PHY-LP2	L2
PHY-LP3	L3
PHY-LP4	L0s
PHY-Down-Retrain	Recovery
All Others	Unused

7.2.5.2. **Physical Layer Link versus Lane Training**

Physical layer link versus lane training shall be as specified in *PCI Express Base Specification*.

7.2.5.3. **Physical Layer (PHY) Sequences**

PHY Sequences for this physical layer clause shall be as specified in *PCI Express Base Specification*.

7.2.5.4. **Alternate Protocol Negotiation**

PHYs that support PCI Express Gen 5 or higher shall request to use the Gen-Z protocol via Alternate Protocol Negotiation. Alternate Protocol Negotiation shall follow the following sequence:

1. Upstream and downstream ports set the “Enhanced Link Behavior Control” field in TS1/2 in Polling/Configuration to “Modified TS1/TS2 Order Sets supported” (11b)
2. Both ports transmit Modified TS1/2s with “Modified TS Usage”=“Alternate Protocol Negotiation” (010b):
 - a. Downstream port requests the Gen-Z Alternate Protocol in Modified TS1s with Alternate Protocol Negotiation Status”=00b and “Alternate Protocol Vendor ID”=“Gen-Z Vendor ID” (0x1E69) in Configuration.Linkwidth.Accept through Configuration.Lanenum.Accept.

- b. Upstream port responds with Modified TS1s with “Alternate Protocol Negotiation Status”=“Requested Protocol Supported” (10b) and “Alternate Protocol Vendor ID”=“Gen-Z Vendor ID” (0x1E69) in Configuration.Linkwidth.Accept through Configuration.Lanenum.Accept.
 - 5 c. Upstream and Downstream ports exchange modified TS2s in Configuration.Complete with “Alternate Protocol Negotiation Status”=00b and “Alternate Protocol Vendor ID”=Gen-Z.
3. After successful protocol negotiation:
- a. Train to Gen1 L0
 - 10 b. Switch to Gen5 and perform equalization
 - c. Reach L0 and send SDS
 - d. Transmit Gen-Z packets encapsulated in TLPs as described in *Framing*.

An example PCI Express Alternate Protocol Negotiation to Gen-Z is shown in *PCIe to Gen-Z Alternate Protocol Negotiation Example*.

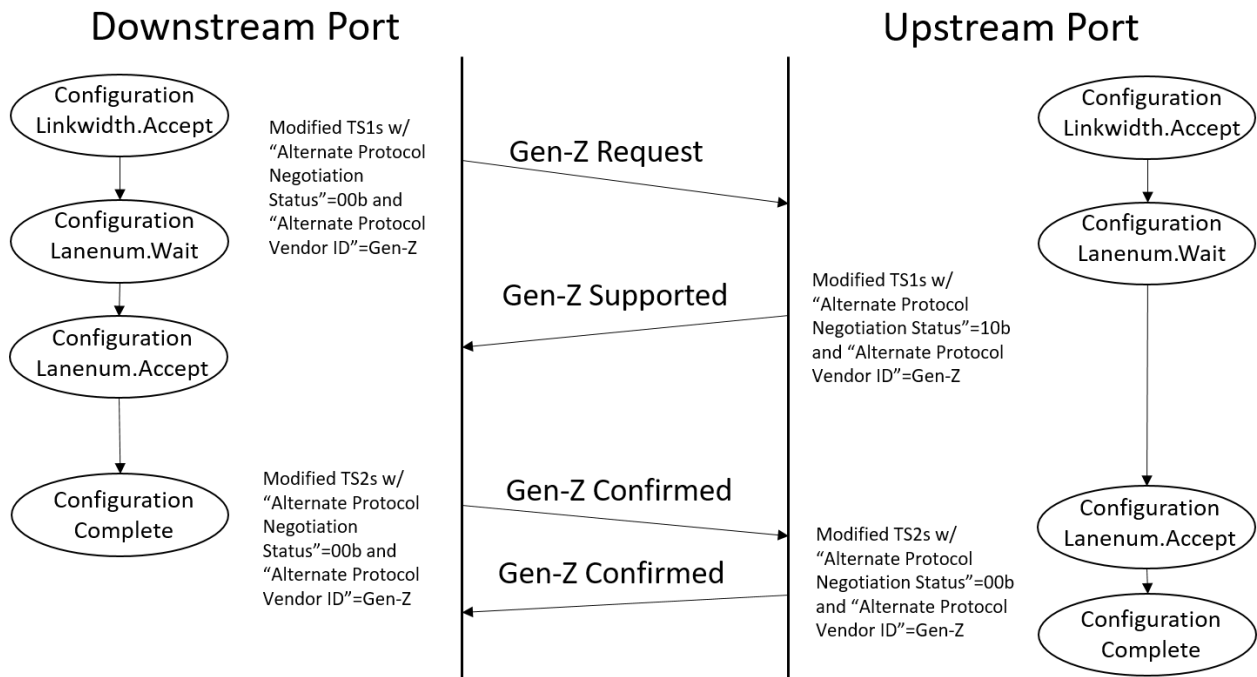


Figure 7-6: PCIe to Gen-Z Alternate Protocol Negotiation Example

7.2.6. Lane Reversal

A physical layer may support lane reversal and if so, shall follow the rules specified in *PCI Express Base Specification*.

7.2.7. Lane Polarity

A physical layer may support lane polarity reversal and if so, shall follow the rules specified in *PCI Express Base Specification*.

7.2.8. Link Width

- 5 This physical layer clause may support symmetric link widths of 1, 2, 4, 8, 12, 16, or 32 lanes as specified in *PCI Express Base Specification*.

7.2.9. Link BIST

This physical layer clause shall not support link BIST.

7.2.10. Physical Layer Retraining

- 10 This physical layer clause shall support physical layer retraining as executed in the Recovery state of the Link Training and Status State Machine (LTSSM) as specified in *PCI Express Base Specification*.

7.2.11. PLTSM

The PLTSM for this physical layer clause shall be as specified in *PCI Express Base Specification*.

7.2.12. Clock Compensation

- 15 This physical layer clause shall support clock compensation as specified in *PCI Express Base Specification*.

7.2.13. Loopback

7.2.13.1. Master Loopback (TX → RX)

This physical layer clause may support physical layer master loopback, transmitter (TX) to receiver (RX).

7.2.13.2. Slave Loopback (Rx → TX)

- 20 Slave loopback, receiver (RX) to transmitter (TX), for this physical layer clause shall be as specified in *PCI Express Base Specification*.

7.2.14. Physical Layer Aggregation

This physical layer clause shall not support dynamic (runtime) aggregation or disaggregation of interfaces.

- 25 An implementation can support static aggregation of multiple physical layers into a single aggregated interface (SAI).

A physical layer shall include its own PLA interface and Interface PHY Structure for each operational Core Interface Structure.

7.2.15. Re-timers

5 This physical layer clause shall support protocol aware re-timers as specified in *PCI Express Base Specification*.

Developer Note: *PCI Express re-timers resolve port orientation during initial link training and have an upstream and downstream port, not both of either type. In Gen-Z topologies, re-timer placement and usage needs to be carefully evaluated to ensure port orientation requirements for this physical layer clause are satisfied.*

10 *In closed systems, manageability can statically configure port orientation if support is provided.*

7.3. Power Management

This physical layer clause shall support power management modes as specified in *PCI Express Base Specification*.

15 This physical layer clause may support the physical layer power management states L0s, L1, and L2/L3 as specified in *PCI Express Base Specification*. Note that some power management states may be optional for Active State Power Management (ASPM) unless specifically required by the form factor.

7.4. Electrical Sublayer

This electrical interface for this physical layer clause shall meet all requirements specified in the *PCI Express Base Specification*.

20 7.5. Management Configuration & Status

7.5.1. Interface Physical Layer Structure

The Interface PHY Structure is specified in the *Gen-Z Core Specification* to support multiple physical layers. The mapping for the Interface PHY Structure associated with the physical layer specified in this clause is detailed below.

25 7.5.1.1. Common Interface Physical Layer Structure

7.5.1.1.1. Common, Lane, and Low Power Interface PHY Structure

The following specifies the features and requirements associated with the Gen-Z-E-NRZ-PCIe specification. The *Gen-Z-E-NRZ-PCIe Interface PHY Structure* indicates restricted or excluded fields or sub-fields from the *Interface PHY Structure Common, Lane, and Low Power Fields*.

Table 7-5: Gen-Z-E-NRZ-PCIe Interface PHY Structure Support Exclusions and Restrictions

Support	Sub-field / Description	Interface PHY Structure Field	Supported Value(s)
Restricted	PHY Clause 7—PCI Express electrical, logical, and LTSSM	PHY Type	0x2
Restricted	Physical Layer Operational Status 0x0—PHY-Down, uninitialized state 0x1—PHY-Up 0x2—PHY-Down-Retrain 0x3-0x6—Reserved 0x7—PHY LP 1 → Mapped to L1 0x8—PHY LP 2 → Mapped to L2 0x9—PHY LP 3 → Mapped to L3 0xA—PHY LP 4 → Mapped to L0s 0xB-0xF—Reserved	PHY Status	See Description
Restricted	Previous Physical Layer Operational Status 0x0—PHY-Down, uninitialized state 0x1—PHY-Up 0x2—PHY-Down-Retrain 0x3-0x6—Reserved 0x7—PHY LP 1 → Mapped to L1 0x8—PHY LP 2 → Mapped to L2 0x9—PHY LP 3 → Mapped to L3 0xA—PHY LP 4 → Mapped to L0s 0xB-0xF—Reserved	PHY Status	See Description
Restricted	Physical Layer Retraining Arc Only Retraining Arc 1 shall be supported (0x0) for this physical layer clause	PHY Control	0x0
Excluded	Retrain Arc 2-4 Support	PHY Cap 1	0b
Excluded	Retrain Arc 2-4 worst case, maximum retraining time	PHY Cap 1	0x0
Restricted	Enable Phit CRC Encoding Phit CRC encoding shall be enabled by default for this physical layer clause	PHY Cap 1 Control	1b (Default)
Restricted	Asymmetric Lane Status Asymmetry shall not be supported for this physical layer clause	PHY Lane Status	0x0

Support	Sub-field / Description	Interface PHY Structure Field	Supported Value(s)
Restricted	Enable Lane Asymmetry	PHY Lane Control	0x0
Restricted	Asymmetric Lane Support	PHY Lane Cap	0x0
Restricted	Reversal Support Only uniform reversal shall be supported for this physical layer clause	PHY Lane Cap	0x0-0x2
Restricted	Asymmetric Lane with Reversal Support	PHY Lane Cap	0b
Restricted	Remote Asymmetric Lane Support	PHY Remote Lane Cap	0x0
Restricted	Remote Reversal Support	PHY Remote Lane Cap	0x0-0x2
Restricted	Remote Asymmetric Lane with Reversal Support	PHY Remote Lane Cap	0b
Restricted	Exit latency from PHY-LP1 to PHY-Up For the physical layer associated with this specification clause, this field should be remapped from the Link Capabilities Register's L1 Exit Latency Field.	PHY Low Power Timing Capability	See Description
Restricted	Exit latency from PHY-LP4 to PHY-Up For the physical layer associated with this specification clause, this field shall be remapped from the Link Capabilities Register's L0s Exit Latency Field.	PHY Low Power Timing Capability	See Description
Excluded	Entry, Exit Latency PHY-Up-LP1-4	PHY-Up Low Power Timing Capability	0x0
Excluded	PHY-Up-LP 1-4 Support	PHY-Up Low Power CAP	0b

7.5.1.1.2. Extended Feature Interface PHY Structure

The following specifies the features and requirements associated with the Interface PHY Structure for Extended Feature fields for the Gen-Z-E-NRZ-PCIe specification.

Table 7-6: Gen-Z-E-NRZ-PCIe Extended Feature Fields

Field Name	Size (bits)	Value / Bit Location	M / O	Access	Description
PHY Extended Status	32	-	M	RO	See <i>Extended Feature Status</i>
PHY Extended Control	32	-	M	RW	See <i>Gen-Z-E-NRZ-PCIe Extended Feature Control</i>
PHY Extended Cap	32	-	M	RO	See <i>Extended Feature Capability</i>
PHY Remote Extended Cap	32	-	M	RW	See <i>Gen-Z-E-NRZ-PCIe Extended Feature Remote Capability</i>

Table 7-7: Gen-Z-E-NRZ-PCIe Extended Feature Status

Bit Location	Access	Description
0	RO	Physical Layer Extended Feature Operational Status 0b—PHY Extended Features Disabled 1b—PHY Extended Features Enabled
2:1	RO	Physical Layer Extended Feature Training Status 0x0—Training with Extended Features has not occurred 0x1—Training with Extended Features succeeded 0x2—Training with Extended Features has failed 0x3—Reserved
5:3	RO	Physical Layer Transfer Generation 0x0—Gen1 0x1—Gen2 0x2—Gen3 0x3—Gen4 0x4—Gen5 0x5-0x7—Reserved
9:6	RO	Physical Layer Transfer Rate 0x0—2.5 GT/s 0x1—5.0 GT/s 0x2—8.0 GT/s

Bit Location	Access	Description
		0x3—16.0 GT/s 0x4—20.0 GT/s 0x5—25.0 GT/s 0x6—32.0 GT/s 0x7-0xF—Reserved
31:10	-	RsvdZ

Table 7-8: Gen-Z-E-NRZ-PCIe Extended Feature Control

Bit Location	Access	Description
2:0	RW	Remap Gen3 Transfer Rate See the physical layer’s capability for remapping both Gen3 and Gen4 transfer rates. When remapping the Gen3 transfer rate, Gen4 remapping should be configured to a higher transfer rate. 0x0—8.0 GT/s (Default) 0x1—16.0 GT/s 0x2—20.0 GT/s 0x3-0x7 Reserved
5:3	RW	Remap Gen4 Transfer Rate See the physical layer’s capability for remapping both Gen3 and Gen4 transfer rates. When remapping the Gen3 transfer rate, Gen4 remapping should be configured a higher transfer rate. 0x0—16.0 GT/s (Default) 0x1—20.0 GT/s 0x2—25.0 GT/s 0x3-0x7 Reserved
6	RW	Gen4 Data Precoding Enable Data precoding can help turn burst errors, caused by the receiver DFE, into two single bit errors and can be useful when remapping Gen4 to a higher transfer rate than the default.
8:7	RW	Remote Gen5 Precoding Request Control if remote transmitter precoding should be requested. If 0x0, PHY hardware is allowed to decide whether or not to request precoding be enabled on the remote transmitter. 0x0— No override. PHY Hardware decides remote precoding (Default) 0x1—Force PHY to request precoding OFF from remote transmitter 0x2— Force PHY to request precoding ON from remote transmitter

Bit Location	Access	Description
		0x3—Reserved
31:9	-	RsvdP

Table 7-9: Gen-Z-E-NRZ-PCIe Extended Feature Capability

Bit Location	Access	Description
0	RO	Remap Gen3 Transfer Rate to 16.0 GT/s Support
1	RO	Remap Gen3 Transfer Rate to 20.0 GT/s Support
3:2	-	RsvdZ
4	RO	Remap Gen4 Transfer Rate to 20.0 GT/s Support
5	RO	Remap Gen4 Transfer Rate to 25.0 GT/s Support
7:6	-	RsvdZ
8	RO	Gen4 Data Precoding Support
31:9	-	RsvdZ

Table 7-10: Gen-Z-E-NRZ-PCIe Extended Feature Remote Capability

Bit Location	Access	Description
0	RO	Remap Gen3 Transfer Rate to 16.0 GT/s Support
1	RO	Remap Gen3 Transfer Rate to 20.0 GT/s Support
3:2	-	RsvdZ
4	RO	Remap Gen4 Transfer Rate to 20.0 GT/s Support
5	RO	Remap Gen4 Transfer Rate to 25.0 GT/s Support
7:6	-	RsvdZ
8	RO	Gen4 Data Precoding Support
31:9	-	RsvdZ

7.5.1.2. Unique Interface Physical Layer Structure

The following specifies the features and requirements associated with the Interface PHY Structure for PHY-specific fields for the Gen-Z-E-NRZ-PCIe specification and are as illustrated in *Gen-Z-E-NRZ-PCIe PHY Specific Fields*. Lane Equalization Control, Margining Lane Control, Margining Lane Status, 16.0 GT/s Lane Equalization Control, and 32.0 GT/s Lane Equalization Control are windowed registers that are accessed by their corresponding Lane Index register as shown in *Gen-Z-E-NRZ-PCIe PHY Specific Windowed Fields*.

+7		+6		+5		+4		+3		+2		+1		+0		
7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	
Remote PCIe PHY Capability				PCIe PHY Capability				PCIe PHY Control				PCIe PHY Status				< Byte 0x60
Link Status Register				Link Control Register				Link Capabilities Register								< Byte 0x68
Link Status 2 Register				Link Control 2 Register				Link Capabilities 2 Register								< Byte 0x70
Link Control 3 Register								Lane Error Status Register								< Byte 0x78
R6				Lane Equalization Control Register				R5				Lane EQC Index				< Byte 0x80
16.0 GT/s Control Register								16.0 GT/s Status Register								< Byte 0x88
16.0 GT/s Local Data Parity Mismatch Status Register								16.0 GT/s Capabilities Register								< Byte 0x90
R9				16.0 GT/s Lane Equalization Control Register				R8				16.0 GT/s Lane EQC Index				< Byte 0x98
R10								Margining Port Status				Margining Port Capabilities				< Byte 0xA0
Margining Lane Status				Margining Lane Control				R11				Margining Lane Index				< Byte 0xA8
16.0 GT/s Second Retimer Data Parity Mismatch Status Register								16.0 GT/s First Retimer Data Parity Mismatch Status Register								< Byte 0xB0
32.0 GT/s Control Register								32.0 GT/s Status Register								< Byte 0xB8
R12								32.0 GT/s Capabilities Register								< Byte 0xC0
Received Modified TS Data 2 Register								Received Modified TS Data 1 Register								< Byte 0xC8
Transmitted Modified TS Data 2 Register								Transmitted Modified TS Data 1 Register								< Byte 0xD0
R14				32.0 GT/s Lane Equalization Control Register				R13				32.0 GT/s Lane EQC Index				< Byte 0xD8
Alternate Protocol Control Register								Alternate Protocol Capabilities Register								< Byte 0xE0
Alternate Protocol Data 2 Register								Alternate Protocol Data 1 Register								< Byte 0xE8
R15								Alternate Protocol Selective Enable Mask Register								< Byte 0xF0

Figure 7-7: Gen-Z-E-NRZ-PCIe PHY Specific Fields

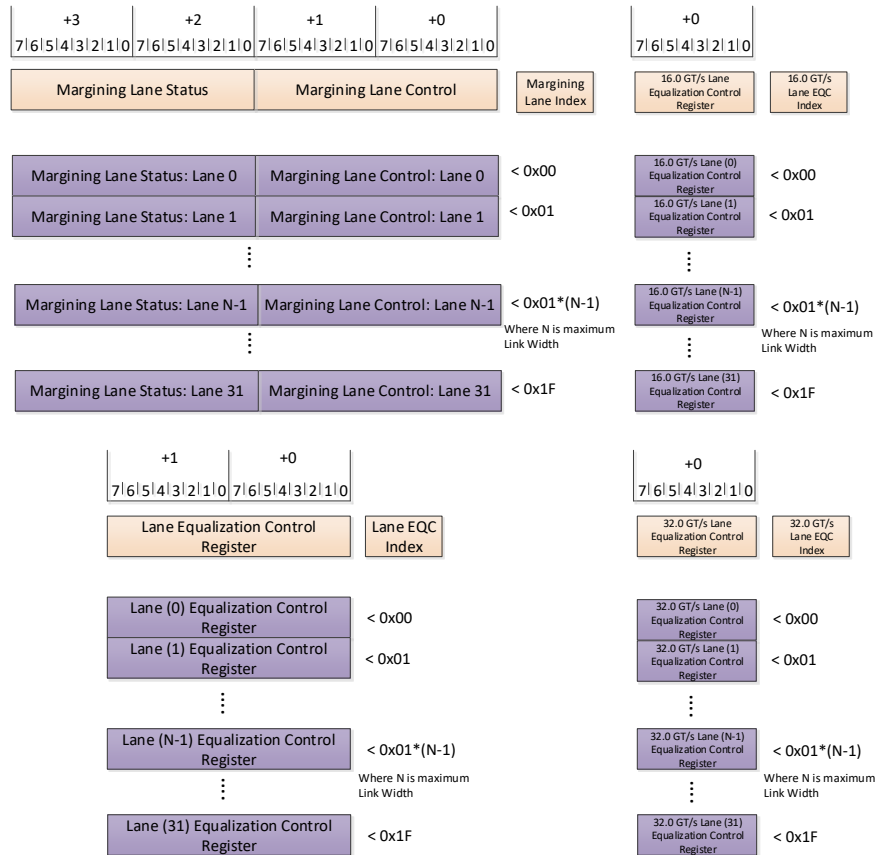


Figure 7-8: Gen-Z-E-NRZ-PCIe PHY Specific Windowed Fields

Table 7-11: Gen-Z-E-NRZ-PCIe PHY-Specific Fields

Field Name	Size (bits)	Value / Bit Location	M / O	Access	Description
PCIe PHY Status	16	-	M	RO	See <i>PCIe PHY Status Fields</i>
PCIe PHY Control	16	-	M	RW	See <i>PCIe PHY Control Fields</i>
PCIe PHY Capability	-	-	M	RO	See <i>PCIe PHY Capability Fields</i>
PCIe PHY Remote Capability	16	-	M	RW	See <i>Remote PCIe PHY Capability Fields</i>
Link Capabilities Register	32	-	M	RO	Shall be as specified by the Link Capabilities Register in the <i>PCI Express Base Specification</i> .

Field Name	Size (bits)	Value / Bit Location	M / O	Access	Description
Link Control Register	16	-	M	RW	Shall be as specified by the Link Control Register in the <i>PCI Express Base Specification</i> .
Link Status Register	16	-	M	RO	Shall be as specified by the Link Status Register in the <i>PCI Express Base Specification</i> .
Link Capabilities 2 Register	32	-	M	RO	Shall be as specified by the Link Capabilities 2 Register in the <i>PCI Express Base Specification</i> .
Link Control 2 Register	16	-	M	RW	Shall be as specified by the Link Control 2 Register in the <i>PCI Express Base Specification</i> .
Link Status 2 Register	16	-	M	RO	Shall be as specified by the Link Status 2 Register in the <i>PCI Express Base Specification</i> .
Lane Error Status Register	32	-	M	RO	Shall be as specified by the Lane Error Status Register in the <i>PCI Express Base Specification</i> .
Link Control 3 Register	32	-	M	RW	Shall be as specified by the Link Control 3 Register in the <i>PCI Express Base Specification</i> .
Lane EQC Index	5	-	M	RW	Determines the lane number for the Lane Equalization Control Register
Lane Equalization Control Register	16	-	M	RW	This interface field is addressed per lane by the Lane EQC Index field and shall be as specified by the Lane Equalization Control Register in the <i>PCI Express Base Specification</i> .
16.0 GT/s Status Register	32	-	M	RO	Shall be as specified by the 16.0 GT/s Status Register in the <i>PCI Express Base Specification</i> .
16.0 GT/s Control Register	32	-	M	RW	Shall be as specified by the 16.0 GT/s Control Register in the <i>PCI Express Base Specification</i> .
16.0 GT/s Capabilities Register	32	-	M	RO	Shall be as specified by the 16.0 GT/s Capabilities Register in the <i>PCI Express Base Specification</i> .

Field Name	Size (bits)	Value / Bit Location	M / O	Access	Description
16.0 GT/s Local Data Parity Mismatch Status Register	32	-	M	RW1C	Shall be as specified by the 16.0 GT/s Local Data Parity Mismatch Register in the <i>PCI Express Base Specification</i> .
16.0 GT/s Lane EQC Index	5	-	M	RW	Determines the lane number for the 16.0 GT/s Lane Equalization Control Register
16.0 GT/s Lane Equalization Control Register	8	-	M	RW	This interface field is addressed per lane by the 16.0 GT/s Lane EQC Index field and shall be as specified by the 16.0 GT/s Lane Equalization Control Register in the <i>PCI Express Base Specification</i> .
Margining Port Capabilities Register	16	-	O	RO	Shall be as specified by the Margining Port Capabilities Register in the <i>PCI Express Base Specification</i> .
Margining Port Status Register	16	-	O	RO	Shall be as specified by the Margining Port Status Register in the <i>PCI Express Base Specification</i> .
Margining Lane Index	5	-	O	RW	Determines the lane number for the Margining Lane Control and Margining Lane Status Register fields
Margining Lane Control Register	16	-	O	RW	This interface field is addressed per lane by the Margining Lane Index field and shall be as specified by the Margining Lane Control Register in the <i>PCI Express Base Specification</i> .
Margining Lane Status	16	-	O	RO	Shall be as specified by the Margining Lane Register in the <i>PCI Express Base Specification</i> .
16.0 GT/s First Retimer Data Parity Mismatch Status Register	32	-	M	RW1C	Shall be as specified by the 16.0 GT/s First Retimer Data Parity Mismatch Register in the <i>PCI Express Base Specification</i> .

Field Name	Size (bits)	Value / Bit Location	M / O	Access	Description
16.0 GT/s Second Retimer Data Parity Mismatch Status Register	32	-	M	RW1C	Shall be as specified by the 16.0 GT/s Second Retimer Data Parity Mismatch Register in the <i>PCI Express Base Specification</i> .
32.0 GT/s Status Register	32	-	M	RO	Shall be as specified by the 32.0 GT/s Status Register in the <i>PCI Express Base Specification</i> .
32.0 GT/s Control Register	32	-	M	RW	Shall be as specified by the 32.0 GT/s Control Register in the <i>PCI Express Base Specification</i> .
32.0 GT/s Capabilities Register	32	-	M	HwInit	Shall be as specified by the 32.0 GT/s Capabilities Register in the <i>PCI Express Base Specification</i> .
Received Modified TS Data 1 Register	32	-	M	RO	Shall be as specified by the Received Modified TS Data 1 Register in the <i>PCI Express Base Specification</i> .
Received Modified TS Data 2 Register	32	-	M	RO	Shall be as specified by the Received Modified TS Data 2 Register in the <i>PCI Express Base Specification</i> .
Transmitted Modified TS Data 1 Register	32	-	M	RO	Shall be as specified by the Transmitted Modified TS Data 1 Register in the <i>PCI Express Base Specification</i> .
Transmitted Modified TS Data 2 Register	32	-	M	RO	Shall be as specified by the Transmitted Modified TS Data 2 Register in the <i>PCI Express Base Specification</i> .
32.0 GT/s Lane EQC Index	5	-	M	RW	Determines the lane number for the 32.0 GT/s Lane Equalization Control Register
32.0 GT/s Lane Equalization	8	-	M	RW	This interface field is addressed per lane by the 32.0 GT/s Lane EQC Index field and shall be as specified by the 32.0 GT/s Lane Equalization

Field Name	Size (bits)	Value / Bit Location	M / O	Access	Description
Control Register					Control Register in the <i>PCI Express Base Specification</i> .
Alternate Protocol Capabilities Register	32	-	M	HwInit	Shall be as specified by the Alternate Protocol Capabilities Register in the <i>PCI Express Base Specification</i> .
Alternate Protocol Control Register	32	-	M	RW	Shall be as specified by the Alternate Protocol Control Register in the <i>PCI Express Base Specification</i> .
Alternate Protocol Data 1 Register	32	-	M	RO	Shall be as specified by the Alternate Protocol Data 1 Register in the <i>PCI Express Base Specification</i> .
Alternate Protocol Data 2 Register	32	-	M	RO	Shall be as specified by the Alternate Protocol Data 2 Register in the <i>PCI Express Base Specification</i> .
Alternate Protocol Selective Enable Mask Register	32	-	M	RW	Shall be as specified by the Alternate Protocol Selective Enable Mask Register in the <i>PCI Express Base Specification</i> .
R5	27	-	-	-	RsvdP
R6	16	-	-	-	RsvdP
R8	27	-	-	-	RsvdP
R9	24	-	-	-	RsvdP
R10	32	-	-	-	RsvdP
R11	27	-	-	-	RsvdP
R12	32	-	-	-	RsvdP
R13	27	-	-	-	RsvdP
R14	24	-	-	-	RsvdP

Field Name	Size (bits)	Value / Bit Location	M / O	Access	Description
R15	32	-	-	-	RsvdP

Table 7-12: PCIe PHY Status Fields

Bit Location	Access	M/O	Description
1:0	RO	M	PHY Port Orientation Status—Indicates the port orientation of the physical layer. 0x0—Upstream facing port (downstream component) 0x1—Downstream facing port (upstream component) Developer Note: For physical layers supporting Crosslink, resolution status of the port orientation may be indicated in the Link Status 2 Register as specified in the PCI Express Base Specification.
15:2	RO	-	RsvdZ

Table 7-13: PCIe PHY Control Fields

Bit Location	Access	M/O	Description
1:0	RW	M	PHY Port Orientation Control—Determines the port orientation of the physical layer in the Equalization sub-state of PCI Express link training. 0x0—Upstream facing port (downstream component) 0x1—Downstream facing port (upstream component) 0x2—Crosslink 0x3—Reserved
15:2	RW	-	RsvdP

Table 7-14: PCIe PHY Capability Fields

Bit Location	Access	M/O	Description
0	RO	M	Upstream facing port (downstream component) support
1	RO	M	Downstream facing port (upstream component) support
2	RO	M	Crosslink support

Bit Location	Access	M/O	Description
15:3			Developer Note: Physical layers supporting Crosslink may also indicate their capability in the Link Capabilities 2 Register as specified in the PCI Express Base Specification.
	RO	-	RsvdZ

Table 7-15: Remote PCIe PHY Capability Fields

Bit Location	Access	M/O	Description
0	RW	M	Upstream facing port (downstream component) support
1	RW	M	Downstream facing port (upstream component) support
2	RW	M	Crosslink support Developer Note: Physical layers supporting Crosslink may also indicate their capability in the Link Capabilities 2 Register as specified in the PCI Express Base Specification.
15:3	RW	-	RsvdP

7.5.2. Physical Layer Management

Refer to the Gen-Z Common Specification clause section on *Physical Layer Management*.

8.Gen-Z-E-PAM4-50G-Fabric Specification

8.1. Introduction

5 Gen-Z-E-PAM4-50G-Fabric defines a Gen-Z physical layer capable of a line-rate at 53.125 GT/s per lane inclusive of the overhead using 4-level Pulse Amplitude Modulation (PAM4) signaling over Fabric media (Medium Reach). One of two different Phit Forward Error Correction (Phit FEC) encodings are used, Phit FEC 288 which uses a BCH(288, 260) codeword or Phit FEC 320 which uses a BCH(320, 260) codeword. Phit FEC 288 has a raw data-rate of 47.222 Gbps and supports a raw BER of 10^{-9} or better. Phit FEC 320 has a raw data-rate of 42.5 Gbps and supports a raw BER of 10^{-7} or better. Both types of Phit FEC have a corrected BER of 10^{-15} or better. A compliant physical layer consists of an electrical interface compatible with the Channel Operating Margin (COM) that has a recommended Insertion Loss¹ (IL) no greater than 20 dB at 13.2813 GHz.

8.2. Logical Sublayer

15 The block diagram in *Gen-Z-E-PAM4-50G Transmit (Tx) Physical Layer Logical Sublayer Block* illustrates the functions and data flow through the physical layer logical sublayer on the transmit side of the data link.

The transmission of a data packet occurs as follows:

1. Transmit data is sent from the Gen-Z core to the Tx Data register over the PLA Interface.
2. The data is then encoded using Phit Forward Error Correction (FEC).
3. The data is striped according to the specifications in *Data Striping*.
- 20 4. The data is scrambled and encoded (gray coding and precoding).
5. The data is transferred to the serializer.
6. The data is transmitted serially by the lane drivers over the electrical or optical physical medium.

The data flow through the PHY logical layer for PHY Up sequences (control packets) follow the same flow as PLA data packets above.

25 The data flow through the PHY logical layer for transmitting training packets occurs as follows:

1. The PLA Tx Data and Data Striping path is bypassed via the Training or Control Sets multiplexer.
2. Some control sets are not scrambled, and in these cases, the scrambler LFSR is disabled (scrambling enable=0) such that it shall not advance and the scrambler function is bypassed.
3. Some control sets are not gray coded or precoded, and in these cases, the encoders are disabled (gray coding enable=0, precoding enable = 0) such that the function is bypassed.
- 30 4. The data is then serialized and transmitted in the same manner as the data sets.

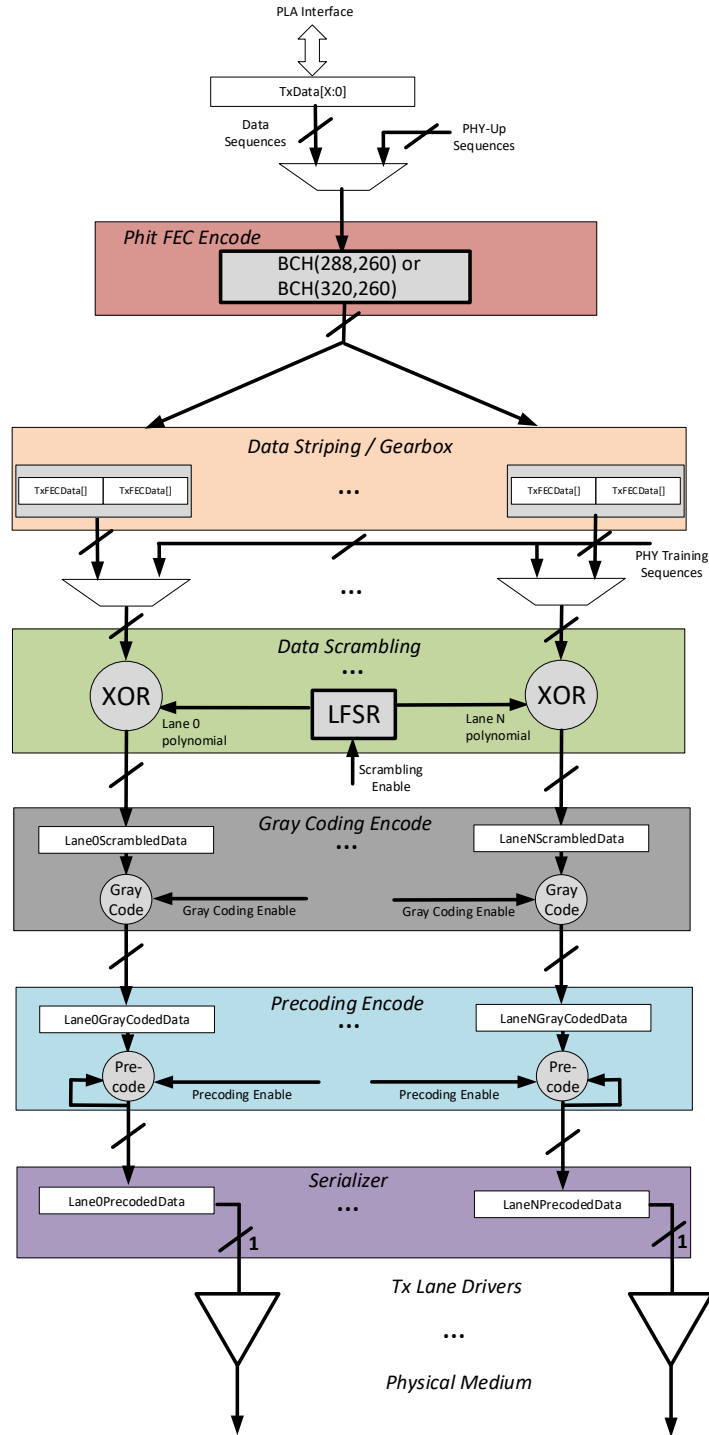


Figure 8-1: Gen-Z-E-PAM4-50G Transmit (Tx) Physical Layer Logical Sublayer Block

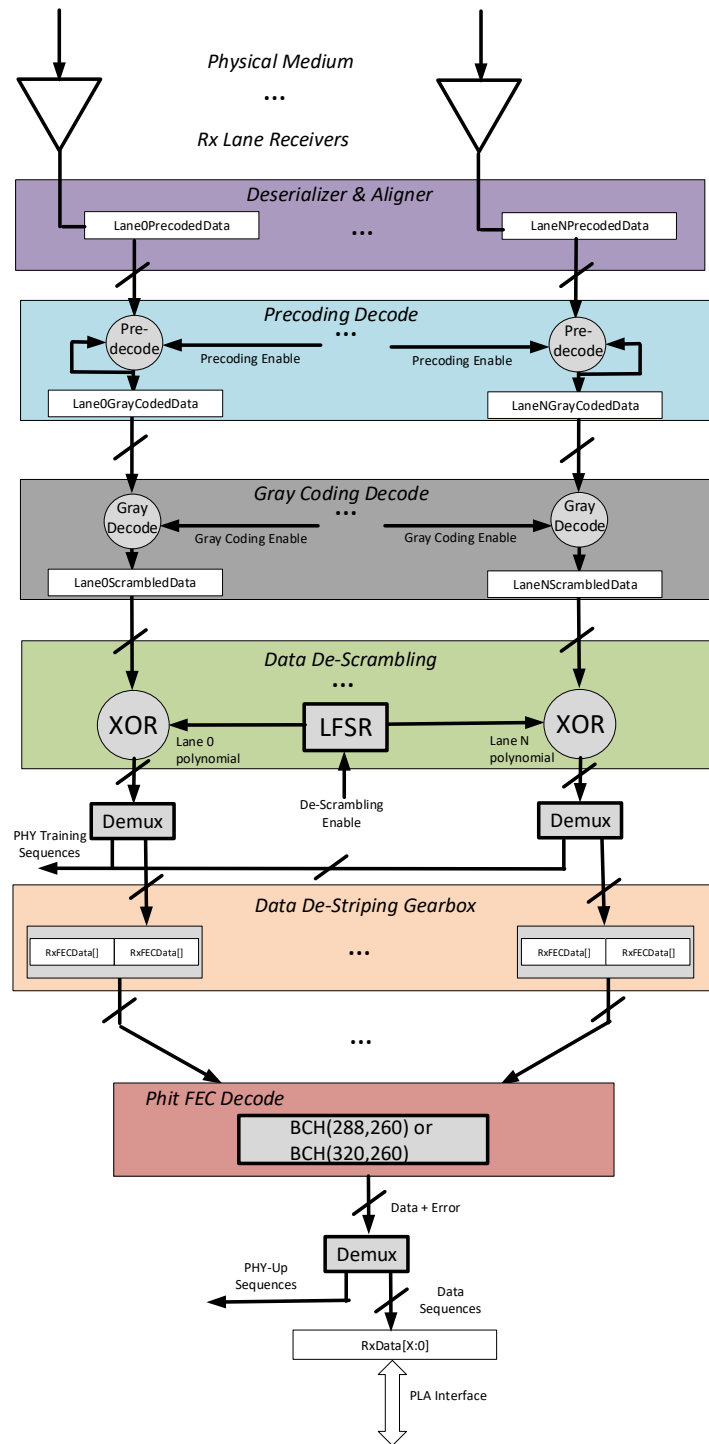


Figure 8-2: Gen-Z-E-PAM4-50G Receive (Rx) Physical Layer Logical Sublayer Block

Gen-Z-E-PAM4-50G Receive (Rx) Physical Layer Logical Sublayer Block illustrates the receiver side of the link and shows how the data is received by the PHY electrical layer and processed by the PHY logical layer in the opposite order of the transmit side.

8.2.1. Link Serialization

Refer to the Common Specification clause on *Link Serialization* with the following modifications detailed below:

- Phit Forward Error Correction (FEC) protection on data and PHY Up sequences.
- Gray Coding to support 4 level Pulse Amplitude Modulation (PAM4).
- Precoding to optionally enhance burst error correction and detection.

The following are the features and requirements associated with encoding /decoding for this physical layer clause:

- PHY Training Sequences shall not be encoded nor protected by Phit FEC.
- Data and PHY-Up Sequences shall be protected by Phit FEC, scrambled, and gray coded and, if negotiated, precoded.
- Data and PHY-Up sequences are not protected by Phit CRC.

8.2.1.1. Encoding / Decoding

Encoding and decoding is not enabled in this physical layer clause, i.e. 64b/66b encoding is not present. Instead, Phy-Up Sequences and Data Sequences are differentiated by the “ctl” bits in the Phit FEC codeword. Training Sequences are differentiated from Phit FEC codewords by the PHY based on the PHY state. Traffic after the Lock Training Sequence is Phit FEC codewords and traffic received after retraining is initiated by Retrain PHY-Up Sequences are Training Sequences.

8.2.1.2. Phit FEC

Phit FEC is a light-weight Forward Error Correction (FEC) that allows correction of most errors, reducing the frequency of transaction retries with higher BER. This is a replacement for the Phit CRC layer described in the Physical Layer Abstraction chapter in the Core specification. With Phit FEC, correctable errors are silently corrected in the receiving PHY before the Phit payload (packet stream) is transferred to the Gen-Z Core. If an uncorrectable error is detected, PHY_RX_PHITCRC_ERR is asserted along with the Phit payload. This causes the Gen-Z Link Layer to force detection of a Packet CRC (PCRC) error for any packet that may be affected (same behavior as with Phit CRC).

Two different Phit FECs are defined based on the raw BER of the link. Links with a raw BER \leq 1E-9 should use the Phit FEC 288, and links with a raw BER \leq 1E-7 and BER $>$ 1E-9 should use the Phit FEC 320. Both Phit FECs result in a corrected BER of 1E-15 or better. If raw BER is not sufficiently random to reach a corrected BER of 1E-15 or better, e.g. due to high crosstalk between lanes, the raw BER should be reduced until a corrected BER of 1E-15 or better can be reached.

8.2.1.2.1. Phit FEC 288

Format

With Phit FEC 288, each Phit is 288 bits in size, consisting of:

- 256 bits of Gen-Z packet stream (not necessarily aligned with packet boundaries).
- 4 copies of a “ctl” bit, which differentiates between Data Sequences and PHY-Up Sequences.
- 28 bits of FEC parity.

The Phit FEC 288 encoding is binary BCH (288,260). This FEC requires the use of Gray Coding and Precoding such that burst errors introduced by DFE are converted to bit errors. The Hamming distance is 8, so this code could potentially be used for correction of 3 bit errors in a Phit. Correction is restricted to 2 bit errors for the following reasons:

- With DFE, the most probable error affecting a Phit is a single burst error, which is converted to 2 bit errors by gray coding and precoding.
- Using a stronger code than strictly required for 2 bit error correction allows stronger detection of uncorrectable Phit errors. This supplements the packet CRC error protection for improved reliability with higher BER (similar to Phit CRC).

The construction of a 288-bit Phit is shown in *Construction of Phit*.

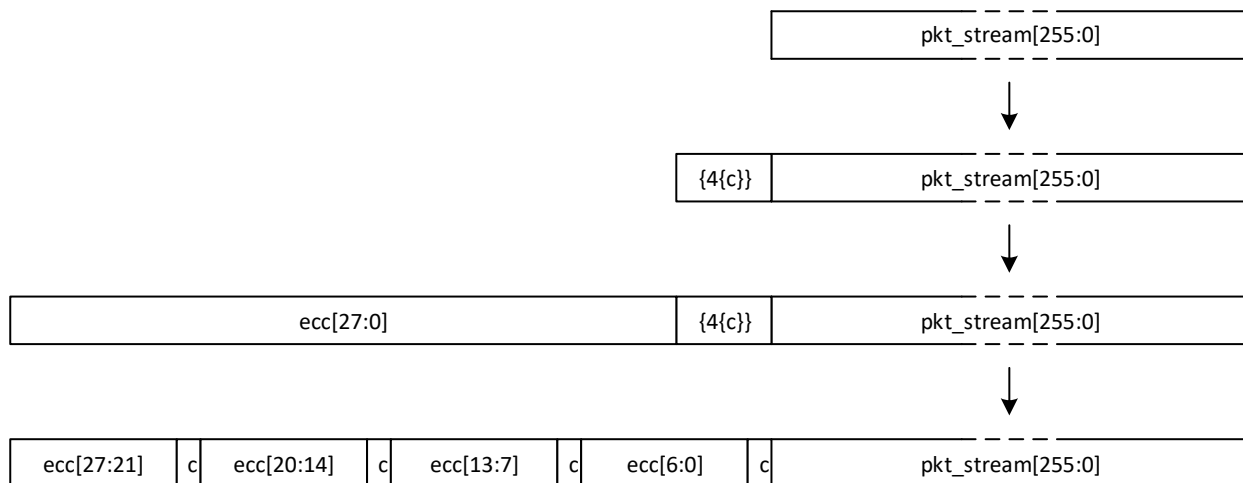


Figure 8-3: Construction of Phit FEC 288

- 1) Start with 256 bits of Gen-Z packet stream.
- 2) Append 4 copies of “ctl” bit, aka “c”.
- 3) Encode as BCH(288,260). Note: this is a binary CRC encoding, and the codeword is bit-reversed for transfer on the link, as is common for CRC codes. The “lowest” bit of the packet stream will travel first on the link and the “highest” bit of the FEC parity will travel last.
- 4) Intersperse the four copies of the “ctl” bit with the FEC parity (aka “ecc”). This reduces the probability that multiple copies will be affected by a single error event, particularly when the copies travel on separate lanes of the link.

The 288-bit Phit is striped across lanes according to the rules described in the Data Striping section in the Physical Layer Abstraction chapter in the Core specification. *Phit FEC 288 on 4-lane link* shows how a Phit is transferred on a four-lane link. In this drawing the bits indicated at the top travel first on the link. For a four-lane link, the Phit is striped across lanes with byte granularity.

When a Phit arrives in a receiving PHY, Phit FEC decode consists of:

- 1) Reverse the final step of Phit construction, so the four copies of the “ctl” bit and the 256 bits of the packet stream are once again a contiguous payload in the codeword.
- 2) Bit-reverse the codeword if necessary (depends on decoder implementation).
- 3) Perform FEC decode, restricting correction to 2 bit errors.

- 4) Check for identical (corrected) copies of the “ctl” bit. If all copies are not identical, initiate link resynchronization.
- 5) Remove “ctl” bits and transfer 256 bits of packet stream (if applicable) to the receiving link layer. If an uncorrectable error was detected by FEC decode, flag affected chunk(s) of packet stream via the PHY_RX_PHITCRC_ERR signal as described in the Physical Layer Abstraction chapter in the Core specification.

5

lane3		lane0	
24	16	8	0
25	17	9	1
26	18	10	2
27	19	11	3
28	20	12	4
29	21	13	5
30	22	14	6
31	23	15	7
56	48	40	32
57	49	41	33
58	50	42	34
59	51	43	35
60	52	44	36
61	53	45	37
62	54	46	38
63	55	47	39
⋮			
248	240	232	224
249	241	233	225
250	242	234	226
251	243	235	227
252	244	236	228
253	245	237	229
254	246	238	230
255	247	239	231
ctl	ctl	ctl	ctl
ecc21	ecc14	ecc7	ecc0
ecc22	ecc15	ecc8	ecc1
ecc23	ecc16	ecc9	ecc2
ecc24	ecc17	ecc10	ecc3
ecc25	ecc18	ecc11	ecc4
ecc26	ecc19	ecc12	ecc5
ecc27	ecc20	ecc13	ecc6

Figure 8-4: Phit FEC 288 on 4-lane link

Encode

The BCH(288,260) code is constructed as follows:

10

- 1) The decoder alphabet is $GF(2^9)$, constructed with primitive polynomial $p(x) = x^9 + x^8 + x^7 + x^6 + x^3 + x^2 + 1$, which may be represented as the binary value 0x3cd.

- 2) Each element of the field may be viewed as a power of the primitive element a , and may be represented as a 9-bit binary value:

$$\text{GF}(2^9) = \{0, a^0, a^1, a^2, a^3, a^4, a^5, a^6, a^7, a^8, a^9, \dots, a^{510}\}$$

$$= \{0x000, 0x001, 0x002, 0x004, 0x008, 0x010, 0x020, 0x040, 0x080, 0x100, 0x1cd, \dots, 0x1e6\}$$

- 3) Minimal polynomials are:

$$m_0(x) = x + a^0$$

$$= x + 1$$

$$m_1(x) = m_2(x) = m_4(x) = (x+a^1)(x+a^2)(x+a^4)(x+a^8)(x+a^{16})(x+a^{32})(x+a^{64})(x+a^{128})(x+a^{256})$$

$$= x^9 + x^8 + x^7 + x^6 + x^3 + x^2 + 1$$

$$m_3(x) = m_6(x) = (x+a^3)(x+a^6)(x+a^{12})(x+a^{24})(x+a^{48})(x+a^{96})(x+a^{192})(x+a^{384})(x+a^{257})$$

$$= x^9 + x^8 + x^6 + x^3 + x^2 + x + 1$$

$$m_5(x) = (x+a^5)(x+a^{10})(x+a^{20})(x+a^{40})(x+a^{80})(x+a^{160})(x+a^{320})(x+a^{129})(x+a^{258})$$

$$= x^9 + x^8 + x^7 + x^3 + x^2 + x + 1$$

- 4) The code generator polynomial is calculated as:

$$g(x) = \text{LCM}(m_0(x), m_1(x), \dots, m_6(x))$$

$$= m_0(x) m_1(x) m_3(x) m_5(x)$$

$$= x^{28} + x^{23} + x^{22} + x^{20} + x^{17} + x^{13} + x^{11} + x^{10} + x^6 + x^5 + x^4 + x^3 + x + 1$$

Given this code generator polynomial, a 260-bit payload is viewed as a polynomial $m(x)$ of degree 259 over $\text{GF}(2)$. A codeword is viewed as a polynomial of degree 287, constructed as:

$$c(x) = x^{28} m(x) - (x^{28} m(x) \bmod g(x))$$

and represented as a 288-bit binary value.

In practice, each bit of FEC parity is calculated as even parity over a set of payload bits, with the parity logic derived from the code generator polynomial.

Decode

The decode process starts with calculating a 6-element power-sum syndrome vector: $\{S_0, S_1, S_2, S_3, S_4, S_5\}$. Each power sum syndrome is calculated by viewing the received codeword as a polynomial $r(x)$ over $\text{GF}(2^9)$ and evaluating $r(x)$ for a particular element of $\text{GF}(2^9)$:

$$S_0 = r(a^0)$$

$$S_1 = r(a^1)$$

:

The power sum syndromes are further decoded with known algorithms to determine whether an error is present, whether it is correctable, and if so, the location(s) of the error(s).

Developer Note: Given the size of the Phit, it will generally be the case that Phits are flowing through the receiving PHY in back-to-back clock cycles (in the deserialized clock domain). An FEC decode of a 2-bit error will likely require several clock cycles to complete. Therefore, it is recommended that the FEC

decoder should be pipelined using a non-iterative method such that it can start a new Phit decode in each clock cycle. As a latency optimization, implementations may choose to bypass correction if the correction pipeline is empty and no error is detected on the Phit.

5 Uncorrectable Error Detection

As previously mentioned, correction is artificially restricted to 2 bit errors. The detection of uncorrectable errors by Phit FEC decode supplements the error detection provided by Gen-Z packet CRC. If an uncorrectable error is detected, PHY_RX_PHITCRC_ERR is asserted along with the Phit payload. To maximize the error detection capability of the Phit FEC, the following sanity checks are recommended (may be some redundancy):

- 1) To confirm decode of a correctable 1 bit error, check the following:
 - Both coefficients of the ELP (s_0 and s_1) are non-zero.
 - Error location < 288 .
 - Calculate an Error Evaluator Polynomial $O(x)$. The syndrome vector may be viewed as coefficients of a degree 5 polynomial $S(x)$ over $GF(2^9)$, Then: $O(x) = S(x) L(x) \bmod x^6$ (where $L(x)$ is the 1-error ELP). Confirm that the degree 0 coefficient of $O(x)$ is non-zero, and all other coefficients are 0.
- 2) To confirm decode of a correctable 2 bit error, check the following:
 - All three coefficients of the ELP are non-zero.
 - Assuming the ad hoc root finding method described above, confirm that the two roots found actually are roots of $L(x)$: $x_0 x_1 == L_0/L_2$ and $x_0 + x_1 == L_1/L_2$.
 - Both error locations < 288 , and not equal to each other.
 - Calculate an Error Evaluator Polynomial: $O(x) = S(x) L(x) \bmod x^6$ (where $L(x)$ is the 2-error ELP). Confirm that the degree 1 coefficient of $O(x)$ is non-zero, and all other coefficients are 0.

With these checks, the probability of error detection for $N = 1 \dots 11$ bit errors in a Phit is illustrated in *Phit FEC 288 Probability of Detection*. Note: This illustrates the error detection capability of the FEC alone. For 6 or more errors, there is a very small probability of FEC mis-correction, but a very high probability that the resulting errors are detected by packet CRC check

Table 8-1: Phit FEC 288 Probability of Detection

N	Probability of detection
1	100% (corrected)
2	100% (corrected)
3	100%
4	100%
5	100%
6	~99.97%
7	~99.9997%

8	~99.97%
9	~99.9998%
10	~99.97%
11	~99.9999%

8.2.1.2.2. Phit FEC 320

Format

With Phit FEC 320, each Phit is 320 bits in size, consisting of:

- 256 bits of Gen-Z packet stream (not necessarily aligned with packet boundaries).
- 4 copies of a “ctl” bit, which differentiates between Data Sequences and PHY-Up Sequences.
- 60 bits of FEC parity.

This encoding uses decoder alphabet $GF(2^{10})$. The Phit FEC 320 encoding is binary BCH (320, 260). This FEC requires the use of Gray Coding and Precoding such that burst errors introduced by DFE are converted to bit errors. The Hamming distance is 13, so this code could potentially be used for correction of 6 bit errors in a Phit. If desired, correction may be limited to fewer errors (e.g. 4), for reduced decoder complexity/area, and increased detection of uncorrectable errors .

The construction of a 320-bit Phit is shown in *Construction of Phit FEC 320*.

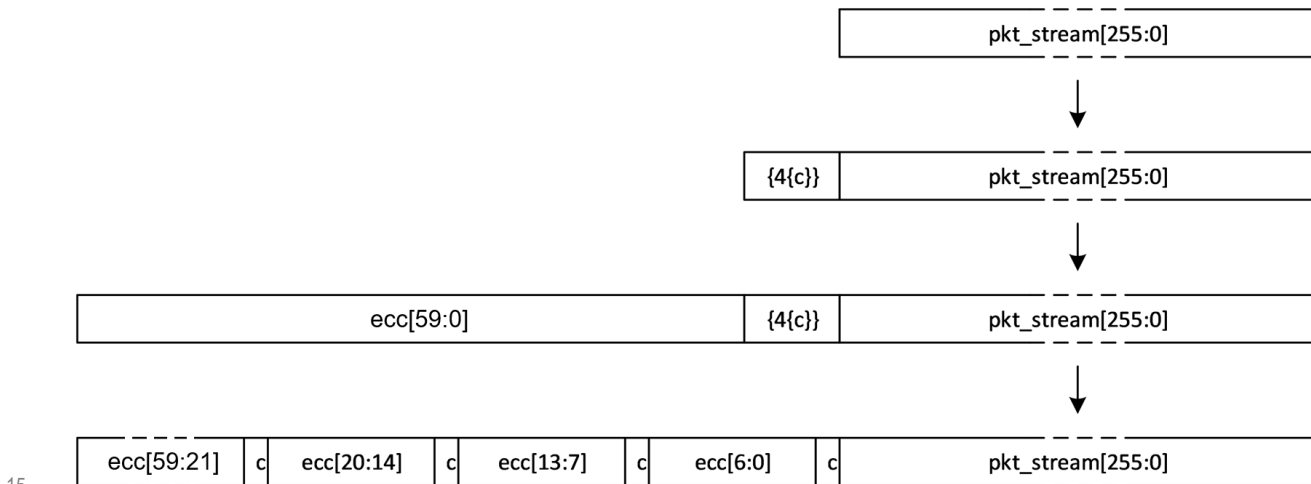


Figure 8-5: Construction of Phit FEC 320

- 1) Start with 256 bits of Gen-Z packet stream.
- 2) Append 4 copies of “ctl” bit, aka “c”.
- 3) Encode as BCH(320,260). Note: this is a binary CRC encoding, and the codeword is bit-reversed for transfer on the link, as is common for CRC codes. The “lowest” bit of the packet stream will travel first on the link and the “highest” bit of the FEC parity will travel last.

- 4) Intersperse the four copies of the “ctl” bit with the FEC parity (aka “ecc”). This reduces the probability that multiple copies will be affected by a single error event, particularly when the copies travel on separate lanes of the link.

- 5) The 320-bit Phit is striped across lanes according to the rules described in the Data Striping section in the Physical Layer Abstraction chapter in the Core specification. *Phit FEC 288 on 4-lane link* shows how a Phit is transferred on a four-lane link. In this drawing the bits indicated at the top travel first on the link. For a four-lane link, the Phit is striped across lanes with byte granularity.

lane3		lane0	
24	16	8	0
25	17	9	1
26	18	10	2
27	19	11	3
28	20	12	4
29	21	13	5
30	22	14	6
31	23	15	7
56	48	40	32
57	49	41	33
58	50	42	34
59	51	43	35
60	52	44	36
61	53	45	37
62	54	46	38
63	55	47	39
⋮			
248	240	232	224
249	241	233	225
250	242	234	226
251	243	235	227
252	244	236	228
253	245	237	229
254	246	238	230
255	247	239	231
ctl	ctl	ctl	ctl
ecc21	ecc14	ecc7	ecc0
ecc22	ecc15	ecc8	ecc1
ecc23	ecc16	ecc9	ecc2
ecc24	ecc17	ecc10	ecc3
ecc25	ecc18	ecc11	ecc4
ecc26	ecc19	ecc12	ecc5
ecc27	ecc20	ecc13	ecc6
ecc52	ecc44	ecc36	ecc28
ecc53	ecc45	ecc37	ecc29
ecc54	ecc46	ecc38	ecc30
ecc55	ecc47	ecc39	ecc31
ecc56	ecc48	ecc40	ecc32
ecc57	ecc49	ecc41	ecc33
ecc58	ecc50	ecc42	ecc34
ecc59	ecc51	ecc43	ecc35

Figure 8-6: Phit FEC 320 on 4-lane link

When a Phit arrives in a receiving PHY, Phit FEC decode consists of:

- 1) Reverse the final step of Phit construction, so the four copies of the “ctl” bit and the 256 bits of the packet stream are once again a contiguous payload in the codeword.
- 2) Bit-reverse the codeword if necessary (depends on decoder implementation).
- 3) Perform FEC decode, restricting correction if desired.
- 4) Check for identical (corrected) copies of the “ctl” bit. If all copies are not identical, initiate link resynchronization.
- 5) Remove “ctl” bits and transfer 256 bits of packet stream (if applicable) to the receiving link layer. If an uncorrectable error was detected by FEC decode, flag affected chunk(s) of packet stream via

5

the PHY_RX_PHITCRC_ERR signal as described in the Physical Layer Abstraction chapter in the Core specification.

Encode

The BCH(320,260) code is constructed as follows:

5

1) The decoder alphabet is $GF(2^{10})$, constructed with primitive polynomial $p(x) = x^{10} + x^3 + 1$, which may be represented as the binary value 0x409.

2) Each element of the field may be viewed as a power of the primitive element a , and may be represented as a 10-bit binary value:

10

$$GF(2^{10}) = \{0, a^0, a^1, a^2, a^3, a^4, a^5, a^6, a^7, a^8, a^9, a^{10}, \dots, a^{1022}\}$$

$$= \{0x000, 0x001, 0x002, 0x004, 0x008, 0x010, 0x020, 0x040, 0x080, 0x100, 0x200, 0x009, \dots, 0x204\}$$

3) Minimal polynomials are:

15

$$m_1(x) = m_2(x) = m_4(x) = m_8(x) = (x+a^1)(x+a^2)(x+a^4)(x+a^8)(x+a^{16})(x+a^{32})(x+a^{64})(x+a^{128})(x+a^{256})(x+a^{512})$$

$$= x^{10} + x^3 + 1$$

$$m_3(x) = m_6(x) = m_{12}(x) = (x+a^3)(x+a^6)(x+a^{12})(x+a^{24})(x+a^{48})(x+a^{96})(x+a^{192})(x+a^{384})(x+a^{768})(x+a^{513})$$

$$= x^{10} + x^3 + x^2 + x + 1$$

$$m_5(x) = m_{10}(x) = (x+a^5)(x+a^{10})(x+a^{20})(x+a^{40})(x+a^{80})(x+a^{160})(x+a^{320})(x+a^{640})(x+a^{257})(x+a^{514})$$

$$= x^{10} + x^8 + x^3 + x^2 + 1$$

20

$$m_7(x) = (x+a^7)(x+a^{14})(x+a^{28})(x+a^{56})(x+a^{112})(x+a^{224})(x+a^{448})(x+a^{896})(x+a^{769})(x+a^{515})$$

$$= x^{10} + x^9 + x^8 + x^7 + x^6 + x^5 + x^4 + x^3 + 1$$

$$m_9(x) = (x+a^9)(x+a^{18})(x+a^{36})(x+a^{72})(x+a^{144})(x+a^{288})(x+a^{576})(x+a^{129})(x+a^{258})(x+a^{516})$$

$$= x^{10} + x^7 + x^5 + x^3 + x^2 + x + 1$$

25

$$m_{11}(x) = (x+a^{11})(x+a^{22})(x+a^{44})(x+a^{88})(x+a^{176})(x+a^{352})(x+a^{704})(x+a^{385})(x+a^{770})(x+a^{517})$$

$$= x^{10} + x^5 + x^4 + x^2 + 1$$

4) The code generator polynomial is calculated as:

$$g(x) = \text{LCM}(m_0(x), m_1(x), \dots, m_{11}(x))$$

$$= m_1(x) m_3(x) m_5(x) m_7(x) m_9(x) m_{11}(x)$$

$$= x^{60} + x^{59} + x^{57} + x^{56} + x^{54} + x^{53} + x^{50} + x^{45} + x^{43} + x^{41} + x^{40} + x^{39} + x^{37} + x^{36} + x^{35} + x^{32} + x^{30} + x^{28} + x^{22} + x^{18} + x^{16} + x^{15} + x^{14} + x^{10} + x^7 + x^5 + x^3 + x^2 + 1$$

30

Given this code generator polynomial, a 260-bit payload is viewed as a polynomial $m(x)$ of degree 259 over $GF(2)$. A codeword is viewed as a polynomial of degree 319, constructed as:

35

$$c(x) = x^{60} m(x) - (x^{60} m(x) \bmod g(x))$$

and represented as a 320-bit binary value.

Decode

The decode process starts with calculating a 12-element power-sum syndrome vector: $\{s_1, s_2, s_3, s_4, s_5, s_6, s_7, s_8, s_9, s_{10}, s_{11}, s_{12}\}$. Each power-sum syndrome is calculated by viewing the received code word as a polynomial $r(x)$ over $GF(2^{10})$ and evaluating $r(x)$ for a particular element of $GF(2^{10})$.

$$\begin{aligned} 5 \quad s_1 &= r(a^1) \\ s_2 &= r(a^2) \\ &: \end{aligned}$$

The power-sum syndromes are further decoded with known algorithms to determine whether an error is present, whether it is correctable, and if so, the location(s) of the error(s). The FEC supports
10 correction of 6 bit errors in the Phit, but correction may be limited to fewer errors (e.g. 4) if desired, to reduce decoder complexity/area.

Developer Note: *In practice each evaluation of $r(x)$ for a (constant) element of $GF(2^{10})$ is best implemented as parity logic. For a 320-bit code word, the parity logic for evaluation of one element corresponds to about $10 \times 320 = 3200$, 2-input XOR gates. Evaluating 12 elements requires $12 \times \sim 3200 = \sim 38,400$, 2-input XOR.*

Given the size of the (320-bit) FEC Phit, it will generally be the case that Phits are flowing through the receiving PHY in back-to-back clock cycles (in the deserialized clock domain). An FEC decode of a multi-bit error will require several clock cycles to complete (e.g. 5..10). Therefore, the FEC decoder should be pipelined such that it can start a new Phit decode in each clock cycle.

20 *The coefficients of the Error Locator Polynomial (ELP) may be calculated with an appropriate Key Equation Solver (KES) algorithm, such as Berlekamp-Massey. Note: only "odd-numbered" iterations of Berlekamp-Massey are required (6 altogether for 6-error correction) which may be implemented as a 6-stage pipeline.*

25 *For finding roots of the ELP (corresponding to error locations), the Chien search is suggested for finding up to 6 roots. The ELP is evaluated for each element of $GF(2^{10})$ corresponding to a valid error location (320 altogether) to find the elements for which it evaluates to 0. For a parallel implementation, each polynomial evaluation of a constant value may be implemented as parity logic. Evaluation of a 6th degree ELP requires the equivalent of about 320, 2-input XOR gates. Parallel evaluation of 320 elements requires about $320 \times 320 = \sim 100,000$ 2-input XOR. After identifying the cases that evaluate to 0 (NOR function), the result is a 320-bit vector with bit(s) set corresponding to error locations.*

30 *If correction is limited to 4 bit errors, 4 ("odd-numbered") iterations (pipeline stages) of Berlekamp-Massey are required to calculate an ELP of degree 4 or less. Note: If more than 4 errors exist in the Phit, the result of this calculation will be an "invalid" ELP which does not have 4 valid roots. In a parallel implementation of the Chien search, each evaluation a 4th degree ELP requires the equivalent of about 200, 2-input XOR gates, resulting in $320 \times 200 = \sim 64,000$, 2-input XOR altogether.*

Uncorrectable Error Detection

The detection of uncorrectable errors by Phit FEC decode supplements the error detection provided by Gen-Z packet CRC. To maximize the error detection capability of the Phit FEC, the following sanity checks are recommended:

- 40 1) The degree of the ELP should be less than or equal to the maximum number of errors being corrected.

- 2) The number of (unique) roots found should correspond to the degree of the ELP.
- 3) The roots should all have corresponding error locations which are valid (0..319).

With these checks, and correcting 6 errors, the probability of detecting N bit errors in a Phit is illustrated in the second column of *Phit FEC 320 Probability of Detection*, and the probability of detecting N bit errors when correction is limited to 4-errors is illustrated in the third column. Note: This illustrates the error detection capability of the FEC alone. When an error is not detected by Phit FEC, there is a very small probability of FEC mis-correction, but a very high probability that the resulting errors are detected by packet CRC check.

Table 8-2: Phit FEC 320 Probability of Detection

N	Probability of detection (6-bit correction)	Probability of detection (4-bit correction)
1	100% (corrected)	100% (corrected)
2	100% (corrected)	100% (corrected)
3	100% (corrected)	100% (corrected)
4	100% (corrected)	100% (corrected)
5	100% (corrected)	100%
6	100% (corrected)	100%
7	~99.9993%	100%
8	~99.9993%	100%
9+	~99.9993%	~100%

8.2.1.3. Framing

Explicit framing beyond Phit FEC shall not be used with this physical layer clause.

8.2.1.4. Gray Coding

Gray Coding, also known as Gray Mapping, is a technique used with PAM4 signaling to limit incorrect decisions to only one bit error per symbol and simplify support for dual-mode PAM2 (NRZ) physical layer implementations.

Gray Coding for this physical layer clause shall be as specified in *IEEE Standard for Ethernet 802.3bs™-2017*, clause 120D.5.7.

8.2.1.5. Precoding

Receiver Decision Feedback Equalization (DFE) mis-corrections can propagate a single error into a burst of back-to-back errors. PAM4 signaling tends to produce longer bursts than NRZ signaling. Precoding is a technique used to convert DFE produced burst errors to only two errors, entry and exit errors on contiguous UI patterns. Precoding shall be as specified in *IEEE Standard for Ethernet 802.3cd™-2018*, clause 135.5.7.2.

The following are the features and requirements associated with precoding for this physical layer clause:

- Precoding in the transmitting direction shall be required.
- Precoding in the receiving direction shall be optional.
- Precoding can be enabled independently per direction and can be non-uniform (enabled in one Tx → Rx direction and disabled in the other Rx ← Tx direction).
- Precoding shall be off during Signal Detection and a local receiving physical layer may request precoding from its remote transmitting physical layer during the Channel Optimize state.
- Precoding previous symbol state shall reset to 2'b00 when the scrambler is reset, e.g. after an Align Training Sequence Header.

Developer Note: *The effect of DFE tap weights on error propagation is well acknowledged in the industry and where tap strengths are significant, the effectiveness of precoding may be minimal. For receiver implementations not requiring a DFE, precoding is not recommended.*

8.2.2. Scrambling / Descrambling

Refer to the Gen-Z Common Specification clause section on *Scrambling / Descrambling*.

8.2.3. Data Striping

Refer to the Gen-Z Common Specification clause section on *Data Striping*.

8.2.4. PLA Data Width

Refer to the Gen-Z Common Specification clause section on *PLA Data Width*.

8.2.5. Physical Layer Initialization and Training

8.2.5.1. PLA State Map

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *PLA State Map*.

8.2.5.2. Physical Layer Link versus Lane Training

Refer to the Gen-Z Common Specification clause section on *Physical Layer Link versus Lane Training*.

8.2.5.3. Physical Layer (PHY) Sequences

Refer to the Gen-Z Common Specification clause section on *Physical Layer (PHY) Sequences* with the following amendments:

The following are the features and requirements associated with PHY training sequences for this physical layer clause:

- The following PHY Sequences shall be defined as PHY training sequences:
 - Signal Detect Training Sequences
 - Align Training Sequences
 - IBIST Training Sequences
 - Config Training Sequences

- Lock Training Sequence
- PHY training sequences shall be transmitted independently on each lane.
- PHY training sequences shall not add encoding headers.
- PHY training sequences shall not be protected by Phit FEC.
- PHY training sequences shall be 128 bits wide.

The following are the features and requirements associated with PHY UP sequences for this physical layer clause

- PHY UP sequences shall be protected by Phit FEC.
- PHY UP sequences shall be 260 bits wide plus FEC parity and striped across lanes.

8.2.5.3.1. Signal Detect Training Sequences

Refer to the Gen-Z Common Specification clause section on *Signal Detect Training Sequences* with following changes illustrated in *Gen-Z-E-PAM4-50G PHY Detect Training Sequences*.

Table 8-3: Gen-Z-E-PAM4-50G PHY Detect Training Sequences

Type	Scrambled, Gray Coded	Word 0	Word 1	Word 2	Word 3
Header	No	0x00000000	0xFFFFFFFF	0x00000000	0xFFFFFFFF
Set	Yes	0x3C3C3C3C	0x0	0x0	0x0

8.2.5.3.2. Align Training Sequences

Refer to the Gen-Z Common Specification clause section on *Align Training Sequences* with following changes illustrated in *Gen-Z-E-PAM4-50G PHY Align Training Sequences*.

Table 8-4: Gen-Z-E-PAM4-50G PHY Align Training Sequences

Type	Scrambled, Gray Coded, Precoded	Word 0	Word 1	Word 2	Word 3
Header	No	0x00000000	0xFFFFFFFF	0x00000000	0xFFFFFFFF
Set	Yes	0x1E1E1E1E	0x0	<Ack>	0x0

8.2.5.3.3. IBIST Training Sequences

Refer to the Gen-Z Common Specification clause section on *IBIST Training Sequences* with following changes illustrated in *Gen-Z-E-PAM4-50G PHY IBIST Training Sequences*.

Table 8-5: Gen-Z-E-PAM4-50G PHY IBIST Training Sequences

Type	Scrambled, Gray Coded, Precoded	Word 0	Word 1	Word 2	Word 3
Header	Yes	0xE5E5E5E5	0x0	0x0	0x0
Pattern	Yes	0x0	0x0	0x0	0x0

8.2.5.3.4. Config Training Sequence

Refer to the Gen-Z Common Specification clause section on *Config Training Sequence* with following changes illustrated in *Gen-Z-E-PAM4-50G PHY Configuration State Training Sequence*.

Table 8-6: Gen-Z-E-PAM4-50G PHY Configuration State Training Sequence

Type	Scrambled, Gray Coded, Precoded	Word 0	Word 1	Word 2	Word 3
Config	Yes	0x2D2D2D2D	<phit_fec>	<width>	<info>

Phit FEC support is advertised in phit_fec field of the Config Training Sequence, and the sent and received phit_fec is used to determine the type of Phit FEC used by the link in PHY-Up. The phit_fec field is encoded as follows where rsvd bits are zero:

- <phit_fec> = {bch_320_supported, bch_288_supported, rsvd[29:16], bch_320_supported, bch_288_supported, rsvd[13:0]}

bch_288_supported can be masked by software by the “Mask Phit FEC 288” field in the *Table 8-15: Gen-Z-E-PAM4-50G-Fabric Extended Feature Control* structure.

Phit FEC selection uses the following priority:

- If local bch_288_supported=1 and bch_288_supported=1, use BCH(288,260)
- Else if bch_320_supported=1 and bch_320_supported=1, use BCH(320,260)
- Else training fails, go back to Signal Detect.

8.2.5.3.5. Lock Training Sequence

Refer to the Gen-Z Common Specification clause section on *Lock Training Sequence* with following changes illustrated in *Gen-Z-E-PAM4-50G PHY Lock State Training Sequence*.

Table 8-7: Gen-Z-E-PAM4-50G PHY Lock State Training Sequence

Type	Scrambled, Gray Coded, Precoded	Word 0	Word 1	Word 2	Word 3
Lock	Yes	0xA4A4A4A4	0x0	0x0	0x0

8.2.5.3.6. PHY-Up and Data Sequences

Refer to the Gen-Z Common Specification clause section on *PHY-Up and Data Sequences* with following changes:

- 5 Data Sequences are defined in the *Data Sequences* table below where “D” represents a stream of Gen-Z packet data.

Table 8-8: Data Sequences

Type	Phit FEC Ctl	Scrambled, Gray Coded, Precoding	Striped	Word 0 Word 4	Word 1 Word 5	Word 2 Word 6	Word 3 Word 7
Data	0b	Yes	Yes	D[31:0] D[159:128]	D[63:32] D[191:160]	D[95:64] D[223:192]	D[127:96] D[255:224]

Each Data Sequence contains 256 bits of Gen-Z Packet Stream Data for both Phit FEC 288 and Phit FEC 320 encodings. The following are the features and requirements associated with the transmission of Data Sequences:

- Data Sequences shall consist of Gen-Z Core data from the PLA, shall be striped, according to the data striping rules, and shall be scrambled, gray coded and precoded.
- Data Sequences shall set the Phit FEC “ctl” bit as specified in *Data Sequences*.

The following are the features and requirements associated with the reception of Data Sequences:

Data Sequences shall be de-precoded, de-gray coded, de-scrambled, shall be re-striped, according to the data striping rules, and FEC decoded.

- Data Sequences shall match the Phit FEC “ctl” bit as specified in *Data Sequences*.
 - If 2 “ctl” bits are 0, the PHY shall assert PHY_RX_TRANSIENT_ERR
 - If 3 “ctl” bits are 0 or an uncorrectable Phit FEC error is detected, the PHY shall send received data across the PLA as Gen-Z Core data to Gen-Z Core and assert PHY_RX_PHITCRC_ERR.
 - If 4 “ctl” bits are 0, the PHY shall send received data across the PLA as Gen-Z Core data to Gen-Z Core and assert PHY_RX_PHITCRC_ERR.

Table 8-9: PHY-Up Sequences

Type	Phit FEC Ctl	Scrambled, Gray Coded, Precoding	Striped	Word 0 Word 4	Word 1 Word 5	Word 2 Word 6	Word 3 Word 7
Idle	1b	Yes	Yes	0xAAAAAAAA	0xAAAAAAAA	0xAAAAAAAA	0xAAAAAAAA
PHY LP	1b	Yes	Yes	0xBBBBBBBB	<LPID>	0xBBBBBBBB	<LPID>
PHY Down	1b	Yes	Yes	0xCCCCCCCC	0xCCCCCCCC	0xCCCCCCCC	0xCCCCCCCC
Retrain	1b	Yes	Yes	0x99999999	<RetrainID>	0x99999999	<RetrainID>

The following are the features and requirements associated with the transmission of PHY-Up Sequences:

- PHY-Up Sequence Words 0-7 shall be encoded in the Phit FEC payload.
- PHY-Up Sequences shall consist of Gen-Z physical layer information, shall be striped across lanes, and shall be scrambled, gray coded, and precoded as specified in *PHY-Up Sequences*.
- PHY-Up Sequences shall set the Phit FEC “ctl” bit as specified in *PHY-Up Sequences*.

The following are the features and requirements associated with the reception of PHY-Up Sequences:

- PHY-Up Sequences shall be de-precoded, de-gray coded, de-scrambled as specified in *PHY-Up Sequences*, FEC decoded, shall be decoded as Gen-Z physical layer information, and shall not be transmitted across the PLA to the Gen-Z Core.
- PHY-Up Sequences shall be considered valid if all they match a valid type, and where applicable, any encoded fields match such as <LPID> and <RetrainID> match within the Phit.
- PHY-Up Sequences shall match the Phit FEC “ctl” as specified in *PHY-Up Sequences*.
- If an invalid PHY-Up Sequence is detected, then:
 - The PHY shall assert PHY_RX_TRANSIENT_ERR and drop the PHY-Up Sequence.

Local and remote physical layers may request conflicting PHY-Up Sequence types that can affect the state of the physical layer link. Reconciliation of these conflicts shall have the following priority:

- PHY Down – highest priority
- Retrain
- LP (Low Power) – lowest priority

The details for field encodings for specific PHY-Up Sequences can be found in the following sections:

- LP (Low Power) <LPID> = *Low Power Operation*
- Retrain <RetrainID> = *Physical Layer Retraining*

8.2.5.3.7. Clock Compensation Sequences

Clock Compensation Sequences are a unique type of PHY Sequence in the Gen-Z-E-PAM4-50G-Fabric Specification and are not a type of PHY-Up Sequence. Clock Compensation Sequences are defined in the *Clock Compensation Sequences* table below.

Table 8-10: Clock Compensation Sequences

Type	Striped, Scrambled	Gray Coded, Precoding	Word 0 Word 1	Word 2 Word 3	Word 4 Word 5	Word 6 Word 7
Clock Comp	No	Yes	0xA0A0A0A0, 0xA0A0A0A0	<sets_remaining>, <sets_remaining>	0x0A00A0AA, 0x0A00A0AA	0x0A00A0AA, 0x0A00A0AA

sets_remaining[31:0] is an encoded value of the remaining double words in the Clock Compensation Sequence, and the decode value ranges from 0-4 double words. sets_remaining[31:0] the uses the follow encoding:

- 5 • 0 double words remaining: 0xAA0082B3
- 1 double words remaining: 0x82B3A0A0
- 2 double words remaining: 0x00AA2828
- 3 double words remaining: 0x2828B30A
- 4 double words remaining: 0xB3820A82

10 Each nibble of the sets_remaining encoded values is unique (0x0, 0x2, 0x3, 0x8, 0xA, or 0xB).

Clock Compensation Sequences may optionally be disabled in the *Unique Interface Physical Layer Structure*. If Clock Compensation Sequences are enabled, the following are the features and requirements associated with the transmission of Clock Compensation Sequences:

- 15 • Clock Compensation Sequences shall be sent in Align, IBIST, Config, and PHY-Up PLTSM States
- Clock Compensation Sequences shall only be sent after every 131,072 UI of non-Clock Compensation Sequences
- The Clock Compensation Sequence UI count shall reset upon entry to Align and reset when a Clock Compensation Sequence is sent
- 20 • Clock Compensation Sequences shall not be striped across lanes. Instead, all lanes shall transmit the same Clock Compensation Sequences like Training Sequences.
- Clock Compensation Sequences shall not be Phit FEC encoded.
- Clock Compensation Sequences may interrupt Phit transmission including being sent in the middle of a Phit.
- The scrambler shall stall during transmission of Clock Compensation Sequences
- 25 • Clock Compensation Sequences shall be gray coded and precoded.
- Clock Compensation Sequence length sent from the link transmitter is 8 Words (32 bytes) per lane and sets_remaining shall be set to the encoded value of 2.
- Precoding state shall reset to 0 after transmission of a Clock Compensation Sequence

30 If Clock Compensation Sequences are enabled, the following are the features and requirements associated with retimer retransmission of Clock Compensation Sequences:

- Retimers may add or delete 2 Words of “0x0A00A0AA” of the Clock Compensation Sequence to compensate for reference clock PPM offsets
- If the length of the Clock Compensation Sequence is changed by the Retimer, the Retimer shall regenerate the <sets_remaining> field and redo precoding and gray coding of the Clock Compensation Sequence

If Clock Compensation Sequences are enabled, the following are the features and requirements associated with the reception of Clock Compensation Sequences:

- The receiver shall count UI between Clock Compensation Sequences starting with the Clock Compensation Sequences identified in the Align PLTSM States
- The receiver shall expect a Clock Compensation Sequence every 131,072 UI of non-Clock Compensation Sequences in the IBIST, Config, and PHY-Up PLTSM States
- The receiver shall support reception of Clock Compensation Sequence that vary from 4 Words to 12 Words in length (support for two retimers).
- Clock Compensation Sequences shall be de-precoded and de-gray coded and the length of the Clock Compensation Sequence shall be decoded from the sets_remaining field
 - The receiver shall perform a majority vote comparison for each of the 16 nibbles of Word 2 and Word 3 to determine remaining double words in the Clock Compensation Sequence (0-4). E.g., reception of the following two sets_remaining fields where “X” nibbles are don’t cares would match two double words remaining: 0x0XXA2XX800AXXX28, since 9 out of 16 nibbles match the 0x00AA282800AA2828 pattern.
 - If the Clock Compensation length cannot be determined, the PHY shall initiate a retrain.
- The descrambler shall stall during Clock Compensation Sequence reception
- De-precoding state shall reset to 0 after transmission of a Clock Compensation Sequence

Clock Compensation Sequences shall not be sent to the Gen-Z Core

8.2.5.4. **Error Recovery**

Refer to the Gen-Z Common Specification clause section on *Error Recovery*.

8.2.5.5. **Lane-to-Lane Deskew**

Refer to the Gen-Z Common Specification clause section on *Lane-to-Lane Deskew*.

The lane-to-lane skew for this physical layer clause shall meet the requirements defined in *Gen-Z-E-PAM4-50G-Fabric Lane-to-Lane Skew***Error! Reference source not found..**

Table 8-11: Gen-Z-E-PAM4-50G-Fabric Lane-to-Lane Skew

Skew Type	Notes	Value	Units
TX lane to lane output skew	Between any two lanes on a single PHY port	1.0 (max)	ns
Rx lane to lane skew	Across all lanes on a single PHY port including lane-to-lane variations due to channel and repeater delays	3.0 (max)	ns

8.2.6. Lane Reversal

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause on *Lane Reversal*.

8.2.7. Lane Polarity

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause on *Lane Polarity*.

5 8.2.8. Link Width

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause on *Link Width*.

8.2.9. Link BIST

Refer to the Gen-Z Common Specification clause section on *Link BIST*.

8.2.10. Physical Layer Retraining

10 Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause on *Physical Layer Retraining*.

8.2.11. PLTSM

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause on *PLTSM*.

8.2.11.1. PHY Idle (Required)

Refer to the Gen-Z Common Specification clause section on *PHY Idle (Required)*.

15 8.2.11.2. Signal Detect (Required)

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Signal Detect (Required)* with the required PHY Sequence modifications detailed in *Signal Detect Training Sequences*.

8.2.11.3. Channel Optimize (Required)

20 Channel Optimize shall be a required state for this individual physical layer clause. Channel Optimize specifies an in-band parameter exchange between physical layers to enable transmitter and receiver equalization adaptation.

8.2.11.3.1. Channel Optimize Entry

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause on *Channel Optimize Entry*.

8.2.11.3.2. Channel Optimize Process

25 Refer to the following topics and clauses from the *IEEE Standard for Ethernet 802.3™* Specification

- CR/KR 50G PAM4 PMD (inband exchange protocol): clause 136.8.11

8.2.11.3.3. Channel Optimize Exit

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause on *Channel Optimize Exit*.

8.2.11.4. *Align (Required)*

Refer to the Gen-Z Common Specification clause section on *Align (Required)* with the required PHY Sequence modifications detailed in *Align Training Sequences*.

Details for precoding negotiation during the Align state will be provided in a later revision.

8.2.11.5. *Interconnect BIST (Optional)*

Refer to the Gen-Z Common Specification clause section on *Interconnect BIST (Optional)* with the required PHY Sequence modifications detailed in *IBIST Training Sequences*.

8.2.11.6. *Config (Required)*

Refer to the Gen-Z Common Specification clause section on *Config (Required)* with the required PHY Sequence modifications detailed in *Config Training Sequence* and *Phit FEC support is advertised in phit_fec* field of the Config Training Sequence, and the sent and received *phit_fec* is used to determine the type of Phit FEC used by the link in PHY-Up. The *phit_fec* field is encoded as follows where *rsvd* bits are zero:

- $\langle \text{phit_fec} \rangle = \{ \text{bch_320_supported}, \text{bch_288_supported}, \text{rsvd}[29:16], \text{bch_320_supported}, \text{bch_288_supported}, \text{rsvd}[13:0] \}$

bch_288_supported can be masked by software by the “Mask Phit FEC 288” field in the *Table 8-15: Gen-Z-E-PAM4-50G-Fabric Extended Feature Control* structure.

Phit FEC selection uses the following priority:

- If local *bch_288_supported*=1 and *bch_288_supported*=1, use BCH(288,260)
- Else if *bch_320_supported*=1 and *bch_320_supported*=1, use BCH(320,260)
- Else training fails, go back to Signal Detect.

Lock Training Sequence.

8.2.11.7. *PHY-Up (Required)*

Refer to the Gen-Z Common Specification clause section on *PHY-Up (Required)*.

8.2.12. Clock Compensation

A maximum clock tolerance offset of 200 ppm (± 100 ppm) is supported in this individual physical layer clause.

Refer to the Gen-Z-E-PAM4-50G-Fabric Specification clause section on *Clock Compensation Sequences* for rules and policies on transmitting and receiving, including addition or removal, of Clock Compensation PHY-Up Sequences.

Developer Note: For closed systems with known clocking topologies, options exist in the PHY Specific Control Fields of the Unique Interface Physical Layer Structure to disable, increase, or decrease the frequency of Clock Compensation Sequences. For example in a system where both the local and remote physical layers share a common reference source, the use of Clock Compensation Sequences can be disabled to reduce physical layer overhead and increase Gen-Z Core data performance.

8.2.13. Loopback

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause on *Loopback*.

8.2.14. Physical Layer Aggregation

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause on *Physical Layer Aggregation*.

8.2.15. Re-timers

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause on *Re-timers*.

8.3. Power Management

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Power Management*.

8.4. Electrical Sublayer

This Gen-Z physical layer electrical sublayer shall meet all single lane/channel requirements specified in the CEI-56G-LR-PAM4 clause in IA # OIF-CEI-04.0 but with the following exceptions:

- The serial baud rate shall be 26.5625 Gsym/s (or 53.125 GT/s) using PAM4 encoding ± 100 ppm.
- The raw Bit Error Rate shall operate at 10^{-9} or better.
- Support for COM with a channel Insertion Loss¹ (IL) recommended to be no greater than 20 dB at 13.28125 GHz.

As a service to the reader, many requirements of particular interest are detailed below and reference the following clauses when related to the operation of a single lane/channel:

- *IEEE Standard for Ethernet 802.3™-2015* for the following clauses:
 - Clause 93. Physical Medium Dependent (PMD) sublayer and baseband medium, type 100GBASE-KR4
 - Annex 93A (normative) Specification methods for electrical channels
 - Annex 93B (informative) Electrical backplane reference model
 - Annex 93C (normative) Receiver interference tolerance
- *IEEE Standard for Ethernet 802.3by™-2016* for the following clauses:

- Clause 111. Physical Medium Dependent (PMD) sublayer and baseband medium, type 25GBASE-KR and 25GBASE-KR-S
- *IEEE Standard for Ethernet 802.3bs™-2017* for the following clauses:
 - Annex 120D (normative) Chip-to-chip 200 Gb/s four-lane Attachment Unit Interface (200GAUI-4 C2C) and 400 Gb/s eight-lane Attachment Unit Interface (400GAUI-8 C2C)
 - Annex 120E (normative) Chip-to-module 200 Gb/s four-lane Attachment Unit Interface (200GAUI-4 C2M) and 400 Gb/s eight-lane Attachment Unit Interface (400GAUI-8 C2M)
- *IEEE Standard for Ethernet 802.3cd™-2018* for the following clauses:
 - Clause 136. Physical Medium Dependent (PMD) sublayer and baseband medium, type 50GBASE-CR, 100GBASE-CR2, and 200GBASE-CR4
 - Clause 137. Physical Medium Dependent (PMD) sublayer and baseband medium, type 50GBASE-KR, 100GBASE-KR2, and 200GBASE-KR4

Developer’s Note: This physical layer clause defines the use of Long Reach (LR) capable transmitters and receivers on restricted Medium Reach (MR) channels as shown in the Gen-Z-E-PAM4-50G-Fabric Electrical . This allows for a light weight, low latency Forward Error Correction (FEC) code at the expense of less correction gain.




Figure 8-7: Gen-Z-E-PAM4-50G-Fabric Electrical Sublayer

8.4.1. Electrical Interface

8.4.1.1. Overview

The electrical interface for this physical layer clause shall be based on high speed, low voltage differential signaling where each connection is point-to-point and signaling is PAM4 and unidirectional.

8.4.1.1.1. Forwarded Clock

This physical layer clause shall not require a forwarded clock.

8.4.1.2. ***Lane Signaling Rate***

The signaling rate for this physical layer clause shall be as specified in *IEEE Standard for Ethernet 802.3bs™-2017*, clause 120D.3.1 at 26.5625 GBd/s \pm 100 Parts Per Million (ppm) per lane using 4-level Pulse Amplitude Modulation (PAM4).

5 8.4.1.3. ***Lane Bit Error Rate***

Each lane associated with this physical layer clause shall operate with a bit error ratio of 10^{-9} or better.

8.4.1.4. ***Lane Equalization***

Tx Equalization coefficient range and step size requirements shall be as specified in *IEEE Standard for Ethernet 802.3™-2018* clause 136.8.11.

10 **8.4.2. High Speed Signaling Specification**

8.4.2.1. ***Transmitter Specification***

The transmitter for this physical layer clause shall be as specified in the CEI-56G-LR-PAM4 clause in IA # OIF-CEI-04.0.

8.4.2.2. ***Receiver Specification***

15 The receiver for this physical layer clause shall be as specified in the CEI-56G-LR-PAM4 clause in IA # OIF-CEI-04.0.

8.4.3. Channel Characteristics

8.4.3.1. ***Channel Operating Margin***

20 The channel for this physical layer clause shall meet compliance with the Channel Operating Margin (COM) for the CEI-56G-LR-PAM4 clause in IA # OIF-CEI-04.0 with the required modifications to Table 21-1:

- Signaling rate (f_b) = 26.5625 Gsym/s
- Target detector error ratio (DER_0) = 10^{-9}

8.4.3.2. ***Insertion Loss***

25 The channel for this physical layer clause is recommended to meet the informative channel insertion loss of the Medium Reach (MR) CEI-56G-MR-PAM4 clause in IA # OIF-CEI-04.0.

8.4.3.3. ***Return Loss***

The Return Loss for this physical layer clause shall be as specified in CEI-56G-LR-PAM4 clause in IA # OIF-CEI-04.0.

8.4.3.4. *Coupling*

The transmitter for this physical layer clause shall be AC-coupled to the receiver as specified in CEI-56G-LR-PAM4 clause in IA # OIF-CEI-04.0.

8.4.4. **Miscellaneous Specification**

5 Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause on *Miscellaneous Specification*.

8.5. Management Control and Status

8.5.1. **Interface Physical Layer Structure**

10 The Interface PHY Structure is defined generically in the Core Specification to cover multiple physical layers. The mapping for the Interface PHY Structure associated with the physical layer defined in this specification clause are detailed below.

8.5.1.1. *Common Interface Physical Layer Structure*

8.5.1.1.1. **Common, Lane, and Low Power Interface PHY Structure**

15 The following defines the features and requirements associated with the Gen-Z-E-PAM4-50G-Fabric specification. The *Gen-Z-E-PAM4-50G-Fabric Interface PHY Structure Support Exclusions and Restrictions* indicates restricted or excluded fields or sub-fields from the *Interface PHY Structure Common, Lane, and Low Power Fields*.

Table 8-12: Gen-Z-E-PAM4-50G-Fabric Interface PHY Structure Support Exclusions and Restrictions

Support	Sub-field / Description	Interface PHY Structure Field	Supported Value(s)
Restricted	PHY Clause 5—50G Fabric	PHY Type	0x4
Restricted	Physical Layer Operational Status 0x0—PHY-Down, uninitialized state 0x1—PHY-Up 0x2—PHY-Down-Retrain 0x3 – 0x6 Reserved 0x7—PHY Low-Power 1 → Mapped to P1 0x8—PHY Low-Power 2 → Mapped to P2 0x9 – 0xA Reserved 0xB – PHY UP Low-Power 1 → Mapped to Dynamic Link Width Reduction 0xC—0xF Reserved	PHY Status	See Description

Support	Sub-field / Description	Interface PHY Structure Field	Supported Value(s)
Restricted	Previous Physical Layer Operational Status 0x0—PHY-Down, uninitialized state 0x1—PHY-Up 0x2—PHY-Down-Retrain 0x3 – 0x6 Reserved 0x7—PHY Low-Power 1 → Mapped to P1 0x8—PHY Low-Power 2 → Mapped to P2 0x9 – 0xA Reserved 0xB – PHY UP Low-Power 1 → Mapped to Dynamic Link Width Reduction 0xC—0xF Reserved	PHY Status	See Description
Restricted	Physical Layer Retraining Arc Only Retraining Arc 1 and Arc 2 are supported for this physical layer clause	PHY Control	0x0 – 0x1
Excluded	Retrain Arc 3-4 Support	PHY Cap 1	0b
Restricted	Enable Phit CRC Encoding Phit CRC encoding should be enabled by default for this physical layer clause	PHY Cap 1 Control	1b (default)
Restricted	Asymmetric Lane Status Asymmetry is not supported for this physical layer clause	PHY Lane Status	0x0
Restricted	Enable Lane Asymmetry	PHY Lane Control	0x0
Restricted	Asymmetric Lane Support	PHY Lane Cap	0x0
Restricted	Reversal Support Only uniform reversal is supported for this physical layer clause	PHY Lane Cap	0x0 – 0x2
Restricted	Asymmetric Lane with Reversal Support	PHY Lane Cap	0b
Restricted	Asymmetric Lane Support	PHY Remote Lane Cap	0x0
Restricted	Reversal Support Only uniform reversal is supported for this physical layer clause	PHY Remote Lane Cap	0x0 – 0x2

Support	Sub-field / Description	Interface PHY Structure Field	Supported Value(s)
Restricted	Asymmetric Lane with Reversal Support	PHY Remote Lane Cap	0b
Excluded	Entry, Exit Latency PHY-LP 3-4	PHY Low Power Timing Capability	0x0
Excluded	PHY-LP 3-4 Support	PHY Low Power CAP	0b
Excluded	Entry, Exit Latency PHY-UP-LP 2-4	PHY Up Low Power Timing Capability	0x0
Excluded	PHY-UP-LP 2-4 Support	PHY Up Low Power CAP	0b

8.5.1.1.2. Extended Feature Interface PHY Structure

The following defines the features and requirements associated with the Gen-Z-E-PAM4-50G-Fabric specification.

Table 8-13: Gen-Z-E-PAM4-50G-Fabric Extended Feature Fields

Field Name	Size (bits)	Value / Bit Location	M / O	Access	Description
PHY Extended Status	32	-	M	RO	See <i>Gen-Z-E-PAM4-50G-Fabric Extended Feature Status</i>
PHY Extended Control	32	-	M	RW	See <i>Gen-Z-E-PAM4-50G-Fabric Extended Feature Control</i>
PHY Extended Cap	32	-	M	RO	See <i>Gen-Z-E-PAM4-50G-Fabric Extended Feature Capability</i>
PHY Remote Extended Cap	32	-	M	RW	See <i>Gen-Z-E-PAM4-50G-Fabric Extended Feature Remote Capability</i>

Table 8-14: Gen-Z-E-PAM4-50G-Fabric Extended Feature Status

Bit Location	Access	Description
0	RO	Physical Layer Operational Status

Bit Location	Access	Description
2:1		0b—PHY Extended features disabled 1b—PHY Extended features enabled
	RO	Physical Layer Extended Feature Training Status 0x0—Training with extended features has not occurred 0x1—Training with extended features succeeded 0x2—Training with extended features has failed 0x3—Reserved
5:3	RO	Current Forward Error Correction (FEC) mode 0x0 – Phit FEC 288 0x1 – Phit FEC 320 0x2 – Phit CRC 0x3 – RS-FEC (802.3 Clause 134) 0x4 – 0x7 – Reserved
8:6	RO	Current Transfer Rate mode 0x0 – 53.125 GT/s 0x1 – 50.0 GT/s 0x2 – 56.25 GT/s 0x3—0x7 – Reserved
31:9	-	RsvdZ

Table 8-15: Gen-Z-E-PAM4-50G-Fabric Extended Feature Control

Bit Location	Access	Description
2:0	RW	Forward Error Correction (FEC) override. If non-zero Phit FEC selection is ignored during Config. 0x0 – Phy Negotiated (Default). Either Phit FEC 288 or Phit FEC 320 are selected during Config. 0x1 – Phit FEC 288 0x2 – Phit FEC 320 0x3 – Phit CRC 0x4 – RS-FEC (802.3 Clause 134) 0x5—0x7 – Reserved

Bit Location	Access	Description
3	RW	Mask Phit FEC 288 0b – Unmasked. PHY advertises Phit FEC 288 Capability normally in PHY Config Training Sequences. 1b – Masked. PHY does not advertise Phit FEC 288 Capability in PHY Config Training Sequences. This bit should be set if the link raw BER is greater than 1E-9.
6:4	RW	Transfer Rate mode 0x0 – 53.125 GT/s (Default) 0x1 – 50.0 GT/s 0x2 – 56.25 GT/s 0x3—0x7 – Reserved
31:7	-	RsvdP

Table 8-16: Gen-Z-E-PAM4-50G-Fabric Extended Feature Capability

Bit Location	Access	Description
0	RO	Phit FEC 288 support 0b—Unsupported 1b—Supported
1	RO	Phit FEC 320 support 0b—Unsupported 1b—Supported
2	RO	Phit CRC support 0b—Unsupported 1b—Supported
3	RO	RS-FEC (802.3 Clause 134) support 0b—Unsupported 1b—Supported
4	RO	50.0 GT/s Transfer Rate support 0b—Unsupported 1b—Supported
5	RO	56.25 GT/s Transfer Rate support 0b—Unsupported

Bit Location	Access	Description
		1b—Supported
31:6	-	RsvdZ

Table 8-17: Gen-Z-E-PAM4-50G-Fabric Extended Feature Remote Capability

Bit Location	Access	Description
0	RW	Phit FEC 288 support 0b—Unsupported 1b—Supported
1	RW	Phit FEC 320 support 0b—Unsupported 1b—Supported
2	RW	Phit CRC support 0b—Unsupported 1b—Supported
3	RW	RS-FEC (802.3 Clause 134) support 0b—Unsupported 1b—Supported
4	RW	50.0 GT/s Transfer Rate support 0b—Unsupported 1b—Supported
5	RW	56.25 GT/s Transfer Rate support 0b—Unsupported 1b—Supported
31:6	-	RsvdZ

8.5.1.2. *Unique Interface Physical Layer Structure*

The Unique Interface Physical Layer structure for the Gen-Z-E-PAM4-50G-Fabric specification references the *Unique Interface Physical Layer Structure* from the Gen-Z-E-NRZ-25G-Fabric specification clause. The *Gen-Z-E-PAM4-50G-Fabric Unique Interface PHY Structure Requirements* specifies the differences.

Table 8-18: Gen-Z-E-PAM4-50G-Fabric Unique Interface PHY Structure Requirements

Differences	Sub-field / Description	Interface PHY Structure Field
Additional Field	New Field: PHY Tx precoding status	<i>Gen-Z-E-PAM4-50G PHY Specific Status Fields</i>
Additional Field	New Field: 50G Clock Compensation Sequence Control	<i>Gen-Z-E-PAM4-50G PHY Specific Control Fields</i>
Additional Field	New Field: PHY Tx precoding New Field: PHY Rx precoding	<i>Gen-Z-E-PAM4-50G PHY Specific Control Fields</i>
Additional Field	New Field: PHY Rx precoding support	<i>Gen-Z-E-PAM4-50G PHY Specific Capability Fields</i>
Additional Field	New Field: Remote PHY Rx precoding support	<i>Gen-Z-E-PAM4-50G Remote PHY Specific Capability Fields</i>
Additional Field + Restructure	New Field: Current pre-cursor (-2) local setting	<i>Gen-Z-E-PAM4-50G PHY TX Lane Status Fields</i>
Additional Field + Restructure	New Field: Pre-cursor (-2) local setting	<i>Gen-Z-E-PAM4-50G PHY TX Lane Control Fields</i>
Additional Field + Restructure	New Field: Requested pre-cursor (-2) remote setting	<i>Gen-Z-E-PAM4-50G PHY Rx Lane Status Fields</i>

Table 8-19: Gen-Z-E-PAM4-50G PHY Specific Status Fields

Bit Location	Access	M/O	Description
3:0	RO	M	Physical layer state Reports current state of the physical layer 0x0—PHY Idle 0x1—Signal Detect 0x2—Channel Optimize 0x3—Align 0x4—Config 0x5—IBIST (optional) 0x6—PHY-Up 0x7—Reserved 0x8—Low Power 1

Bit Location	Access	M/O	Description
7:4			0x9—Low Power 2 0xA—0xB—Reserved 0xC—PHY-UP Low Power 1 0xD—0x1F—Reserved
	RO	O	Physical layer sub-state See <i>PHY Sub-state Mapping</i> for optional physical layer sub-state mappings.
11:8	RO	M	Previous Physical layer state When automatic re-initialization control is enabled, this field shall report the final state of the physical layer prior to reattempting physical layer training. The values are equivalent to those of the Physical layer state field.
15:12	RO	O	Previous Physical layer sub-state When automatic re-initialization control is enabled, this field shall report the final sub-state of the physical layer prior to reattempting physical layer training. The values are equivalent to those of the Physical layer sub-state field.
16	RO	O	PHY Tx precoding status PHY Tx precoding is negotiated and requested by the remote PHY Rx during physical layer training. 0b—disabled 1b—enabled
31:17	RO	-	RsvdZ

Table 8-20: Gen-Z-E-PAM4-50G PHY Specific Control Fields

Bit Location	Access	M/O	Description
0	RW	M	Automatic re-initialization control (AutoReinit) If enabled, the PHY will automatically attempt to restart training on errors during physical layer training. If disabled, the PHY will remain in the failing state to aid in debug and/or error recovery. 0b—disabled 1b—enabled
2:1	RW	M	25G PHY Up Clk Comp sequence count

Bit Location	Access	M/O	Description
			<p>This field shall specify the number of consecutive PHY Up Clk Comp sequences to be sent to allow a physical layer to autonomously compensate for clock differences with the remote PHY.</p> <p>0x0—1 0x1—2 0x2—3 (default) 0x3—4</p>
4:3	RW	M	<p>25G PHY Up Clk Compensation sequence frequency</p> <p>This field shall specify the minimum frequency for PHY Up Clk Comp sequences to be sent to allow a physical layer to autonomously compensate for clock differences with the remote PHY.</p> <p>0x0—1 per 3200 blocks 0x1—1 per 1600 blocks (default) 0x2—1 per 800 blocks 0x3—1 per 400 blocks</p> <p>A block shall be specified as 66 UI for 64b/66b and 130 UI for 128b/130b.</p>
6:5			<p>50G Clock Compensation Sequence Control</p> <p>This field controls disabling Clock Compensation Sequences or modifies the frequency that they are sent.</p> <p>0x0—Disabled 0x1— Enabled, 1 per 262,144 UI 0x2— Enabled, 1 per 131,072 UI (default) 0x3— Enabled, 1 per 65,536 UI</p>
7	RW	-	RsvdP
8	RW	M	<p>PHY Tx precoding</p> <p>PHY Tx precoding is negotiated and requested by the remote PHY Rx during physical layer training. This control field allows manual control of the Tx and when enabled must be supported and enabled on the remote PHY Rx. PHY Tx precoding shall be supported in this physical layer clause.</p> <p>0b—disabled or automatic 1b—enabled</p>
9	RW	O	PHY Rx precoding

Bit Location	Access	M/O	Description
			0b—disabled 1b—enabled
15:10	RW	-	RsvdP

Table 8-21: Gen-Z-E-PAM4-50G PHY Specific Capability Fields

Bit Location	Access	M/O	Description
0	RO	M	Automatic re-initialization control (AutoReinit) support 0b—Not supported 1b—Supported
1	RW	O	PHY Rx precoding support 0b—Not supported 1b—Supported
31:2	RO	-	RsvdZ

Table 8-22: Gen-Z-E-PAM4-50G Remote PHY Specific Capability Fields

Bit Location	Access	M/O	Description
0	RW	M	Remote PHY automatic re-initialization control (AutoReinit) support 0b—Not supported 1b—Supported
1	RW	O	Remote PHY Rx precoding support 0b—Not supported 1b—Supported
31:2	RW	-	RsvdP

Table 8-23: Gen-Z-E-PAM4-50G PHY TX Lane Status Fields

Bit Location	Access	M/O	Description
1:0	RO	M	Cursor status mode 0x0—PMD final adjustment count 0x1—Explicit ratio settings

Bit Location	Access	M/O	Description
5:2			0x2-0x3—Reserved
	RO	M	<p>Current pre-cursor (-2) local setting</p> <p>If cursor status mode is PMD final adjustment count:</p> <p>0x0—Reserved</p> <p>0x1—0x7—Final positive adjustments from current initialization mode (Initialize or Preset). Example, 0x3 → 3 positive adjustments.</p> <p>0x8—Reserved</p> <p>0x9—0xF—Final negative (minus 8) adjustments from current initialization mode (Initialize or Preset). Example, 0xA → 10-8 = 2 negative adjustments.</p> <p>If cursor status mode is explicit ratio settings:</p> <p>0x0—c(-2) ratio 0</p> <p>0x1—c(-2) ratio -0.05</p> <p>0x2—c(-2) ratio -0.1</p> <p>0x3—c(-2) ratio -0.15</p> <p>0x4-0xF—Reserved</p>
9:6	RO	M	<p>Current pre-cursor (-1) local setting</p> <p>If cursor status mode is PMD final adjustment count:</p> <p>0x0—Reserved</p> <p>0x1—0x7—Final positive adjustments from current initialization mode (Initialize or Preset). Example, 0x3 → 3 positive adjustments.</p> <p>0x8—Reserved</p> <p>0x9—0xF—Final negative (minus 8) adjustments from current initialization mode (Initialize or Preset). Example, 0xA → 10-8 = 2 negative adjustments.</p> <p>If cursor status mode is explicit ratio settings:</p> <p>0x0—c(-1) ratio 0</p> <p>0x1—c(-1) ratio -0.05</p> <p>0x2—c(-1) ratio -0.1</p> <p>0x3—c(-1) ratio -0.15</p> <p>0x4-0xF—Reserved</p>

Bit Location	Access	M/O	Description
13:10	RO	M	<p>Current post-cursor local setting</p> <p>If cursor status mode is PMD final adjustment count:</p> <p>0x0—Reserved</p> <p>0x1—0x7—Final positive adjustments from current initialization mode (Initialize or Preset). Example, 0x3 → 3 positive adjustments.</p> <p>0x8—Reserved</p> <p>0x9—0xF—Final negative (minus 8) adjustments from current initialization mode (Initialize or Preset). Example, 0xA → 10-8 = 2 negative adjustments.</p> <p>If cursor status mode is explicit ratio settings:</p> <p>0x0—c(1) ratio 0</p> <p>0x1—c(1) ratio -0.05</p> <p>0x2—c(1) ratio -0.1</p> <p>0x3—c(1) ratio -0.15</p> <p>0x4—c(1) ratio -0.2</p> <p>0x5—c(1) ratio -0.25</p> <p>0x6-0xF—Reserved</p>
17:14	RO	M	<p>Current output margin / primary-cursor local setting</p> <p>If cursor status mode is PMD final adjustment count:</p> <p>0x0—Reserved</p> <p>0x1—0x7—Final positive adjustments from current initialization mode (Initialize or Preset). Example, 0x3 → 3 positive adjustments.</p> <p>0x8—Reserved</p> <p>0x9—0xF—Final negative (minus 8) adjustments from current initialization mode (Initialize or Preset). Example, 0xA → 10-8 = 2 negative adjustments.</p> <p>If cursor status mode is explicit ratio settings:</p> <p>0x0—Normal output swing</p> <p>0x1—~87.5% output swing</p> <p>0x2—~75% output swing</p> <p>0x3—~62.5% output swing</p>

Bit Location	Access	M/O	Description
19:18			0x4—~50% output swing 0x5—~37.5% output swing 0x6—~25% output swing 0x7—0xF—Reserved
	RO	M	Current initialization mode 0x0—PMD Initialize settings 0x1—PMD Preset settings 0x2—0x3—Reserved
21:20	RO	M	Local equalization status 0x0—local PMD equalization not run 0x1—local PMD equalization in progress 0x2—local PMD equalization complete with no errors 0x3—local PMD equalization complete with errors
23:22	RO	M	Remote equalization status 0x0—remote PMD equalization not run 0x1—remote PMD equalization in progress 0x2—remote PMD equalization complete with no errors 0x3—remote PMD equalization complete with errors
31:24	RO	-	RsvdZ

Table 8-24: Gen-Z-E-PAM4-50G PHY TX Lane Control Fields

Bit Location	Access	M/O	Description
1:0	RW	M	Cursor control mode 0x0—Automatic PMD inband exchange (default) 0x1—Direct Interface PHY Structure control 0x2—0x3—Reserved
5:2	RW	M	Pre-cursor (-2) local setting If cursor mode is automatic PMD inband exchange, the setting shall be invalid and set to 0x0. If cursor control mode is Direct Interface PHY Structure control: 0x0—c(-2) ratio 0

Bit Location	Access	M/O	Description
9:6			0x1—c(-2) ratio -0.05 0x2—c(-2) ratio -0.1 0x3—c(-2) ratio -0.15 0x4-0xF—Reserved
	RW	M	Pre-cursor (-1) local setting If cursor mode is automatic PMD inband exchange, the setting shall be invalid and set to 0x0. If cursor control mode is Direct Interface PHY Structure control: 0x0—c(-1) ratio 0 0x1—c(-1) ratio -0.05 0x2—c(-1) ratio -0.1 0x3—c(-1) ratio -0.15 0x4-0xF—Reserved
13:10	RW	M	Post-cursor local setting If cursor mode is automatic PMD inband exchange, the setting shall be invalid and set to 0x0. If cursor control mode is Direct Interface PHY Structure control: 0x0—c(1) ratio 0 0x1—c(1) ratio -0.05 0x2—c(1) ratio -0.1 0x3—c(1) ratio -0.15 0x4—c(1) ratio -0.2 0x5—c(1) ratio -0.25 0x6-0xF—Reserved
17:14	RW	M	Output Margin / Primary-cursor local setting If cursor mode is automatic PMD inband exchange, the setting shall be invalid and set to 0x0. If cursor control mode is Direct Interface PHY Structure control: 0x0—Normal output swing 0x1—~87.5% output swing 0x2—~75% output swing 0x3—~62.5% output swing

Bit Location	Access	M/O	Description
19:18			0x4—~50% output swing 0x5—~37.5% output swing 0x6—~25% output swing 0x7—0xF—Reserved
	RW	M	Initialization mode 0x0—PMD Initialize settings 0x1—PMD Preset settings 0x2—0x3—Reserved
31:20	RW	-	RsvdP

Table 8-25: Gen-Z-E-PAM4-50G PHY Rx Lane Status Fields

Bit Location	Access	M/O	Description
0	RO	M	Signal detected 0b—signal not detected 1b—signal detected
1	RO	M	Current lane polarity 0b—polarity normal 1b—polarity reversed
2	RO	O	CDR status 0b—CDR not locked 1b—CDR locked
3	RO	O	Rx equalization status 0b—Equalization not run 1b—Equalization complete
4	RO	M	Rx symbol alignment status 0b—Symbol alignment not achieved 1b—Symbol alignment achieved
5	RO	M	Lane error status 0b—No error detected 1b—Errors detected

Bit Location	Access	M/O	Description
7:6	RO	-	RsvdZ
11:8	RO	O	Rx IBIST status 0x0—IBIST not run 0x1—IBIST run; Bit Error Ratio (BER) < ~1e-15 0x2—IBIST run; BER = ~1e-14 0x3—IBIST run; BER = ~1e-13 0x4—IBIST run; BER = ~1e-12 0x5—IBIST run; BER = ~1e-11 0x6—IBIST run; BER = ~1e-10 0x7—IBIST run; BER = ~1e-9 0x8—IBIST run; BER = ~1e-8 0x9—0xF—Reserved
15:12	RO	O	Requested pre-cursor (-2) remote setting 0x0—c(-2) ratio 0 0x1—c(-2) ratio -0.05 0x2—c(-2) ratio -0.1 0x3—c(-2) ratio -0.15 0x4-0xF—Reserved
19:16	RO	O	Requested pre-cursor (-1) remote setting 0x0—c(-1) ratio 0 0x1—c(-1) ratio -0.05 0x2—c(-1) ratio -0.1 0x3—c(-1) ratio -0.15 0x4-0xF—Reserved
23:20	RO	O	Requested post-cursor remote setting 0x0—c(1) ratio 0 0x1—c(1) ratio -0.05 0x2—c(1) ratio -0.1 0x3—c(1) ratio -0.15 0x4—c(1) ratio -0.2 0x5—c(1) ratio -0.25

Bit Location	Access	M/O	Description
27:24			0x6-0xF—Reserved
	RO	O	Requested output margin / primary-cursor remote setting 0x0—Normal output swing 0x1—~87.5% output swing 0x2—~75% output swing 0x3—~62.5% output swing 0x4—~50% output swing 0x5—~37.5% output swing 0x6—~25% output swing 0x7—0xF—Reserved
25	RO	O	Change in remote equalization requested 0b—No request 1b—Request for remote equalization setting change
31:26	RO	-	RsvdZ

8.5.2. Physical Layer Management

Refer to the Gen-Z Common Specification clause section on *Physical Layer Management*.

9. Gen-Z-E-PAM4-50G-Local Specification

9.1. Introduction

5 Gen-Z-E-PAM4-50G-Local defines a Gen-Z physical layer capable of a line-rate at 53.125 GT/s per lane inclusive of the overhead with Forward Error Correction (raw data-rate of 47.222 Gbps) using 4 level Pulse Amplitude Modulation (PAM4) signaling over Local mediums (Very Short Reach) providing a raw BER of 10^{-9} or better and a corrected BER of 10^{-15} or better. A compliant physical layer consists of an electrical interface compatible with the Channel Operating Margin (COM) that has a recommended Insertion Loss¹ (IL) no greater than 10 dB and no less than 4 dB at 13.2813 GHz.

9.2. Logical Sublayer

10 Refer to the Gen-Z-E-PAM4-50G-Fabric Specification clause section on *Logical Sublayer*.

9.2.1. Link Serialization

Refer to the Gen-Z-E-PAM4-50G-Fabric Specification clause section on *Link Serialization* with the modification that Phit FEC 288 is always used and Phit FEC 320 is never used.

9.2.2. Scrambling / Descrambling

15 Refer to the Gen-Z Common Specification clause section on *Scrambling / Descrambling*.

9.2.3. Data Striping

Refer to the Gen-Z Common Specification clause section on *Data Striping*.

9.2.4. PLA Data Width

Refer to the Gen-Z Common Specification clause section on *PLA Data Width*.

9.2.5. Physical Layer Initialization and Training

20 Refer to the Gen-Z-E-PAM4-50G-Fabric Specification clause section on *Physical Layer Initialization and Training*.

9.2.6. Lane Reversal

Refer to the Gen-Z Common Specification clause section on *Lane Reversal*.

9.2.7. Lane Polarity

25 Refer to the Gen-Z-E-NRZ-25G-Local Specification clause section on *Lane Polarity*.

9.2.8. Link Width

Refer to the Gen-Z-E-NRZ-25G-Local Specification clause section on *Link Width*.

9.2.9. Link BIST

Refer to the Gen-Z Common Specification clause section on *Link BIST*.

5 9.2.10. Physical Layer Retraining

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Physical Layer Retraining*.

9.2.11. PLTSM

Refer to the Gen-Z-E-NRZ-25G-Local Specification clause section on *PLTSM*.

9.2.11.1. PHY Idle (Required)

10 Refer to the Gen-Z Common Specification clause section on *PHY Idle (Required)*.

9.2.11.2. Signal Detect (Required)

Refer to the Gen-Z-E-PAM4-50G-Fabric Specification clause section on *Signal Detect (Required)*.

9.2.11.3. Channel Optimize (Required)

Refer to the Gen-Z-E-PAM4-50G-Fabric Specification clause section on *Channel Optimize (Required)*.

15 9.2.11.4. Align (Required)

Refer to the Gen-Z-E-PAM4-50G-Fabric Specification clause section on *Align (Required)*.

9.2.11.5. Interconnect BIST (Optional)

Refer to the Gen-Z-E-PAM4-50G-Fabric Specification clause section on *Interconnect BIST (Optional)*.

9.2.11.6. Config (Required)

20 Refer to the G Gen-Z-E-PAM4-50G-Fabric Specification clause section on *Config (Required)*.

9.2.11.7. PHY Up (Required)

Refer to the Gen-Z Common Specification clause section on *PHY-Up (Required)*.

9.2.12. Clock Compensation

Refer to the Gen-Z-E-PAM4-50G-Fabric Specification clause section on *Clock Compensation*.

9.2.13. Loopback

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Loopback*.

9.2.14. Physical Layer Aggregation

Refer to the Gen-Z-E-NRZ-25G-Local Specification clause section on *Physical Layer Aggregation*.

5 9.2.15. Re-timers

Re-timers for this this physical layer clause shall not be allowed.

Developer Note: *If a system topology using this specification clause exceeds its channels requirements, developers should consider using a physical layer capable of supporting the Gen-Z-E-PAM4-50G-Fabric specification.*

10 9.3. Power Management

Refer to the Gen-Z-E-PAM4-50G-Fabric Specification clause section on *Power Management*.

9.4. Physical Layer Electrical Sublayer

This Gen-Z physical layer electrical sublayer shall meet all single lane/channel requirements specified in Gen-Z-E-PAM4-50G-Fabric clause section *Electrical Interface* but with the following exceptions:

- 15 • Support for COM with a channel Insertion Loss¹ (IL) recommended to be no greater than 10 dB at 13.28125 GHz.
- Link equalization shall not require Physical Medium Dependent (PMD) inband exchange of transmit coefficients.

20 **Developer's Note:** *This physical layer clause defines the use of Medium Reach (MR) capable transmitters and receivers on restricted Very Short Reach (VSR) channels as shown in the Gen-Z-E-PAM4-50G-Local Electrical . This allows for a light weight, low latency Forward Error Correction (FEC) code at the expense of less correction gain.*




Figure 9-1: Gen-Z-E-PAM4-50G-Local Electrical Sublayer

9.4.1. Electrical Interface

9.4.1.1. Overview

5 The electrical interface for this physical layer clause shall be based on high speed, low voltage differential signaling where each connection is point-to-point and signaling is PAM4 and unidirectional.

9.4.1.1.1. Forwarded Clock

This physical layer clause shall not require a forwarded clock.

9.4.1.2. Lane Signaling Rate

10 The signaling rate for this physical layer clause shall be as specified in *IEEE Standard for Ethernet 802.3bs™-2017*, clause 120D.3.1 at 26.5625 GBd/s \pm 100 Parts Per Million (ppm) per lane using 4-level Pulse Amplitude Modulation (PAM4).

9.4.1.3. Lane Bit Error Rate

Each lane associated with this physical layer clause shall operate with a bit error ratio of 10^{-9} or better.

9.4.1.4. Lane Equalization

15 Tx Equalization coefficient range and step size requirements shall be as specified in *IEEE Standard for Ethernet 802.3bs™-2017* clause Annex 120D.3.1.5. The interface physical layer structure should provide fields to configure the transmit equalization coefficients, which can be manipulated by manageability prior to the Align physical layer training state.

9.4.2. High Speed Signaling Specification

9.4.2.1. Transmitter Specification

The transmitter for this physical layer clause shall be as specified in the CEI-56G-MR-PAM4 clause in IA # OIF-CEI-04.0.

5 9.4.2.2. Receiver Specification

The receiver for this physical layer clause shall be as specified in the CEI-56G-MR-PAM4 clause in IA # OIF-CEI-04.0.

9.4.3. Channel Characteristics

9.4.3.1. Channel Operating Margin

10 The channel for this physical layer clause shall meet compliance with the Channel Operating Margin (COM) for the CEI-56G-MR-PAM4 clause in IA # OIF-CEI-04.0 with the required modifications to Table 17-1:

- Signaling rate (f_b) = 26.5625 Gsym/s
- Target detector error ratio (DER_0) = 10^{-9}

15 9.4.3.2. Insertion Loss

The channel insertion loss ($IL(f)$ in dB) for this physical layer clause should meet the equation illustrated in *Gen-Z-E-PAM4-50G-Local Channel Insertion Loss*.

$$IL(f) \leq 0. + 7.5 \sqrt{\frac{f}{f_b}} + 2.38 \frac{f}{f_b} + 13.56 \left(\frac{f}{f_b}\right)^2 \quad (\text{dB})$$

$$IL(f) \geq \begin{cases} 0 & 0.5 \leq f \leq 1 \\ 0.303(f - 1) & 1 \leq f \leq 17.5 \\ 5.0 & 17.5 < f \leq f_b \end{cases} \quad (\text{dB})$$

Where,

f is the frequency in GHz

f_b is the signaling rate (26.5625) in Gb/s

$IL(f)$ is the insertion loss at frequency f




Figure 9-2: Gen-Z-E-PAM4-50G-Local Channel Insertion Loss

9.4.3.3. Return Loss

The Return Loss for this physical layer clause shall be as specified in CEI-56G-MR-PAM4 clause in IA # OIF-CEI-04.0.

9.4.3.4. Coupling

The transmitter for this physical layer clause shall be AC-coupled to the receiver as specified in CEI-56G-MR-PAM4 clause in IA # OIF-CEI-04.0.

9.4.4. Miscellaneous Specification

Refer to the Gen-Z-E-NRZ-25G-Fabric Specification clause section on *Miscellaneous Specification*.

9.5. Management Control and Status

9.5.1. Interface Physical Layer Structure

9.5.1.1. Common Interface Physical Layer Structure

9.5.1.1.1. Common, Lane, and Low Power Interface PHY Structure

The following defines the features and requirements associated with the Gen-Z-E-PAM4-50G-Local specification. The *Gen-Z-E-PAM4-50G-Local Interface PHY Structure Support Exclusions and Restrictions* indicates restricted or excluded fields or sub-fields from the *Interface PHY Structure Common, Lane, and Low Power Fields*.

Table 9-1: Gen-Z-E-PAM4-50G-Local Interface PHY Structure Support Exclusions and Restrictions

Support	Sub-field / Description	Interface PHY Structure Field	Supported Value(s)
Restricted	PHY Clause 6—50G Local	PHY Type	0x5
Restricted	Physical Layer Operational Status 0x0—PHY-Down, uninitialized state 0x1—PHY-Up 0x2—PHY-Down-Retrain 0x3 – 0x6 Reserved 0x7—PHY Low-Power 1 → Mapped to P1 0x8—PHY Low-Power 2 → Mapped to P2 0x9 – 0xA Reserved 0xB – PHY UP Low-Power 1 → Mapped to Dynamic Link Width Reduction 0xC—0xF Reserved	PHY Status	See Description
Restricted	Previous Physical Layer Operational Status 0x0—PHY-Down, uninitialized state 0x1—PHY-Up 0x2—PHY-Down-Retrain 0x3 – 0x6 Reserved 0x7—PHY Low-Power 1 → Mapped to P1 0x8—PHY Low-Power 2 → Mapped to P2 0x9 – 0xA Reserved 0xB – PHY UP Low-Power 1 → Mapped to Dynamic Link Width Reduction 0xC—0xF Reserved	PHY Status	See Description
Restricted	Physical Layer Retraining Arc Only Retraining Arc 1 and Arc 2 are supported for this physical layer clause	PHY Control	0x0 – 0x1
Excluded	Retrain Arc 3-4 Support	PHY Cap 1	0b
Restricted	Asymmetric Lane with Reversal Support	PHY Lane Cap	0b
Restricted	Asymmetric Lane Support	PHY Remote Lane Cap	0x0

Support	Sub-field / Description	Interface PHY Structure Field	Supported Value(s)
Excluded	Entry, Exit Latency PHY-LP 3-4	PHY Low Power Timing Capability	0x0
Excluded	PHY-LP 3-4 Support	PHY Low Power CAP	0b
Excluded	Entry, Exit Latency PHY-UP-LP 2-4	PHY Up Low Power Timing Capability	0x0
Excluded	PHY-UP-LP 2-4 Support	PHY Up Low Power CAP	0b

9.5.1.1.2. Extended Feature Interface PHY Structure

The Extended Interface Physical Layer structure for the Gen-Z-E-PAM4-50G-Local specification references the *Extended Feature Interface PHY Structure* from the Gen-Z-E-PAM4-50G-Fabric specification clause.

5

9.5.1.2. Unique Interface Physical Layer Structure

The Unique Interface Physical Layer structure for the Gen-Z-E-PAM4-50G-Local specification references the *Unique Interface Physical Layer Structure* from the Gen-Z-E-PAM4-50G-Fabric specification clause. The *Gen-Z-E-PAM4-50G-Local Unique Interface PHY Structure Requirements* specifies the differences.

Table 9-2: Gen-Z-E-PAM4-50G-Local Unique Interface PHY Structure Requirements

Differences	Sub-field / Description	Interface PHY Structure Field
Excluded	Current pre-cursor (-2) local setting Pre-cursor (-2) is not supported for this physical layer clause	<i>Gen-Z-E-PAM4-50G PHY TX Lane Status Fields</i>
Excluded	Pre-cursor (-2) local setting Pre-cursor (-2) is not supported for this physical layer clause	<i>Gen-Z-E-PAM4-50G PHY TX Lane Control Fields</i>
Default Value	Cursor control mode 0x0—Automatic PMD inband exchange 0x1—Direct Interface PHY Structure control (default) 0x2—0x3—Reserved	<i>Gen-Z-E-PAM4-50G PHY TX Lane Control Fields</i>

Differences	Sub-field / Description	Interface PHY Structure Field
Excluded	Requested pre-cursor (-2) remote setting Pre-cursor (-2) is not supported for this physical layer clause	<i>Gen-Z-E-PAM4-50G PHY Rx Lane Status Fields</i>
Default Value	Channel Optimize state control The Channel Optimize state is optional for this physical layer clause. 0b—Channel Optimize state bypassed (default)	<i>PHY Channel Optimize State Control Fields</i>

9.5.2. Physical Layer Management

Refer to the Gen-Z Common Specification clause section on *Physical Layer Management*.

Appendix A Gen-Z AIC Specification

A.1. AIC Introduction

Gen-Z AIC specifies the loss budgets supported with Gen-Z physical layers for Add-in Cards (AIC). *Gen-Z AIC Reference* shows an example of a one connector AIC topology, where A defines the routing length (reach) on the host Multi-Layered Board (MLB) and B defines the routing length on the AIC.




Figure A-0-1: Gen-Z AIC Reference

A.2. AIC Loss Budget

AIC loss budgets for each Gen-Z physical layer clause are divided into two sections to accompany different measurement locations—the die pad interface and the package to board interface as illustrated in *Gen-Z AIC Reference*.

A.2.1 Die Pad Interface Loss Budgets

The Insertion Loss (IL) budgets in this section follow PCI-SIG standard conventions with test points at the die pad interface and inclusive of the package. Total AIC insertion loss shall comply with following requirements:

- For Gen-Z PECFF and PCIe CEM AICs, the Gen-Z-E-NRZ-PCIe (16G) AIC total loss (including package) shall not exceed 8 dB up to 8 GHz.
- For Gen-Z PECFF and PCIe CEM AICs, the Gen-Z-E-NRZ-PCIe (32G) AIC total loss (including package) shall not exceed 9 dB up to 16 GHz.

When designing the channel pertaining to the topology in *Gen-Z AIC Reference* at the die pad interface, PECFF and CEM AIC developers should follow the guidelines specified in *PECFF and CEM AIC Die Pad Interface Design Guidelines*.

Table A-1: PECFF and CEM AIC Die Pad Interface Design Guidelines

Gen-Z Physical Layer Specification	Channel	Add-in Card		Remaining Host MLB	
	Loss (dB)	Loss (dB)	Reach B (inch)	Loss (dB)	Reach A (inch)
Gen-Z-E-NRZ-PCIe (16G)	28	8	4.1	14	11.6
Gen-Z-E-NRZ-PCIe (32G)	36	9	4.1	17	14.1

Developer Note:

- Full channel and host MLB loss budgets in PECFF and CEM AIC Die Pad Interface Design Guidelines factor in 1dB insertion loss for the AIC connector.
- Gen-Z-E-NRZ-PCIe (16G) host MLB reach was computed assuming a package insertion loss of 5 dB at 8 GHz and is not included in the Remaining Host MLB loss budget.
- Gen-Z-E-NRZ-PCIe (16G) AIC reach was computed assuming a package insertion loss of 3 dB at 8 GHz and is included in the Add-in Card loss budget.
- The trace reach for 16G in PECFF and CEM AIC Die Pad Interface Design Guidelines on the host MLB (A) and the AIC (B) was computed based on materials capable of 1.2 dB/inch transmission line losses at 8 GHz.
- Gen-Z-E-NRZ-PCIe (32G) host MLB reach was computed assuming a package insertion loss of 9 dB at 16 GHz and is not included in the Remaining Host MLB loss budget.
- Gen-Z-E-NRZ-PCIe (32G) AIC reach was computed assuming a package insertion loss of 4 dB at 16 GHz and is included in the Add-in Card loss budget.
- The trace reach for 32G in PECFF and CEM AIC Die Pad Interface Design Guidelines on the host MLB (A) and the AIC (B) was computed based on materials capable of 1.2 dB/inch transmission line losses at 16 GHz.

A.2.2 Package to Board Interface Loss Budgets

The Insertion Loss (IL) budgets in this section follow 802.3 and OIF-CEI standard conventions with test points at the package to board interface. Total PECFF and CEM AIC insertion loss shall comply with following requirements:

- Gen-Z-E-NRZ-25G-Fabric AIC board loss shall not exceed 4 dB up to 12.89 GHz
- Gen-Z-E-NRZ-25G-Local AIC board loss shall not exceed 4 dB up to 12.89 GHz
- Gen-Z-E-PAM4-50G-Fabric AIC board loss shall not exceed 4 dB up to 13.2813 GHz
- Gen-Z-E-PAM4-50G-Local AIC board loss shall not exceed 4 dB up to 13.2813 GHz

When designing the channel pertaining to the topology in *Gen-Z AIC Reference* at the package to board interface, PECFF and CEM AIC developers should follow the guidelines specified in *Gen-Z PECFF and CEM AIC Package to Board Interface Design Guidelines*.

Table A-2: Gen-Z PECFF and CEM AIC Package to Board Interface Design Guidelines

Gen-Z Physical Layer Specification	Channel	Add-in Card		Remaining Host MLB	
	Loss (dB)	Loss (dB)	Reach B (inch)	Loss (dB)	Reach A (inch)
Gen-Z-E-NRZ-25G-Fabric	30	4	4.4	25	27.8
Gen-Z-E-NRZ-25G-Local	10	4	4.4	5	5.6
Gen-Z-E-PAM4-50G-Fabric	20	4	4	15	15
Gen-Z-E-PAM4-50G-Local	10	4	4	5	5

Developer Note:

- Full channel and host MLB loss budgets in Gen-Z PECFF and CEM AIC Package to Board Interface Design Guidelines factor in 1dB insertion loss for the AIC connector.
- The informative trace reach for 25G in Gen-Z PECFF and CEM AIC Package to Board Interface Design Guidelines on the host MLB (A) and the AIC (B) was computed based on materials capable of 0.9 dB/inch transmission line losses assumption at 12.89 GHz.
- The informative trace reach for 50G in Gen-Z PECFF and CEM AIC Package to Board Interface Design Guidelines on the host MLB (A) and the AIC (B) was computed based on materials capable of 1.0 dB/inch transmission line losses assumption at 13.2813 GHz.

5