

Errata for the PCI Express Base Specification, Revision 1.0

15 Apr 2003

Protocol WG Errata:

C1 – Release Date 11 Feb 2003
C2 – Release Date 11 Feb 2003
C3 – Release Date 11 Feb 2003
C4 – Release Date 11 Feb 2003
C5 – Release Date 11 Feb 2003
C6 – Release Date 11 Feb 2003
C7 – Release Date 11 Feb 2003
C8 – Release Date 11 Feb 2003
C9 – Release Date 11 Feb 2003
C10 – Release Date 11 Feb 2003
C11 – Release Date 11 Feb 2003
C12 – Release Date 11 Feb 2003
C13 – Release Date 11 Feb 2003
C14 – Release Date 11 Feb 2003
C15 – Release Date 11 Feb 2003
C16 – Release Date 11 Feb 2003
C17 – Release Date 11 Feb 2003
C18 – Release Date 11 Feb 2003
C19 – Release Date 11 Feb 2003
C20 – Release Date 11 Feb 2003
C21 – Release Date 11 Feb 2003
C22 – Release Date 11 Feb 2003
C23 – Release Date 11 Feb 2003
C24 – Release Date 11 Feb 2003
C25 – Release Date 11 Feb 2003
C26 – Release Date 11 Feb 2003
C27 – Release Date 11 Feb 2003
C28 – Release Date 11 Feb 2003
C29 – Release Date 11 Feb 2003
C30 – Release Date 11 Feb 2003
C31 – Release Date 11 Feb 2003
C32 – Release Date 11 Feb 2003
C33 – Release Date 11 Feb 2003
C34 – Release Date 11 Feb 2003
C35 – Release Date 11 Feb 2003
C36 – Release Date 11 Feb 2003
C37 – Release Date 11 Feb 2003
C38 – Release Date 11 Feb 2003
C39 – Release Date 11 Feb 2003
C40 – Release Date 11 Feb 2003
C41 – Release Date 11 Feb 2003
C42 – Release Date 11 Feb 2003
C43 – Release Date 11 Feb 2003
C44 – Release Date 11 Feb 2003
C45 – Release Date 11 Feb 2003

C46 – Release Date 11 Feb 2003
C47 – Release Date 11 Feb 2003
C48 – Release Date 11 Feb 2003
C49 – Release Date 11 Feb 2003
C50 – Release Date 11 Feb 2003
C51 – Retracted
C52 – Release Date 11 Feb 2003
C53 – Release Date 27 Mar 2003
C54 – Release Date 27 Mar 2003
C55 – Release Date 27 Mar 2003
C56 – Release Date 27 Mar 2003
C57 – Release Date 27 Mar 2003
C58 – Release Date 27 Mar 2003
C59 – Release Date 27 Mar 2003
<C58-C66 below>
C67 – Release Date 15 Apr 2003

Electrical WG Errata:

C60 – Release Date 27 Mar 2003
C61 – Release Date 27 Mar 2003
C62 – Release Date 27 Mar 2003
C63 – Release Date 27 Mar 2003
C64 – Release Date 27 Mar 2003
C65 – Release Date 27 Mar 2003
C66 – Release Date 27 Mar 2003

E1 (p.246 & p.33) – Release Date 11 Feb 2003
E2.1 through E2.5 - Release Date 11 Feb 2003
E2.6 through E2.20 - Release Date 11 Feb 2003
E3.1 - Release Date 11 Feb 2003
E4.1 - Release Date 11 Feb 2003
E4.2 through E4.15 - Release Date 11 Feb 2003
E2.21 through E2.23 - Release Date 27 Mar 2003
E4.16 - Release Date 27 Mar 2003
E2.24, E3.2, E4.17 - Release Date 15 Apr 2003

Electrical WG Errata:

Chapter 4 Electrical (EWG) Content and Editorial
errata – Release Date 11 Feb 2003

REVISION	REVISION HISTORY	DATE
A	Initial revision of errata document	20 Dec 2002
A	Addendum	14 Feb 2003
A	<p>Final for 1.0 to 1.0a --</p> <p>Errata C66 added based on feedback to earlier errata from PCISIG membership. Summary of C66:</p> <ol style="list-style-type: none">1. Clarifications to common sense rules (e.g.: fundamental reset only applies when power is on)2. Clarifications regarding test loads and how they relate to eye diagrams. Additional wording added to point out that test loads are placed at package pins and that the transmit and receive eye diagrams apply to the package pins with 50 ohm loads. We also point out that the eye at the package pins will be different than the eye at the silicon pads dependent primarily on the package design.3. Notes added to clarify LTSSM text - no content/functional changes. L0/L0s interactions clarified. Added some explanatory text to improve spec understanding. <p>Errata C67 added added to clarify definition of Hwlnit.</p> <p>Small modifications to C65 and C53 (indicated in document)</p> <p>Noted in Editorial Errata section that number of capitalization, grammatical, spelling and formatting editorial errata were fixed at the discretion of the technical writer.</p>	15 Apr 2003

Table of Contents

Table of Contents	3
Content Errata – Chaps 1-3, 5-7	5
C1. Max_Payload_Size	5
C2. 4K Request Restriction	6
C3. TD/EP Definition Errata	7
C4. Switch Support for INTx Messages	8
C5. Missing Sub-Row for PM_PME	9
C6. Byte Count Calculation Table	10
C7. Secondary Bus Reset is Hot Reset	12
C8. RCRB Optional Clarification	12
C9. Extended Tag Field Clarification	12
C10. Size of MM Cfg Space Access Operations	13
C11. Improved SW Model Support for Advanced Error Reporting	14
C12. Slot Power Limit Message also uses MsgD	16
C13. Clarify Error Tables	16
C14. Byte Count and Lower Address for Non-Mem Completions	18
C15. Starting Point for REPLAY_TIMER	19
C16. Use of Cpl vs. CplD	21
C17. Error in Table for Hot Plug Message	21
C18. Inconsistency in Header Log Register Diagram	22
C19. Clarify use of “Bridge”	23
C20. Switch Unlocking Rule	23
C21. Clarify Multi-Function Devices with Differing Error Severities	25
C22. Clarify BME Bit Bridge vs. Device Operation	25
C23. Clarify Link Retrain and Link Training Bits	26
C24. Clarify Extended Sync Bit	27
C25. Changing Common Clock Config Bit	28
C26. Clarify Length as Reserved for Cpl	29
C27. Reset Conditions for Replay_Timer	29
C28. FC Updates for Infinite FC Advertisements	31
C29. Size of RCRB Access Operations	32
C30. Clarify Requirements for Sticky Registers	33
C31. Replay Timing L0s Correction Factor	35
C32. Wrong Units in Figure 5-7	38
C33. Ack/Nak Rule Clarification	39
C34. Error Logging Clarification	40
C35. Poisoned Write Clarification	41
C36. Remove Bridge Spec Material	42
C37. Remove Unnecessary Root Complex Coherence Reference	44
C38. Read/Write Fragmentation	44
C39. Clarify Device/Bridge Behavior for Config Bits	46
C40. In Msg Tables, All ‘B’ Should Be ‘BD’	48
C41. PCI-X Bridges can also set BCM	51
C42. Configuration Address Mapping & Rules Clarification	51

C43. Clarify Slot Power Limit Register Operation	52
C44. Clarify use of Max Read Request Size	53
C45. Byte Enable Valid Usage Model	54
C46. Reset/"Power Good" – Protocol Aspects.....	54
C47. Wakeup Mechanisms – Protocol Aspects.....	60
C48. Device 0 Decode Clarification	70
C49. Remove Advanced Switching References.....	70
C50. Optional DLLP Timeout Mechanism.....	76
C51. Retracted.....	76
C52. Clarify Configuration Retry	77
C53. Assorted Errata	77
C54. Non-Infinite FC permission for RC	96
C55. Slot/Common clock config clarification.....	97
C56. L0s Invocation and Exit Policy	97
C57. Aux Power for Beacon	99
C58. Legacy Endpoint Requirements.....	99
C59. Integrated device.....	100
C60 Receiver FTS Timeout	100
C61 - Add the terms TTX-IDLE-MIN and TTX-IDLE-SET-TO-IDLE to the L2.Idle state.	101
C62 - Detect must always change the voltage towards ground.	101
C63 - Loopback in 10 bit domain.	102
C64. - Scrambling LFSR Equation Change	104
C65 Assorted Errata	115
C66 Assorted Errata	137
C67 Add Item to HwInit Definition	150
Editorial Errata – Chaps 1-3, 5-7	151
E1. Section Cross References.....	151
E2. Typographical Errors.....	151
E3. Formatting Errors.....	156
E4. Minor Wording Errors.....	157
Content and Editorial Errata – Chapter 4	161
Appendix A – Chapter 4 w/Changebars 22Jul to 11 Dec (Released 11 Feb 2003 – Does Not Include C54 onward)	232
Physical Layer Specification (Released 11 Feb 2003 – Does Not Include C54 onward).....	233
Appendix B – New and Redrawn Figures.....	325

Content Errata – Chaps 1-3, 5-7

This section corrects errors that affect the technical meaning of the specification and clarifies ambiguous or incomplete areas of the specification. These corrections will be included in a future revision of the specification.

C1. Max_Payload_Size

Release Date: 2/11/03

Clarify operation and requirements of Max_Payload_Size.

Section 2.2.2, p. 51, change text as shown:

- ❑ The **Transmitter of a TLP with a data payload** ~~of a TLP~~ must not **allow the data payload length to** exceed the length specified by the value in the Max_Payload_Size field of the **Transmitter's** Device Control Register (see Section <7.8.4>).
 - Note: Max_Payload_Size applies only to TLPs with data payloads; Memory Read Requests are not restricted in length by Max_Payload_Size. The size of the Memory Read Request is controlled by the Length field
- ❑ **The data payload of a Received TLP must not exceed the length specified by the value in the Max_Payload_Size field of the Receiver's Device Control Register (see Section <7.8.4>).**
 - Receivers must check for violations of this rule. If a Receiver determines that a TLP violates this rule, the TLP is a Malformed TLP
 - ? This is a reported error associated with the Receiving Port (see Section <6.2>)

Section 2.3.1.1, p. 90, change text as shown:

- ❑ Completions must not include more data than permitted by the Max_Payload_Size parameter. ~~calculated as a naturally aligned boundary.~~
 - Receivers must check for violations of this rule – TLPs in violation are Malformed TLPs
 - ? This is a reported error associated with the Receiving Port (see Section <6.2>)

Note: This is simply a special case of the rules which apply to all TLPs with data payloads

Add this implementation note following Table 7-12 "Device Control Register" (currently on p. 342).



IMPLEMENTATION NOTE

Use of Max_Payload_Size

The Max_Payload_Size mechanism allows software to control the maximum payload in packets sent by Endpoints to balance latency versus bandwidth trade-offs, particularly for isochronous traffic.

If software chooses to program the Max_Payload_Size of various system elements to non-default values, it must take care to ensure that each packet does not exceed the Max_Payload_Size parameter of any system element along the packet's path. Otherwise, the packet will be rejected by the system element whose Max_Payload_Size parameter is too small.

Discussion of specific algorithms used to configure Max_Payload_Size so that this requirement is met is beyond the scope of this specification, but software should base its algorithm upon factors such as the following:

- ☐ the Max_Payload_Size capability of each system element within a hierarchy
- ☐ awareness of when system elements are added or removed through hot-plug operations
- ☐ knowing which system elements send packets to each other, what type of traffic is carried, what type of transactions are used, or if packet sizes are constrained by other mechanisms

For the case of firmware that configures system elements in preparation for running legacy OS environments, the firmware may need to avoid programming a Max_Payload_Size above the default of 128B, which is the minimum supported by Endpoints.

For example, if the OS environment does not comprehend PCI Express, firmware probably should not program a non-default Max_Payload_Size for a hierarchy that supports hot-plug operations. Otherwise, if no software is present to manage Max_Payload_Size settings when a new element is added, improper operation may result. Note that a newly added element may not even support a Max_Payload_Size setting as large as the rest of the hierarchy, in which case software may need to deny enabling the new element or reduce the Max_Payload_Size settings of other elements.

C2. 4K Request Restriction

Release Date: 2/11/03

Clarify the application and requirements of the restriction on Requests against passing a 4K memory address boundary.

Section 2.2.2, p.51 - The current placement of the following bullet implies the rule applies only to writes.

- ❑ Requests must not specify an Address/Length combination which causes a Memory Space access to cross a 4K boundary.
 - Receivers may optionally check for violations of this rule. If a Receiver implementing this check determines that a TLP violates this rule, the TLP is a Malformed TLP
 - ? If checked, this is a reported error associated with the Receiving Port (see Section <6.2>)

This rule (highlighted) should be moved to Section 2.2.7, and edits made as shown (rule itself is unmodified):

For Memory Requests, the following rules ~~applies~~ apply:

- ❑ Memory Requests route by address, using either 64 bit or 32 bit Addressing (see Figure <2 13> and <Figure 2 14>)
- ❑ Requests must not specify an Address/Length combination which causes a Memory Space access to cross a 4K boundary.
 - Receivers may optionally check for violations of this rule. If a Receiver implementing this check determines that a TLP violates this rule, the TLP is a Malformed TLP
 - ? If checked, this is a reported error associated with the Receiving Port (see Section <6.2>)

C3. TD/EP Definition Errata

Release Date: 2/11/03

Fix errata caused by edits missed in earlier change to the TD/EP bit definitions

Section 2.2.3, p.52, edit as shown:

- ❑ For any TLP, a value of 1b in the TD field indicates the presence of the TLP Digest field ~~including an ECRC value~~ at the end of the TLP
 - ~~• The presence or absence of the TLP Digest field must be checked for all TLPs~~
 - A TLP where the TD field value does not correspond with the observed size (accounting for the data payload, if present) ~~with a 1b in the TD field but without a TLP Digest, or a TLP with a TLP Digest but without a 1b in the TD field,~~ is a Malformed TLP
 - ? This is a reported error associated with the Receiving Port (see Section <6.2>)
- ~~❑ For any TLP with a TLP Digest field, a value of 1b in the EP field indicates that the TLP Digest field is used for an end-to-end CRC (ECRC)~~
 - ~~• The presence or absence of the ECRC must be checked for all TLPs~~

☐ If the device at the ultimate destination of the TLP-

⊕☐ ~~supports neither data poisoning nor does not support~~ ECRC checking, the device must ignore the TLP Digest

~~• supports data poisoning but not ECRC checking, the device interprets the value in the TLP Digest field according to Section 2.7.2~~

- If the device at the ultimate destination of the TLP supports ECRC checking, the device interprets the value in the TLP Digest field as an ECRC value, according to the rules in Section <2.7.1>

Section 2.7, p.115, edit as shown:

The basic data reliability **mechanism** in PCI Express is ~~achieved-contained~~ within the Data Link Layer, which uses a 32-bit CRC ... On all other TLPs a Switch must preserve the ECRC (forward it untouched) as an integral part of the TLP.

In some cases, the data in a TLP payload is known to be corrupt at the time the TLP is generated, or may become corrupted while passing through an intermediate component, such as a Switch. In these cases, error forwarding, also known as data poisoning, can be used to indicate the corruption to the device consuming the data.

Section 2.7.2, p.119, edit as shown:

Error Forwarding (also known as data poisoning), is ~~enabled-indicated in PCI Express by either modifying the value placed in the TLP Digest field or by setting the proper value in the TD and EP field to 1bs~~. The rules for doing this are specified below.

C4. Switch Support for INTx Messages

Release Date: 2/11/03

Switches are required to support the reception and transmission of INTx Messages. This is described in the text, but not shown in the table of INTx Messages.

Table 2-12, p.52, edit as shown (note: footnotes remain – they are not shown here for clarity):

Name	Code[7:0]	Routing r[2:0]	Support				Req ID	Description/Comments
			R C	E p	S w	B r		
Assert_INTA	0010 0000	100	All:				B	Assert INTA virtual wire Note: These Messages are used for PCI 2.3 compatible INTx emulation
			r		tr			
			As Required:					
				t		t		
Assert_INTB	0010 0001	100	All:				B	Assert INTB virtual wire
			r		tr			
			As Required:					

Name	Code[7:0]	Routing r[2:0]	Support				Req ID	Description/Comments
			R C	E p	S w	B r		
				t		t		
Assert_INTC	0010 0010	100	All:				B	Assert INTC virtual wire
			r		tr			
			As Required:					
				t		t		
Assert_INTD	0010 0011	100	All:				B	Assert INTD virtual wire
			r		tr			
			As Required:					
				t		t		
Deassert_INTA	0010 0100	100	All:				B	De-assert INTA virtual wire
			r		tr			
			As Required:					
				t		t		
Deassert_INTB	0010 0101	100	All:				B	De-assert INTB virtual wire
			r		tr			
			As Required:					
				t		t		
Deassert_INTC	0010 0110	100	All:				B	De-assert INTC virtual wire
			r		tr			
			As Required:					
				t		t		
Deassert_INTD	0010 0111	100	All:				B	De-assert INTD virtual wire
			r		tr			
			As Required:					
				t		t		

C5. Missing Sub-Row for PM_PME

Release Date: 2/11/03

There is a missing sub-row for PM_PME in the table for Power Management Messages.

Table 2-14, p.72, edit as shown:

Table 2-14: Power Management Messages

Name	Code[7:0]	Routing r[2:0]	Support				Req ID	Description/Comments
			R C	E p	S w	B r		

PM_Active_State_Nak	0001 0100	100	t	r	tr	r	B	Terminate at Receiver
PM_PME	0001 1000	000	All:				BDF	Sent Upstream by PME-requesting component. Propagates Upstream.
			r		tr	t		
			If PME supported:					
				t				
PME_Turn_Off	0001 1001	011	t	r		r	BDF	Broadcast Downstream
PME_TO_Ack	0001 1011	101	r	t		t	BDF	Sent Upstream by Endpoint. Sent Upstream by Switch when received on all Downstream Ports.

C6. Byte Count Calculation Table

Release Date: 2/11/03

There is an error in one row of the Table 2-21, and a missing row for the case of no bytes enabled (zero length read, a.k.a. “flush”).

Table 2-21, p.93, edit as shown (Note: existing footnotes not shown here for clarity):

Table 2-21: Calculating Byte Count from Length and Byte Enables

1 st DW BE[3:0]	Last DW BE[3:0]	Total Byte Count
1xx1	0000	4
01x1	0000	3
1x10	0000	3
0011	0000	2
0110	0000	2
1100	0000	2
0001	0000	1
0010	0000	1
0100	0000	1
1000	0000	1
0000	0000	1
xxx1	1xxx	Length * 4
xxx1	01xx	(Length * 4) -1
xxx1	001x	(Length * 4) -2
xxx1	0001	(Length * 4) -3
xx10	1xxx	(Length * 4) -1
xx10	01xx	(Length * 4) -2
xx10	001x	(Length * 4) -3

1 st DW BE[3:0]	Last DW BE[3:0]	Total Byte Count
0001 x10	10000001	(Length * 4) -4
x100	1xxx	(Length * 4) -2
x100	01xx	(Length * 4) -3
x100	001x	(Length * 4) -4
x100	0001	(Length * 4) -5
1000	1xxx	(Length * 4) -3
1000	01xx	(Length * 4) -4
1000	001x	(Length * 4) -5
1000	0001	(Length * 4) -6

C7. Secondary Bus Reset is Hot Reset

Release Date: 2/11/03


Table 7-6 incorrectly states that setting the Secondary Bus Reset bit of the Bridge Control Register causes a Warm Reset – it should say Hot Reset. Additionally, the term “hierarchy domain” is used incorrectly – fix by deleting this text.

Section 7.5.3.3, Table 7-6, edit as shown:

6	Secondary Bus Reset – Setting this bit triggers a warm-hot reset on the corresponding PCI Express Port and the PCI Express hierarchy domain subordinate to the Port. Default value of this field is 0.	RW
---	--	----

C8. RCRB Optional Clarification

RCRBs are optional as indicated in Section 7.2.3 of the specification. Update Figure 7-1 appropriately to match with Section 7.2.3:



DM14256

Figure 7-1: PCI Express Root Complex Device Mapping

C9. Extended Tag Field Clarification

Release Date: 2/11/03

Clarify that the extended tag support bit applies to the device as a Requester

Section 7.8.3, Table 7-11, edit as shown:

5	<p>Extended Tag Field Supported – This field indicates the maximum supported size of the Tag field as a Requester. Defined encodings are:</p> <table><tr><td>0b</td><td>5-bit Tag field supported</td></tr><tr><td>1b</td><td>8-bit Tag field supported</td></tr></table> <p>Note that 8-bit Tag field support must be enabled by the corresponding control field in the Device Control register.</p>	0b	5-bit Tag field supported	1b	8-bit Tag field supported	RO
0b	5-bit Tag field supported					
1b	8-bit Tag field supported					

C10. Size of MM Cfg Space Access Operations

Release Date: 2/11/03

In an earlier revision, we had some text in the specification concerning the permitted sizes of reads and writes to the enhanced config access mechanism using memory mapped config space. This text was not accurate and was removed, but was not replaced, leaving this area unspecified. The following change restricts the permitted access sizes to DW and smaller and forbids accesses that cross DW boundaries. Additionally, locked requests to this space are expressly prohibited.

Section 7.2.2, p.314, add as shown:

The enhanced PCI Express configuration access mechanism utilizes a flat memory-mapped address space to access device configuration registers. In this case, the memory address determines the configuration register accessed and the memory data returns the contents of the addressed register. The mapping from memory address A[27:0] to PCI Express configuration space address is defined in Table 7-1. The base address A[63:28] is allocated in an implementation specific manner and reported by the system firmware to the operating system.

In some processor architectures, it is possible to generate memory space requests that cannot be expressed in a single Configuration Request, for example due to crossing a DW aligned boundary, or because a locked access is used. A Root Complex implementation is not required to support the translation to Configuration Requests of such memory space requests.

Table 0-1: Configuration Address Mapping

Memory Address	PCI Express Configuration Space
A[27:20]	Bus[7:0]
A[19:15]	Device[4:0]
A[14:12]	Function[2:0]
A[11:8]	Extended Register[3:0]
A[7:0]	Register[7:0]



IMPLEMENTATION NOTE

Generating Configuration Requests

Because Root Complex implementations are not required to support the generation of Configuration Requests from memory space accesses that cross DW boundaries, or that use locked semantics, software should take care not to cause the generation of such requests when using the memory-mapped configuration access mechanism unless it is known that the Root Complex implementation being used will support the translation.

C11. Improved SW Model Support for Advanced Error Reporting

Release Date: 2/11/03

Advanced Error Reporting provides a new mechanism for detailed error reporting. In the process of reviewing the software usage model it was determined that additional information was required in the Root Error Status Register to support models where different handlers are used for Fatal and Non-Fatal errors. The least disruptive way to include this additional information appears to be the addition of three bits: one to indicate Fatal/Non-Fatal for the 1st error message received, one to indicate that one or more Fatal error messages were received, and one to indicate that one or more Non-Fatal error messages were received.

Section 7.10.10, p.370, add as shown:

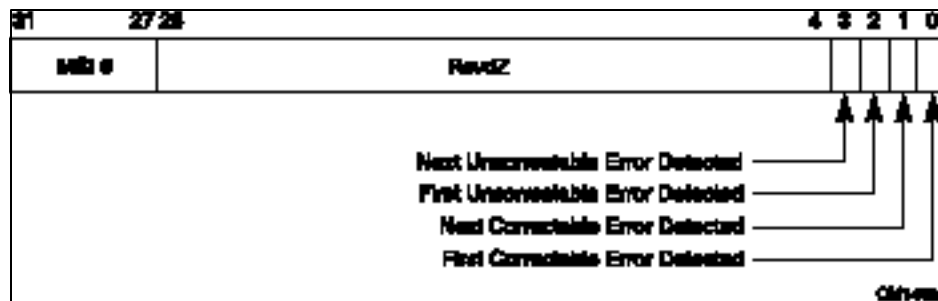


Figure 0-1: Root Error Status Register <!!MODIFY – See table>

The Root Error Status register reports status of error messages (ERR_COR), ERR_NONFATAL, and ERR_FATAL) received by the root complex, and of errors detected by the Root Port itself (which are treated conceptually as if the Root Port had sent an error message to itself). ERR_NONFATAL and ERR_FATAL messages are grouped together as uncorrectable. Each correctable and uncorrectable (non-fatal and fatal) error source has a first error bit and a next error bit associated with it respectively. When an error is received by a Root Complex, the respective first error bit is set and the Requestor ID is logged in the Error Source Identification register. A set individual error status bit indicates that a particular error category occurred; software may clear an error status by writing a 1 to the respective bit. If software does not clear the first reported error before another error message is received of the same category (correctable or uncorrectable), the corresponding

next error status bit will be set but the Requestor ID of the subsequent error message is discarded. The next error status bits may be cleared by software by writing a 1 to the respective bit as well. Refer to Section ~~<6.2>~~6.2 for further details. This register is updated regardless of the settings of the Root Control register and the Root Error Command register.

Table ~~7-327-32~~: Root Error Status Register

Bit Location	Description	Register Attribute
0	ERR_COR Received – Set when a correctable error message is received and this bit is not already set. Default value of this field is 0.	RW1CS
1	Multiple ERR_COR Received – Set when a correctable error message is received and ERR_COR Received is already set. Default value of this field is 0.	RW1CS
2	ERR_FATAL/NOFATAL Received – Set when either a fatal or a non-fatal error message is received and this bit is not already set. Default value of this field is 0.	RW1CS
3	Multiple ERR_FATAL/NOFATAL Received – Set when either a fatal or a non-fatal error is received and ERR_FATAL/NOFATAL Received is already set. Default value of this field is 0.	RW1CS
4	First Uncorrectable Fatal – Set to 1b when the first Uncorrectable error message received is for a fatal error. Default value of this field is 0.	RW1CS
5	Non-Fatal Error Messages Received – Set to 1b when one or more Non-Fatal Uncorrectable error messages have been received. Default value of this field is 0.	RW1CS
6	Fatal Error Messages Received – Set to 1b when one or more Fatal Uncorrectable error messages have been received. Default value of this field is 0.	RW1CS

Bit Location	Description	Register Attribute
31:27	<p>Advanced Error Interrupt Message Number – If this function is allocated more than one MSI interrupt number, this register is required to contain the offset between the base Message Data and the MSI Message that is generated when any of the status bits of this capability are set.</p> <p>Hardware is required to update this field so that it is correct if the number of MSI Messages assigned to the device changes.</p>	RO

C12. Slot Power Limit Message also uses MsgD

Release Date: 2/11/03

Intro text in Section 2.2.8, Message Request Rules, is incorrect in that it states that the Vendor Defined messages are the only ones that use MsgD. This is incorrect – the Set_Slot_Power_Limit Message also uses MsgD. In addition, there is a typo.

Section 2.2.8, p. 65, edit as shown:

- ☐ All Message Requests use the Msg Type field encoding, except for the Vendor_Defined messages, which can use either Msg or MsgD, and the Set_Slot_Power_Limit message, which uses MsgD.

C13. Clarify Error Tables

Release Date: 2/11/03

The tables for Physical, Data Link and Transaction Layer errors are in the same format, and for all the detecting agent action is given as if all enable bits are set to “enable” and, for Advanced Error Handling, mask bits are disabled and severity bits are set to their default values. The footnote refs and table comments are potentially confusing, however, and should be improved.

Section 6.2.6, p.270, edit as shown:

Table 6-2: Physical Layer Error List

Error Name	Default Severity	Detecting Agent Action ²
------------	------------------	-------------------------------------

² For these tables, detecting agent action is given as if all enable bits are set to “enable” and, for Advanced Error Handling, mask bits are disabled and severity bits are set to their default values. Actions must be modified according to the actual settings of these bits.

Error Name	Default Severity	Detecting Agent Action ²
Receiver Error	Correctable	<i>Receiver (if checking):</i> Send ERR_COR to Root Complex unless masked.
Training Error	Uncorrectable (Fatal)	If checking, send ERR_FATAL/ ERR_NONFATAL to Root Complex ³ unless masked

Table 6-3: Data Link Layer Error List

Error Name	Severity	Detecting Agent Action [Add ref to existing footnote: "For these..."]
Bad TLP	Correctable	<i>Receiver:</i> Send ERR_COR to Root Complex.
Bad DLLP		<i>Receiver:</i> Send ERR_COR to Root Complex.
Replay Timeout		<i>Transmitter:</i> Send ERR_COR to Root Complex.
REPLAY NUM Rollover		<i>Transmitter:</i> Send ERR_COR to Root Complex.
Data Link Layer Protocol Error	Uncorrectable (Fatal)	If checking, send ERR_FATAL to Root Complex.

Table 6-4: Transaction Layer Error List

Error Name	Severity	Detecting Agent Action [Add ref to existing footnote: "For these..."]
Poisoned TLP Received	Uncorrectable (Non-Fatal)	<i>Receiver (if data poisoning is supported):</i> Send ERR_NONFATAL to Root Complex. Log the header of the poisoned TLP. ⁴
ECRC Check Failed		<i>Receiver:</i> Send ERR_NONFATAL to Root Complex. Log the header of the TLP that encounter the ECRC error.
Unsupported Request (UR)		<i>Request Receiver:</i> Send ERR_NONFATAL to Root Complex. Log the header of the TLP that caused the error.

³ Only the component closer to the Root Complex is typically capable of sending the error Message.⁴ Advanced Error Handling only.

Error Name	Severity	Detecting Agent Action [Add ref to existing footnote: “For these...”]
Completion Timeout		<i>Requester:</i> Send ERR_NONFATAL/ ERR_FATAL to Root Complex.
Completer Abort		<i>Completer:</i> Send ERR_NONFATAL to Root Complex. Log the header of the Completion that encountered the error.
Unexpected Completion		<i>Receiver:</i> Send ERR_NONFATAL to Root Complex. Log the header of the Completion that encountered the error. Note that if the Unexpected Completion is a result of misrouting, the Completion Timeout mechanism will be triggered at the corresponding Requester.
Receiver Overflow	Uncorrectable (Fatal)	<i>Receiver (if checking):</i> Send ERR_FATAL to Root Complex.
Flow Control Protocol Error		<i>Receiver (if checking):</i> Send ERR_FATAL to Root Complex.
Malformed TLP		<i>Receiver:</i> Send ERR_FATAL to Root Complex. Log the header of the TLP that encountered the error.

C14. Byte Count and Lower Address for Non-Mem Completions

Release Date: 2/11/03

The requirements for the Byte Count and Lower Address fields are included in Section “2.3.1.1. Data Return for Read Requests” even though they apply to IO and Config Completions. The text should be edited as shown to relocate these rules to the correct section. Note that there is no actual change to the rules themselves.

Edit as shown:

2.2.9 Completion Rules

...

- Byte Count[11:0] – ~~For Memory Read Completions, the~~ remaining byte count for Request ~~(see Section <2.3.1.1>)~~

- ? The Byte Count value is specified as a binary number, with 0000 0000 0001b indicating 1 byte, 1111 1111 1111b indicating 4095 bytes, and 0000 0000 0000b indicating 4096 bytes
- ? For Memory Read Completions, Byte Count[11:0] is set according to the rules in Section <2.3.1.1>.
- ? For all other types of Completions, the Byte Count field must be 4
- ...
- Lower Address[6:0] – lower byte address for starting byte of Completion
 - ? For Memory Read Completions, the value in this field is the byte address for the first enabled byte of data returned with the Completion (~~discussed in more detail~~ see rules in Section <2.3.1.1>)
 - ? This field is set to all ‘0’s for all ~~other~~ types of Completions other than Memory Read Completions
- ...

2.3.1.1. Data Return for Read Requests

- ❑ For each Memory Read Completion, ...
 - ~~❑ For all other types of Completions, the Byte Count field must be 4.~~
- ...
- ❑ For all Memory Read Completions, ...
 - ~~❑ For types of Completions other than Memory Read Completions, the Lower Address field must be all ‘0’s~~
- ...

C15. Starting Point for REPLAY_TIMER

Release Date: 2/11/03

In Section 3.5.2.1, there is an inconsistency in the spec regarding the starting condition for REPLAY_TIMER.

Edit as shown (related section included for context):

- ❑ The following timer is used:
 - REPLAY_TIMER - Counts time since last Ack or Nak DLLP received
 - ? Started at the ~~start-last~~ Symbol of any TLP transmission or retransmission, if not already running
- ...

<Table 3-4> defines the timeout count values for the REPLAY_TIMER. The values are specified according to the largest TLP payload size and Link width.

The values are measured at the Port of the TLP Transmitter, from last Symbol of TLP to First Symbol of TLP retransmission. The values are calculated using the formula (note – this is simply three times the Ack Latency value – see Section <3.5.3.1>):

...

C16. Use of Cpl vs. CplD**Release Date: 2/11/03**

In Table 2-3, the description of Cpl incorrectly implies that unsuccessful IO and Configuration Read Completions use CplD.

In Table 2-3, “Fmt[1:0] and Type[4:0] Field Encodings”, edit as shown:

Cpl	00	0 1010	Completion without Data – used for I/O and Configuration Write Completions, and Memory Read Completions (IO, Configuration or Memory) with Completion Status other than Successful Completion
-----	----	--------	--

C17. Error in Table for Hot Plug Message**Release Date: 2/11/03**

In Table 2-19, the support requirements for the Attention_Button_Pressed Message are incorrect.

In Table 2-19, p.78, change table entry:

Attention_Button_Pressed	0100 1000	100	t	r	tr	r	BDF	This message is issued by a device in a slot that implements an Attention Button on the card to signal the Switch/Root Port to generate the Attention Button Pressed Event. The Switch Switch/Root Port terminates the message and sets the Attention Button Pressed register to 1b which may result in an interrupt being generated.
--------------------------	-----------	-----	---	---	----	---	-----	---

to this (RC & Sw as receiver must support; Ep, Sw & Br support as transmitter if needed):

Attention_Button_Pressed	0100 1000	100	All:		BDF	This message is issued by a device in a slot that implements an Attention Button on the card to signal the Switch/Root
			r	r		
			As Required:			
			t	t		

					Port to generate the Attention Button Pressed Event. The Switch Switch/Root Port terminates the message and sets the Attention Button Pressed register to 1b which may result in an interrupt being generated.
--	--	--	--	--	--

C18. Inconsistency in Header Log Register Diagram

Release Date: 2/11/03

In Figure 7-34, “Header Log Register”, the arrangement of the bytes of the header is at odds with the labels for each DW. It is proposed to resolve this in favor of the DW view, as this makes the header fields easier for software to process.

7.10.8. Header Log Register (Offset 1Ch)

The header log register captures the header for the TLP corresponding to a detected error; refer to Section 6.2 for further details. Section 6.2 also describes the conditions where the packet header is logged. This register is 16 bytes and adheres to the format of the headers defined throughout this specification. **The header is captured such that the fields of the header read by software in the same way the headers are presented in this document, when the register is read using DW accesses. Therefore, byte 0 of the header is located in byte 3 of the Header Log Register, byte 1 of the header is in byte 2 of the Header Log Register and so forth. For 12 byte headers, only bytes 0 through 11 of the Header Log Register are used, and values in bytes 12 through 15 are undefined.**

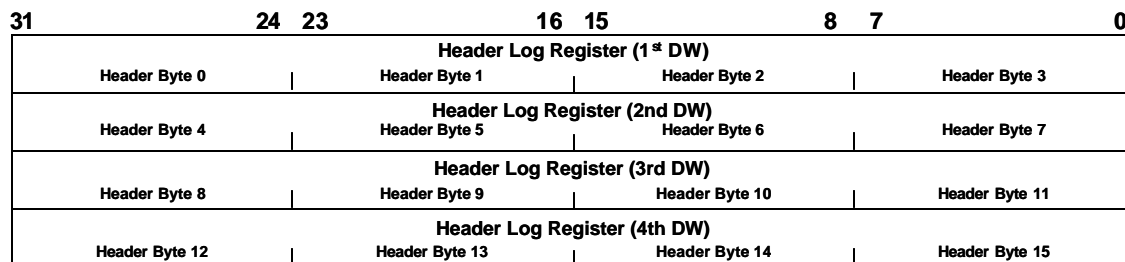


Figure 0-2: Header Log Register **<!note new fig>**

Table 0-2: Header Log Register

Bit	Description	Register	Default
-----	-------------	----------	---------

Location		Attribute	Value
127:0	Header of TLP associated with error	ROS	0

C19. Clarify use of “Bridge”

Release Date: 2/11/03

In the Message tables in Section 2.2.8.1, the term “Bridge” is used ambiguously. It should be clarified that the use of “Bridge” here refers only to the case where the primary side of the bridge connects to PCI Express and the secondary side connects to conventional PCI or PCI-X.

Section 2.8.1.1, p.68, edit footnote as shown:

Abbreviations:

RC = Root Complex

Sw = Switch (only used with “Link” routing)

Ep = Endpoint

Br = PCI Express (primary) to /PCI/PCI-X (secondary) Bridge

...

C20. Switch Unlocking Rule

Release Date: 2/11/03

In the case of a serious protocol error, there is an unintended potential deadlock exposure in the spec’s rules for locked sequences. The exposure results when a locked sequence that was established legally is disrupted by an unsuccessful completion to a subsequent read of the locked sequence. As written, a Switch would unlock as a result of this unsuccessful completion (intended as a safety mechanism to minimize the length of the lock & the chance of being stuck in it). However, once unlocked, a write could be sent up the locked path, potentially passing the unsuccessful Completion, and getting stuck at the RC (preventing the Completion from reaching the RC, and thereby preventing the RC from reacting to the unsuccessful Completion Status).

The proposed resolution is to remove the clause requiring the Switch to unlock due to the unsuccessful completion, and (for additional safety) to add text clarifying the behavior of lock sequences.

In Sections 6.5.2 and 6.5.3, Edit as shown:

6.5.2 Initiation and Propagation of Locked Transactions - Rules

Locked **transaction** sequences are generated by the Host CPU(s) as one or more reads followed by a ~~n-equal~~ number of writes to the same location(s). When a lock is established,

all other traffic is blocked from using the path between the Root Complex and the locked Legacy Endpoint or Bridge.

- ❑ ~~Lock~~-A locked transaction sequence or attempted locked transaction sequence is initiated on PCI Express using the “lock”-type Read Request/Completion (MRdLk/CplDLk) and terminated with the Unlock Message
 - Locked Requests which are completed with a status other than Successful Completion do not establish lock (explained in detail in the following sections)
 - Regardless of the status of any of the Completions associated with a locked sequence, all locked sequences and attempted locked sequences must be terminated by the transmission of an Unlock Message.
 - MRdLk, CplDLk and Unlock semantics are allowed only for the default Traffic Class (TC0)
 - Only one locked transaction sequence attempt may be in progress at a given time within a single hierarchy domain
- ❑ The Unlock Message is sent from the Root Complex down the locked transaction path to the Completer, and may be broadcast from the Root Complex to all Endpoints and Bridges
 - Any device which is not involved in the locked sequence must ignore this Message

The initiation and propagation of a locked transaction sequence through PCI Express is performed as follows:

- ❑ A locked transaction sequence is started with a MRdLk Request
 - Any successive reads for the locked transaction sequence must also use MRdLk Requests
 - The Completions for any successful MRdLk Request use the CplDLk Completion type, or the CplLk Completion type for unsuccessful Requests
- ❑ If any read associated with a locked sequence is completed unsuccessfully, the Requester must assume that the atomicity of the lock is no longer assured, and that the path between the Requester and Completer is no longer locked
- ❑ All writes for the locked sequence use MWr Requests
- ❑ The Unlock Message is used to indicate the end of a locked sequence
 - A Switch propagates Unlock Messages to the locked Egress Port
 - A PCI Express/PCI Bridge may propagate the semantic of Unlock by deasserting LOCK# on its PCI interface
- ❑ Upon receiving an Unlock Message, a Legacy Endpoint or Bridge must unlock itself if it is in a locked state
 - If not locked, or if the Receiver is a PCI Express Endpoint or Bridge which does not support lock, the Unlock Message is ignored and discarded

...

6.5.3 Switches and Lock - Rules

...

- ❑ When the CplDLk for the first MRdLk Request is returned, if the Completion indicates a Successful Completion status, the Switch must block all Requests from all other Ports from being propagated to either of the Ports involved in the locked access, except for Requests which map to non-VC0 on the Egress Port

~~• If at any time during a locked operation there is a Completion with completion status other than Successful Completion, the Switch must unblock the Ingress and Egress Ports~~

C21. Clarify Multi-Function Devices with Differing Error Severities

Release Date: 2/11/03

In a multi-function device, when an error occurs that affects multiple functions, only one error message is sent (Section 6.2.4). It is not stated, however, what action is taken when the functions implement Advanced Error Handling, and different functions map the error to different severity levels. The following clarification is proposed.

In Section 6.2.4, edit as shown:

6.2.4. Error Logging

...

In a multi-function device, ...

On the detection of one of these errors, a multi-function device should generate at most one error reporting message **of a given severity**, where the message must report the Requestor ID of a function of the device that is enabled to report that specific type of error. If no function is enabled to send a reporting message, then the device does not send a reporting message. **If all reporting-enabled functions have the same severity level set for the error, then only one error message is sent. If all reporting-enabled functions do not have the same severity level set for the error, then one error message for each severity level is sent.** Software is responsible for scanning all functions in a multi-function device when it detects one of those errors.

C22. Clarify BME Bit Bridge vs. Device Operation

Release Date: 2/11/03

The current text does not clearly distinguish between the operation of the BME bit in the case of a bridge (real or virtual) and the case of a device (Endpoint).

In Section 7.5.1.1. Table 7-3, edit as shown:

2	<p>Bus Master Enable – Controls the ability of a PCI Express agent-Endpoint to issue memory-Memory and I/O readRead/write-Write requestsRequests, and the ability of a Root or Switch Port to forward Memory and IO Read/Write Requests in the upstream direction.</p> <p><i>Endpoints:</i></p> <p>–Disabling this bit prevents a PCI Express agent from issuing any memory-Memory or I/O read/write requests. Note that as MSI interrupt messages are in-band memory writes, disabling the bus master enable bit disables MSI interrupt messages as well.</p> <p>Requests other than memory or I/O requests are not controlled by this bit.</p> <p>Default value of this field is 0b.</p> <p>This bit is hardwired to 0b if a device does not generate memory or I/O Requests.</p> <p><i>Root and Switch Ports:</i></p> <p>This bit also-controls forwarding of memory-Memory or I/O Requests by a switch-Switch or root-Root port-Port in the upstream directionfrom the secondary interface to the primary interface of the PCI-PCI bridge structure representing the respective switch or root port. When this bit is not set, Memory and I/O Requests received at a Root Port or the downstream side of a Switch Port must be handled as Unsupported Requests (UR), and for Non-Posted Requests a Completion with UR completion status must be returned. This bit does not affect forwarding of Completions from-in either the primary interface to the secondary interfaceupstream or downstream direction.</p> <p>The forwarding of Requests other than memory or I/O requests are not controlled by this bit.</p> <p>Default value of this field is 0b.</p>	RW
---	---	----

C23. Clarify Link Retrain and Link Training Bits

Release Date: 2/11/03

Clarify operation of Link Retrain and Link Training bits (+ fix minor typo).

7.8.7 Link Control Register (Offset 10h)

...

Table 0-3: Link Control Register

Bit Location	Register Description	Attributes
...		
4	<p>Link Disable – This bit disables the Link when set to 1b; this field is not applicable and reserved for endpoint devices and Upstream Ports of a-Switches.</p> <p>Writes to this bit are immediately reflected in the value read from the bit, regardless of actual Link state.</p> <p>Default value of this field is 0b.</p>	RW
5	<p>Retrain Link – This bit A write of 1b to this bit initiates Link retraining when set by directing the Physical Layer LTSSM to the Recovery state. The LTSSM must enter the Recovery state before the Completion is returned for the Write Request to this bit. Reads of this bit always return 0b.;</p> <p>†This field is not applicable and reserved for endpoint devices and Upstream Ports of a-Switches.</p> <p>This bit always returns 0b when read.</p>	RW
...		

7.8.8. Link Status Register (Offset 12h)

...

Table 0-4: Link Status Register

Bit Location	Register Description	Attributes
...		
11	<p>Link Training – This read-only bit indicates that Link training is in progress (Physical Layer LTSSM in Configuration or Recovery state); hardware clears this bit once Link training is complete.</p> <p>This field is not applicable and reserved for endpoint devices and Upstream Ports of Switches.</p>	RO
...		

C24. Clarify Extended Sync Bit

Release Date: 2/11/03

Clarify/correct description of Extended Sync bit.

7.8.7 Link Control Register (Offset 10h)

...

7	Extended Synch – This bit when set forces the transmission of 4096 FTS ordered sets in the L0s state followed by a single SKP Ordered Set (see Section <4.2.4.3>) prior to entering L0 state, and the transmission of 1024 TS1 ordered sets in the L1 state prior to entering the Recovery state. This mode provides external devices monitoring the link time to achieve bit and symbol lock before the link enters L0 or Recovery states and resumes communication. Default value for this bit is 0b	RW
---	--	----

C25. Changing Common Clock Config Bit

Release Date: 2/11/03

Clarify that the reported N_FTS can be changed at initialization/configuration time.

Note: See revised Ch4 text

4.2.4.3 Fast Training Sequence (FTS)

...

~~After initial power up~~ In the Polling, Configuration, and Recovery states, the N_FTS value is exchanged in the TS1/TS2 training sequence. “N_FTS defines the number of FTS ordered-set that must be transmitted when transitioning from L0s to L0. For the data rate in this specification, the value that can be requested by a component corresponds to a bit lock time of 16 ns to 4 μ s, except when the Extended Synch bit is set forcing the transmission of 4096 FTS ordered sets resulting in a bit lock time of 64 μ s. Note that the N_FTS value reported by a component may change, for example due to software modifying the value in the Common Clock Configuration bit (Section <7.8.7>).

...

Clarify need to trigger link retrain following a change to the setting of the Common Clock Configuration bit.

7.8.7 Link Control Register (Offset 10h)

...

6	<p>Common Clock Configuration – This bit when set indicates that this component and the component at the opposite end of this Link are operating with a distributed common reference clock.</p> <p>A value of 0 indicates that this component and the component at the opposite end of this Link are operating with asynchronous reference clock.</p> <p>Components utilize this common clock configuration information to report the correct L0s and L1 Exit Latencies.</p> <p>After changing the value in this bit in both components on a Link, software must trigger the Link to retrain by writing a 1b to the Retrain Link bit.</p> <p>Default value of this field is 0.</p>	RW
---	---	----

C26. Clarify Length as Reserved for Cpl

Release Date: 2/11/03

Clarify that the Length field is reserved for Cpl type TLPs

2.2.1. Common Packet Header Fields

...

- ☐ Length[9:0] – Length of data payload in DW (see <>) – bits 1:0 of Byte 2 concatenated with bits 7:0 of Byte 3
 - TLP data must be four-byte naturally aligned and in increments of four-Byte Double Words (DW).
 - Reserved for TLPs that do not contain or refer to data payloads, including Cpl, CplLk, and Messages (except as specified)

C27. Reset Conditions for Replay_Timer

Release Date: 2/11/03

Clarify reset conditions for REPLAY_TIMER

3.5.2.1. LCRC and Sequence Number Rules (TLP Transmitter)

...

- ☐ The following timer is used:
 - REPLAY_TIMER - Counts time since last Ack or Nak DLLP received
 - ? Started at the ~~start~~-last Symbol[**This change in C15] of any TLP transmission or retransmission, if not already running

- ? For each replay, reset and restart **REPLAY_TIMER** when sending the last Symbol of the first TLP to be retransmitted
- ? Restarts for each Ack/~~Nak~~ DLLP received while there are unacknowledged TLPs outstanding, if, and only if, the received Ack ~~or Nak~~-DLLP acknowledges some TLP in the retry buffer
 - | Note: This ensures that **REPLAY_TIMER** is reset only when forward progress is being made
- ? Reset and hold until restart conditions are met for each Nak received or when the **REPLAY_TIMER** expires
- ? Resets and holds when there are no outstanding unacknowledged TLPs

...

When a replay is initiated, either due to reception of a Nak or due to **REPLAY_TIMER** expiration, the following rules describe the sequence of operations that must be followed:

- ☐ If all TLPs transmitted have been acknowledged (the Retry Buffer is empty), terminate replay, otherwise continue.
- ☐ Increment **REPLAY_NUM**.
 - If **REPLAY_NUM** rolls over from 11b to 00b, the Transmitter signals the Physical Layer to retrain the Link. This is a reported error associated with the Port (see Section <>).

Note that Data Link Layer state, including the contents of the Retry Buffer, are not reset by this action unless the Physical Layer reports Physical LinkUp = 0 (causing the Data Link Control and Management State Machine to transition to the DL_Inactive state).
- If **REPLAY_NUM** does not roll over from 11b to 00b, continue.
- ☐ Block acceptance of new TLPs from the Transmit Transaction Layer.
- ☐ Complete transmission of any TLP currently being transmitted.
- ☐ Retransmit unacknowledged TLPs, starting with the oldest unacknowledged TLP and continuing in original transmission order
 - Reset and restart **REPLAY_TIMER** when sending the last Symbol of the first TLP to be retransmitted
 - Once all unacknowledged TLPs have been re-transmitted, return to normal operation.
 - If any Ack or Nak DLLPs are received during a replay, the transmitter is permitted to complete the replay without regard to the Ack or Nak DLLP(s), or to skip retransmission of any newly acknowledged TLPs.
 - ? Once the transmitter has started to resend a TLP, it must complete transmission of that TLP in all cases.
 - Ack and Nak DLLPs received during a replay must be processed, and may be collapsed

- ? Example: If multiple Acks are received, only the one specifying the latest Sequence Number value must be considered – Acks specifying earlier Sequence Number values are effectively “collapsed” into this one
- ? Example: During a replay, Nak is received, followed by an Ack specifying a later Sequence Number – the Ack supercedes the Nak, and the Nak is ignored
- Note: Since all entries in the Retry Buffer have already been allocated space in the Receiver by the Transmitter’s Flow Control gating logic, no further flow control synchronization is necessary.
- ❑ Re-enable acceptance of new TLPs from the Transmit Transaction Layer.
- ...
- ❑ For Ack and Nak DLLPs, the following steps are followed (see <>):
 - If the ~~Sequence Number specified by the~~ AckNak_Seq_Num does not ~~specify the Sequence Number of~~ correspond either to an unacknowledged TLP, or to of the most recently acknowledged TLP, the DLLP is discarded
 - ? If the DLLP is an Ack DLLP, this is a DL Layer Protocol Error which is a reported error associated with the Port (see Section <>).
 - If the AckNak_Seq_Num does not specify the Sequence Number of the most recently acknowledged TLP, then the DLLP acknowledges some TLPs in the retry buffer:
 - ? Purge from the retry buffer all TLPs from the oldest to the one corresponding to the AckNak_Seq_Num
 - ? Load ACKD_SEQ with the value in the AckNak_Seq_Num field
 - ? Reset REPLAY_NUM ~~and REPLAY_TIMER~~

C28. FC Updates for Infinite FC Advertisements

Release Date: 2/11/03

Intent is that FC Updates are permitted but not required following an initial infinite advertisement. Clarify this rule and related checking.

2.6.1. Flow Control Rules

...

- ❑ If an Infinite Credit advertisement (value of 00h or 000h) has been made during initialization, no Flow Control updates are required following initialization.
 - If UpdateFC DLLPs are sent, the credit value fields must be set to zero and must be ignored by the receiver. The receiver may optionally check for non-zero update values (in violation of this rule). If a component implementing this check determines a violation of this rule, the violation is a Flow Control Protocol Error (FCPE)

- ? If checked, this is a reported error associated with the Receiving Port (see Section <6.2>)
- ❑ If only the Data or Header advertisement (but not both) for a given type (N, NP, or CPL) has been made with infinite credits during initialization, then the transmission of UpdateFC DLLPs is still required, but the credit field corresponding to the Data/Header (advertised as infinite) must be set to zero and must be ignored by the receiver.
- The receiver may optionally check for non-zero update values (in violation of this rule). ~~Components may optionally check for violations of this rule.~~ If a ~~component~~ receiver implementing this check determines a violation of this rule, the violation is a Flow Control Protocol Error (FCPE)
- ? If checked, this is a reported error associated with the Receiving Port (see Section <6.2>)

C29. Size of RCRB Access Operations

Release Date: 2/11/03

RCRBs behave similarly to configuration space. Update text for Root Complex Register Block to match configuration space clarification C10.

Section 7.2.3, p.314, add as shown:

A root port may be associated with an optional 4096 byte block of memory mapped registers referred to as the Root Complex Register Block (RCRB). These registers are used in a manner similar to configuration space and can include PCI Express extended capabilities and other implementation specific registers that apply to the root complex. The structure of the RCRB is described in Section 7.9.2.

System firmware communicates the base address of the RCRB for each Root Port or internal device in the Root Complex to the operating system. Multiple Root Ports or internal devices may be associated with the same RCRB. The RCRB memory-mapped registers must not reside in the same address space as the memory-mapped configuration space.

A Root Complex implementation is not required to support memory space requests to a Root Complex Register Block that cross DWORD aligned boundaries or that use locked semantics.



IMPLEMENTATION NOTE

Accessing Root Complex Register Block

Because Root Complex implementations are not required to support memory space requests to a Root Complex Register Block that cross DWORD boundaries, or that use locked

semantics, software should take care not to cause the generation of such requests when accessing a Root Complex Register Block unless it is known that the Root Complex implementation being used will support the request.

C30. Clarify Requirements for Sticky Registers

Release Date: 2/11/03

The usage models for the PM sticky bits and the Advanced Error Reporting sticky bits are different, and place different requirements on the bits themselves, but this is not reflected in the text.

5.5.1. Auxiliary Power Enabling

...
The Aux Power PM Enable bit is sticky (see Section <7.4>) so ~~meaning that~~ its state is preserved in the D3cold state, and is not affected by the transitions from the D3cold state to the D0Uninitilaized state.

7.4 Configuration Register Types

...

Table 7-2: Register (and Register Bit-Field) Types

Register Attribute	Description
...	
ROS	<p>Sticky - Read-only register: Registers are read-only and cannot be altered by software. Registers are not initialized or modified by hot reset.</p> <p>Where noted, Devices that consume AUX power must preserve sticky register values when AUX power consumption (either via AUX power or PME Enable) is enabled. In these cases, registers are not initialized or modified by hot, warm or cold reset (see Section <6.6>).</p>
RWS	<p>Sticky - Read-Write register: Registers are read-write and may be either set or cleared by software to the desired state. Bits are not initialized or modified by hot reset.</p> <p>Where noted, Devices that consume AUX power must preserve sticky register values when AUX power consumption (either via AUX power or PME Enable) is enabled. In these cases, registers are not initialized or modified by hot, warm or cold reset (see Section <6.6>).</p>

Register Attribute	Description
RW1CS	Sticky - Read-only status, Write-1-to-clear status register: Registers indicate status when read, a set bit indicating a status event may be cleared by writing a 1. Writing a 0 to RW1CS bits has no effect. Bits are not initialized or modified by hot reset. Where noted, D devices that consume AUX power must preserve sticky register values when AUX power consumption (either via AUX power or PME Enable) is enabled. In these cases, registers are not initialized or modified by hot, warm or cold reset (see Section <6.6>).
...	

...

7.6 PCI Power Management Capability Structure

...

Table 7-8: Power Management Status/Control Register

Bit Location	Register Description	Attributes
1:0	Power State	RW
8	PME Enable Note: Devices that consume AUX power must preserve the value of this sticky register when AUX power is available. In such devices this register value is not modified by hot, warm or cold reset.	RWS
...		
15	PME Status Note: Devices that consume AUX power must preserve the value of this sticky register when AUX power is available. In such devices this register value is not modified by hot, warm or cold reset.	RW1CS
...		

...

7.8.4 Device Control Register (Offset 08h)

...

Table 7-12: Device Control Register

Bit Location	Register Description	Attributes
...		
10	<p>Auxiliary (AUX) Power PM Enable – This bit when set enables a device to draw AUX power independent of PME AUX power. Devices that require AUX power on legacy operating systems should continue to indicate PME AUX power requirements. AUX power is allocated as requested in the AUX_Current field of the Power Management Capabilities Register (PMC), independent of the PME_En bit in the Power Management Control/Status Register (PMCSR) (see Chapter <5>). For multi-function devices, a component is allowed to draw AUX power if at least one of the functions has this bit set.</p> <p>Note: Devices that consume AUX power must preserve the value of this sticky register when AUX power is available. In such devices this register value is not modified by hot, warm or cold reset. Default value of this field is 0.</p> <p>Devices that do not implement this capability hardwire this bit to 0.</p>	RWS
...		

C31. Replay Timing L0s Correction Factor

Release Date: 2/11/03

The timing equations and tables in the spec do not account for delays cause by L0s. These delays can make satisfying the specified timing values impossible, and therefore create a contradiction. The following text shows changes that resolve this contradiction.

In 3.5.2.1:

Table 3-4 The following formula defines the timeout count values for the REPLAY_TIMER. The values are specified according to the largest TLP payload size and Link width.

The values are measured at the Port of the TLP Transmitter, from last Symbol of TLP to First Symbol of TLP retransmission. The values are calculated using the formula (note – this is simply three times the Ack Latency value – see Section 3.5.3.1):

$$\left(\frac{(Max_Payload_Size + TLPOverhead) * AckFactor}{LinkWidth} + InternalDelay \right) * 3 + Rx_L0s_Adjustment$$
~~$$\left(\frac{(Max_Payload_Size + TLPOverhead) * AckFactor}{LinkWidth} + InternalDelay \right) * 3$$~~

where

Max_Payload_Size	is the value in the Max_Payload_Size field of the Device Control Register
TLP Overhead	represents the additional TLP components which consume Link bandwidth (Header, LCRC, framing Symbols) and is treated here as a constant value of 28 Symbols
AckFactor	represents the number of maximum size TLPs which that can be received before an Ack is sent, and is used to balance Link bandwidth efficiency and retry buffer size – the value varies according to Max_Payload_Size and Link width, and is included in Table 3-5
LinkWidth	is the operating width of the Link
InternalDelay	represents the internal processing delays for received TLPs and transmitted DLLPs, and is treated here as a constant value of 19 Symbol Times

Rx_L0s_Adjustment equals the time required by the component's receive circuits to exit from L0s to L0 (as to receive an Ack DLLP from the other component on the Link) expressed in Symbol Times (see Section <ref 4.2.6.6.1>)

The values in Table <ref 3-7> do not include this adjustment offset.

Thus, the timeout value for REPLAY_TIMER is the value given in Table <ref 3-7> plus the receiver L0s adjustment offset, described above.

Table 3-4: Unadjusted[ADD FOOTNOTE] REPLAY_TIMER Limits by Link Width and Max_Payload_Size (Symbol Times) Tolerance: -0%/+100%

		Link Operating Width						
		x1	x2	x4	x8	x12	x16	x32
Max_Payload_Size	128B	711	384	219	201	174	144	99
	256B	1248	651	354	321	270	216	135
	512B	1677	867	462	258	327	258	156
	1024B	3213	1635	846	450	582	450	252
	2048B	6285	3171	1614	834	1095	834	444
	4096B	12429	6243	3150	1602	2118	1602	828

FOOTNOTE: The values in this table are determined using the formula shown above minus the "Rx_L0s_Adjustment" term

In 3.5.3.1:

<> defines the threshold values for the AckNak_LATENCY_TIMER timer, which for any specific case is called the Ack Latency. The values are specified according to the largest TLP payload size and Link width. The values are measured at the Port of the TLP Receiver, starting with the time the last Symbol of a TLP is received to the first Symbol of the Ack/Nak DLLP being transmitted. The values are calculated using the formula:

$$\frac{(Max_Payload_Size + TLPOverhead) * AckFactor}{LinkWidth} + InternalDelay + Tx_L0s_Adjustment$$

$$\frac{(Max_Payload_Size + TLPOverhead) * AckFactor}{LinkWidth} - InternalDelay$$

where

Max_Payload_Size	is the value in the Max_Payload_Size field of the Device Control Register
TLP Overhead	represents the additional TLP components which consume Link bandwidth (Header, LCRC, framing Symbols) and is treated here as a constant value of 28 Symbols
AckFactor	represents the number of maximum size TLPs which can be received before an Ack is sent, and is used to balance Link bandwidth efficiency and retry buffer size – the value varies according to Max_Payload_Size and Link width, and is defined in Error! Reference source not found.
LinkWidth	is the operating width of the Link
InternalDelay	represents the internal processing delays for received TLPs and transmitted DLLPs, and is treated here as a constant value of 19 Symbol Times
Tx_L0s_Adjustment	if L0s is enabled, the time required for the transmitter to exit L0s (see Section <ref 4.2.6.6.2>), expressed in Symbol Times, or 0 if L0s is not enabled The values in Table <ref 3-9> do not include this adjustment offset.

Thus, the Ack Latency is the value given in Table <ref 3-9> plus the transmitter L0s adjustment offset, described above.

Table 3-5: Unadjusted[ADD FOOTNOTE] Ack Transmission Latency Limit and AckFactor by Link Width and Max Payload (Symbol Times)

		Link Operating Width						
		x1	x2	x4	x8	x12	x16	x32
Max_Payload_Size	128B	237 AF = 1.4	128 AF = 1.4	73 AF = 1.4	67 AF = 2.5	58 AF = 3.0	48 AF = 3.0	33 AF = 3.0
	256B	416 AF = 1.4	217 AF = 1.4	118 AF = 1.4	107 AF = 2.5	90 AF = 3.0	72 AF = 3.0	45 AF = 3.0
	512B	559 AF = 1.0	289 AF = 1.0	154 AF = 1.0	86 AF = 1.0	109 AF = 2.0	86 AF = 2.0	52 AF = 2.0

	1024B	1071 AF = 1.0	545 AF = 1.0	282 AF = 1.0	150 AF = 1.0	194 AF = 2.0	150 AF = 2.0	84 AF = 2.0
	2048B	2095 AF = 1.0	1057 AF = 1.0	538 AF = 1.0	278 AF = 1.0	365 AF = 2.0	278 AF = 2.0	148 AF = 2.0
	4096B	4143 AF = 1.0	2081 AF = 1.0	1050 AF = 1.0	534 AF = 1.0	706 AF = 2.0	534 AF = 2.0	276 AF = 2.0

FOOTNOTE: The values in this table are determined using the formula shown above minus the “Tx_L0s_Adjustment” term



IMPLEMENTATION NOTE

Retry Buffer Sizing

The Retry Buffer should be large enough to ensure that under normal operating conditions, transmission is never throttled because the retry buffer is full. In determining the optimal buffer size, one must consider the Ack Latency value ([Table 3-5](#)), any differences between the actual implementation and the internal processing delay used to generate these values, and the delays caused by the physical Link interconnect.

The receiver L0s exit latency (see Section 4.6.2.2.1) should also be accounted for, as is demonstrated with the following example using components A and B:

- φ A exits L0s on its transmit path to B and starts transmitting a long burst of write Requests to B
- φ B initiates L0s exit on its transmit path to A, but the L0s exit time required by A's receiver is large
- φ Meanwhile, B is unable to send Ack DLLPs to A, and A stalls due to lack of Retry Buffer space
- φ The transmit path from B to A returns to L0, B transmits an Ack DLLP to A, and the stall is resolved

This stall can be avoided by matching the size of a component's Retry Buffer to the L0s exit latency of the components receiver, or, conversely, matching the receiver L0s exit latency to the desired size of the Retry Buffer.

AckFactor values were chosen to allow implementations to achieve good performance without requiring an uneconomically large retry buffer. To enable consistent performance across a general purpose interconnect with differing implementations and applications, it is necessary to set the same requirements for all components without regard to the application space of any specific component. If a component does not require the full transmission bandwidth of the Link, it may reduce the size of its retry buffer below the minimum size required to maintain available retry buffer space with the Ack Latency values specified.

Note that the Ack Latency values specified ensure that the range of permitted outstanding Sequence Numbers will never be the limiting factor causing transmission stalls.

C32. Wrong Units in Figure 5-7

Release Date: 2/11/03

An apparent font bug changed the units shown in Figure 5-7 to be milliseconds instead of microseconds. The corresponding text is correct.

Fix figure as shown (note that original is correct – problem is in conversion):




Figure 5-7: Example of L1 Exit Latency Computation

C33. Ack/Nak Rule Clarification

Release Date: 2/11/03

The rules for checking the validity of a received Ack or Nak DLLP are unclear and potentially conflicting.

In Section 3.5.2.1, edit as shown:

- ❑ For Ack and Nak DLLPs, the following steps are followed (see <>):
 - If the AckNak_Seq_Num does not specify the Sequence Number of an unacknowledged TLP, ~~or of the most recently acknowledged TLP~~ and is not equal to the value in ACKD_SEQ, the DLLP is discarded
 - ? ~~If the DLLP is an Ack DLLP, this~~ This is a DL Layer Protocol Error which is a reported error associated with the Port (see Section <>)
- Note that it is not an error to receive an Ack DLLP when there are no outstanding unacknowledged TLPs, including the time between reset and the

first TLP transmission, so long as the specified Sequence Number matches the value in ACKD_SEQ.

...

edit figure as shown:

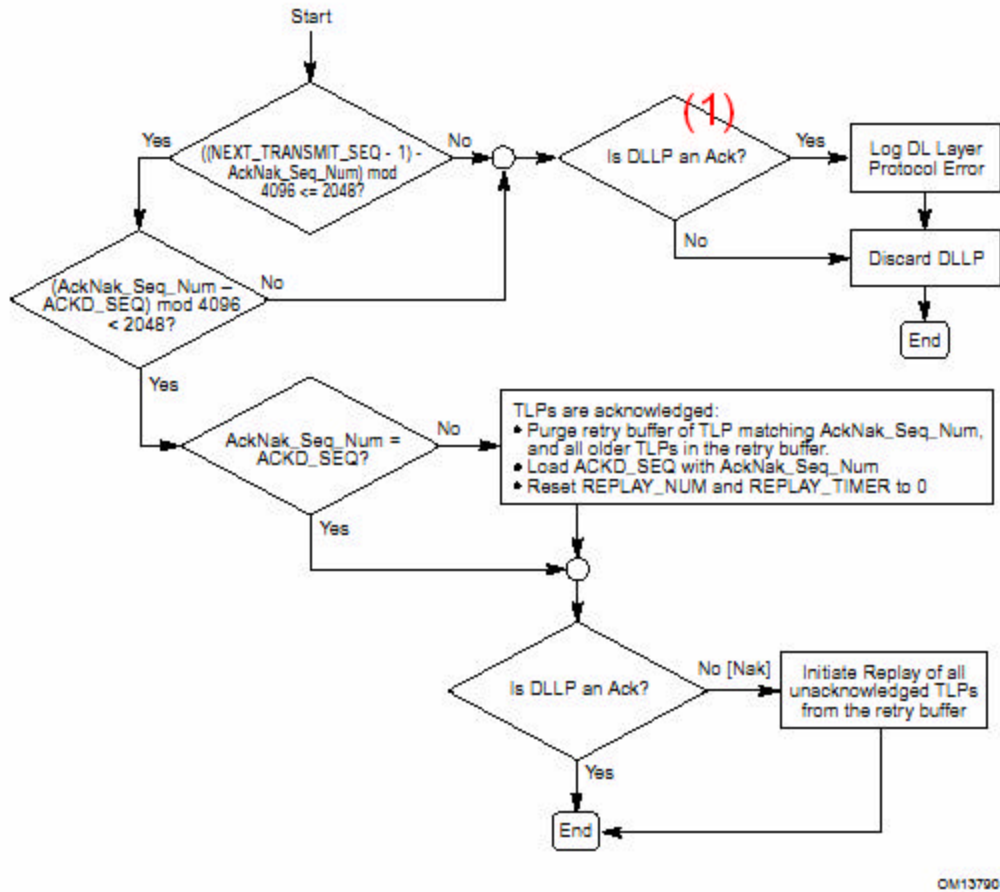


Figure 3-17: Ack/Nak DLLP Processing Flowchart **!!MODIFY DIAGRAM – ELIM (1), GO TO “Log DL ...”**

C34. Error Logging Clarification

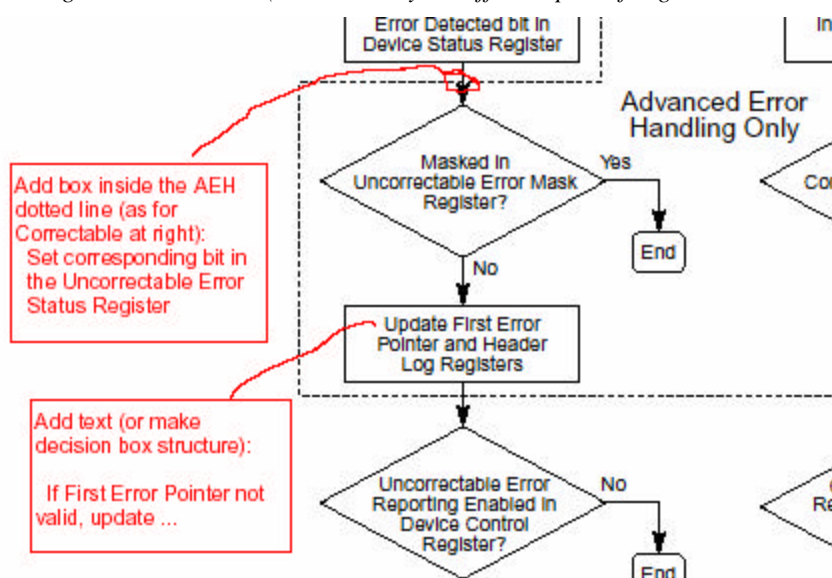
Release Date: 2/11/03

Figure 6-2 is missing a step described in the text. Additionally, one step is not completely described in the figure; the figure and text should agree.

In 6.2.4.2, edit as shown:

... The First Error Pointer register is valid when ~~remains valid until~~ the corresponding bit of the Uncorrectable Error Status register is set. The First Error Pointer value is not meaningful when the corresponding bit of the Uncorrectable Error Status register is not set, or is an unimplemented or undefined bit. ~~reset by software. This will cause the value in the First Error Pointer register to become undefined until a subsequent unmasked uncorrectable error is detected.~~ ...

Edit Figure 6-2 as shown (note that only the affected part of Figure 6-2 is shown here):



C35. Poisoned Write Clarification

Release Date: 2/11/03

Handling of poisoned writes to control spaces needs clarification.

In 2.7.2.2, edit as shown:

...

- ❑ A poisoned I/O or Memory Write Request, or a Message with data (except for Vendor_Defined Messages), that ~~which~~ addresses ~~an I/O or Memory mapped~~ a control space-register or control structure in the Completer must be ~~discarded~~ handled as an Unsupported Request (UR) by the Completer. ~~For such I/O requests, a Completion with a Completion Status of UR is returned~~ (see Section <2.2.9>).
- A Switch must route a poisoned I/O or Memory Write Request or Message with data in the same way it would route a ~~non-poisoned I/O or Memory Write Request~~ the same Request if it were not poisoned, unless the Request addresses a control register or control structure space of the Switch itself, in which case the Switch is the Completer for the Request and must follow the above rule.

For some applications it may be desirable for the Completer to use poisoned data in Write Requests which ~~do not target control registers or control structures non-control spaces~~

such use is not forbidden. Similarly, it may be desirable for the Requester to use data marked poisoned in Completions - such use is also not forbidden. The appropriate use of poisoned information is application specific, and is not discussed in this document.

...

C36. Remove Bridge Spec Material

Release Date: 2/11/03

The following three changes correct text that is incorrect or incomplete in the base spec but will be covered fully and correctly in the bridge spec.

In 2.4, p.99, edit as shown:

...

- ❑ PCI Express Switches are permitted to allow a Memory Write or Message Request with the Relaxed Ordering bit set to pass any previously posted Memory Write or Message Request moving in the same direction. Switches must forward the Relaxed Ordering attribute unmodified. The Root Complex is also permitted to allow data bytes within the Request to be written to system memory in any order. (The bytes must be written to the correct system memory locations. Only the order in which they are written is unspecified). ~~PCI Express PCI-X Bridge devices must forward the Relaxed Ordering attribute unmodified but must treat all transactions as if the Relaxed Ordering attribute bit is not set.~~

~~Note: This maintains compatibility with PCI-X relaxed ordering usage models and corresponding rules. For more details, refer to the PCI-X Addendum to the PCI Local Bus Specification, Rev 2.0.~~

In Table 2-27, p.109, edit as shown (note footnote added):

Table 2-27: Minimum Initial Flow Control Advertisements [footnote ref]

Credit Type	Minimum Advertisement
PH	1 unit – credit value of 01h
PD	Largest possible setting of the Max_Payload_Size for the component divided by FC Unit Size. Example: If the largest Max_Payload_Size value supported is 1024B, the smallest permitted initial credit value would be 040h.
NPH	1 unit – credit value of 01h

Credit Type	Minimum Advertisement
NPD	1 unit – credit value of 01h
CPLH	Switch and PCI Express to PCI-X Bridge (PCI-X mode only) : 1 FC unit – credit value of 01h Root Complex and , Endpoint, and PCI Express to PCI Bridge : “infinite” FC units – initial credit value of all ‘0’s ⁵
CPLD	Switch and PCI Express to PCI-X Bridge (PCI-X mode only) : Largest possible setting of the Max_Payload_Size for the component divided by FC Unit Size, or the size of the largest Read Request the component will ever generate, whichever is smaller. Root Complex and, Endpoint, and PCI Express to PCI Bridge : “infinite” FC units – initial credit value of all ‘0’s

[footnote] Note that PCI Express to PCI/PCI-X Bridges requirements are addressed in the <PCI Express Bridge Specification>.

Edit Section 6.5.4 as shown:

6.5.4 PCI Express/PCI Bridges and Lock - Rules

The requirements for PCI Express/PCI Bridges are similar to those for Switches, except that, because PCI Express/PCI Bridges use only the default Virtual Channel and Traffic Class, all other traffic is blocked during the locked access. The requirements on the PCI bus side of the PCI Express/PCI Bridge match the requirements for a PCI/PCI Bridge (see *PCI-to-PCI Bridge Architecture Specification, Rev. 1.1* ~~and <PCI Express Bridge Specification>~~).

- ~~□ When the Locked Completion for the first Locked Read Request is returned, the Bridge must block all Requests not associated with the locked access from flowing through the Bridge in either direction~~
- ~~□ The Bridge must remain blocked as described above until the Bridge is unlocked from the side which initiated the locked access~~
 - ~~• The Bridge unlocks itself and propagates the unlock to the other side of the Bridge~~

The following changes improve synergy with the bridge spec.

Section 2.2.8.1.5, p. 74, edit as shown:

The component on the other side of the Link (Endpoint, Switch, or ~~PCI Express-PCI~~ Bridge) that receives Set_Slot_Power_Limit message must copy ...

Section 6.5.2, p. 287, edit as shown:

⁵ This value is interpreted as infinite by the Transmitter, which will, therefore, never throttle.

- ❑ The Unlock Message is used to indicate the end of a locked sequence
 - A Switch propagates Unlock Messages to the locked Egress Port
 - ~~A PCI Express/PCI Bridge may propagate Unlock by deasserting LOCK# on its PCI interface~~

C37. Remove Unnecessary Root Complex Coherence Reference

Release Date: 2/11/03

This change was requested by Sun. The following justification was provided:

The requirement seems to refer to the coherency of the memory attached to the Root Complex in this architecture. The sentence should be removed; there are many other possible system architectures and coherency schemes for which this sentence does not apply.

Section 2.4, p. 100, delete as shown:

~~❑ Main memory writes from the CPU accepted by the Root Complex are architecturally part of the system memory image; the Root Complex must ensure coherency for subsequent device reads from main memory.~~

C38. Read/Write Fragmentation

Release Date: 2/11/03

Restructure Section 2.4 and add text as shown:

2.4 Transaction Ordering

2.4.1 Transaction Ordering Rules

Table 2-23 defines the ordering ...

...

2nd Completion returned: Data from 1000h to 107Fh.

2.4.2 Update Ordering and Granularity Observed by a Read Transaction

If a Requester using a single transaction reads a block of data from a Completer, and the Completer's data buffer is concurrently being updated, the ordering of multiple updates and granularity of each update reflected in the data returned by the read is outside the scope of

this specification. This applies both to updates performed by PCI Express write transactions and updates performed by other mechanisms such as host CPUs updating host memory.

As an example of update ordering, assume that the block of data is in host memory, and a host CPU writes first to location A and then to a different location B. A Requester reading that data block with a single read transaction is not guaranteed to observe those updates in order. In other words, the Requester may observe an updated value in location B and an old value in location A, regardless of the placement of locations A and B within the data block. Unless a Completer makes its own guarantees (outside the PCI Express specification) with respect to update ordering, a Requester that relies on update ordering must observe the update to location B via one read transaction before initiating a subsequent read to location A to return its updated value.

As an example of update granularity, if a host CPU writes a QWORD to host memory, a Requester reading that QWORD from host memory may observe a portion of the QWORD updated and another portion of it containing the old value.

While not required by this specification, it is strongly recommended that host platforms guarantee that when a host CPU writes aligned DWORDs or aligned QWORDS to host memory, the update granularity observed by a PCI Express read will not be smaller than a DWORD.



IMPLEMENTATION NOTE

No Ordering Required Between Cachelines

A Root Complex serving as a Completer to a single Memory Read that requests multiple cachelines from host memory is permitted to fetch multiple cachelines concurrently, to help facilitate multi-cacheline completions, subject to `Max_Payload_Size`. No ordering relationship between these cacheline fetches is required.

2.4.2 Update Ordering and Granularity Provided by a Write Transaction

If a single write transaction containing multiple DWORDs and the Relaxed Ordering bit clear is accepted by a Completer, the observed ordering of the updates to locations within the Completer's data buffer must be in increasing address order. This semantic is required in case a PCI or PCI-X bridge along the path combines multiple write transactions into the single one. However, the observed granularity of the updates to the Completer's data buffer is outside the scope of this specification.

While not required by this specification, it is strongly recommended that host platforms guarantee that when a PCI Express write updates host memory, the update granularity observed by a host CPU will not be smaller than a DWORD.

As an example of update ordering and granularity, if a Requester writes a QWORD to host memory, in some cases a host CPU reading that QWORD from host memory could observe the first DWORD updated and the second DWORD containing the old value.

C39. Clarify Device/Bridge Behavior for Config Bits**Release Date: 2/11/03**

Several config bit descriptions need improvement to better distinguish between the behavior for an Endpoint vs. for a Switch or Root Complex.

In Table 7-3 (Section 7.5.1.1), edit as shown:

8	SERR Enable – See Section <>. <p>This bit, when set, enables reporting of non-fatal and fatal errors detected by the device to the Root Complex. Note that errors are reported if enabled either through this bit or through the PCI-Express specific bits in the Device Control Register (see Section <>).</p> <p>Default value of this field is 0.</p>	RW
...		
10	Interrupt Disable - Controls the ability of a PCI Express device to generate INTx interrupt messages. When set, devices are prevented from generating INTx interrupt messages. <p>Any INTx emulation interrupts already asserted by the device must be deasserted when this bit is set.</p> <p>Note that INTx emulation interrupts forwarded by Root and Switch Ports from devices downstream of the Root or Switch Port are not affected by this bit.</p> <p>Default value of this field is 0.</p>	RW

In Table 7-4 (Section 7.5.1.2), edit as shown:

3	Interrupt Status – Indicates that an INTx interrupt message is pending internally to the device. <p>Note that INTx emulation interrupts forwarded by Root and Switch Ports from devices downstream of the Root or Switch Port are not reflected in this bit.</p> <p>Default value of this field is 0.</p>	RO
---	---	----

In Table 7-7 (Section 7.6), edit as shown:

31:27	<p>PME Support – For a device, this 5-bit field indicates the power states in which the device may generate a PME.</p> <p>Bits 31, 30, and 27 must be set to 1b for PCI-PCI bridge structures representing ports on Root Complexes/Switches to indicate that the bridge will forward PME Messages.</p>	RO
-------	---	----

In Table 7-13 (Section 7.8.5), edit as shown:

5	<p>Transactions Pending –</p> <p><i>Endpoints:</i></p> <p>This bit when set indicates that a-the device has issued Non-Posted Requests which have not been completed. A device reports this bit cleared only when all outstanding Non-Posted Requests have completed or have been terminated by the Completion Timeout mechanism. Completions for any outstanding Non-Posted Requests have been received.</p> <p><i>Root and Switch Ports:</i></p> <p>This bit when set indicates that a Port has issued Non-Posted Requests on its own behalf (using the Port's own Requester ID) which have not been completed. The Port reports this bit cleared only when all such outstanding Non-Posted Requests have completed or have been terminated by the Completion Timeout mechanism. Note that Root and Switch Ports implementing only the functionality required by this document do not issue Non-Posted Requests on their own behalf, and therefore are not subject to this case. Root and Switch Ports that do not issue Non-Posted Requests on their own behalf hardwire this bit to 0b.</p>	RO
---	---	----

In Table 7-15 (Section 7.8.7), edit as shown:

3	<p><i>Root Ports:</i></p> <p>Read Completion Boundary (RCB) – Indicates the RCB value for the Root Port. Refer to Section <> for the definition of the parameter RCB.</p> <p>Encodings are:</p> <table><tr><td>0b</td><td>64 byte</td></tr><tr><td>1b</td><td>128 byte</td></tr></table> <p>This field is hardwired for a Root Port and returns its RCB support capabilities.</p> <p><i>Other Than Root PortsEndpoints:</i></p> <p>Read Completion Boundary (RCB) – May be set by configuration software to indicate the RCB value of the Root Port upstream from the Endpoint. Refer to Section <> for the definition of the parameter RCB.</p> <p>Encodings are:</p> <table><tr><td>0b</td><td>64 byte</td></tr><tr><td>1b</td><td>128 byte</td></tr></table> <p>Devices that do not implement this feature must hardwire the field to 0b.</p> <p><i>Switch Ports:</i></p> <p>Not applicable - must hardwire the field to 0b.</p>	0b	64 byte	1b	128 byte	0b	64 byte	1b	128 byte	<p><i>Root Ports:</i></p> <p>RO</p> <p><i>Other Than Root PortsEndpoints:</i></p> <p>:</p> <p>RW</p> <p><i>Switch Ports:</i></p> <p>RO</p>
0b	64 byte									
1b	128 byte									
0b	64 byte									
1b	128 byte									

C40. In Msg Tables, All 'B' Should Be 'BD'

Release Date: 2/11/03

Since there is only (and exactly) one logical device associated with an Upstream Port, and since in the future the Device Number for this Device might not always be 0, messages (like other TLPs) must always use the correct (captured) Device Number for their Requester IDs.

In Section 2.2.8.1 (and subsections), edit as shown (note that Sw column entries were changed in C4):

Table 2-12: INTx Mechanism Messages

Name	Code[7:0]	Routing r[2:0]	Support [Ⓢ]				Req ID ⁸	Description/Comments
			R C	E p	S w	B r		

Name	Code[7:0]	Routing r[2:0]	Support ⁸				Req ID ⁸	Description/Comments
			R C	E p	S w	B r		
Assert_INTA	0010 0000	100	All:				BD	Assert INTA virtual wire Note: These Messages are used for PCI 2.3 compatible INTx emulation
			r		tr			
			As Required:					
				t		t		
Assert_INTB	0010 0001	100	All:				BD	Assert INTB virtual wire
			r		tr			
			As Required:					
				t		t		
Assert_INTC	0010 0010	100	All:				BD	Assert INTC virtual wire
			r		tr			
			As Required:					
				t		t		
Assert_INTD	0010 0011	100	All:				BD	Assert INTD virtual wire
			r		tr			
			As Required:					
				t		t		
Deassert_INTA	0010 0100	100	All:				BD	De-assert INTA virtual wire
			r		tr			
			As Required:					
				t		t		
Deassert_INTB	0010 0101	100	All:				BD	De-assert INTB virtual wire
			r		tr			
			As Required:					
				t		t		
Deassert_INTC	0010 0110	100	All:				BD	De-assert INTC virtual wire
			r		tr			
			As Required:					
				t		t		
Deassert_INTD	0010 0110	100	All:				BD	De-assert INTD virtual wire
			r		tr			
			As Required:					
				t		t		

[in footnote:]

⁸ The Requester ID includes sub-fields for Bus Number, Device Number and Function Number. Some Messages are not associated with specific ~~Devices or~~ Functions in a component, and for such Messages

this ~~ese~~ field is ~~are~~-Reserved; this is shown in this column using a code. Some messages can be used in more than one context, and therefore more than one code may be listed. The codes in this column are:

~~B~~ — = ~~Bus Number included; Device Number and Function Number are Reserved~~
 BD = Bus Number and Device Number included; Function Number is Reserved
 BDF = Bus Number, Device Number, and Function Number are included

Table 2-14: Power Management Messages

Name	Code[7:0]	Routing r[2:0]	Support				Req ID	Description/Comments
			R C	E p	S w	B r		
PM_Active_State_Nak	0001 0100	100	t	r	tr	r	BD	Terminate at Receiver
...								

Table 2-15: Error Signaling Messages

Name	Code[7:0]	Routing r[2:0]	Support				Req ID	Description/Comments
			R C	E p	S w	B r		
ERR_COR	0011 0000	000	r	t		t	B BD BDF	This Message is issued when the component or device detects a correctable error on the PCI Express interface.
ERR_NONFATAL	0011 0001	000	r	t		t	B BD BDF	This Message is issued when the component or device detects a non-fatal, uncorrectable error on the PCI Express interface.
ERR_FATAL	0011 0011	000	r	t		t	B BD BDF	This Message is issued when the component or device detects a fatal, uncorrectable error on the PCI Express interface.

In Section 5.6, p.258, edit as shown:

☐ Requester ID

- PM_PME message
 - ? Endpoints report their upstream Link bus number and the device and function number where the PME originated.

- All other ~~devices messages~~ report their upstream Link bus ~~and device~~ numbers, and ~~device and the~~ function number must ~~both~~ be zero.

C41. PCI-X Bridges can also set BCM**Release Date: 2/11/03***In Section 2.3.1.1, edit as shown:***IMPLEMENTATION NOTE****BCM Bit Usage**

To satisfy certain PCI-X protocol constraints, a **PCI-X Bridge** or PCI-X Completer for a PCI-X burst read in some cases will set the Byte Count field in the first PCI-X transaction of the Split Completion sequence to indicate the size of just that first transaction instead of the entire burst read. When this occurs, the **PCI-X Bridge/PCI-X** completer will also set the BCM bit in that first PCI-X transaction, to indicate that the Byte Count field has been modified from its normal usage. Refer to the PCI-X 2.0 specification for further details.

A PCI Express Memory Read Requester needs to correctly handle the case when a **PCI-X Bridge/PCI-X** Completer sets the BCM bit. When this occurs, the first Read Completion packet returned to the Requestor will have the BCM bit set, indicating that the Byte Count field reports the size of just that first packet instead of the entire remaining byte count. The Requester should not conclude at this point that other packets of the Read Completion are missing.

The BCM bit will never be set in subsequent packets of the Read Completion, so the Byte Count field in those subsequent packets will always indicate the remaining byte count in each instance. Thus, the Requester can use the Byte Count field in these packets to determine if other packets of the Read Completion are missing.

PCI Express Completers will never set the BCM bit.

C42. Configuration Address Mapping & Rules Clarification**Release Date: 2/11/03***In Section 7.2.2, edit as shown:*

... The mapping from memory address A[27:0] to PCI Express configuration space address is defined in Table 7-1. The base address A[63:28] is allocated in an implementation specific manner and reported by the system firmware to the operating system.

Table 7-1: Enhanced Configuration Address Mapping

Memory Address ^[footnote]	PCI Express Configuration Space
A[27:20]	Bus Number [7:0]
A[19:15]	Device Number [4:0]
A[14:12]	Function Number [2:0]
A[11:8]	Extended Register Number [3:0]
A[7:20]	Register Number [7:0]
A[1:0]	Along with size of the access, used to generate Byte Enables

[FOOTNOTE: This address refers to the byte-level address from a software point of view.]

In Section 7.3.2, edit as shown:

- ☐ Extended Register Number and Register Number – Specify the configuration space address of the register being accessed (concatenated such that Extended Register Number forms the more significant bits).

In Section 7.3.3, edit as shown:

...

For Root Ports, Switches and PCI Express-PCI Bridges, the following rules apply:

...

Additional rules ~~for specific to~~ Root Complexes:

- ☐ Configuration Requests addressing bus numbers assigned to devices within the Root Complex are processed by the Root Complex
 - The assignment of bus numbers to the logical devices within a Root Complex may be done in an implementation specific way.

~~☐ Configuration Requests addressing other buses are processed according to the rules for Switches (above)~~

C43. Clarify Slot Power Limit Register Operation

Release Date: 2/11/03

In Section 7.8.9, Table 7-17, edit as shown:

Table 7-17: Slot Capabilities Register

Bit Location	Register Description	Attributes
--------------	----------------------	------------

Bit Location	Register Description	Attributes
...		
14:7	<p>Slot Power Limit Value – In combination with the Slot Power Limit Scale value, specifies the upper limit on power supplied by slot (see Section 6.9).</p> <p>Power limit (in Watts) calculated by multiplying the value in this field by the value in the Slot Power Limit Scale field.</p> <p>This register must be implemented if the Slot Implemented bit is set.</p> <p>Writes to this register also cause the Port to send the Set_Slot_Power_Limit message.</p> <p>The default value prior to hardware/firmware initialization is 0000 0000b.</p>	HwInit
16:15	<p>Slot Power Limit Scale – Specifies the scale used for the Slot Power Limit Value (see Section 6.9).</p> <p>Range of Values:</p> <p>00b = 1.0x</p> <p>01b = 0.1x</p> <p>10b = 0.01x</p> <p>11b = 0.001x</p> <p>This register must be implemented if the Slot Implemented bit is set.</p> <p>Writes to this register also cause the Port to send the Set_Slot_Power_Limit message.</p> <p>The default value prior to hardware/firmware initialization is 00b.</p>	HwInit

C44. Clarify use of Max Read Request Size

Release Date: 2/11/03

In Section 2.2.7, add clarification and reference as shown:

For Memory Requests, the following rules apply:

- ☐ Memory Requests route by address, using either 64 bit or 32 bit Addressing (see <> and <>)
- ☐ For Memory Read Requests, Length must not exceed the value specified by Max_Read_Request_Size (see Section 7.8.4)

C45. Byte Enable Valid Usage Model

Release Date: 2/11/03

In Section 2.2.5, edit as shown

- ❑ Non-contiguous byte enables (enabled bytes separated by non-enabled bytes) are permitted in ~~the 1st DW BE field for all Requests with length of 1 DW~~~~both Byte Enable fields.~~
- ~~Non-contiguous byte enable examples~~~~Examples:~~ 1010b, 0101b, 1001b, 1011b, 1101b
- ❑ Non-contiguous byte enables are permitted in both Byte Enable fields for QW aligned Memory Requests with length of 2 DW (1 QW).
- ❑ All Memory Requests with length of 3 DW or more must enable only bytes that are contiguous with the data between the first and last DW of the Request
 - Contiguous byte enable examples:
 - 1st DW BE: 1100, Last DW BE: 0011
 - 1st DW BE: 1000, Last DW BE: 0111

C46. Reset/"Power Good" – Protocol Aspects

Release Date: 2/11/03

Note: For this change, there are corresponding changes for Chapter 4 (Physical Layer) and in the EM specs. These changes are covered in other entries.

In Terms and Acronyms, edit as shown:

...	
cold reset	A " Power Good Fundamental Reset" following the application of power.
...	
hot reset	A reset propagated in-band across a Link using a Physical Layer mechanism.
...	
warm reset	A "Fundamental Reset" reset caused by driving the "Power Good" signal inactive and then active without cycling the supplied power.
...	

In "2.9.1. Transaction Layer Behavior in DL_Down Status", edit as shown:

...

For a Port on an Endpoint, and the Port on a Switch or Bridge that is closest to the Root Complex, DL_Down status is handled as a ~~Link~~-reset by:

...

❑ (for Switch and Bridge) propagating ~~Link-hot Reset-reset~~ to all other Ports

...

In “5.3.1.4.2. D3cold State”, edit as shown:

A function transitions to the D3cold state when its power is removed. A power-on sequence ~~with its associated cold reset~~ transitions a function from the D3cold state to the D0Uninitialized state. At this point software must perform a full initialization of the function in order to re-establish all functional context, completing the restoration of the function to its D0active state.

...

An auxiliary power source must be used to support PME event detection, Link reactivation, and to preserve PME context from within D3cold. Note that once the I/O Hierarchy has been brought back to a fully communicating state, as a result of the Link reactivation, the waking agent then propagates a PME message to the root of the Hierarchy indicating the source of the PME event. Refer to Section <5.3.3> for further PME specific detail. ~~Exit from D3cold is accomplished with assertion of PWRGOOD, (either provided as an auxiliary signal or internally generated by the component), followed by the Link training sequence.~~

...

In “5.3.3.5. PM_PME Delivery State Machine”, edit as shown:

Note: This section also affected by Beacon/WAKE# write-up (C47).

The following diagram conceptually outlines the PM_PME delivery control state machine. This state machine determines ability of a Link to service PME events by issuing PM_PME immediately vs. requiring initial Link reactivation.

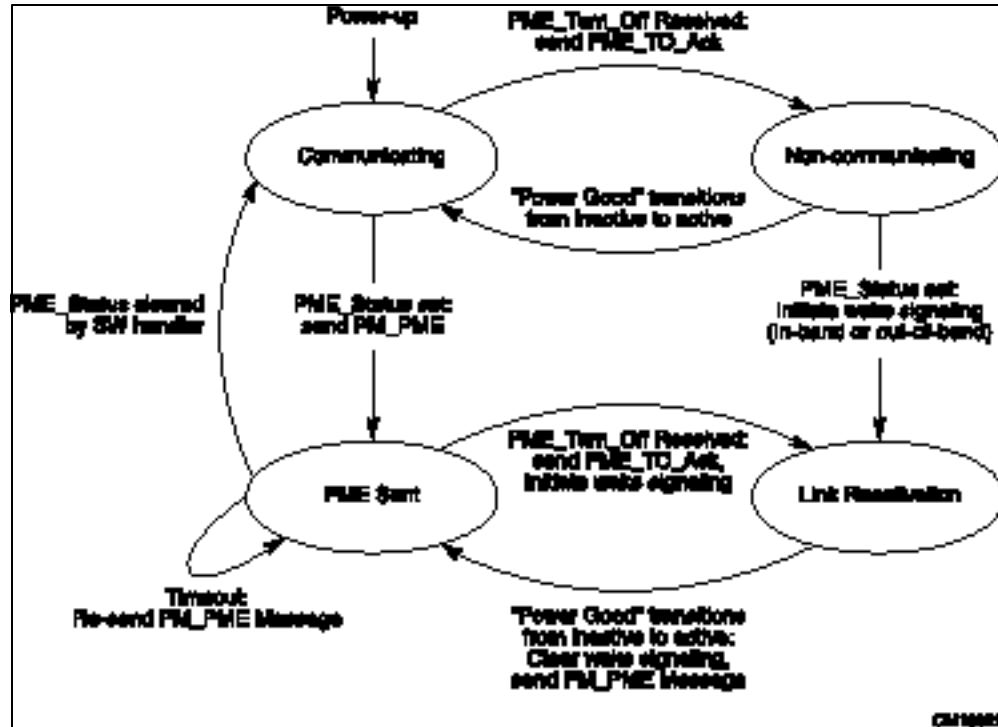


Figure 0-3: A Conceptual PME Control State Machine <!! Change figure to match text>

Communicating State:

At initial power-up **and associated reset**, the Upstream Link enters the “Communicating” state

- ❑ If PME_Status is asserted (assuming PME delivery is enabled), a PM_PME Message will be issued upstream, terminating at the root of the PCI Express Hierarchy. The next state is the “PME Sent” state
- ❑ If a PME_Turn_Off Message is received, the Link enters the “Non-Communicating” state following its acknowledgment of the message and subsequent entry into the L2/L3 Ready state.

Non-communicating State:

- ❑ ~~If a Power Good signal transitions from inactive to active state (an indication of power and clock restoration), the next state is the “Communicating” state.~~
Following the restoration of ~~at~~ power and clock, ~~and the associated reset have been restored~~, the next state is the “Communicating” state.
- ❑ If PME_Status is asserted, the Link will transition to “Link Reactivation” state, and activate the wake mechanism.

PME Sent State

...

Link Reactivation State

- ❑ ~~If a Power Good signal transitions from inactive to active st~~Following the restoration of power and clock, and the associated ~~reset~~, the Link resumes a transaction-capable state. The device clears the wake signaling, issues a PM_PME Upstream and transitions into the “PME Sent” state.

...

In “6.2.2.2.1. Fatal Errors”, edit as shown:

Fatal errors are uncorrectable error conditions which render the particular PCI Express Link and related hardware unreliable. For fatal errors, a reset of the **components on the** Link may be required to return to reliable operation. Platform handling of fatal errors, and any efforts to limit the effects of these errors, is platform implementation specific.

Comparing with PCI/PCI-X, reporting a fatal error is somewhat analogous to asserting SERR#.

In “6.2.2.2.2. Non-Fatal Errors”, edit as shown:

Non-fatal errors are uncorrectable errors which cause a particular transaction to be unreliable but the Link is otherwise fully functional. Isolating non-fatal from fatal errors provides system management software the opportunity to recover from the error without resetting the **components on the** Link(s) and disturbing other transactions in progress. Devices not associated with the transaction in error are not impacted by the error.

...

In “6.6. PCI Express Reset – Rules”, edit as shown:

This section specifies the behavior of PCI Express **Link**-reset. ~~This document covers the relationship between the architectural mechanisms defined in this document and the reset mechanisms defined in this document. The reset can be generated by the platform or on the component, but a~~Any relationship between the PCI Express **Link**-reset and component or platform reset is component or platform specific (respectively).

In all form factors and system hardware configurations, there must at some level

~~? There must~~ be a hardware mechanism for setting or returning all Port states to the initial conditions specified in this document— this mechanism is called “~~Power Good~~**Fundamental Reset**” . ~~This mechanism can take the form of an auxiliary signal provided by the system to a component or add-in card, in which case the signal must be called PERST#, and must conform to the rules specified in Section <4.2.4.5.1>. When PERST# is provided to a component or add-in card, this signal must be used by the component or add-in card to cause a “Fundamental Reset”. When PERST# is not provided to a component or add-in card, “Fundamental Reset” is generated autonomously by the component or an add-in card, and the details of how this is done are outside the scope of this document. If “Fundamental Reset” is generated autonomously by the component or add-in card, and if power is supplied by the platform to the component/add-in card, then the component/add-in card must generate a “Fundamental Reset” to itself if the supplied power goes outside of the limits specified for the form-factor or system.~~

- ❑ This document describes three distinct types of reset: cold, warm, and hot.

- A “~~Power Good~~**Fundamental Reset**” ~~will~~**must** occur following the application of power to the component. This is called a “cold” reset.

- In some cases, it may be possible for the “~~Power Good~~Fundamental Reset” mechanism to be triggered by hardware without the removal and re-application of power to the component. This is called a “warm” reset. ~~Note that this document does not specify a means for generating a warm reset.~~

- ~~Note that t~~There is ~~also~~ an in-band mechanism for propagating reset across a Link. This is called a “hot” reset and is described in Section <4.2.4.5>. ~~Note that this mechanism is automatically triggered as a part of the Link initialization process that follows a cold or warm reset.~~

~~?~~ Note also that ~~entering the DL_Inactive state~~ the Data Link Layer reporting ~~DL_Down~~ is in some ways identical to a “hot” reset – see Section <2.9>.

- ❑ On exit from any type of reset (cold, warm, or hot), all Port registers and state machines must be set to their initialization values as specified in this document, ~~except for sticky registers (Section <7.4>)~~
- ❑ On exit from a “~~Power Good~~Fundamental Reset”, the Physical Layer will attempt to bring up the Link (see Section <4.2.5>). Once both components on a Link have entered the initial Link Training state, they will proceed through Link initialization for the Physical Layer and then through Flow Control initialization for VC0, making the Data Link and Transaction Layers ready to use the Link
 - Following Flow Control initialization for VC0, it is possible for TLPs and DLLPs to be transferred across the Link

Following a reset, some devices may require additional time before they are able to respond to Requests they receive. Particularly for Configuration Requests it is necessary that components and devices behave in a deterministic way, which the following rules address. The first set of rules address requirements for components and devices:

- ❑ A component must enter the initial active Link Training state within 80 ms of the end of “~~Power Good~~Fundamental Reset” (Link Training is described in Section <4.2.4>)
 - Note: In some systems, it is possible that the two components on a Link may exit “~~Power Good~~Fundamental Reset” at different times. Each component must observe the requirement to enter the initial active Link Training state within 80 ms of the end of “~~Power Good~~Fundamental Reset” from its own point of view.
- ❑ On the completion of Link Training (entering the DL_Active state, see Section <3.2>), a component must be able to receive and process TLPs and DLLPs

The second set of rules address requirements placed on the system:

- ❑ To allow components to perform internal initialization, system software must wait for at least 100 ms from the end of a reset ~~of one or more device (cold/warm/hot)~~ before it is permitted to issue Configuration Requests to ~~PCI Express those~~ devices
 - A system must guarantee that all components intended to be software visible at boot time are ready to receive Configuration Requests within 100 ms of the end of “~~Power Good~~Fundamental Reset” ~~at the Root Complex~~ – how this is done is beyond the scope of this specification

- ❑ The Root Complex and/or system software must allow 1.0s (+50%/-0%) after a reset ~~(hot/warm/cold), of a device~~ before it may determine that a device which fails to return a Successful Completion status for a valid Configuration Request is a broken device

~~2~~Note: This delay is analogous to the Trhfa parameter specified for PCI/PCI-X, and is intended to allow an adequate amount of time for devices which require self initialization.

- ❑ When attempting a Configuration access to devices on a PCI or PCI-X segment behind a PCI Express/PCI(-X) Bridge, the timing parameter T_{rhfa} must be respected

When a Link is in normal operation, the following rules apply:

- ❑ If, for whatever reason, a normally operating Link goes down, the Transaction and Data Link Layers will enter the DL_Inactive state (see Sections <2.9> and <3.2.1>)

- ❑ For any ~~virtual or actual PCI Bridge~~Root or Switch downstream Port, setting the Secondary Bus Reset bit of the Bridge Control register associated with the Port ~~any of the following~~ must cause a hot reset ~~of the secondary side of the Bridge using the Physical Layer mechanism for communicating Link Reset~~ to be sent (see Section <4.2.4.5>):

- ~~Setting the Secondary Bus Reset bit of the Bridge Control register~~

- ❑ For a Switch, the following must cause a hot reset to be sent on all downstream Ports:

- Setting the Secondary Bus Reset bit of the Bridge Control register associated with the upstream Port
- The Data Link Layer of the Upstream Port reporting DL_Down~~Entering DL_Inactive on the primary side of the Bridge~~
- Receiving a hot ~~Primary side Link~~ reset ~~using the Physical Layer mechanism for communicating Link Reset~~ on the upstream Port

Certain aspects of “~~Power Good Fundamental~~ Reset” are specified in this document and others are specific to a platform, form factor and/or implementation. Specific platforms, form factors or application spaces may require the additional specification of the timing and/or sequencing relationships between the components of the system for “~~Power Good Fundamental~~ Reset”. For example, it might be required that all PCI Express components within a chassis observe the assertion and deassertion of “~~Power Good Fundamental~~ Reset” at the same time (to within some tolerance). In a multi-chassis environment, it might be necessary to specify that the chassis containing the Root Complex be the last to exit “~~Power Good Fundamental~~ Reset.”

In all cases where power ~~and~~ PERST# ~~are~~ supplied, the following parameters must be defined:

- ❑ $T_{pvppl} - T_{pvperl}$ – PERST# “~~Power Good~~” must remain ~~in~~ active at least this long after power becomes valid
- ❑ $T_{pwgpd} - T_{perst}$ – When ~~de~~asserted, PERST# “~~Power Good~~” must remain ~~de~~asserted at least this long

- ❑ T_{fail} – When power becomes invalid, ~~PERST#~~“Power Good” must be ~~de~~asserted within this time

Additional parameters may be specified.

In all cases where a reference clock is supplied, the following parameter must be defined:

- ❑ $T_{perst\cancel{wgd}\text{-clk}}$ – ~~PERST#~~“Power Good” must remain ~~in~~active at least this long after any supplied reference clock ~~is~~ stable

Additional parameters may be specified.

...

In “7.4. Configuration Register Types”, edit as shown:

...

Table 7-2: Register (and Register Bit-Field) Types

Register Attribute	Description
HwInit	Hardware Initialized: Register bits are initialized by firmware or hardware mechanisms such as pin strapping or serial EEPROM. Bits are read-only after initialization and can only be reset (for write-once by firmware) with “ Power GoodFundamental Reset” (see Section <6.6>).
...	

...

C47. Wakeup Mechanisms – Protocol Aspects

Release Date: 2/11/03

Note: For this change, there are corresponding changes for Chapter 4 (Physical Layer) and in the EM specs. These changes are covered in other entries.

In Terms and Acronyms, edit as shown:

...

Beacon An optional 30 kHz–500 MHz inband signal used to exit the L2 Link power management state. One of two defined mechanisms for waking up a Link in L2 (see also wakeup)

...

inband signaling A method for signaling events and conditions using the Link between two components, as opposed to the use of separate physical (sideband) signals. All mechanisms defined in this document can be implemented using inband signaling, although in some form factors sideband signaling may be used instead.

...

sideband signaling A method for signaling events and conditions using physical signals separate from the signals forming the Link between two components. All mechanisms defined in this document can be implemented using inband signaling, although in some form factors sideband signaling may be used instead.

...

wakeup An optional mechanism used by a component to request the reapplication of main power when in the L2 Link state. Two such mechanisms are defined: Beacon (using inband signaling), and WAKE# (using sideband signaling).

...

In “1.1. A Third Generation I/O Interconnect”, edit as shown:

...

- Power management and budgeting
 - ? Ability to identify power management capabilities of a given function
 - ? Ability to transition a function into a specific power state
 - ? Ability to receive notification of the current power state of a function
 - ? Ability to ~~propagate~~ generate an event request to wakeup the system from a main power-off state of the main power supply
 - ? Ability to sequence device power-up to allow graceful platform policy in power budgeting.

...

In “5.1. Overview”, edit as shown:

PCI Express-PM provides the following services:

- ☐ A mechanism to identify power management capabilities of a given function
- ☐ The ability to transition a function into a certain power management state
- ☐ Notification of the current power management state of a function
- ☐ The option to wakeup the system on a specific event

...

PCI Express components are permitted to ~~wake~~wakeup the system ~~from any supported power management state using a wakeup mechanism followed by through the request of a~~ power management event (PME) message. PCI Express systems may provide the optional auxiliary power supply (Vaux) needed for ~~PME-wakeup~~ operation from the “off” system states where the main power supplies are off. PCI Express-PM extends beyond ~~its~~ the PME mechanism defined in conventional PCI-PM ~~predecessor in this regard~~ as PCI Express PME “messages” include the ~~geographical location (Requestor ID) within the Hierarchy~~ of the requesting agent. These PME messages are in-band TLPs routed from the requesting device ~~to a towards the~~ Root Complex.

Another distinction of the PCI Express-PM PME mechanism is in its separation of the following two tasks that are associated with PME:

- ❑ Reactivation (~~wake~~wakeup) of the ~~I/O Hierarchy-associated~~ resources (i.e., re-establishing reference clocks and main power rails to the PCI Express components)
- ❑ Sending ~~the actual a~~ PME Message (~~vector~~) to the Root Complex

...

In “5.2. Link State Power Management”, edit as shown:

...

- ❑ L2 – Auxiliary powered Link deep energy saving state.

L2 support is optional, and dependent upon platform support of Vaux.

L2 – The downstream component’s main power supply inputs and reference clock inputs are shut off.

- When in L2, all ~~PME detection logic~~, Link reactivation ~~“Beacon”~~wakeup logic (~~Beacon or WAKE#~~), PME context, and any other “keep alive” logic is powered by Vaux.

TLP and DLLP communication over a Link that is in L2 is ~~prohibited~~not possible.

Exiting the L2 state is accomplished by reestablishing main power and reference clocks to all components within the domain of the power manager, followed by full Link training and initialization. Once a given Link has completed Link training and initialization it is then in the L0 state and may begin sending and receiving TLPs and DLLPs.

...

In “5.3.1.4.2. D3cold State”, edit as shown:

...

An auxiliary power source must be used to support PME event detection ~~within a device~~, Link reactivation, and to preserve PME context from within D3cold...

In “5.3.3. Power Management Event Mechanisms’ and subsections, edit as shown:

5.3.3.1. Motivation

...

Power management software may transition a PCI Express Hierarchy into a low power state, and transition the upstream links of these devices into the non-communicating L2 state⁶. The PCI Express PME generation mechanism is therefore broken into two components:

- ❑ Waking a non-communicating Hierarchy (~~wake~~up). This step is required only if the upstream Link of the device originating the PME is in the non-communicating L2 state, since in that state the device cannot send a PM_PME message upstream.

⁶ The L2 state is defined as “non-communicating” since component reference clock and main power supply are removed in that state.

❑ Sending a PM_PME message to the root of the PCI Express Hierarchy

...

5.3.3.2. Link ~~Reactivation~~Wakeup

The Link ~~reactivation~~-wakeup mechanisms provides a means of signaling the platform to re-establish power and reference clocks to the components within its domain. ~~There are two defined wakeup mechanisms: Beacon and WAKE#. The Beacon mechanism uses inband signaling to implement wakeup functionality, and is described in Refer to Ssection <4.3.2.4>. For components that support wakeup functionality, Beacon is the required mechanism for all components, except for components designed exclusively for the following form factors: PCI Express Card Electromechanical Specification, and PCI Express Mini-Card Electromechanical Specification [Ed note: Need to double-check this name and add Mini-Card to the "Reference Documents" section at front of spec]. [See C53 for additional clarifying statement]~~

The WAKE# mechanism uses sideband signaling to implement wakeup functionality. WAKE# is an "open drain" signal asserted by components requesting wakeup and observed by the associated power controller. WAKE# is only defined for certain form factors, and the detailed specifications for WAKE# are included in the relevant form factor specifications. ~~for details on the in-band mechanism for Link reactivation. Refer to the PCI Express Card Electromechanical Specification for details on the out-of-band mechanism for Link reactivation.~~ Specific form factor specifications may require the use of either Beacon or WAKE# as the wakeup mechanism. All form factors that require WAKE# as the wakeup mechanism must permit components to also generate Beacon, although the Beacon may not be observed.

When WAKE# is used as a wakeup mechanism, ~~Once~~once WAKE# has been asserted, the asserting PCI Express function must continue to drive the signal low until ~~software explicitly clears the PME_Status bit in that function's PMCSR~~main power has been restored to the component as indicated by "power good".

WAKE# is not intended to be used as an input by any ~~PCI Express function other than the PM controller~~Endpoint, and the system is not required to route or buffer it in such a way that ~~an PCI Express function Endpoint~~ is guaranteed to be able to detect that the signal has been asserted by another function.

Before using any wakeup mechanism, ~~PCI Express~~functions must be enabled by software ~~before to do so asserting WAKE#~~ by setting the device's PME_En bit in the Power Management Control/Status Register (PMCSR). ~~Devices~~The PME_Status bit is sticky, and ~~devices~~ must maintain the value of the PME_Status bit through reset if ~~+3.3V~~aux power is available and they are enabled for ~~wake~~wakeup events.

Systems that allow PME generation from D3cold state must provide auxiliary power to support Link ~~reactivation~~-wakeup when the main system power rails are off.

~~The reactivation period ends when the upstream-directed Link of a device enters the initialization phase as a result of Link transition from the L2 state to the L0 state—a power-on sequence transitions the device from the D3cold state to the D0uninitialized state.~~

~~The downstream device shall cease requesting Link reactivation (either in-band or auxiliary out-of-band) once it has entered the D0uninitialized state.~~

Regardless of the wakeup mechanism used, Once the Link has been re-activated and trained, the requesting agent then propagates a PM_PME message upstream to the Root Complex. From a power management point of view, the two wakeup mechanisms provide the same functionality, and are not distinguished elsewhere in this chapter.



IMPLEMENTATION NOTE

Example of WAKE# to Beacon Translation

Switch products targeting applications that connect “Beacon domains” and “WAKE# domains” must translate the wakeup mechanism appropriately. Figure <!!XX> shows two example systems, each including slots that use the WAKE# wakeup mechanism. In Case 1, WAKE# is input directly to the Power Management Controller, and no translation is required. In Case 2, WAKE# is an input to the Switch, and in response to WAKE# being asserted the Switch must generate a Beacon that is propagated to the Root Complex/Power Management Controller.

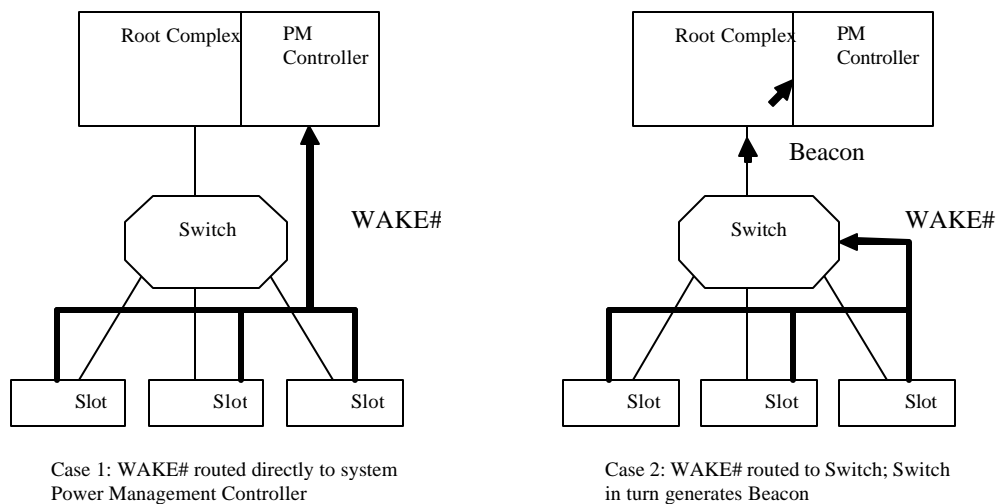


Figure <!!XX>: Conceptual Diagrams Showing Two Example Cases of WAKE# Routing

...

5.3.3.4. PME Rules

- ❑ All PCI Express components supporting PCI Express-PM must implement the PCI-PM PMC and PMCSR registers in accordance with the PCI-PM specification. These registers reside in the PCI-PM compliant PCI Capability List format.

- PME capable functions must implement the PME_Status bit, and underlying functional behavior, in their PMCSR configuration register.
- When a function initiates Link ~~reactivation~~wakeup, or issues a PM_PME Message, it must set its PME_Status bit.

...

5.3.3.5. PM_PME Delivery State Machine

The following diagram conceptually outlines the PM_PME delivery control state machine. This state machine determines ability of a Link to service PME events by issuing PM_PME immediately vs. requiring ~~initial~~-Link ~~reactivation~~wakeup.

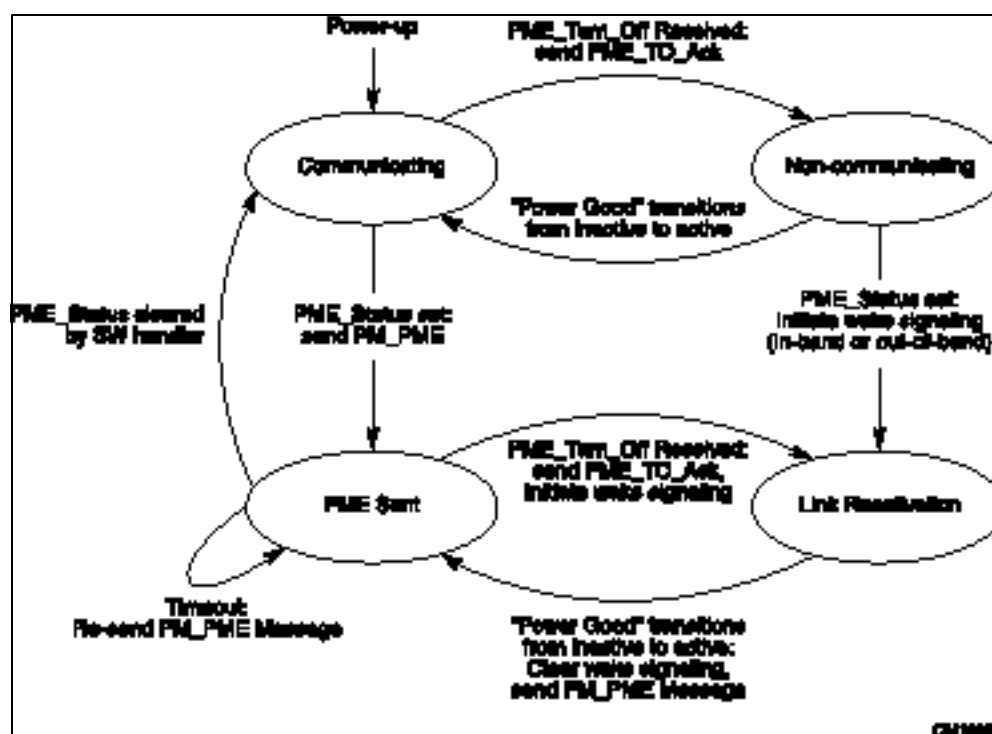


Figure 0-4: A Conceptual PME Control State Machine <!! Change “wake” to “wakeup” – see text(!)>

Communicating State:

At initial power-up, the Upstream Link enters the “Communicating” state

- ❑ If PME_Status is asserted (assuming PME delivery is enabled), a PM_PME Message will be issued upstream, terminating at the root of the PCI Express Hierarchy. The next state is the “PME Sent” state
- ❑ If a PME_Turn_Off Message is received, the Link enters the “Non-Communicating” state following its acknowledgment of the message and subsequent entry into the L2/L3 Ready state.

Non-communicating State:

- ❑ *[ed note: This is modified by C46]* If a “Power Good” signal transitions from inactive to active state (an indication that power and clock have been restored), the next state is the “Communicating” state.
- ❑ If PME_Status is asserted, the Link will transition to “Link Reactivation” state, and activate the ~~wake~~wakeup mechanism.

PME Sent State

- ❑ If PME_Status is cleared, the function becomes PME Capable again. Next state is the “Communicating” state.
- ❑ If the PME_Status bit is not cleared by the time the PME service timeout expires, a PM_PME message is re-sent upstream. See Section <> for an explanation of the timeout mechanism.
- ❑ If a PME message has been issued but the PME_Status has not been cleared by software when the Link is about to be transitioned into a messaging incapable state (a PME_Turn_Off Message is received), the Link transitions into “Link Reactivation” state after sending a PME_TO_Ack message. The device also activates the ~~wake~~wakeup mechanism.

Link Reactivation State

- ❑ If a Power Good signal transitions from inactive to active state, the Link resumes a transaction-capable state. The device clears the ~~wake~~up signaling, *if necessary*, and issues a PM_PME Upstream and transitions into the “PME Sent” state.

...

In “6.1.6. Native PME Software Model”, edit as shown:

...

The software handler for this interrupt can determine which device ~~caused the wake~~ sent the ~~PME message event~~ by reading the PME Requester ID field in the Root Status Register in a Root Port. ...

...

In “6.1.9. PME Routing Between PCI Express and PCI Hierarchies”, edit as shown:

PME-capable ~~conventional~~ PCI and PCI-X devices assert the PME# pin to signal a power management event. The PME# signal from PCI devices may either be converted to a PCI Express in-band PME message by a PCI Express-PCI bridge or routed directly to the Root Complex.

...

In “6.7.3.1. Attention Button Registers”, edit as shown:

...

Attention Button Pressed Enable – This bit when set enables the generation of the hot plug interrupt or a ~~wake~~wakeup ~~signal-event~~ on an Attention Button Pressed event.

...

In “6.7.3.4. Power Controller Registers”, edit as shown:

...

Power Fault Detected Enable – This bit when set enables the generation of the hot plug interrupt or a ~~wake~~wakeup event ~~signal~~ on a power fault event.

In “6.7.3.5. Presence Detect Register”, edit as shown:

...

Presence Detect Changed Enable – This bit when set enables the generation of the hot plug interrupt or a ~~wake~~wakeup event ~~signal~~ on a presence detect changed event

...

In “6.7.3.6. MRL Sensor Registers”, edit as shown:

...

MRL Sensor Changed Enable – This bit when set enables the generation of the hot plug interrupt or a ~~wake~~wakeup event ~~signal~~ on a MRL Sensor changed event.

...

In Section 6.7.7, edit as shown:

6.7.7. PCI Express Hot Plug Interrupt/~~Wake~~Wakeup Signal Logic

...

When the system is in a sleep state or if the hot plug capable port is in a device state D1, D2, or D3hot, the above hot plug controller events generate a ~~wake~~wakeup event ~~message~~ (using PME mechanism) instead of a hot plug interrupt. ~~Note that the hot plug controller generates the wakeup on behalf of the hot plugged device, and it is not necessary for that device to have auxiliary (or main) power.~~

A hot plug capable port also supports generation of hot plug interrupt when the hot plug control logic completes an issued command. However, a command completed event will not cause a ~~wake~~wakeup event if the system is in a sleep state or if the hot plug capable port is in a device state D1, D2, or D3hot.

Figure <> shows the logical connection between the hot plug event logic and the system interrupt/~~wake~~wakeup event generation logic.

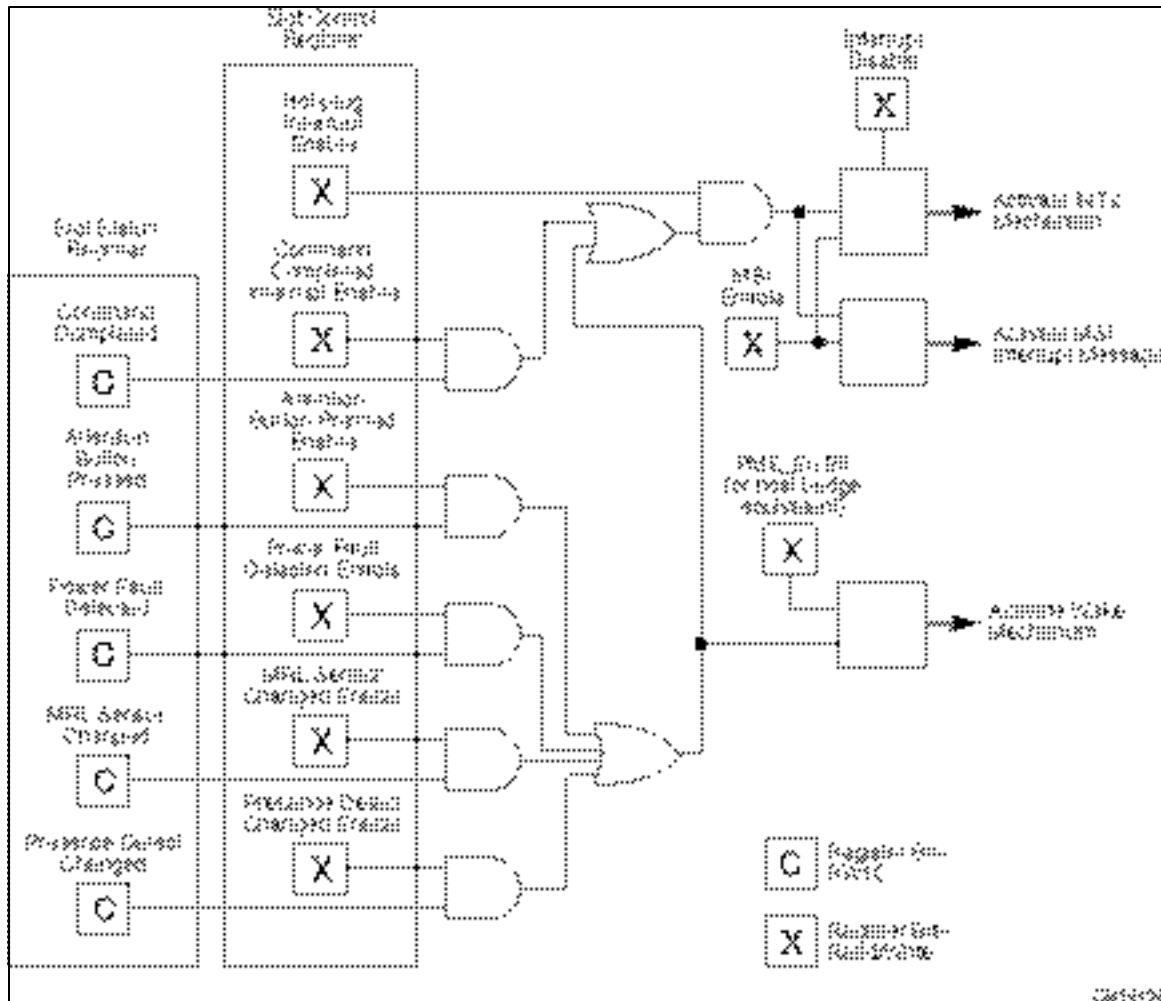


Figure 6-9: Hot Plug Logic *<!! Change “wake” to “wakeup”>*

In “6.7.8. The Operating System Hot Plug Method”, edit as shown:

...

```
Device(PPB1){
```

```
...
```

```
Method(OSHP, 0) {
```

```
    // Disable firmware access to SHPC and restore
```

```
    // the normal System Interrupt and Wakeup mechanismSignal
```

```

        // connection-
    }

    ...
}

```

In “7.6. PCI Power Management Capability Structure”, edit as shown:

This structure is required for all PCI Express devices. Figure <7-7> details allocation of the PCI PM Capability Structure register fields in a PCI Express Context. PCI Express devices are required to support D0 and D3 device states (refer to Section <5.1.1>); PCI-PCI bridge structures representing PCI Express ports as described in Section <7.1> are required to indicate PME ~~wake-message passing~~ capability due to the in-band nature of PME messaging for PCI Express.

The PME status bit for the PCI-PCI bridge structure representing PCI Express ports, however, is only set when the PCI-PCI bridge function is itself generating a PME. The PME status bit is not set when the bridge is propagating a PME ~~message~~ but the PCI-PCI bridge function itself is not internally ~~asserting-generating a~~ PME.

...

Edit Table 7-18 as shown:

Table 7-18: Slot Control Register

Bit Location	Register Description	Attributes
0	Attention Button Pressed Enable – This bit when set enables the generation of hot plug interrupt or wake wakeup message event on an attention button pressed event. Default value of this field is 0.	RW
1	Power Fault Detected Enable – This bit when set enables the generation of hot plug interrupt or wake wakeup message -event on a power fault event. Default value of this field is 0.	RW
2	MRL Sensor Changed Enable – This bit when set enables the generation of hot plug interrupt or wake wakeup message -event on a MRL sensor changed event. Default value of this field is 0.	RW
3	Presence Detect Changed Enable – This bit when set enables the generation of hot plug interrupt or wake wakeup message event on a presence detect changed event. Default value of this field is 0.	RW

Bit Location	Register Description	Attributes
...		

C48. Device 0 Decode Clarification

Release Date: 2/11/03

Clarify text so that device implementations do not wrongly assume a fixed device number of 0.

In Section 7.1, edit as shown:

...

A PCI Express endpoint is mapped into configuration space as a single logical device (~~Device 0~~) with one or more logical functions.

...

In Section 7.3.1, move footnote into body text and edit as shown:

...

Switches and Root Complexes must associate only Device 0 with the device attached to the logical bus representing the link from a Switch Downstream Port or a Root Port. Configuration Requests targeting the Bus Number associated with a link specifying Device Number 0 are delivered to the device attached to the link; Configuration Requests specifying all other Device Numbers (1-31) must be terminated by the Switch Downstream Port or the Root Port with an Unsupported Request Completion Status (equivalent to Master Abort in PCI).⁷ ~~Devices must not assume that Device Number 0 is associated with their upstream port, but must capture their assigned Device Number as discussed in Section <2.2.6.2>.~~

...

C49. Remove Advanced Switching References

Release Date: 2/11/03

Note: This errata has no material effect on the base specification document, but is written as a “C” item to keep the material together.

In Section 1.1:

- ~~• Multi-hierarchy and advanced peer-to-peer communications~~
- ~~? Ability to support vendor-specific and PCI Express-standard peer-to-peer communications messaging~~

⁷ ~~Future switch components that are implemented as a single switch device (without the PCI PCI Bridges) that is not limited by legacy compatibility requirements may not have this restriction. To accommodate such future implementations, devices may not assume that device 0 is associated with their upstream port.~~

~~? Ability to connect multiple hierarchies to support peer-to-peer communications across large fabric topologies~~

In Section 1.3.3:

A Switch is defined as a logical assembly of multiple virtual PCI-to-PCI bridge devices as illustrated in Figure 1-3. All Switches are governed by the following base rules ~~(advanced Switch components will support additional capabilities beyond those described below).~~


In Section 1.5.4.1:

~~□ Transaction-level support for the switching and advanced communication applications~~

Delete all of Section 1.6:

~~1.6 Advanced Peer-to-Peer Communication Overview~~

~~Advanced peer-to-peer communication is an optional functionality used to support peer-to-peer communications across one or more hierarchies that constitute a single fabric instance. Figure 1-6 shows an example of a fabric with multiple hierarchies.~~



~~Figure 1-6: Advanced Peer-to-Peer Communication~~

The primary attributes/requirements are:

- ~~□ A communications paradigm used to exchange control and data packets.~~
- ~~□ Optional support for multicast packet replication within Advanced Switch components. Multicast allows an Endpoint to inject a single packet into the fabric targeted at a multicast group identifier and have the Advanced Switch components replicate this packet to all participating Endpoints. This eliminates the need for the injecting Endpoint having to know all of the participating Endpoints within the multicast group.~~

- ~~❑ Packets are routed from source to destination based on a flat node addressable ID.~~
- ~~❑ Each PCI Express Base Specification-compatible hierarchy effectively defines an individual partition within the advanced switching fabric.~~
 - ~~• At any given time, one Root Complex (RC) controls a partition.~~
 - ~~• Multiple partitions may be collapsed into a single partition by assigning ownership of the partition to one of the active RC. For example, dual-redundant fabrics that are inter-linked such that either RC may take over for the other should one RC fail.~~
- ~~❑ Communication between two independent partitions is not allowed until both the partitions are initialized and configured. At this point, configuration software can configure the Links between the partitions (contained within advanced Switch components). During this step, fabric routing is established between components attached to the fabric.~~

~~The scope of the information provided in this document is limited to providing definition of basic primitives required to support advanced PCI Express packet switching applications. A detailed description of typical usage models and operation of the capabilities enabled by these optional features is beyond the scope of this document. This information will be provided in the PCI Express Advanced Packet Switching Specification.~~

In Section 2.1.1, Table 2-1:

Table 2-1: Transaction Types for Different Address Spaces

Address Space	Transaction Types	Basic Usage
...		
Message	Baseline (including Vendor-defined) Advanced Switching	From event signaling mechanism to general purpose messaging

In Section 2.1.1.4:

2.1.1.4 Message Transactions

~~The Message Transactions, or simply Messages, are used to support two primary usage models and corresponding groups of messages:~~

- ~~❑ In-band communication of events between PCI Express devices~~
- ~~❑ Peer-to-peer communication between PCI Express devices (associated with Advanced Switching)~~

~~These two usage models map to two different groups of Messages in PCI Express.~~

~~In addition to the specified Messages, PCI Express provides support for vendor-defined messages using [See E4.12] specific-specified message reserved codes given in this document. The definition of specific vendor-defined messages is outside the scope of this document.~~

This specification establishes a standard framework within which vendors can specify their own vendor-defined Messages tailored to fit the specific requirements of their platforms (see Sections **Error! Reference source not found.** and **Error! Reference source not found.**).

Note that these vendor-defined messages are not guaranteed to be interoperable with components from different vendors.

In Section 2.2.1, Table 2-3:

Table 2-3: Fmt[1:0] and Type[4:0] Field Encodings

TLP Type	Fmt [1:0]	Type [4:0]	Description
...			
MsgAS	01	11n ₂ n ₁ n ₀	Message for Advanced Switching—The sub-field n[2:0] specifies the message type: 1n₂n₁n₀—Signaling Packet Messages A detailed description of message types and message headers will be presented in a separate document entitled Advanced PCI Express Packet Switching Specification. This is a companion specification to the PCI Express Base Specification.
MsgASD	11	11c ₂ c ₁ c ₀	Message for Advanced Switching—The sub-field c[2:0] specifies the message type: 1c₂c₁c₀—Data Packet Messages A detailed description of message types and message headers will be presented in a separate document entitled Advanced PCI Express Packet Switching Specification. This is a companion specification to the PCI Express Base Specification.

In Section 2.2.8:

2.2.8 Message Request Rules

This document defines the following groups of Messages:

~~Baseline Message Group—includes:~~

- ~~⊕~~ INTx Interrupt Signaling
- ~~⊕~~ Power Management
- ~~⊕~~ Error Signaling
- ~~⊕~~ Locked Transaction Support

⊕□ Slot Power Limit Support

⊕□ Vendor Defined Messages

⊕□ Hot Plug Signaling

~~□ Advanced Switching Support Message Group—includes:~~

~~• Data Packet Messages~~

~~• Signal Packet Messages~~

The following rules apply to all ~~Baseline~~ Message ~~Group~~ Requests. Additional rules specific to each type of Message ~~Request~~ follow.

...

Delete section heading and title:

~~2.2.8.1—Baseline Message Group~~

For Sections 2.2.8.1.1 through 2.2.8.1.7, promote the section headings as shown:

2.2.8.1.1 becomes 2.2.8.1

2.2.8.1.2 becomes 2.2.8.2

2.2.8.1.3 becomes 2.2.8.3

2.2.8.1.4 becomes 2.2.8.4

2.2.8.1.5 becomes 2.2.8.5

2.2.8.1.6 becomes 2.2.8.6

2.2.8.1.7 becomes 2.2.8.7

Delete section 2.2.8.2:

~~2.2.8.2.—Advanced Switching Support Message Group~~

~~The Messages in this group can be divided into the following two types:~~

~~□ Data Packet Messages (MsgASD—see Figure 2-19):~~

~~• Unicast, Data Packet~~

~~• Multicast, Data Packet~~

~~□ Signaling Packet Messages (MsgAS—see Figure 2-20):~~

~~• Signaling Packet, without interrupt~~

~~• Null signaling Packet, interrupt to Host in the destination Hierarchy~~

~~• Null signaling Packet, interrupt to destination device~~

~~• Signaling Packet, with interrupt to Host in the destination Hierarchy~~

• **Signaling Packet, with interrupt to destination device**

A detailed description of message types and message headers will be presented in a separate document entitled Advanced PCI Express Packet Switching Specification. This is a companion specification to the PCI Express Base Specification.




Figure 2-19: Request Header Format for MsgAS Request

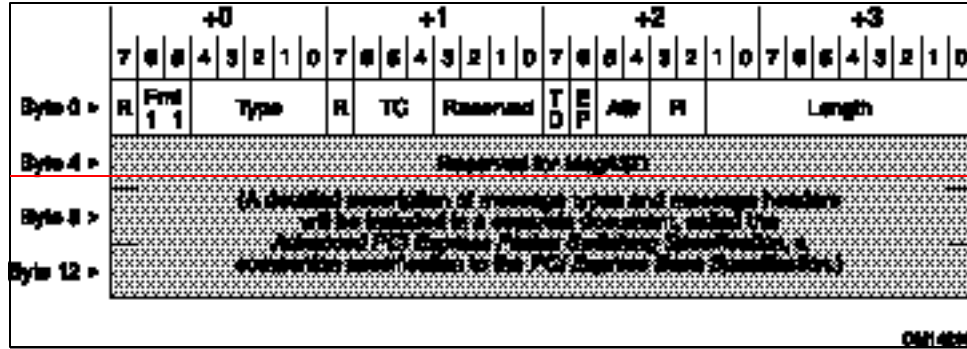


Figure 2-20: Request Header Format for MsgASD Request

❑ Components which are compliant to this specification but not to the Advanced PCI Express Packet Switching Specification must treat a received MsgAS or MsgASD Request as an Unsupported Request

• **This is a reported error associated with the Receiving Port (see Section 7.2)**

In 7.11.5, Table 7-39:

Table 7-39: VC Resource Capability Register

Bit Location	Description	Attribute
...		

Bit Location	Description	Attribute
14	<p>Advanced Packet Switching – Indicates the VC resource only supports transactions optimized for Advanced Packet Switching (AS). This field is valid for all PCI Express ports and RCRB.</p> <p>When this field is set to 0, it indicates that the VC resource is capable of supporting all transactions defined by this specification (including AS transport packets).</p> <p>When this field is set to 1, it indicates that the VC resource only supports transactions optimized for Advanced Packet Switching, and, therefore, must not be used for non-AS packet traffic.</p> <p>Detailed operation of this type of VC resource will be described in the PCI Express Advanced Packet Switching Specification.</p>	RO
...		

C50. Optional DLLP Timeout Mechanism

Release Date: 2/11/03

There is a requirement to send FC Updates at a minimum periodic frequency, implying that a check may/should be implemented. Such a check is not currently specified – this change clarifies how such a check could optionally be implemented.

In Section 2.6.1.2, add as shown:

- ❑ When the Link is in the L0 or L0s Link state, Update FCPs for each enabled type of non-infinite FC credit must be scheduled for transmission at least once every 30 μ s (-0%/+50%)
 - A timeout mechanism may optionally be implemented – If implemented, such a mechanism must:
 - ? be active only when the Link is in the L0 or L0s Link state
 - ? use a timer with a limit of 100 μ s (-0%/+50%), where the timer is reset by the receipt of any Init or Update FCP; Alternately, the timer may be reset by the receipt of any DLLP (See Section <3.4>)
 - ? upon timer expiration, instruct the Physical Layer to retrain the Link (via the LTSSM Recovery state, Section <4.2.6???)

NOTE: [see C53 for added text]

C51. Retracted

C52. Clarify Configuration Retry**Release Date:** 2/11/03*Clarify operation of Configuration Retry.**In Section 2.8, edit/add as shown:*

...

~~The~~ PCI Express ~~elements~~ ~~devices~~ that ~~are capable of initiating~~ ~~issue~~ Requests ~~that invoke requiring~~ Completions must implement ~~the~~ Completion Timeout mechanism. An exception is made for Configuration Requests (see below). ~~This~~ ~~The Completion Timeout~~ mechanism is activated for each Request that requires ~~one or more~~ Completions when the Request is transmitted. Since PCI Express Switches do not autonomously initiate Requests ~~(that need Completion)~~ ~~s~~, the requirement for Completion Timeout support is limited only to Root Complexes, PCI Express-PCI Bridges, and Endpoint devices.

The PCI Express Specification defines the following range for the min/max acceptable timer values for the Completion Timeout mechanism:

- ☐ The Completion Timeout timer must not expire (i.e., cause a timeout event) in less than 50 μ s. It is strongly recommend that unless an application requires this level of timer granularity the ~~the~~ minimum time should not expire in less than 10 ms
- ☐ The Completion Timeout timer must expire if a Request is not completed in 50 ms.

A Completion Timeout is a reported error associated with the Requestor device/function (see Section <6.2>).

Note: A Memory Read Request for which there are multiple Completions must be considered completed only when all Completions have been received by the Requester. If some, but not all, requested data is returned before the Completion Timeout timer expires, the Requestor is permitted to keep or to discard the data which was returned prior to timer expiration.

Configuration Requests have special requirements (see Sections <2.3.1> and <6.6>). Because of these special requirements, the support and timer values for a Completion Timeout for Configuration Requests are implementation specific. ~~Furthermore, PCI Express to PCI/PCI-X bridges, by default, are not enabled to return Configuration Request Retry Status (CRS) for Configuration Requests to a PCI/PCI-X device behind the bridge. This may result in lengthy completion delays that must be comprehended by the Completion Timeout value in the Root Complex. System software may enable PCI Express to PCI/PCI-X bridges to return Configuration Request Retry Status by setting the Bridge Configuration Retry Enable bit in the Device Control register, subject to the restrictions noted in the PCI Express to PCI/PCI-X Bridge specification.~~

C53. Assorted Errata**Release Date:** 3/27/03

Add bullet to 1.3.2.1, as shown:

- ❑ A Legacy Endpoint may implement extended configuration space capabilities, but such capabilities may be ignored by software

In 1.3.2.2, edit as shown:

- ❑ A PCI Express Endpoint must not ~~depend on operating system allocation of~~ require I/O resources claimed through BAR(s).

In 2.6.1., edit as shown:

...

- ❑ The unit of Flow Control credit is ~~4DW 16 Bytes~~ for Data

...

Table 2-26: TLP Flow Control Credit Consumption

TLP	Credit Consumed
...	
Memory Write Request	1 PH + n PD units ^[FOOTNOTE]
...	

[FOOTNOTE] For all cases where “n” appears, $n = \text{Roundup}(\text{LengthDataLen}/\text{FC unit size})$.

In 3.5.2.1, edit/add as shown:

...

- **REPLAY_TIMER** - Counts time since last Ack or Nak DLLP received
 - ? Started at the last symbol of any TLP transmission or retransmission, if not already running
 - ? For each replay, reset and restart REPLAY_TIMER when sending the last Symbol of the first TLP to be retransmitted
 - ? Restarts for each Ack DLLP received while there are unacknowledged TLPs outstanding, if, and only if, the received Ack DLLP acknowledges some TLP in the retry buffer

| Note: This ensures that REPLAY_TIMER is reset only when forward progress is being made

- ? Reset and hold until restart conditions are met for each Nak received (**except during a replay**) or when the REPLAY_TIMER expires
- ? **Not advanced during link retraining (holds value when LTSSM in Recovery <ref LTSSM Recovery state>)**
- ? Resets and holds when there are no outstanding unacknowledged TLPs

...

- If REPLAY_NUM rolls over from 11b to 00b, the Transmitter signals the Physical Layer to retrain the Link, **and waits for the completion of retraining before proceeding**. This is a reported error associated with the Port (see Section <6.2>).

...

Note: The following reverts a change made in Errata C27

- ? Reset REPLAY_NUM **and REPLAY_TIMER**

In 6.6, remove text inserted by C46:

- ❑ There is an in-band mechanism for propagating reset across a Link. This is called a hot reset and is described in Section 4.2.4.5. **~~Note that this mechanism is automatically triggered as a part of the Link initialization process that follows a cold or warm reset.~~**

[Modifying C50,] In Section 2.6.1.2, add as shown:

- ❑ When the Link is in the L0 or L0s Link state, Update FCPs for each enabled type of non-infinite FC credit must be scheduled for transmission at least once every 30 μ s (-0%/+50%), **except when the Extended Synch bit of the Link Control Register is set, in which case the limit is 120 μ s (-0%/+50%).**
 - A timeout mechanism may optionally be implemented – If implemented, such a mechanism must:
 - ? be active only when the Link is in the L0 or L0s Link state
 - ? use a timer with a limit of 200 μ s (-0%/+50%), **~~except when the Extended Synch bit of the Link Control Register is set, in which case the limit is 400 μ s (-0%/+50%).~~** where the timer is reset by the receipt of any Init or Update FCP; Alternately, the timer may be reset by the receipt of any DLLP (See Section <3.4>)
 - ? upon timer expiration, instruct the Physical Layer to retrain the Link (via the LTSSM Recovery state, Section <4.2.6???)

NOTE: The implementation of this optional mechanism is strongly encouraged. It is anticipated that future revisions of this specification may change this mechanism from optional to required.

In 3.5.3.1, edit as shown:

Tx_L0s_Adjustment	if L0s is enabled, the time required for the transmitter to exit L0s (see Section 4.2.6.6.2), expressed in Symbol Times, or 0 if L0s is not enabled
-------------------	---

Note that the setting of the Extended Synch bit of the Link Control register affects the exit time from L0s to L0, and must be taken into account in this adjustment

The values in Table 3-5 do not include this adjustment offset.

In 2.2.6.2:

- ☐ Prior to the initial Configuration Write to a device, the device is not permitted to initiate **Non-Posted** Requests.

Exception: Logical devices within a Root Complex are permitted to initiate Requests prior to software-initiated configuration for accesses to system boot device(s).

Note that this rule and the exception are consistent with the existing PCI model for system initialization and configuration.

[note: this section affected by earlier errata – included here] In Section 7.3.1, edit as shown:

7.3.1. Device Number

... Configuration Requests targeting the Bus Number associated with a link specifying Device Number 0 are delivered to the device attached to the link; Configuration Requests specifying all other Device Numbers (1-31) must be terminated by the Switch Downstream Port or the Root Port with an Unsupported Request Completion Status (equivalent to Master Abort in PCI). Devices must not assume that Device Number 0 is associated with their upstream port, but must capture their assigned Device Number as discussed in Section <2.2.6.2>. **Devices must respond to all Type 0 Configuration Read Requests, regardless of the Device Number specified in the Request.**

In 2.6.1, Insert as shown:

- ☐ During FC initialization for any Virtual Channel...

<Table>

- ❑ A Receiver must never cumulatively issue more than 2048 outstanding unused credits to the transmitter for data payload or 128 for header.
 - Components may optionally check for violations of this rule. If a component implementing this check determines a violation of this rule, the violation is a Flow Control Protocol Error (FCPE)
 - ? If checked, this is a reported error associated with the Receiving Port (see Section <6.2>)
- ❑ If an Infinite Credit advertisement (value of 00h or 000h) has been made during ...

In Section 3.5.2.1, pg 142:

- ❑ Copies of Transmitted TLPs must be stored in the Data Link Layer Retry Buffer, **except for nullified TLPs (see Section <3.5.2.1>)**

In 5.2, edit as shown:

- ❑ L0s – A low resume latency, energy saving “standby” state.

...

It is possible for the transmit side of one component on a Link to be in L0s while the transmit side of the other component on the Link is in L0.

- ❑ L1 – Higher latency, lower power “standby” state.

...

- ❑ L2/L3 Ready – Staging point **for L2 or L3 entry**~~removal of main power~~ **Lowest Power Link state, resume latency same as power on transition**

L2/L3 Ready transition protocol support is required

The L2/L3 Ready state is related to PCI-PM D-state transitions. L2/L3 Ready is ~~the~~ **a transitional pseudo-state (not corresponding directly to the state of the LTSSM)** that a given Link enters into when the platform is actively executing the process to remove power and clocks from the downstream component or both attached components. ~~preparing to enter its system sleep state. Following the completion of the L2/L3 Ready state transition protocol for that Link, the Link is then ready for either L2 or L3, but not actually in either of those states until main power has been removed.~~ This process is initiated after PM software transitions a device into a D3 state, and subsequently calls power management software to initiate the removal of power and clocks. After the link enters the L2/L3 Ready state the component(s) are ready for power removal. Depending upon the platform’s implementation choices with respect to providing a Vaux supply, after main power has been removed the Link will either settle into L2 (i.e., Vaux is provided), or it will settle into a zero power “off” state (see L3). **Note that these are PM pseudo-states for the Link; under these conditions, the LTSSM will in general operate only on main power, and so will power off with main power removal.**

The L2/L3 Ready state entry transition process must begin as soon as possible following the acknowledgment of a PM_Turn_Off message, (i.e., the injection of a PME_TO_Ack TLP). The Downstream component initiates L2/L3 Ready entry by injecting a PM_Enter_L23 DLLP onto its transmit Port. Refer to Section <5.6> for further detail on power management system messages.

TLP and DLLP communication over a Link that is in L2/L3 Ready is ~~not possible~~ prohibited.

~~Exit from L2/L3 Ready back to L0 may only be initiated by an upstream initiated transaction targeting the Downstream component in the same manner that an upstream initiated transaction would trigger the transition from L1 back to L0. The case where an upstream initiated exit from L2/L3 Ready would occur corresponds to the scenario where, sometime following the transition of the Link to L2/L3 Ready but before main power is removed, the platform power manager decides not to enter the system sleep state. Once the PME_Turn_Off/PME_TO_Ack sequence is initiated, the downstream component is guaranteed that a Fundamental Reset will occur (with or without the removal of power). Exit from L2/L3 Ready back to L0 must only be initiated through the power controller, and must follow entry into D3_{cold} or Fundamental Reset.~~

~~A Link's transition into the L2/L3 Ready state is one of the final stages involving PCI Express protocol for transitioning a device into a D3 state in preparation for an unused/low power device state, or as a prelude to a system or subsystem "Sleep" state.~~

❑ L2 – Auxiliary powered Link deep energy saving state.

L2 support is optional, and dependent upon platform support of Vaux.

L2 – The ~~Downstream~~ component's main power supply inputs and reference clock inputs are shut off.

- When in L2 (with main power removed), any Link reactivation wakeup logic (Beacon or WAKE#), PME context, and any other "keep alive" logic is powered by Vaux.

TLP and DLLP communication over a Link that is in L2 is not possible.

~~Exiting the L2 state is accomplished by reestablishing main power and reference clocks to all components within the domain of the power manager, followed by full Link training and initialization. Once a given Link has completed Link training and initialization it is then in the L0 state and may begin sending and receiving TLPs and DLLPs.~~

❑ L3 – Link Off state.

Zero power state.

Refer to Section <4.2> for further detail relating to entering and exiting each of the PCI Express L-states between L0 and L2/L3 Ready (L2.Idle from the chapter 4 perspective). The L2 state is an abstraction for PM purposes distinguished by the presence of auxiliary power, and should not be construed to imply a requirement that the LTSSM remains active. The L3 state is not defined or discussed in the electrical section (as the component is completely un-powered) with the exception of defining un-powered electrical driver and receiver specifications.

<Figure 5-1> highlights the legitimate L-state transitions that may occur during the course of Link operation.




Figure 0-5: Link Power Management State Transitions

Note that these states and state transitions do not correspond directly to the actions of the Physical Layer LTSSM. For example, the LTSSM is typically powered by main power (not Vaux), and so in both the L2 and L3 states will be unpowered. The arc noted in <Figure 5-1> indicates the case where the platform does not provide Vaux. In this case, the L2/L3 Ready state transition protocol results in a state of readiness for loss of main power, and once main power is removed the Link settles into the L3 state.

...

In 5.3.1.4.2, edit as shown:

5.3.1.4.2 D3cold State

A function transitions to the D3_{cold} state when its power is removed. A power-on sequence transitions a function from the D3_{cold} state to the D0_{Uninitialized} state. At this point software must perform a full initialization of the function in order to re-establish all functional context, completing the restoration of the function to its D0_{active} state. Functions that support wakeup functionality from D3_{cold} must maintain their PME context (in the Power Management Status/Control Register) for inspection by PME service routine software during the course of the resume process. ~~Functions that support PME assertion from D3_{cold} must maintain their PME context for inspection by PME service routine software during the course of the resume process. Functions may only generate PME messages from D3_{cold} if the platform supplies them with a Vaux supply or if they have an independent source of power.³¹ PME context consists of all information relating to the function's assertion of PME.~~

In Section 5.3.2.3, add as shown:

- ❑ L2/L3 Ready entry transition protocol uses the PM_Enter_L23 DLLP.

~~Note that the PM_Enter_L23 DLLPs are sent continuously until an acknowledgement is received. The L1 entry protocol uses the PM_Enter_L1 DLLP.~~

In 5.3.3.2, edit as shown:

5.3.3.2. Link Wakeup

... Refer to the ~~PCI Express Card~~ ~~Relevant~~ electromechanical ~~Spec~~ ~~ification~~ for details on the out-of-band mechanism ...

Edit figures as indicated below each caption:

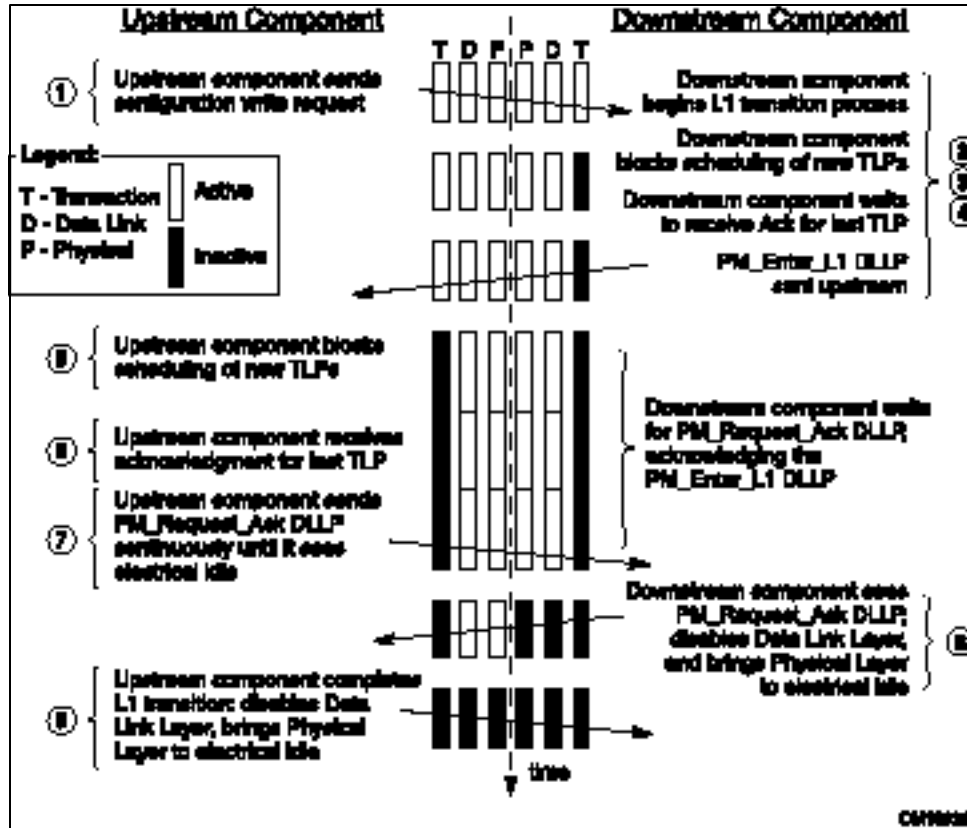


Figure 5-3: Entry into L1 Link State

Fig 5-2: step 4 - change from "PM_Enter_L1 DLLP sent upstream" to "PM_Enter_L1 DLLPs sent continuously"

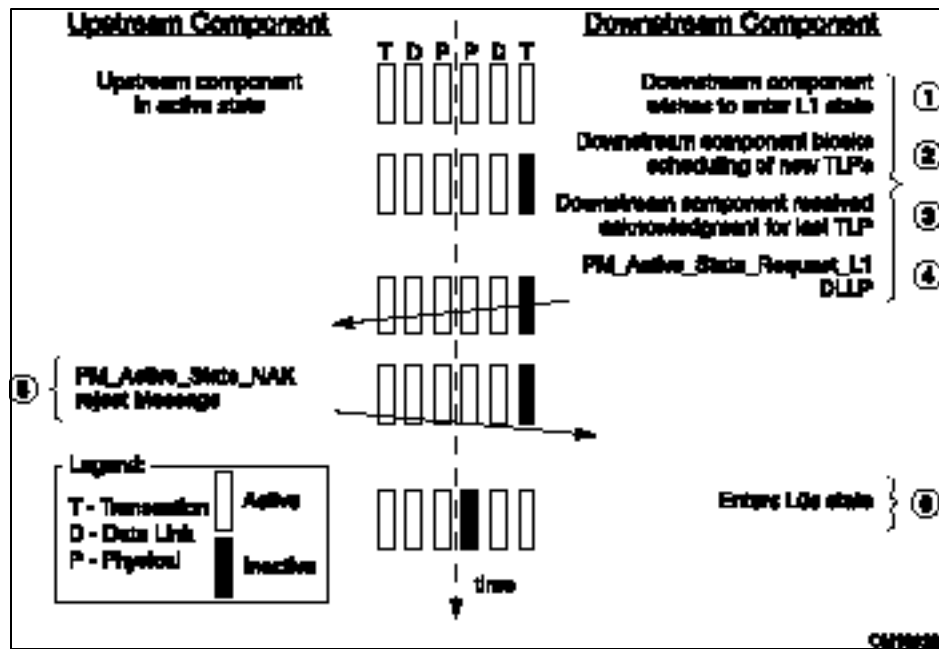


Figure 5-5: L1 Transition Sequence Ending with a Rejection (L0s Enabled)

Fig 5-5: step 4 - change from "PM_Active_State_Request_L1 DLLP" to "PM_Active_State_Request_L1 DLLPs sent continuously"

Fig 5-5: between step 4 and 6, add comment similar to Fig 5-2: "Component waits for a response to the PM_Active_State_Request_L1 DLLPs"

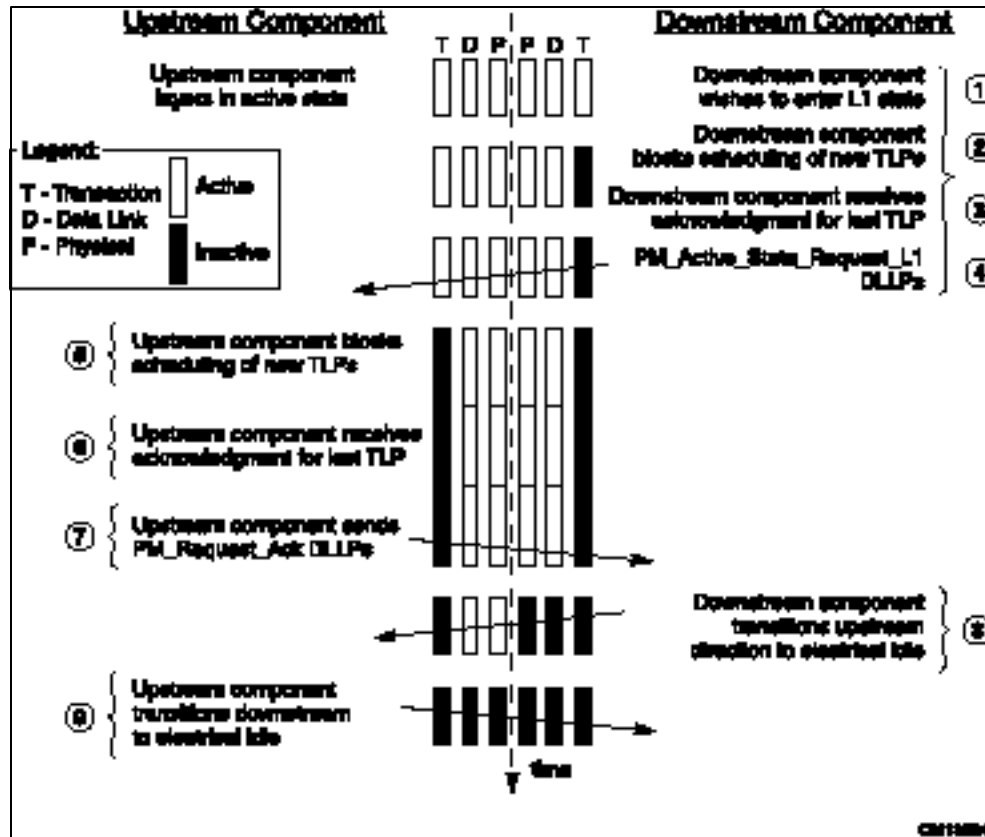


Figure 5-6: L1 Successful Transition Sequence

Fig 5-6: step 4 - change from "PM_Active_State_Request_L1 DLLPs" to "PM_Active_State_Request_L1 DLLPs sent continuously"

Fig 5-6: between step 4 and 6, add comment similar to Fig 5-2: "Component waits for a response to the PM_Active_State_Request_L1 DLLPs"

Fig 5-6: step 7 - change "Upstream component sends ..." to "Upstream component sends PM_Request_Ack DLLPs continuously"

In 5.3.3.2, add as shown:

... For components that support wakeup functionality, Beacon is the required mechanism for all components, except for components designed exclusively for the following form factors: PCI Express Card Electromechanical Specification, and PCI Express Mini- Card Electromechanical Specification. **Switch products targeting applications where Beacon is used on some ports of the Switch and WAKE# is used for other ports must translate the wakeup mechanism appropriately (see Implementation Note "Example of WAKE# to Beacon Translation"). In applications where WAKE# is the only wakeup mechanism used, the Root Complex is not required to support the receipt of Beacon.**

In Section 5.3.3.2.1, edit as shown:

PME_Turn_Off Broadcast Message

Before main component power and reference clocks are turned off the Root Complex or ~~Hot Plug controller within a~~ Switch Downstream Port, must issue a broadcast message that instructs all agents downstream of that point within the hierarchy to cease initiation of any subsequent PM_PME messages, effective immediately upon receipt of the PME_Turn_Off message.

Each PCI Express agent is required to respond with a TLP “acknowledgement” Packet, PME_TO_Ack that is, as in the case of a PME Message, always routed upstream. In all cases, the PME_TO_Ack message must terminate at the PME_Turn_Off message’s point of origin.³⁵

Footnotes:

35 Point of origin for the PME_Turn_Off message could be all of the Root Ports for a given Root Complex (full platform sleep state transition), an individual ~~hot plug capable~~ Root Port, or a ~~hot plug capable~~ Switch Downstream Port.

37 Power delivery control within this context relates to control over the entire PCI Express Link hierarchy, or over a subset of PCI Express links ranging down to a single PCI Express Link ~~and associated endpoint~~ for sub hierarchies ~~supporting independently managed power and clock distribution residing downstream of a Hot Plug controller managed interconnect.~~

In Section 5.4.1, on first occurrence- Active State Power Management (ASPM)

Use “ASPM” in the rest of the section.

Further down in 5.4.1, edit:

The L0s exit latency may differ significantly if the reference clock for opposing sides of a given Link is provided from the same source, or delivered to each component from a different source. PCI Express-PM software informs each PCI Express device of its clock configuration via the “common clock configuration” bit in their PCI Express Capability Structure’s Link Control Register. This bit serves as the determining factor in the L0s exit latency value reported by the device. ~~ASPM may be enabled or disabled by default depending on implementation specific criteria and/or the requirements of the associated form-factor specification(s). All PCI Express devices power on with Active State Link Power Management turned off by default.~~ Software can enable or disable ASPM~~active state Link power management~~ using a process described in Section 5.4.1.3.1.

In 5.4.1.3.1, edit:

5.4.1.3.1. Software Flow for Enabling ~~or Disabling~~ Active State Link Power Management

Following is an example software algorithm that highlights how to enable ~~or disable~~ active state Link power management in a PCI Express component.

...

- PCI Express software updates the “Common Clock Configuration” bits in the components on both ends of each Link to indicate if those devices share the same

reference clock, and triggers Link retraining by writing '1' to the Retrain Link bit in the Link Control register of one of the components.

...

In Table 6-4:

Completer Abort		<i>Completer:</i> Send ERR_NONFATAL to Root Complex. Log the header of the Request Completion that encountered the error.
-----------------	--	---

in 2.9.1, edit as shown:

...

For a Root Complex, or any Port on a Switch other than the one closest to the Root Complex, DL_Down status is handled by:

- ☐ ~~returning all PCI Express-specific registers, state machines and externally observable state to the specified default or initial conditions (except for registers defined as sticky—see Section 0)~~
- ☐ initializeing back to their default state any buffers or internal state associated with outstanding requests transmitted downstream
 - ? Note: Port configuration registers must not be affected, except as required to update status associated with the transition to DL_Down

In 7.5.3.3, add as shown:

Table 0-5: Bridge Control Register

Bit Location	Register Description	Attributes
...		
6	Secondary Bus Reset – Setting this bit triggers a hot reset on the corresponding PCI Express Port. Port configuration registers must not be affected, except as required to update Port status. Default value of this field is 0.	RW

Add text following Table 6-4 (as shown):

Table 6-4: Transaction Layer Error List

Error Name	Severity	Detecting Agent Action
...		Error! Bookmark not defined.

For all errors listed above, the appropriate status bit(s) must be set upon detection of the error. For Unsupported Request (UR), the logging and signaling control/status mechanisms are more complicated than for the other types of errors, and so are clarified below. Note that upon receiving an Non-Posted (NP) Request that is UR, the return of a Completion with UR status is unconditional, and the reporting settings in the Device Control (and Uncorrectable Error Mask register, if Advanced Error Reporting (AER) is implemented) register(s) do not affect the return of this completion.

In the Device Status Register:

- ☐ Set Unsupported Request Detected any time a received Request is determined to be UR
- ☐ Set Non-Fatal Error Detected any time a received Request is determined to be UR, and either AER is not implemented or AER Unsupported Request Error Severity is set to Non-Fatal
- ☐ Set Fatal Error Detected any time a received Request is determined to be UR, and AER is implemented, and AER Unsupported Request Error Severity is set to Fatal

UR is also reported through the AER structure, if implemented.

If the Unsupported Request Reporting Enable bit in the Device Control Register is not set, then the device must not send an error message in the case of a UR, although the device will set status bits as described above (regardless of AER settings, if AER implemented).

If AER is not implemented, and if the Unsupported Request Reporting Enable bit in the Device Control Register is set, then an ERR_NONFATAL error message will be sent to the Root Complex (or signaled within the Root Complex, in the case of a Root Port) on detection of a UR.

If AER is implemented and the Unsupported Request Reporting Enable bit in the Device Control Register is set, then an error message is sent if the Unsupported Request Error Mask bit is not set. The severity of the error message is determined by the setting of the Unsupported Request Error Severity bit.

---**From JoeC's collected text email**---

Page 32 summarizes requirements for Legacy and non-legacy Endpoints. Under "1.3.2.1 Legacy Endpoint Rules", add the following bullets to the end of the list:

- ☐ A Legacy Endpoint operating as the Requester of a Memory Transaction is not required to be capable of generating addresses 4 GB or greater.

- ❑ A Legacy Endpoint is required to support MSI if an interrupt resource is requested, but is permitted to support either the 32-bit or 64-bit Message Address version of the MSI capability structure.
- ❑ A Legacy Endpoint is permitted to support 32-bit addressing for Base Address registers that request memory resources.

Under "1.3.2.2 PCI Express Endpoint Rules", add the following bullets to the end of the list:

- ❑ A PCI Express Endpoint operating as the Requester of a Memory Transaction is required to be capable of generating addresses greater than 4 GB.
- ❑ A PCI Express Endpoint is required to support MSI if an interrupt resource is requested, and must support the 64-bit Message Address version of the MSI capability structure.
- ❑ A PCI Express Endpoint requesting memory resources through a BAR must set the BAR's Prefetchable bit unless the range contains locations with read side-effects or locations in which the device does not tolerate write merging.
- ❑ For a PCI Express Endpoint, 64b addressing must be supported for all BARs that have the prefetchable bit set. 32b addressing is permitted for all BARs that do not have the prefetchable bit set.
- ❑ The minimum memory address range requested by a BAR is 128 bytes.

Add implementation note in section 2.2.7 "Memory, I/O, and Configuration Request Rules" following Figure 2-14 on page 63:

IMPLEMENTATION NOTE: Generation of 64b Addresses

It is strongly recommended that PCI Express Endpoints be capable of generating the full range of 64-bit addresses. However, if a PCI Express Endpoint supports a smaller address range, and is unable to reach the full address range required by a given platform environment, the corresponding Device Driver must ensure that all Memory Transaction target buffers fall within the address range supported by the Endpoint. The exact means of ensuring this is platform and OS specific, and beyond the scope of this specification.

Page 260 contains the requirements for MSI. Insert the following paragraph following the existing second paragraph:

A Legacy Endpoint is required to support either the 32-bit or 64-bit Message Address version of the MSI capability structure. A PCI Express Endpoint is required to support the 64-bit Message Address version of the MSI capability structure.

Page 325 shows the Type 0 Configuration Space Header details that are unique in Express for devices. Add a new sub-section within section 7.5.2 prior to 7.5.2.1:

7.5.2.x Base Address Registers (Offset 10h - 24h)

A PCI Express Endpoint requesting memory resources through a BAR must set the BAR's Prefetchable bit unless the range contains locations with read side-effects or locations in which the device does not tolerate write merging. It is strongly encouraged that memory-mapped resources be designed as prefetchable whenever possible. PCI Express devices other than Legacy Endpoints must support 64-bit addressing for any Base Address register that requests prefetchable memory resources. The minimum memory address range requested by a BAR is 128 bytes.

Page 326 shows the Type 1 Configuration Space Header details that are unique in Express for virtual bridges. Add a new sub-section within section 7.5.3 prior to 7.5.3.1:

7.5.3.x Base Address Registers (Offset 10h/14h)

A PCI Express Endpoint requesting memory resources through a BAR must set the BAR's Prefetchable bit unless the range contains locations with read side-effects or locations in which the device does not tolerate write merging. It is strongly encouraged that memory-mapped resources be designed as prefetchable whenever possible. PCI Express devices other than Legacy Endpoints must support 64-bit addressing for any Base Address register that requests prefetchable memory resources. The minimum memory address range requested by a BAR is 128 bytes.

Add a new subsection between existing subsections 7.5.3.2 and 7.5.3.3:

7.5.3.y. Prefetchable Memory Base/Limit (Offset 24h)

The Prefetchable Memory Base and Prefetchable Memory Limit registers must indicate that 64-bit addresses are supported, as defined in the PCI Bridge 1.1 specification.

---From Suparna's Ch5 and Ch7 emails --

In section 5.4.1.2.1 pg 248. 249, Add following lines in sections below:

pg 248: Active State Link PM L1 Negotiation Rules:

2nd paragraph:

.The downstream component must wait until it receives a Link layer acknowledgement for the last TLP it had previously sent (**i.e. the retry buffer is empty**). The component may retransmit a TLP if required by the Data Link layer rules

5th paragraph:

.During this waiting period, the downstream component must not initiate any Transaction Layer transfers. It must still accept TLPs and DLLPs from the upstream component. It also responds with DLLPs as needed by the Link layer protocol. **All received TLP's must be acknowledged before entering the L1 state(i.e., DLLP ACK/NAK Latency timers are reset before entering the L1 state).**

In Section 5.4.1.2.1 Entry to L1 state, Implementation Note 2nd paragraph (pg 247) change:

"One possible approach would be for the downstream device to initiate a transition to the L1 state **once the device has both its Receiver and Transmitter in L0s state (RxL0s and TxL0s)** for a set amount of time."

In Section 5.3.2.1 pg 233, change:

disable Link Layerdisable DLLP, TLP transmission

In Section 5.4.1.2.1, p 249, change:

disable Link Layerdisable DLLP, TLP transmission

In Section 5.2 page 224 add note:

Exit from L2/L3 ready cannot exit directly to L0 state. As per the PCI Express Base Specification, Revision 1.0 Errata, L2/L3 ready exit is through detect state. **Note: Exit from L2/L3 ready back to L0 will be through intermediate LTSSM states. Refer to chapter 4 for details.**

In Section 7.8.7 add:

This mode provides external devices monitoring the link **(e.g., logic analyzers)**, time to achieve bit and symbol lock before the link enters L0 ...

In Section 7.8.7 Active state Link PM Control, add the following at end of Active State Link PM Control in pg 347

Note: "L0s entry supported" indicates the transmitter entering L0s is supported. Receiver must be capable of entering L0s even when the field is disabled (00b) .

Default value of this field is (00b) or (01b) depending on form factor.

In Table 7-15, edit as shown:

Table 7-15: Link Control Register

Bit Location	Register Description	Attributes
--------------	----------------------	------------

Bit Location	Register Description	Attributes
...		
5	<p>Retrain Link – A write of 1b to this bit initiates Link retraining by directing the Physical Layer LTSSM to the Recovery state. The LTSSM must enter the Recovery state before the Completion is returned for the Write Request to this bit. Reads of this bit always return 0b.</p> <p>This field is not applicable and reserved for endpoint devices and Upstream Ports of Switches.</p> <p>This bit always returns 0b when read.</p>	RW
...		

In Table 7-16, edit as shown:

Table 7-16: Link Status Register

Bit Location	Register Description	Attributes
...		
11	<p>Link Training – This read-only bit indicates that Link training is in progress (Physical Layer LTSSM in Configuration or Recovery state), or that 1b was written to the Retrain Link bit, but Link training has not yet begun; hardware clears this bit once Link training is complete.</p> <p>This field is not applicable and reserved for endpoint devices and Upstream Ports of Switches.</p>	RO
...		

In 2.6.1.2, in implementation note “Flow Control Update Frequency”, add as shown:

... <Table 2-28> shows recommended values for the frequency of transmission **during normal operation** based on Link Width and Max_Payload_Size values.

...

<Table 0-6: UpdateFC Transmission Latency Guidelines by Link Width and Max Payload (Symbol Times)>

Note that the values given in this table do not account for any delays caused by the receiver or transmitter being in L0s. For improved performance and/or power-saving, it may be desirable to use a Flow Control update policy that is more sophisticated than a simple timer. Any such policy is implementation specific, and beyond the scope of this document.

---Allow L1 enabled w/o L0s enabled

In 5.4.1.2.1, “Entry to L1 State”, edit as shown:

... If the endpoint receives a negative acknowledgement in response to its issuance of a PM_Active_State_Request_L1 DLLP, then the endpoint must enter the L0s state as soon as possible [FOOTNOTE: Assuming **L0s is enabled and** that the conditions for L0s entry are met]. Note that the component on the upstream side of the Link may not support L1 Active State Link PM, or it may be disabled and so could legitimately respond to the L1 entry request with a negative acknowledgement.

...

If the Switch’s Upstream Port receives a negative acknowledgement in response to its issuance of a PM_Active_State_Request_L1 DLLP, then the Switch’s Upstream Port transmit Lanes must instead transition to the L0s state as soon as possible [FOOTNOTE: Assuming **L0s is enabled and** that the conditions for L0s entry are met].

...

- ❑ If the request was rejected, the Downstream component must immediately transition its transmit Lanes into the L0s state, provided **L0s is enabled and** that conditions for L0s entry are met.

...

In Section 5.4.1.3 Active State Link PM Configuration, Table 5-10, change:

Table 5-10: Encoding of the Active State Link PM Control Field

Field	Description
Active State Link PM Control	00b – Disabled
	01b – L0s Entry Enabled
	10b – L1 Entry Enabled Reserved
	11b – L0s and L1 Entry enabled

and add:

Active State Link PM Control = 10

Port’s transmitter must not enter L0s

Ports connected to the Downstream end of the Link may issue PM_Active_State_Request_L1 DLLPs.

Ports connected to the Upstream end of the Link must respond with positive acknowledgement to L1 request and transition into L1 if conditions for Root Complex Root Port or Switch Downstream Port in Section <5.4.1.2.1> are met.

In 7.8.6, edit as shown (also make corresponding change in Section 5.4.1.3, Table 5-6):

Table 7-14: Link Capabilities Register

Bit Location	Register Description	Attributes																
...																		
14:12	<p>L0s Exit Latency – This field indicates the L0s exit latency for the given PCI Express Link. The value reported indicates the length of time this Port requires to complete transition from L0s to L0. Defined encodings are:</p> <table><tr><td>000b</td><td>Less than 64 ns</td></tr><tr><td>001b</td><td>64 ns to less than 128 ns</td></tr><tr><td>010b</td><td>128 ns to less than 256 ns</td></tr><tr><td>011b</td><td>256 ns to less than 512 ns</td></tr><tr><td>100b</td><td>512 ns to less than 1 μs</td></tr><tr><td>101b</td><td>1 μs to less than 2 μs</td></tr><tr><td>110b</td><td>2 μs-4 μs</td></tr><tr><td>111b</td><td>ReservedMore than 4 μs</td></tr></table> <p>Note that exit latencies may be influenced by PCI Express reference clock configuration depending upon whether a component uses a common or separate reference clock.</p>	000b	Less than 64 ns	001b	64 ns to less than 128 ns	010b	128 ns to less than 256 ns	011b	256 ns to less than 512 ns	100b	512 ns to less than 1 μ s	101b	1 μ s to less than 2 μ s	110b	2 μ s-4 μ s	111b	Reserved More than 4 μ s	RO
000b	Less than 64 ns																	
001b	64 ns to less than 128 ns																	
010b	128 ns to less than 256 ns																	
011b	256 ns to less than 512 ns																	
100b	512 ns to less than 1 μ s																	
101b	1 μ s to less than 2 μ s																	
110b	2 μ s-4 μ s																	
111b	Reserved More than 4 μ s																	
...																		

In Section 7.8.7, *Link Control Register*, Table 7-15, change:

Table 7-15: Link Control Register

Bit Location	Register Description	Attributes								
1:0	<p>Active State Link PM Control – This field controls the level of active state PM supported on the given PCI Express Link. Defined encodings are:</p> <table><tr><td>00b</td><td>Disabled</td></tr><tr><td>01b</td><td>L0s Entry EnabledSupported</td></tr><tr><td>10b</td><td>L1 Entry EnabledReserved</td></tr><tr><td>11b</td><td>L0s and L1 Entry EnabledSupported</td></tr></table>	00b	Disabled	01b	L0s Entry Enabled Supported	10b	L1 Entry Enabled Reserved	11b	L0s and L1 Entry Enabled Supported	RW
00b	Disabled									
01b	L0s Entry Enabled Supported									
10b	L1 Entry Enabled Reserved									
11b	L0s and L1 Entry Enabled Supported									

[At the request of the EWG] At the end of Section 6.6, add the following implementation note:

Implementation Note: Use of PERST#

The deassertion of PERST# may not be sufficient in all implementations to say that “fundamental reset” is exited. Such implementations may require other additional conditions to be met, for example the IO PLL locking, before exiting “fundamental reset.”

C54. Non-Infinite FC permission for RC

Release Date: 3/27/03

In Section 2.6.1 Flow Control Rules, Table 2-27:

Entry: CPLH: Change

~~“Switch and PCI Express to PCI X Bridge (PCI X mode only): 1 FC unit - credit value of 01h”~~

~~“Root Complex, Endpoint, and PCI Express to PCI Bridge: infinite FC units - initial credit value of all 0s”~~

To

~~“Root Complex (supporting peer to peer traffic between all Root Ports) and Switch, and PCI Express to PCI X Bridge (PCI X mode only): 1 FC unit - credit value of 01h”~~

~~“Root Complex (not supporting peer to peer traffic between all Root Ports) and Endpoint and PCI Express to PCI Bridge: infinite FC units - initial credit value of all 0s”~~

Entry: CPLD: Change

~~“Switch and PCI Express to PCI X Bridge (PCI X mode only): Largest possible”~~

~~“Root Complex, Endpoint, and PCI Express to PCI Bridge: infinite FC units - initial credit value of all 0s”~~

To

~~“Root Complex (supporting peer to peer traffic between all Root Ports) and Switch, and PCI Express to PCI X Bridge (PCI X mode only): Largest possible”~~

~~“Root Complex (not supporting peer to peer traffic between all Root Ports) and Endpoint and PCI Express to PCI Bridge: infinite FC units - initial credit value of all 0s”~~

In Section 2.6.1 Flow Control Rules.

After Table 2-27 add the following bullets:

- o A Root Complex that does not support peer to peer traffic between all Root Ports must advertise infinite Completion credits.

o A Root Complex that supports peer to peer traffic between all Root Ports may optionally advertise non-infinite Completion credits. In this case, the Root Complex must ensure that deadlocks are avoided and forward progress is maintained for completions directed towards the RC. Note that temporary stalls of completion traffic (due to a temporary lack of credit) are possible since Non-Posted requests forwarded by the RC may not have explicitly allocated completion buffer space.

C55. Slot/Common clock config clarification

Release Date: 3/27/03

Add the following implementation note in Chapter 7 at the end of Section 7.8.7:

Implementation Note: Use of the Slot Clock Configuration and Common Clock Configuration bits

In order to accurately determine the common clocking configuration of components on opposite ends of a link that crosses a connector there are two pieces of information that are required. The following description defines these requirements.

The first necessary piece of information is whether the port that connects to the system side of the connector uses a clock that has a common source and therefore constant phase relationship to the clock signal provided on the connector. This information is provided by the system side component through a hardware initialized bit (Slot Clock Configuration) in its Link Status register. Note that some electromechanical form-factor specifications may require the port that connects to the system side of the connector use a clock that has a common source to the clock signal provided on the connector.

The second necessary piece of information is whether the device on the option card uses the clock supplied on the connector or one generated locally on the card. It is the card design and layout that will determine whether or not the endpoint component is connected to the clock source provided by the slot connector. An endpoint going onto this card should have some hardware initialized method for the card design/designer to indicate the configuration used for this particular card design. This information is reported by the endpoint device in bit 12 (Slot Clock Configuration) of its Link Status register. Note that some electromechanical form-factor specifications may require the port that connects to the card side of the connector use the clock signal provided on the connector.

System firmware or software will read this value from the components on both ends of a physical link. If both devices report the use of a common clock connection this firmware/software will program bit 6 (Common Clock Configuration) of the Link Control register to a one on both components connected to the Link. Each component uses this bit to determine the length of time required to re-synch its receiver to the opposing component's transmitter when exiting L0s.

This value is reported as a time value in bits 12-14 of the Link Capabilities register (offset 0Ch) and is sent to the opposing transmitter as part of the initialization process as N_FTS. Components would be expected to require much longer synch times without common clocking and would therefore report a longer L0s exit latency in bits 12-14 of the Link Capabilities register and would send a larger number for N_FTS during training. This forces a requirement that whatever software changes this bit should force a link retrain in order to get the correct N_FTS set for the receivers at both ends of the link.

C56. L0s Invocation and Exit Policy

Release Date: 3/27/03

In 5.4.1.1.1. edit:

L0s Invocation Policy

PCI Express Ports that are enabled for L0s entry must transition their transmit Lanes to the L0s state if the defined idle conditions (below) are met for a period of time not to exceed 7 μ s. Within this time period, the policy used by the Port to determine when to enter L0s is implementation specific. ~~are met for a specified period of time. The Port may choose this period of time to be anywhere within the range of:~~

~~(opposing Port's reported L0s exit latency)/4 ? t ? opposing Port's reported L0s exit latency~~

~~Defining the invocation time as a range enables the tuning of Active State Power Management behavior, balancing power savings with performance.~~

Definition of Idle

The definition of "idle" varies with device category

An Endpoint Port, ~~or Root Complex~~ Root Port ~~or Switch Downstream Port~~ is determined to be idle if the following conditions are met:

- ☐ No TLP is pending to transmit over the Link, or no FC credits are available to transmit ~~any TLPs~~ anything

<addition previously located here moved to 5.4.1.1.2 (below)>

- ☐ No DLLPs are pending for transmission

A Switch's Upstream Port is determined to be idle if the following conditions are met:

- ☐ All of the Switch's Downstream Port receive Lanes are in the L0s state
- ☐ No pending TLPs to transmit, or no FC credits are available to transmit anything
- ☐ No DLLPs are pending for transmission

<the following had been marked deleted – restore (i.e. do not change existing published text)>

A Switch's Downstream Port is determined to be idle if the following conditions are met:

- ☐ The Switch's Upstream Port's receive Lanes are in the L0s state
- ☐ No pending TLPs to transmit on this Link, or no FC credits are available
- ☐ No DLLPs are pending for transmission

See Section <4.2> for details on L0s entry by the Physical Layer.

5.4.1.1.2. Exit from L0s State

~~Components from either end of a PCI Express Link may initiate an exit from the L0s low power Link state.~~

A component with its transmitter in L0s must initiate L0s exit when it has a TLP or DLLP to transmit across the Link. Note that a transition from the L0s Link state ~~does not should never~~ depend on the status (or availability) of FC credits. The Link must be able to reach the L0 Link Active state, and to exchange FC credits across the Link. For example, if all credits of some type were consumed when the Link entered L0s, then any component on either side of the Link must still be able to transition the Link to the L0 state when new credits need to be sent across the Link. ~~Note that it may be appropriate for a component to anticipate the end of the idle condition and initiate L0s transmit exit, for example, when a NP request is received.~~

...

C57. Aux Power for Beacon

Release Date: 3/27/03

In 5.2, description of L2 state - add as shown:

- ❑ L2 – Auxiliary powered Link deep energy saving state.

L2 support is optional, and dependent upon platform support of Vaux. A component may only consume Vaux power if enabled to do so as described in Section 5.5.1.

In Section 5.3.3.2, “Link Wakeup”:

...

Systems that allow PME generation from D3cold state must provide auxiliary power to support Link wakeup when the main system power rails are off. A component may only consume auxiliary power if software has enabled it to do so as described in Section 5.5.1. Software is required to enable auxiliary power consumption in all components that participate in Link wakeup, including all components that must propagate the Beacon signal. In the presence of legacy system software, this is the responsibility of system firmware.

At the end of Section 5.5.1 (“5.5.1. Auxiliary Power Enabling”), add:

Typically, Aux Power is required to support the Beacon wakeup mechanism, and components supporting Beacon must not consume Aux Power unless enabled by software to consume Aux Power. To enable finer granularity of power consumption, Root Ports and Switch Ports should support independent control of Aux Power consumption for each port. However, it is permitted to logically OR the Aux Power enables for multiple ports and combine the Aux Power control for those ports so long as the combined Aux Power consumption does not exceed the sum of the amounts enabled by software.

C58. Legacy Endpoint Requirements

Release Date: 3/27/03

In Table 7-10 (PCI Express Capabilities Register), Bit Location 7:4 (Device/Port Type), change last paragraph from:

Legacy PCI Express Endpoint devices are not allowed to implement extended configuration space capabilities; PCI Express extended configuration capabilities for such devices are ignored by software.

To:

Extended configuration space capabilities, if implemented on Legacy PCI Express Endpoint devices, may be ignored by software.

C59. Integrated device

Release Date: 3/27/03

Group 1: Changes required to clarify integrated device limitations w/o adding support for RC integrated devices.

1.3.2.1, p.32, add to bulleted list

- A Legacy Endpoint must appear within one of the Hierarchy Domains originated by the Root Complex.

1.3.2.2, p.32, add to bulleted list

- A PCI Express Endpoint must appear within one of the Hierarchy Domains originated by the Root Complex.

1.3.3, p.33, add to bulleted list - Endpoint devices (represented by Type 00h Configuration Space headers) may not appear to configuration software on the switch's internal bus as peers of the Virtual PCI-to-PCI Bridges representing the Switch Downstream Ports.

7.1, p.311, add at end of 2nd paragraph - Only the PCI-PCI Bridges representing the Switch Downstream Ports may appear on the internal bus. Endpoints, represented by Type 0 configuration space headers, may not appear on the internal bus.

7.3.1, p.315, 2nd paragraph - Switches, and components wishing to incorporate more than eight functions at their upstream Port, may implement one or more "virtual switches," represented by multiple Type 1 (PCI-to-PCI Bridge) configuration space headers as illustrated in Figure 7-2. These virtual switches serve to allow fan-out beyond eight functions. Since switch downstream ports may appear on any device number, in this case all address information fields (bus, device and function numbers) must be completely decoded to access the correct register. Any configuration access targeting an unimplemented bus, device or function must return a Completion with Unsupported Request Completion Status.

C60 Receiver FTS Timeout

Release Date: 3/27/03

In section 4.2.6.6.1.3. (Rx_L0s.FTSRX_L0s.FTS)

change line from:

"Note: The N_FTS timeout is approximately $40 \times [N_FTS + 1 \text{ (SKP)}] \times \text{UI}$. This time would be exact except for that SKPs may be longer than a 4 symbol ordered set."

To:

"Note: The N_FTS timeout shall be no shorter than $40 \times [N_FTS + 3] \times \text{UI}$ (The $3 \times 40 \text{ UI}$ is derived from 6 symbols to cover a maximum SKP ordered set + 4 symbols for a possible extra FTS ordered set + 2 symbols of design margin), and no longer than twice this amount. When the

extended synch bit is set the receiver N_FTS timeout must be adjusted to no shorter than 40* [2048] * UI (2048 FTS ordered sets) and no longer than 40* [4096] * UI (4096 FTS ordered sets)."

C61 - Add the terms TTX-IDLE-MIN and TTX-IDLE-SET-TO-IDLE to the L2.Idle state.

Release Date: 3/27/03

In section 4.2.6.8.1

Change line from:

☐ All configured transmitters are in Electrical Idle.

- Note: The DC common mode voltage does not have to be within specification.⁸

To (changes in red):

☐ All configured transmitters are in Electrical Idle for a minimum time of $T_{TX-IDLE-MIN}$ (Table 0-4).

- Note: The DC common mode voltage must be within specification by $T_{TX-IDLE-SET-TO-IDLE}$.⁹
- Note: The DC common mode voltage does not have to be within specification.¹⁰

Additional note – the footnotes are found at the bottom of this page.

C62 - Detect must always change the voltage towards ground.

Release Date: 3/27/03

Replace section 4.3.1.8:

The Receiver Detection circuit is performed by a transmitter and must correctly detect whether a load impedance of Z_{RX-DC} or lower is present.

The behavior of the receiver detection sequence is described below.

Step 1. Transmitter is in a stable Electrical Idle state.¹¹

⁸ The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).

⁹ The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).

¹⁰ The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).

¹¹ The common mode being driven does not have to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).

- Step 2.** A transmitter changes the common mode voltage on D+ and D- to a different value.
- a. A receiver is detected based on the rate that the lines change to the new voltage.
 - i. The receiver is not present if the voltage at the transmitter charges at a rate dictated only by the transmitter impedance, and capacitance of the interconnect, and series capacitor.
 - ii. The receiver is present if the voltage at the transmitter charges at a rate dictated by the transmitter impedance, the series capacitor, the interconnect capacitance, and the receiver termination.

Replace with...

The Receiver Detection circuit is performed by a transmitter and must correctly detect whether a load impedance of Z_{RXDC} or lower is present.

The recommended behavior of the receiver detection sequence is described below.

- Step 1.** Transmitter must start at a stable voltage prior to the detect common mode shift. This voltage can be VDD, GND, or some other stable common-mode voltage between VDD and GND.
- Step 2.** Transmitter changes the common mode voltage on D+ and D- to a different value.
- a. If the common-mode voltage is equal to VDD then the change must be towards ground.
 - b. If the common-mode voltage is equal to GND then the change must be towards VDD.
 - c. If the common-mode is set to a stable value between VDD and GND, then the direction that the voltage moved to get to this stable initial Electrical Idle voltage is important and the detect common mode shift must be in the opposite direction.
- Step 3.** A receiver is detected based on the rate that the lines change to the new voltage.
- iii. The receiver is not present if the voltage at the transmitter charges at a rate dictated only by the transmitter impedance, and capacitance of the interconnect, and series capacitor.
 - iv. The receiver is present if the voltage at the transmitter charges at a rate dictated by the transmitter impedance, the series capacitor, the interconnect capacitance, and the receiver termination.

C63 - Loopback in 10 bit domain.

Release Date: 3/27/03

Section 4.2.6.10.2

Replace:

- ☐ A Loopback Slave is required to retransmit each 10b data and control symbol exactly as received, without applying scrambling/descrambling or disparity corrections, with three important exceptions:

- If a received 10b symbol is determined to be an invalid 10b code (i.e., no legal translation to a control or data value possible) then the slave must instead transmit the EDB symbol in the corresponding time slot of the invalid symbol. Either a positive or negative disparity can be chosen for the EDB symbol.
- If a SKP ordered set retransmission requires adding a SKP symbol to accommodate timing tolerance correction, the SKP symbol is inserted in the retransmitted symbol stream anywhere in the SKP ordered set following the COM symbol. Either a positive or negative disparity can be chosen for the inserted SKP symbol.
- If a SKP ordered set retransmission requires dropping a SKP symbol to accommodate timing tolerance correction, the SKP symbol is simply not retransmitted and transmission continues with the next received symbol or an EDB, as defined above.

As result of these rules, received valid 10b codes are retransmitted even if they fail to match expected disparity in the receiver and result in retransmission violating normal disparity rules.

With:

- ❑ A Loopback Slave is required to retransmit ~~each the received 10 bit data and control symbol information~~ exactly as received, ~~without applying scrambling/descrambling or disparity corrections, with three important exceptions:~~ while continuing to perform clock tolerance compensation:
 - SKPs must be added or deleted on a per Lane basis as outlined in <section 4.2.7> with the exception that SKPs don't have to be simultaneously added or removed across Lanes of a configured Link.
 - ~~If a received 10b symbol is determined to be an invalid 10b code (i.e., no legal translation to a control or data value possible) then the slave must instead transmit the EDB symbol in the corresponding time slot of the invalid symbol. Either a positive or negative disparity can be chosen for the EDB symbol.~~
 - If a SKP ordered set retransmission requires adding a SKP symbol to accommodate timing tolerance correction, the SKP symbol is inserted in the retransmitted symbol stream anywhere ~~adjacent to a SKP symbol~~ in the SKP ordered set following the COM symbol. ~~Either a positive or negative disparity can be chosen for the inserted SKP symbol.~~ The inserted SKP symbol must be of the same disparity as the received SKPs symbol(s) in the SKP ordered set.
 -
 - ⊕▪ If a SKP ordered set retransmission requires dropping a SKP symbol to accommodate timing tolerance correction, the SKP symbol is simply not retransmitted ~~and transmission continues with the next received symbol or an EDB, as defined above.~~
 - Note: No modifications of the received 10 bit data is allowed by the Loopback Slave even if it is determined to be an invalid 10b code (i.e., no legal translation to a control or data value possible).

~~As result of these rules, received valid 10b codes are retransmitted even if they fail to match expected disparity in the receiver and result in retransmission violating normal disparity rules.~~

C64. - Scrambling LFSR Equation Change

Release Date: 3/27/03

Section 4.2.3

Replace polynomial:

$$G(X)=X^{16}+X^{15}+X^{13}+X^4+1$$

With:

$$G(X)=X^{16}+X^5+X^4+X^3+1$$


Second Bullet in this section, Replace:

“The LFSR value is advanced eight serial shifts for each character except the SKP.”


With:

“The LFSR value is advanced eight serial shifts for each character **including K-codes** except the SKP.”

Replace Figure 4-10: LFSR with Scrambling Polynomial:



With:



Appendix C.1

Fully replace the appendix to match the new polynomial, including:

- Sample code must be largely rewritten
- Table showing the LSFR feedback values must be updated
- Table showing the “scramble” byte values must be updated
- The EMI spectrum is essentially unchanged

The replacement text is appended to the end of this SCN.

Updated Appendix C:

C. Physical Layer Appendix

C.1. Data Scrambling

The following subroutines encode and decode an eight-bit value contained in “inbyte” with the LFSR. This is presented as one example only; there are many ways to obtain the proper output. This example demonstrates how to advance the LFSR eight times in one operation and how to XOR the data in one operation. Many other implementations are possible but they must all produce the same output as that shown here.

The following algorithm uses the “C” programming language conventions, where “<<” and “>>” represent the shift left and shift right operators, “>” is the compare greater than operator, and “^” is the exclusive or operator, and & is the logical “AND” operator.

```
/*
    this routine implements the serial descrambling algorithm in parallel form
    for the LSFR polynomial:  x^16+x^5+x^4+x^3+1
    this advances the LSFR 8 bits every time it is called
    this requires fewer than 25 xor gates to implement (with a static register)
```

```
    The XOR required to advance 8 bits/clock is:

    bit    0    1    2    3    4    5    6    7    8    9   10   11   12   13   14   15

           8    9   10   11   12   13   14   15    0    1    2    3    4    5    6    7

                8    9   10   11   12   13   14   15

                    8    9   10   11   12   13   14   15

                        8    9   10   11   12   13   14   15
```

```
    The serial data is just the reverse of the upper byte:

    bit    0    1    2    3    4    5    6    7
```

```
    15  14  13  12  11  10   9   8
*/

int scramble_byte(int inbyte)
{

static int scrambit[16];
static int bit[16];
static int bit_out[16];
static unsigned short lfsr = 0xffff; // 16 bit short for polynomial
int i, outbyte;

    if (inbyte == COMMA)    // if this is a comma
    {
        lfsr = 0xffff;      // reset the LFSR
        return (COMMA);     // and return the same data
    }

    if (inbyte == SKIP)     // don't advance or encode on skip
        return (SKIP);

    for (i=0; i<16;i++)     // convert LFSR to bit array for legibility
        bit[i] = (lfsr >> i) & 1;

    for (i=0; i<8; i++)     // convert byte to be scrambled for legibility
        scrambit[i] = (inbyte >> i) & 1;

    // apply the xor to the data
    if (! (inbyte & 0x100) &&    // if not a KCODE, scramble the data
        ! (TrainingSequence == TRUE)) // and if not in the middle of
    {
        // a training sequence
```

```
        scrambit[0] ^= bit[15];
        scrambit[1] ^= bit[14];
        scrambit[2] ^= bit[13];
        scrambit[3] ^= bit[12];
        scrambit[4] ^= bit[11];
        scrambit[5] ^= bit[10];
        scrambit[6] ^= bit[9];
        scrambit[7] ^= bit[8];
    }

    // Now advance the LFSR 8 serial clocks
    bit_out[ 0] = bit[ 8];
    bit_out[ 1] = bit[ 9];
    bit_out[ 2] = bit[10];
    bit_out[ 3] = bit[11] ^ bit[ 8];
    bit_out[ 4] = bit[12] ^ bit[ 9] ^ bit[ 8];
    bit_out[ 5] = bit[13] ^ bit[10] ^ bit[ 9] ^ bit[ 8];
    bit_out[ 6] = bit[14] ^ bit[11] ^ bit[10] ^ bit[ 9];
    bit_out[ 7] = bit[15] ^ bit[12] ^ bit[11] ^ bit[10];
    bit_out[ 8] = bit[ 0] ^ bit[13] ^ bit[12] ^ bit[11];
    bit_out[ 9] = bit[ 1] ^ bit[14] ^ bit[13] ^ bit[12];
    bit_out[10] = bit[ 2] ^ bit[15] ^ bit[14] ^ bit[13];
    bit_out[11] = bit[ 3]           ^ bit[15] ^ bit[14];
    bit_out[12] = bit[ 4]           ^ bit[15];
    bit_out[13] = bit[ 5];
    bit_out[14] = bit[ 6];
    bit_out[15] = bit[ 7];

    lfsr = 0;

    for (i=0; i <16; i++) // convert the LFSR back to an integer
        lfsr += (bit_out[i] << i);

    outbyte = 0;
```

```
    for (i= 0; i<8; i++) // convert data back to an integer
        outbyte += (scrambit[i] << i);

    return outbyte;
}

/* NOTE THAT THE DESCRAMBLE ROUTINE IS IDENTICAL TO THE SCRAMBLE ROUTINE
   this routine implements the serial descrambling algorithm in parallel form
   this advances the lfsr 8 bits every time it is called
   this uses fewer than 25 xor gates to implement (with a static register)
   The XOR tree is the same as the scrambling routine
*/

int unscramble_byte(int inbyte)
{

    static int descrambit[8];
    static int bit[16];
    static int bit_out[16];
    static unsigned short lfsr = 0xffff; // 16 bit short for polynomial
    int outbyte, i;

    if (inbyte == COMMA) // if this is a comma
    {
        lfsr = 0xffff; // reset the LFSR
        return (COMMA); // and return the same data
    }

    if (inbyte == SKIP) // don't advance or encode on skip
        return (SKIP);

    for (i=0; i<16;i++) // convert the LFSR to bit array for legibility
```

```
    bit[i] = (lfsr >> i) & 1;

for (i=0; i<8; i++)    // convert byte to be de-scrambled for legibility
    descrambit[i] = (inbyte >> i) & 1;

// apply the xor to the data
if (! (inbyte & 0x100) &&    // if not a KCODE, scramble the data
    ! (TrainingSequence == TRUE))    // and if not in the middle of
{
    // a training sequence

    descrambit[0] ^= bit[15];
    descrambit[1] ^= bit[14];
    descrambit[2] ^= bit[13];
    descrambit[3] ^= bit[12];
    descrambit[4] ^= bit[11];
    descrambit[5] ^= bit[10];
    descrambit[6] ^= bit[9];
    descrambit[7] ^= bit[8];
}

// Now advance the LFSR 8 serial clocks
bit_out[ 0] = bit[ 8];
bit_out[ 1] = bit[ 9];
bit_out[ 2] = bit[10];
bit_out[ 3] = bit[11] ^ bit[ 8];
bit_out[ 4] = bit[12] ^ bit[ 9] ^ bit[ 8];
bit_out[ 5] = bit[13] ^ bit[10] ^ bit[ 9] ^ bit[ 8];
bit_out[ 6] = bit[14] ^ bit[11] ^ bit[10] ^ bit[ 9];
bit_out[ 7] = bit[15] ^ bit[12] ^ bit[11] ^ bit[10];
bit_out[ 8] = bit[ 0] ^ bit[13] ^ bit[12] ^ bit[11];
bit_out[ 9] = bit[ 1] ^ bit[14] ^ bit[13] ^ bit[12];
bit_out[10] = bit[ 2] ^ bit[15] ^ bit[14] ^ bit[13];
bit_out[11] = bit[ 3]          ^ bit[15] ^ bit[14];
```

```

    bit_out[12] = bit[ 4]                ^ bit[15];

    bit_out[13] = bit[ 5];

    bit_out[14] = bit[ 6];

    bit_out[15] = bit[ 7];

    lfsr = 0;

    for (i=0; i <16; i++) // convert the LFSR back to an integer
        lfsr += (bit_out[i] << i);

    outbyte = 0;

    for (i= 0; i<8; i++) // convert data back to an integer
        outbyte += (descrambit[i] << i);

    return outbyte;
}

```

The initial 16 bit values of the LFSR for the first 128 LFSR advances following a reset are listed below:

	0, 8	1, 9	2, A	3, B	4, C	5, D	6, E	7, F
00	FFFF	E817	0328	284B	4DE8	E755	404F	4140
08	4E79	761E	1466	6574	7DBD	B6E5	FDA6	B165
10	7D09	02E5	E572	673D	34CF	CB54	4743	4DEF
18	E055	40E0	EE40	54BE	B334	2C7B	7D0C	07E5
20	E5AF	BA3D	248A	8DC4	D995	85A1	BD5D	4425
28	2BA4	A2A3	B8D2	CBF8	EB43	5763	6E7F	773E
30	345F	5B54	5853	5F18	14B7	B474	6CD4	DC4C
38	5C7C	70FC	F6F0	E6E6	F376	603B	3260	64C2
40	CB84	9743	5CBF	B3FC	E47B	6E04	0C3E	3F2C
48	29D7	D1D1	C069	7BC0	CB73	6043	4A60	6FFA
50	F207	1102	01A9	A939	2351	566B	6646	4FF6

58	F927	3081	85B0	AC5D	478C	82EF	F3F2	E43B
60	2E04	027E	7E72	79AE	A501	1A7D	7F2A	2197
68	9019	0610	1096	9590	8FCD	D0E7	F650	46E6
70	E8D6	C228	3AB2	B70A	129F	9CE2	FC3C	2B5C
78	5AA3	AF6A	70C7	CDF0	E3D5	C0AB	B9C0	D9C1

An 8 bit value of 0 repeatedly encoded with the LFSR after reset produces the following consecutive 8 bit values:

	00	01	02	03	04	05	06	07	08	09	0A	0B	0C	0D	0E	0F
00	FF	17	C0	14	B2	E7	02	82	72	6E	28	A6	BE	6D	BF	8D
10	BE	40	A7	E6	2C	D3	E2	B2	07	02	77	2A	CD	34	BE	E0
20	A7	5D	24	B1	9B	A1	BD	22	D4	45	1D	D3	D7	EA	76	EE
30	2C	DA	1A	FA	28	2D	36	3B	3A	0E	6F	67	CF	06	4C	26
40	D3	E9	3A	CD	27	76	30	FC	94	8B	03	DE	D3	06	52	F6
50	4F	88	80	95	C4	6A	66	F2	9F	0C	A1	35	E2	41	CF	27
60	74	40	7E	9E	A5	58	FE	84	09	60	08	A9	F1	0B	6F	62
70	17	43	5C	ED	48	39	3F	D4	5A	F5	0E	B3	C7	03	9D	9B
80	8B	0D	8E	5C	33	98	77	AE	2D	AC	0B	3E	DA	0B	42	7A
90	7C	D1	CF	A8	1C	12	EE	41	C2	3F	38	7A	0D	69	F4	01
A0	DA	31	72	C5	A0	D7	93	0E	DC	AF	A4	55	E7	F0	72	16
B0	68	D5	38	84	DD	00	CD	18	9E	CA	30	59	4C	75	1B	77
C0	31	C5	ED	CF	91	64	6E	3D	FE	E8	29	04	CF	6C	FC	C4
D0	0B	5E	DA	62	BA	5B	AB	DF	59	B7	7D	37	5E	E3	1A	C6
E0	88	14	F5	4F	8B	C8	56	CB	D3	10	42	63	04	8A	B4	F7
F0	84	01	A0	01	83	49	67	EE	3E	2A	8B	A4	76	AF	14	D5

100	4F	AC	60	B6	79	D6	62	B7	43	E7	E5	2A	40	2C	6E	7A
110	56	61	63	20	6A	97	4A	38	05	E5	DD	68	0D	78	4C	53
120	8B	D6	86	57	B2	AA	1A	80	18	DC	BA	FC	03	A3	4B	30

Scrambling produces the power spectrum (in the 10b domain) shown in Figure C-1.




Figure C-1: Scrambling Spectrum for a Data Value of 0

C65 Assorted Errata

Release Date: 3/27/03

Section 4.3.3.2 , replace the last paragraph:

"CTX is an optional portion of the measurement test load. The measurement should be taken on the opposite side of the capacitor from the package, and the value of the CTX must be in the range of 75 nF to 200 nF."

With the following:

"When measuring transmitter output parameters CTX is an optional portion of the Test/Measurement load. When used, the value of CTX must be in the range of 75 nF to 200 nF. CTX must not be used when the Test/Measurement load is placed in the receiver package reference plane."

Replace Figure 4-51 (page 247):





Figure 4-~~51~~⁵⁷⁴⁻²⁹: BEACON, Which Includes a 2 ns Pulse Through a 75 nF Capacitor

with:



Section 4.2.7.1. Rules for Transmitters

Insert the following two bullets after the last bullet in this section. The bullet reads:

" SKIP ordered sets do not count as an interruption when monitoring for consecutive characters or ordered set (i.e., eight consecutive TS1 ordered sets in Polling.Active)"

After the above bullet insert the following two bullets:

"Any and all time spent in any lower power link state (L0s, L1, L2), does not count in the 1180 to 1538 symbol interval used to schedule the transmission of SKIP ordered sets. During all lower power link states any counter(s) or other mechanisms used to schedule SKIP ordered sets must be reset."

Section 4.2.4.3. Fast Training Sequence (FTS)

[Second paragraph]

Replace:

"SKP ordered sets must be scheduled and transmitted between FTS ordered sets as necessary to meet the definitions in Section 4.2.7 with the exception that no SKP ordered sets can be scheduled during the first 255 FTS ordered sets. A single SKP ordered set is always sent after the last FTS is transmitted."

With:

"SKP ordered sets must be scheduled and transmitted between FTS ordered sets as necessary to meet the definitions in Section 4.2.7 with the exception that no SKP ordered sets can be **transmitted** during the first 255 FTS ordered sets. A single SKP ordered set is always sent

after the last FTS is transmitted. **Note that it is possible that two SKP ordered-sets can be transmitted back to back (one SKP ordered set to signify the completion of the 4096 FTSs and one scheduled and transmitted to meet the definitions described in section 4.2.7)"**

>>>>Note to tech writer: In the above, references to section 4.2.7 should be cross references to those sections. If other spec changes cause the section number to move, the number 4.2.7 may change – the correct section is called "Clock Tolerance Compensation"

Section 4.5.9 - Disable

Replace:

Disabled uses Bit 1 (Disable Link) in the Training Control Register (see Table 4-2 and Table 4-3) which is sent within the TS1 and TS2 training ordered set.

With:

Disabled uses Bit 1 (Disable Link) in the Training Control **field** (see Table 4-2 and Table 4-3) which is sent within the TS1 and TS2 training ordered set.

4.2.5.10. Loopback

Replace:

Loopback uses Bit 2 (Loopback) in the Training Control Register (see Table 4-2 and Table 4-3) which is sent within the TS1 and TS2 training ordered set.

With:

Loopback uses Bit 2 (Loopback) in the Training Control **field** (see Table 4-2 and Table 4-3) which is sent within the TS1 and TS2 training ordered set.

4.2.5.11. Hot Reset

Replace:

Hot Reset uses Bit 0 (Hot Reset) in the Training Control Register (see Table 4-2 and Table 4-3) which is sent within the TS1 and TS2 training ordered set.

With:

Hot Reset uses Bit 0 (Hot Reset) in the Training Control **field** (see Table 4-2 and Table 4-3) which is sent within the TS1 and TS2 training ordered set.

In section 4.6.3.5.1

Change line from

- ❑ The next state is Configuration.Idle immediately after all Lanes that are transmitting TS2 ordered sets receive eight consecutive TS2 ordered sets with matching Lane and Link numbers (non-Pad) and 16 TS2 ordered sets are sent after receiving one TS2 ordered set.
- Note: All remaining Lanes that are not part of the configured Link are no longer associated with the LTSSM in progress and must:
 - i. Be associated with a new LTSSM if this optional feature is supported.
 - or
 - ii. All Lanes that cannot be associated with an optional new LTSSM must transition to Electrical Idle.¹²
 - ‡ Note: In the case of an optional crosslink, the receiver terminations are required to meet $Z_{RX-HIGH-IMP-DC}$ (see Table **Table 0-5**).
 - ‡ Note: These Lanes must be re-associated with the LTSSM immediately after the LTSSM in progress transitions back to Detect.

To (changes in Red)

- ❑ The next state is Configuration.Idle immediately after all Lanes that are transmitting TS2 ordered sets receive eight consecutive TS2 ordered sets with matching Lane and Link numbers (non-Pad) and 16 TS2 ordered sets are sent after receiving one TS2 ordered set.
- Note: All remaining Lanes that are not part of the configured Link are no longer associated with the LTSSM in progress and must:
 - iii. Be associated with a new LTSSM if this optional feature is supported.
 - or
 - iv. All Lanes that cannot be associated with an optional new LTSSM must transition to Electrical Idle.¹³

¹² The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).

¹³ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode

- | Note: In the case of an optional crosslink, the receiver terminations are required to meet $Z_{RX-HIGH-IMP-DC}$ (see Table **Table 0-5**).
- | Note: These Lanes must be re-associated with the LTSSM immediately after the LTSSM in progress transitions back to Detect.
- | **Note: An Electrical Idle ordered-set does not need to be sent before transitioning to Electrical Idle.**

In section 4.6.3.5.2

Change line from

- ❑ The next state is Configuration.Idle immediately after all Lanes that are transmitting TS2 ordered sets receive eight consecutive TS2 ordered sets with matching Lane and Link numbers (non-Pad) and 16 consecutive TS2 ordered sets are sent after receiving one TS2 ordered set.
 - Note: All remaining Lanes that are not part of the configured Link are no longer associated with the LTSSM in progress and must:
 - i. Optionally be associated with a new crosslink LTSSM if this feature is supported.
 - or
 - ii. All remaining Lanes that are not associated with a new crosslink LTSSM must transition to Electrical Idle¹⁴, and receiver terminations are required to meet $Z_{RX-HIGH-IMP-DC}$ (see Table 0-5).
 - | Note: These Lanes must be re-associated with the LTSSM immediately after the LTSSM in progress transitions back to Detect.

To (changes in Red)

- ❑ The next state is Configuration.Idle immediately after all Lanes that are transmitting TS2 ordered sets receive eight consecutive TS2 ordered sets with matching Lane and Link numbers (non-Pad) and 16 consecutive TS2 ordered sets are sent after receiving one TS2 ordered set.
 - Note: All remaining Lanes that are not part of the configured Link are no longer associated with the LTSSM in progress and must:

During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).

¹⁴ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).

- iii. Optionally be associated with a new crosslink LTSSM if this feature is supported.
- or
- iv. All remaining Lanes that are not associated with a new crosslink LTSSM must transition to Electrical Idle¹⁵, and receiver terminations are required to meet $Z_{RX-HIGH-IMP-DC}$ (see Table 0-5).
 - | Note: These Lanes must be re-associated with the LTSSM immediately after the LTSSM in progress transitions back to Detect.
 - | **Note: An Electrical Idle ordered-set does not need to be sent before transitioning to Electrical Idle.**

In section 4.2.6.1.2

change line from:

- ☐ If at least one but not all un-configured Lanes detect a receiver, then:
 - 1. Wait for 12 ms.
 - 2. The transmitter performs a Receiver Detection sequence on all un-configured Lanes that can form one or more Links (see Section 0 for more information),
 - i) The next state is Polling if exactly the same Lanes detect a receiver as the first Receiver Detection sequence.
 - Note: Lanes that did not detect a receiver must:
 - i) Be associated with a new LTSSM if this optional feature is supported.
 - or
 - ii) All Lanes that cannot be associated with an optional new LTSSM must transition to Electrical Idle.¹⁶
 - Note: These Lanes must be re-associated with the LTSSM immediately after the LTSSM in progress transitions back to Detect.
 - ii) Otherwise, the next state is Detect.Quiet.

To (changes in red):

- ☐ If at least one but not all un-configured Lanes detect a receiver, then:

¹⁵ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).

¹⁶ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see **Table 0-4**).

3. Wait for 12 ms.
4. The transmitter performs a Receiver Detection sequence on all un-configured Lanes that can form one or more Links (see Section 0 for more information),
 - i) The next state is Polling if exactly the same Lanes detect a receiver as the first Receiver Detection sequence.
 - Note: Lanes that did not detect a receiver must:
 - i) Be associated with a new LTSSM if this optional feature is supported.
 - or
 - ii) All Lanes that cannot be associated with an optional new LTSSM must transition to Electrical Idle.¹⁷
 - Note: These Lanes must be re-associated with the LTSSM immediately after the LTSSM in progress transitions back to Detect.
 - **Note: An Electrical Idle ordered-set does not need to be sent before transitioning to Electrical Idle.**
 - ii) Otherwise, the next state is Detect.Quiet.

Section 4.2.6.5, change:

This is the normal operational state.

- ☐ LinkUp = 1 (status is set true).
- ☐ Next state is **Recovery** if a TS1 or TS2 ordered set is received on any configured Lane.
- ☐ Next state is **Recovery** if directed to this state or if Electrical Idle is detected without receiving an Electrical Idle ordered set.
 - Note: “if directed” applies to a Port that is instructed by a higher Layer to transition to Recovery.
 - Note: The transmitter may complete any TLP or DLLP in progress.
- ☐ Next state of transmitter is **L0s** if directed to this state.
 - Note: “if directed” applies to a Port that is instructed by a higher Layer to initiate L0s.
- ☐ Next state of receiver is **L0s** if receiver detects **Electrical Idle** ordered set and is not directed to L1 or L2 states.
- ☐ Next state is **L1**:
 - i. If directed
 - and

¹⁷ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see **Table 0-4**).

- ii. an Electrical Idle ordered set is received
and

- iii. an Electrical Idle ordered set is transmitted.

- Note: “if directed” is defined as both ends of the Link having agreed to enter L1 immediately after the condition of both the receipt and transmission of an Electrical Idle ordered set is met.

- Note: When directed by a higher Layer one side of the Link always initiates and exits to L1 by transmitting an Electrical Idle ordered set, followed by a transition to Electrical Idle.¹⁸ The same Port then waits for the receipt of an Electrical Idle ordered set, and then immediately transitions to L1. Conversely, the side of the Link that first receives an Electrical Idle ordered set must send an Electrical Idle ordered set and immediately transition to L1.

☐ Next state is **L2**:

- i.If directed

and

- ii.an Electrical Idle ordered set is received

and

- iii.an Electrical Idle ordered set is transmitted.

- Note: “if directed” is defined as both ends of the Link having agreed to enter L2 immediately after the condition of both the receipt and transmission of an Electrical Idle ordered set is met.

- Note: When directed by a higher Layer, one side of the Link always initiates and exits to L2 by transmitting an Electrical Idle ordered set followed by a transition to Electrical Idle.¹⁹ The same Port then waits for the receipt of an Electrical Idle ordered set, and then immediately transitions to L2. Conversely, the side of the Link that first receives an Electrical Idle ordered set must send an Electrical Idle ordered set and immediately transition to L2.

To (changes in red):

This is the normal operational state.

☐ LinkUp = 1 (status is set true).

¹⁸ The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).

¹⁹ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).

- ❑ Next state is **Recovery** if a TS1 or TS2 ordered set is received on any configured Lane.
- ❑ Next state is **Recovery** if directed to this state or if Electrical Idle is detected without receiving an Electrical Idle ordered set.
 - Note: “if directed” applies to a Port that is instructed by a higher Layer to transition to Recovery.
 - Note: The transmitter may complete any TLP or DLLP in progress.
- ❑ Next state of transmitter is **L0s** if directed to this state.
 - Note: “if directed” applies to a Port that is instructed by a higher Layer to initiate L0s (see section <5.4.1.1.1> for more details).
- ❑ Next state of receiver is **L0s** if receiver detects **Electrical Idle** ordered set and is not directed to L1 or L2 states.
- ❑ Next state is **L1**:
 - iv. If directed
 - and
 - v. an Electrical Idle ordered set is received
 - and
 - vi. an Electrical Idle ordered set is transmitted.
 - Note: “if directed” is defined as both ends of the Link having agreed to enter L1 immediately after the condition of both the receipt and transmission of an Electrical Idle ordered set is met (see section <5.3.2.1> for more details).
 - Note: When directed by a higher Layer one side of the Link always initiates and exits to L1 by transmitting an Electrical Idle ordered set, followed by a transition to Electrical Idle.²⁰ The same Port then waits for the receipt of an Electrical Idle ordered set, and then immediately transitions to L1. Conversely, the side of the Link that first receives an Electrical Idle ordered set must send an Electrical Idle ordered set and immediately transition to L1.
- ❑ Next state is **L2**:
 - iv.If directed
 - and
 - v.an Electrical Idle ordered set is received
 - and
 - vi.an Electrical Idle ordered set is transmitted.

²⁰ The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).

- Note: “if directed” is defined as both ends of the Link having agreed to enter L2 immediately after the condition of both the receipt and transmission of an Electrical Idle ordered set is met (see section <4.3.2.3> for more details).
- Note: When directed by a higher Layer, one side of the Link always initiates and exits to L2 by transmitting an Electrical Idle ordered set followed by a transition to Electrical Idle.²¹ The same Port then waits for the receipt of an Electrical Idle ordered set, and then immediately transitions to L2. Conversely, the side of the Link that first receives an Electrical Idle ordered set must send an Electrical Idle ordered set and immediately transition to L2.

In table 4.5

Change line from

T _{TX-IDLE-MIN}	Minimum time spent in Electrical Idle	50			UI	Minimum time a transmitter must be in Electrical Idle
T _{TX-IDLE-SET-TO-IDLE}	Maximum time to transition to a valid Electrical Idle after sending an Electrical Idle ordered set			20	UI	After sending an Electrical Idle ordered set, the transmitter must meet all Electrical Idle specifications within this time.

To (changes in red):

T _{TX-IDLE-MIN}	Minimum time spent in Electrical Idle	50			UI	Minimum time a transmitter must be in Electrical Idle. Utilized by the receiver to start looking for an Electrical Idle Exit after successfully receiving an Electrical Idle ordered-set.
T _{TX-IDLE-SET-TO-IDLE}	Maximum time to transition to a valid Electrical Idle after sending an Electrical Idle ordered set			20	UI	After sending an Electrical Idle ordered set, the transmitter must meet all Electrical Idle specifications within this time. This is considered a debounce time for the TX to meet Electrical Idle after transitioning from L0.

²¹ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle (V_{TX-CM-DC-ACTIVE-IDLE-DELTA}) specification (see Table 0-4).

In section 4.2.6.8.1

Change line from:

☐ For Downstream Lanes:

- For a Root Port, the next state is Detect if a Beacon is received on at least Lane 0.
 - ? Note: Main power must be restored before entering Detect.

To (changes in red):

☐ For Downstream Lanes:

- For a Root Port, the next state is Detect if a Beacon is received on at least Lane 0 **or “if directed”**.
 - ? Note: Main power must be restored before entering Detect.
 - ? **Note: “if directed” is defined as a higher layer decides to exit to Detect.**

Replace the following text on page 217:

“An appropriate average TX UI must be used as the interval for measuring the eye diagram.”

With:

A recovered TX UI is calculated over 3500 consecutive unit intervals of sample data. The eye diagram is created using all edges of the 250 consecutive UI in the center of the 3500 UI used for calculating the TX UI.

Change the following text on page 218 in the comments section of the table:

“Jitter is defined as the measurement variation of the crossing points ($V_{RX-DIFFp-p} = 0$ V) in relation to an appropriate average TX UI.”

With:

Jitter is defined as the measurement variation of the crossing points ($V_{RX-DIFFp-p} = 0$ V) in relation to a recovered TX UI. A recovered TX UI is calculated over 3500 consecutive unit intervals of sample data. Jitter is measured using all edges of the 250 consecutive UI in the center of the 3500 UI used for calculating the TX UI.

Change the following text on page 219 in note 7:

“If the clocks to the RX and TX are not derived from the same clock chip the TX UI must be used as a reference for the eye diagram.”

With:

If the clocks to the RX and TX are not derived from the same reference clock, the TX UI recovered from 3500 consecutive UI must be used as the reference for the eye diagram.

Change the following text on page 219 in note 8:

“If the clocks to the RX and TX are not derived from the same clock chip the TX UI must be used as a reference for the eye diagram.”

With:

If the clocks to the RX and TX are not derived from the same reference clock, the TX UI recovered from 3500 consecutive UI must be used as the reference for the eye diagram.

Change the following text on page 220:

“An appropriate average TX UI must be used as the interval for measuring the eye diagram.”

With:

A recovered TX UI is calculated over 3500 consecutive unit intervals of sample data. The eye diagram is created using all edges of the 250 consecutive UI in the center of the 3500 UI used for calculating the TX UI.

Add the following text as an additional note on page 220:

13. It is recommended that the recovered TX UI is calculated using all edges in the 3500 consecutive UI interval with a fit algorithm using a minimization merit function. Least squares and median deviation fits have worked well with experimental and simulated data.

In section 4.2.6.3.1.2

Replace:

- ☐ Immediately after all Upstream Lanes receive two consecutive TS1 ordered sets with Link numbers that are different than PAD (K23.7), a single Link number is selected and transmitted on all Lanes that can form a Link. Any left over Lanes that detected a receiver during Detect must transmit TS1 ordered sets with the Link and Lane number set to PAD (K23.7). The next state is Configuration.Linkwidth.Accept.

With(changes in red):

- ❑ ~~If any Lanes first received at least one or more TS1 ordered sets with a Link and Lane number set to PAD (K23.7) Immediately after all Upstream Lanes~~ receive two consecutive TS1 ordered sets with Link numbers that are different than PAD (K23.7) ~~and Lane Number equal to PAD~~, a single Link number is selected and transmitted on all Lanes that can form a Link. Any left over Lanes that detected a receiver during Detect must transmit TS1 ordered sets with the Link and Lane number set to PAD (K23.7). The next state is Configuration.Linkwidth.Accept.

Section 4.2.6.3.1.1, second bullet, replace:

Next state is Loopback if directed to this state and the transmitter is capable of being a Loopback Master.

With:

Next state is Loopback if directed to this state, and the transmitter is capable of being a Loopback Master, **which is determined by implementation specific means.**

Section 4.2.6.3.1.2, second bullet, replace:

Next state is Loopback if directed to this state and the transmitter is capable of being a Loopback Master.

With:

Next state is Loopback if directed to this state, and the transmitter is capable of being a Loopback Master, **which is determined by implementation specific means.**

Section 4.2.6.4.3, fourth bullet, replace:

Next state is Loopback if directed to this state and the transmitter is capable of being a Loopback Master.

With:

Next state is Loopback if directed to this state, and the transmitter is capable of being a Loopback Master, **which is determined by implementation specific means.**

Replace in section 4.2.6.11

- ❑ Lanes that were not directed by a higher Layer to initiate Hot Reset (i.e., received two consecutive TS1 ordered sets with the Hot Reset bit (Bit 0) asserted on any configured Lanes):

- LinkUp = 0 (False)
- If any Lane of an Upstream Port of a Switch receives a training sequence with the Hot Reset bit asserted, all configured Downstream Ports must transition to Hot Reset as soon as possible.
- All Lanes in the configured Link transmit TS1 ordered sets with the Hot Reset bit (Bit 0) asserted and the configured Link and Lane numbers.
- After a 2 ms timeout:
 - ? If two consecutive TS1 ordered sets were received with the Hot Reset bit (Bit 0) asserted and the configured Link and Lane numbers, the next state is Hot Reset.
 - ? Otherwise, the next state is Detect.

With: (changes in Red)

- ☐ Lanes that were not directed by a higher Layer to initiate Hot Reset (i.e., received two consecutive TS1 ordered sets with the Hot Reset bit (Bit 0) asserted on any configured Lanes):
 - LinkUp = 0 (False)
 - If any Lane of an Upstream Port of a Switch receives a training sequence with the Hot Reset bit asserted, all configured Downstream Ports must transition to Hot Reset as soon as possible.
 - ? **Note: Any optional crosslinks on the switch are an exception to this rule and the behavior is system specific.**
 - All Lanes in the configured Link transmit TS1 ordered sets with the Hot Reset bit (Bit 0) asserted and the configured Link and Lane numbers.
 - After a 2 ms timeout:
 - ? If two consecutive TS1 ordered sets were received with the Hot Reset bit (Bit 0) asserted and the configured Link and Lane numbers, the next state is Hot Reset.
 - ? Otherwise, the next state is Detect.

Change in Table 4-2 and Table 4-3

Training Control

Bit 0 – Hot Reset

Bit 0 = 0, De-assert

Bit 0 = 1, Assert

Bit 1 – Disable Link

Bit 1 = 0, De-assert

Bit 1 = 1, Assert

Bit 2 – Loopback

Bit 2 = 0, De-assert

Bit 2 = 1, Assert

Bit 3 – Disable Scrambling

Bit 3 = 0, De-assert

Bit 3 = 1, Assert Bit 4:7, Reserved, set to 0

Replace with

Training Control

Bit 0 – Hot Reset

Bit 0 = 0, De-assert

Bit 0 = 1, Assert

Bit 1 – Disable Link

Bit 1 = 0, De-assert

Bit 1 = 1, Assert

Bit 2 – Loopback

Bit 2 = 0, De-assert

Bit 2 = 1, Assert

Bit 3 – Disable Scrambling

Bit 3 = 0, De-assert

Bit 3 = 1, Assert

Bit 4:7, Reserved, set to 0

Change text in Section 4.2.1.3

The symbol tables for the valid 8b/10b codes are given in Appendix B. These tables have one column for the positive disparity and one column for the negative disparity.

If a received symbol is found in the proper column corresponding to the current running disparity, then that symbol is valid.

If a received symbol is found in the column corresponding to the incorrect running disparity or if the symbol does not correspond to either column, the Physical Layer must notify the Data Link

Layer that the received symbol is invalid. This is a Receiver Error, and is a reported error associated with the Port (see Section **Error! Reference source not found.**).

Replace with

The symbol tables for the valid 8b/10b codes are given in Appendix B. These tables have one column for the positive disparity and one column for the negative disparity.

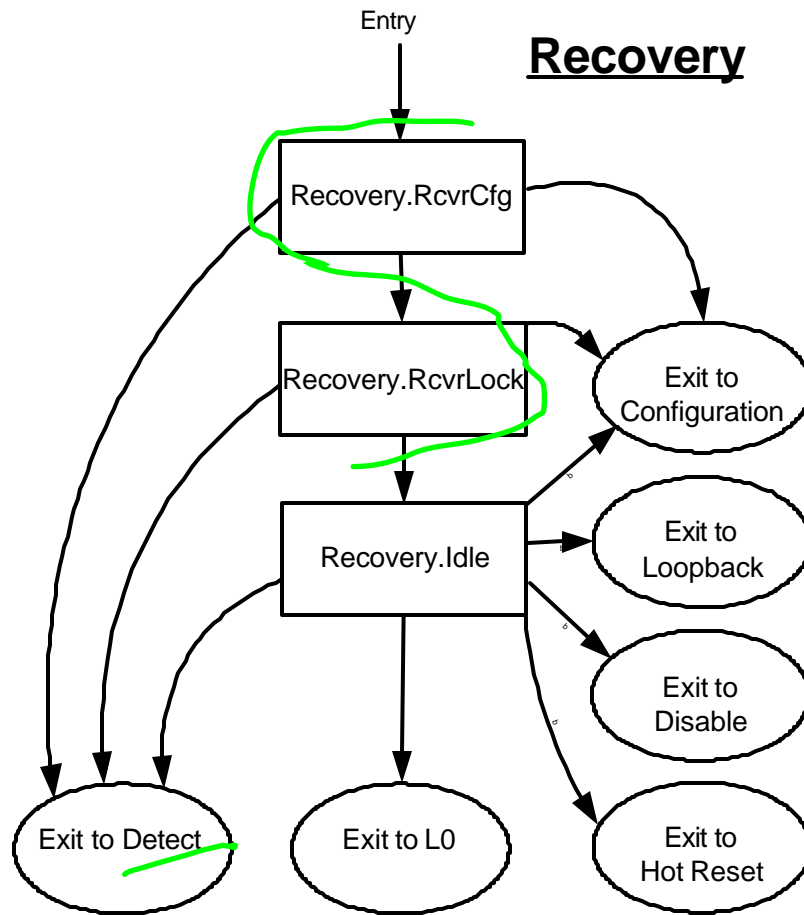
~~If a received symbol is found in the proper column corresponding to the current running disparity, then that symbol is valid.~~

~~A Transmitter may pick any disparity when first transmitting differential data after being in an Electrical Idle state. The Transmitter must then follow proper 8b/10b encoding rules until the next Electrical Idle state is entered.~~

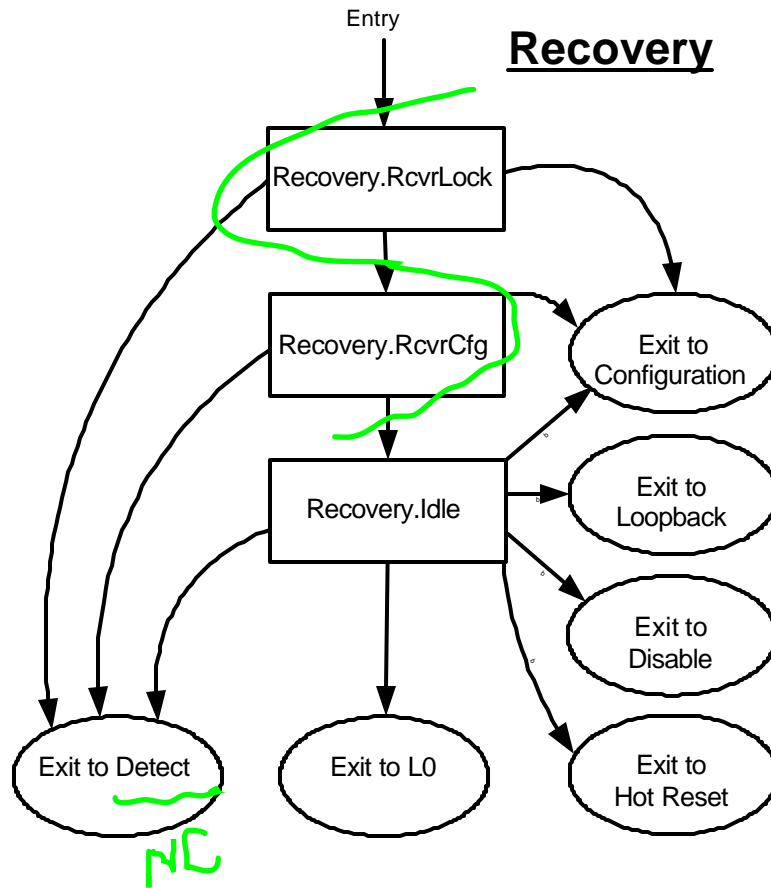
~~The initial disparity for a Receiver that detects an exit from Electrical Idle is set to the disparity of the first character used to obtain symbol lock. Disparity may also be reinitialized if symbol lock is lost and regained during the transmission of differential information due to an implementation specific number of errors. All following received symbols after the initial disparity is set must be in the found in the proper column corresponding to the current running disparity.~~

~~If a received symbol is found in the column corresponding to the incorrect running disparity or if the symbol does not correspond to either column, the Physical Layer must notify the Data Link Layer that the received symbol is invalid. This is a Receiver Error, and is a reported error associated with the Port (see Section **Error! Reference source not found.**).~~

Replace Figure 4-35:



Replace with



Replace the following text in Chapter 4, Section 4.3.2.1, Page 213:

“The worst-case operational loss budget is calculated by taking the minimum output voltage ($V_{TX-DIFFp-p} = 800$ mV) divided by the minimum input voltage to the receiver ($V_{RX-DIFFp-p} = 175$ mV), which results in 13.2 dB.”

with:

“An approximate way to understand the worst-case operational loss budget at 1.25 GHz is calculated by taking the minimum output voltage ($V_{TX-DIFFp-p} = 800$ mV) divided by the minimum input voltage to the receiver ($V_{RX-DIFFp-p} = 175$ mV), which results in a maximum loss of 13.2 dB. The approximate way to understand the worst-case operational loss budget at 625 MHz is calculated by taking the minimum de-emphasized output voltage ($V_{TX-DIFFp-p} = 505$ mV) divided by the minimum input voltage to the receiver ($V_{RX-DIFFp-p} = 175$ mV), which results in a maximum loss of 9.2 dB. Although loss vs. frequency is useful in understanding how to design an effective interconnect the timing and voltage margin measured in the TX and RX eye diagrams end up being the ultimate constraints of insertion Loss”

Add to the end of Section 4.3.1.9, as the second to last paragraph:

When the transmitter transitions from Electrical Idle to a valid differential signal level it must meet the output return loss specifications in Table 4-5 and transition from zero differential to full differential voltage consistent with the requirements of the Transmitter Compliance Eye Diagram of Figure 4-24 after the allotted debounce time of $T_{TX-IDLE-TO-TO-DIFF-DATA}$ (See Table 4-5).

Add to Table 4-5 - Addition is in RED

$T_{TX-IDLE-SET-TO-IDLE}$	Maximum time to transition to a valid Electrical Idle after sending an Electrical Idle ordered set			20	UI	After sending an Electrical Idle ordered set, the transmitter must meet all Electrical Idle specifications within this time. This is considered a debounce time for the TX to meet Electrical Idle after transitioning from L0.
$T_{TX-IDLE-TO-TO-DIFF-DATA}$	Maximum time to transition to valid TX Specifications after leaving an Electrical Idle condition			20	UI	Maximum time to meet all TX specifications when transitioning from Electrical Idle to sending differential data. This is considered a debounce time for the TX to meet all TX specifications after leaving Electrical Idle.

Add to 4.2.6.6.2.3 - Addition is in RED

❑ Transmitter sends N_FTS Fast Training Sequences on all configured Lanes.

- Note: Up to one full FTS ordered-set may be sent before the N_FTS FTS ordered-sets are sent.
- Note: No SKP ordered sets can be inserted before all FTS ordered sets as defined by the agreed upon N_FTS parameter are transmitted.
- Note: If the Extended Synch bit is set, the Transmitter sends 4096 Fast Training Sequences.

Section 4.2.6.3.1.1**Replace:**

A single FTS training sequence is an ordered set composed of one K28.5 (COM) symbol and three K28.1 symbols. The maximum number of FTS ordered sets (N_FTS) that a component can request is 255, providing a bit time lock of $4 * 255 * 10 * UI$. 4096 FTS ordered sets must be sent when the Extended Synch bit is set in order to provide external Link monitoring tools with enough time to achieve bit and framing synchronization. SKP ordered sets must be scheduled and

transmitted between FTS ordered sets as necessary to meet the definitions in Section 04.2.7 with the exception that no SKP ordered sets can be scheduled during the first 255 FTS ordered sets. A single SKP ordered set is always sent after the last FTS is transmitted.

With:

A single FTS training sequence is an ordered set composed of one K28.5 (COM) symbol and three K28.1 symbols. The maximum number of FTS ordered sets (N_FTS) that a component can request is 255, providing a bit time lock of $4 * 255 * 10 * UI$. 4096 FTS ordered sets must be sent when the Extended Synch bit is set in order to provide external Link monitoring tools with enough time to achieve bit and framing synchronization. SKP ordered sets must be scheduled and transmitted between FTS ordered sets as necessary to meet the definitions in Section 04.2.7 with the exception that no SKP ordered sets can be scheduled during the first ~~255~~ N_FTS FTS ordered sets. A single SKP ordered set is always sent after the last FTS is transmitted.

Section 4.2.6.3.1.1

Replace:

- ❑ Optionally, if crosslinks are supported and all Downstream Lanes initially receive two consecutive TS1 ordered sets with a Link number different than PAD (K23.7), then the Downstream Lanes are now designated as Upstream Lanes and a new random cross Link timeout is chosen (see $T_{\text{crosslink}}$ in Table 0-4). The next state is Configuration.Linkwidth.Start as Upstream Lanes.

With:

- ❑ Optionally, if crosslinks are supported and all Downstream Lanes initially receive two consecutive TS1 ordered sets with a Link number different than PAD (K23.7) **and a Lane Number set to PAD**, then the Downstream Lanes are now designated as Upstream Lanes and a new random cross Link timeout is chosen (see $T_{\text{crosslink}}$ in Table 0-4). The next state is Configuration.Linkwidth.Start as Upstream Lanes.

Replace the third paragraph of section 4.2.6.4.2:

"Next state is configuration if 8 consecutive TS1 order-sets are received on any configured Lanes with Link and Lane number that don't match what is being transmitted on those same lanes and 16 TS2 order-sets are sent after receiving one TS2 order-set ."

With :

"Next state is configuration if 8 consecutive TS1 order-sets are received on any configured Lanes with Link and Lane number that don't match what is being transmitted on those same lanes and 16 TS2 order-sets are sent after receiving one **TS1** order-set ."

In section 4.2.4.5.1 add the comments in red

- ❑ Fundamental Reset only applies when MAIN Power is present.
- ❑ Fundamental Reset does not apply when no power or only AUX power is present.

When Fundamental Reset is asserted:

- ❑ The receiver terminations are required to only meet $Z_{RX-HIGH-IMP-DC}$ (see Table 0-5).
- ❑ The transmitter terminations is required to only meet Z_{TX-DC} (see Table 0-5).
- ❑ The transmitter holds a constant DC common mode voltage.²²

When Fundamental Reset is de-asserted:

- ❑ The Port LTSSM (see Section **Error! Reference source not found.**) is initialized (i.e., Detect is immediately entered).

In table 4-5 make the following edits as listed.

Replace:

$V_{TX-CM-ACp}$	AC Peak Common Mode Output Voltage			20	mV	$V_{TX-CM-ACp} = V_{TX-D+} + V_{TX-D-} /2 - V_{TX-CM-DC}$ $V_{TX-CM-DC} = DC_{(avg)}$ of $ V_{TX-D+} + V_{TX-D-} /2$ during L0 See Note 2.
-----------------	------------------------------------	--	--	----	----	---

With:

$V_{TX-CM-ACp}$	AC Peak Common Mode Output Voltage			20	mV	$V_{TX-CM-ACp} = V_{TX-D+} + V_{TX-D-} /2 - V_{TX-CM-DC}$ $V_{TX-CM-DC} = DC_{(avg)}$ of $ V_{TX-D+} + V_{TX-D-} /2$ See Note 2.
-----------------	------------------------------------	--	--	----	----	---

Replace:

$V_{TX-CM-DC-LINE-DELTA}$	Absolute Delta of DC Common Mode Voltage between D+ and D-	0		25	mV	$ V_{TX-CM-DC-D+} [during L0] - V_{TX-CM-DC-D-} [During L0.] \leq 25 \text{ mV}$ $V_{TX-CM-DC-D+} = DC_{(avg)}$ of $ V_{TX-D+} $ [during L0] $V_{TX-CM-DC-D-} = DC_{(avg)}$ of $ V_{TX-D-} $
---------------------------	--	---	--	----	----	---

²² The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode during L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 4-5).

						[during L0] See Note 2.
--	--	--	--	--	--	----------------------------

With:

$V_{TX-CM-DC-LINE-DELTA}$	Absolute Delta of DC Common Mode Voltage between D+ and D-	0		25	mV	$ V_{TX-CM-DC-D+} - V_{TX-CM-DC-D-} \leq 25 \text{ mV}$ $V_{TX-CM-DC-D+} = DC_{(avg)} \text{ of } V_{TX-D+} $ $V_{TX-CM-DC-D-} = DC_{(avg)} \text{ of } V_{TX-D-} $ See Note 2.
---------------------------	--	---	--	----	----	---

Replace the term in Table 4-6:

$V_{RX-CM-ACp}$	AC Peak Common Mode Input Voltage			150	mV	$V_{RX-CM-AC} = V_{RX-D+} + V_{RX-D-} /2 - V_{RX-CM-DC}$ $V_{RX-CM-DC} = DC_{(avg)} \text{ of } V_{RX-D+} + V_{RX-D-} /2 \text{ during L0}$ See Note 7.
-----------------	-----------------------------------	--	--	-----	----	---

With:

$V_{RX-CM-ACp}$	AC Peak Common Mode Input Voltage			150	mV	$V_{RX-CM-AC} = V_{RX-D+} + V_{RX-D-} /2 - V_{RX-CM-DC}$ $V_{RX-CM-DC} = DC_{(avg)} \text{ of } V_{RX-D+} + V_{RX-D-} /2$ See Note 7.
-----------------	-----------------------------------	--	--	-----	----	---

Section 4.3.3 – Table 4-5, second page, line 11

Change line from:

Z_{TX-DC}	Transmitter DC Impedance	40 k			Ω	Required TX D+ as well as D- DC impedance during all states
-------------	--------------------------	------	--	--	----------	---

To (changes in red):

Z_{TX-DC}	Transmitter DC Impedance	40 Ohms			Ω	Required TX D+ as well as D- DC impedance during all states
-------------	--------------------------	------------	--	--	----------	---

C66 Assorted Errata

Release Date: 4/15/03

This section corrects errors that affect the technical meaning of the specification and clarifies ambiguous or incomplete areas of the specification.

In Section 4.3.3.2, edit as shown:

The test load is shown at the transmitter package reference plane. The same test/Measurement load is applicable to the receiver package reference plane **at the end of system interconnect in place of the receiver silicon and package.**

In Section 4.3.3.1, add footnote:

A recovered TX UI is calculated over 3500 consecutive unit intervals of sample data. The eye diagram is created using all edges of the 250 consecutive UI in the center of the 3500 UI used for calculating the TX UI²³.

Change the following text in Table 4-5 (section 4.3.4) in the comments section of TRX-EYE-MEDIAN-to-MAX-JITTER:

Jitter is defined as the measurement variation of the crossing points ($V_{RX-DIFFp-p} = 0$ V) in relation to ~~an appropriate average TX UI.~~ **recovered TX UI. A recovered TX UI is calculated over 3500 consecutive unit intervals of sample data. Jitter is measured using all edges of the 250 consecutive UI in the center of the 3500 UI used for calculating the TX UI.**

In 4.2.6.11, delete the note that was added in C65:

~~?—Note: Any optional crosslinks on the switch are an exception to this rule and the behavior is system specific.~~

In Section 4.3.3 – Table 4-5, second page, line 11, change edit made in C65:

Z_{TX-DC}	Transmitter DC Impedance	40 Ohms			Ω	Required TX D+ as well as D- DC impedance during all states
-------------	--------------------------	-----------------------	--	--	----------	---

In table 4-5 make the following edits as listed (shown as edits to C65 edits already made).

²³ **It is recommended that the recovered TX UI is calculated using all edges in the 3500 consecutive UI interval with a fit algorithm using a minimization merit function (i.e. Least squares and median deviation fits).**

$V_{TX-CM-ACp}$	RMS AC Peak Common Mode Output Voltage			20	mV	$V_{TX-CM-ACp} = \text{RMS}(V_{TX-D+} + V_{TX-D-} /2 - V_{TX-CM-DC})$ $V_{TX-CM-DC} = DC_{(avg)} \text{ of } V_{TX-D+} + V_{TX-D-} /2$ See Note 2.
...						
$V_{TX-CM-DC-LINE-DELTA}$	Absolute Delta of DC Common Mode Voltage between D+ and D-	0		25	mV	$ V_{TX-CM-DC-D+} - V_{TX-CM-DC-D-} \leq 25 \text{ mV}$ $V_{TX-CM-DC-D+} = DC_{(avg)} \text{ of } V_{TX-D+} $ $V_{TX-CM-DC-D-} = DC_{(avg)} \text{ of } V_{TX-D-} $ See Note 2.

In Table 4-6, edit as shown (shown as edits to C65 edits already made):

$V_{RX-CM-ACp}$	RMS AC Peak Common Mode Input Voltage			150	mV	$V_{RX-CM-AC} = \text{RMS}(V_{RX-D+} + V_{RX-D-} /2 - V_{RX-CM-DC})$ $V_{RX-CM-DC} = DC_{(avg)} \text{ of } V_{RX-D+} + V_{RX-D-} /2$ See Note 7.
-----------------	--	--	--	-----	----	---

In section 4.2.4.5.1 add the comments in red

- ☐ Fundamental Reset only applies when MAIN Power is present.
- ☐ Fundamental Reset does not apply when no power or only AUX power is present.

When Fundamental Reset is asserted:

- ☐ The receiver terminations are required to only meet $Z_{RX-HIGH-IMP-DC}$ (see <>).
- ☐ The transmitter terminations is required to only meet Z_{TX-DC} (see <>).
- ☐ The transmitter holds a constant DC common mode voltage.²⁴

In Section 4.2.5.9, edit as shown:

The intent of the Disabled state is to allow a **configured Link-group-of-Lanes** to be disabled until directed or Electrical Idle is exited (i.e., due to a hot removal and insertion) after entering Disabled

In Section 4.2.8, edit as shown:

²⁴ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode during L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 4-5).

During **Polling** the **Polling.Compliance** sub-state must be entered based on the presence of **passive** test equipment being attached to one Lane of a possible Link **and being Detected during the Detect**.

The compliance pattern consists of the sequence of 8b/10b symbols K28.5, D21.5, K28.5, and D10.2 repeating. ~~Current running disparity must be set to negative before sending the first symbol.~~

In Section 4.3.1.2, edit as shown:

The requirement for the inclusion of AC coupling capacitors on the interconnect media is specified ~~at~~by the transmitter.

In Section 4.3.1.2, edit as shown:

The output Beacon voltage level can range between ~~the pre-emphasized and a~~ specified voltage level (see $V_{TX-DIFFP-P}$ in <Table 4-5>) and a corresponding -3.5 dB de-emphasized voltage levels for Beacon pulses smaller than 500 ns.

In Section 4.2.6.2.1, edit as shown:

...

☐ Otherwise, after a 24 ms timeout the next state is:

1. Polling.Configuration if,
 - a) Any Lane, **which detected a receiver during Detect**, received eight consecutive TS1 or TS2 ordered sets...

...

2. Polling.Compliance if at least one Lane's receiver **which detected a receiver during Detect**in **Polling** has never detected an exit from **Electrical Idle** since entering Polling.Active.

In Section 4.2.6.2.3, edit as shown:

- ☐ The next state is Configuration after eight consecutive TS2 ordered sets **with Link and Lane numbers set to PAD (K23.7)** are received on any Lanes that detected a receiver during Detect ...
- ☐ The next state is Polling.Speed after eight consecutive TS2 ordered sets **with Link and Lane numbers set to PAD (K23.7)** are received on ~~any~~all Lanes that detected a receiver during Detect ...

In Section 4.2.6.3.1.1, edit as shown:

- ☐ If any Lanes first received at least one or more TS1 ordered sets with a Link and Lane number set to PAD (K23.7), the next state is Configuration.Linkwidth.Accept immediately after **any of those same Downstream Lanes receive two consecutive TS1 ordered sets with a non-PAD Link number that matches any of the transmitted Link numbers, and with a Lane number set to all those same Downstream Lanes receive two consecutive TS1 ordered sets with any Link number that was transmitted and different than** PAD (K23.7).

In Section 4.2.6.4.1, edit as shown:

The next state is Configuration if ~~any of all~~ the configured Lanes that are receiving a TS1 or TS2 ordered set have received at least one TS1 or TS2 with Link and Lane numbers that match what is being transmitted on those same Lanes.

In Section 4.2.6.4.2, edit as shown:

- ☐ Next state is Recovery. Idle if eight consecutive TS2 ordered sets ...

In Section 4.2.6.4.2, the Eye Diagram in figure 4-24 section 4.3.3.1 is shown as 0.07 UI but it should be 0.7 UI.

In Section 4.2.4.1, edit as shown:

~~SKIP ordered sets (see Section <=>) must be transmitted during but never interrupt a TS1 or TS2 ordered set. Training sequences (TS1 or TS2) are transmitted consecutively and can only be interrupted by SKIP order sets (see Section <4.2.7>).~~

In Section 4.2.7.2, edit as shown:

- ☐ Receivers shall be tolerant to receive and process consecutive SKIP ordered sets.
 - Note: Receivers shall be tolerant to receive and process SKIP ordered sets that have a maximum separation dependent on the Max_payload_size a component supports. The formula for the max number of symbols (N) between SKIP ordered-sets is: $N = 1538 + (\text{Max_payload_size_byte} + 26)$ ~~separated from each other at most 5664 symbol times—measured as the distance between the leading COM symbols.~~
 - e.g if Max_payload_size is 4096B , $N = 1538 + 4096 + 26 = 5660$

In Section 4.2.6.3.1.1, edit as shown:

- ☐ In the optional case where a crosslink is supported, the next state is **Disable** after all Lanes that detected a receiver during **Detect**, ~~that are also receiving TS1 ordered-sets, receive the Disable and are receiving TS1 ordered-sets with the Disable Link~~ bit asserted in two consecutive TS1 ordered sets.
- ☐ Next state is **Loopback** ~~after if~~ all Lanes that detected a receiver during **Detect** ~~that are also receiving receive the Loopback bit asserted in two consecutive TS1 ordered-sets receive the Loopback bit asserted in two consecutive on all Lanes receiving a TS1 ordered sets.~~

In Section 4.2.6.3.1.2, edit as shown:

- ☐ Next state is **Disable** after any Lanes that detected a receiver during **Detect** and are receiving TS1 ordered sets with the Disable Link bit asserted in two consecutive TS1 ordered sets.

- Note: In the optional case where a crosslink is supported, the next state is **Disable** only after all Lanes that detected a receiver during **Detect**, that are also receiving TS1 ordered sets, receive the ~~and are receiving TS1 ordered sets with the~~ Disable Link bit asserted in two consecutive TS1 ordered sets.
- ❑ Next state is **Loopback** after ~~if~~ all Lanes that detected a receiver during **Detect**, that are also receiving TS1 ordered sets, receive the Loopback bit asserted in two consecutive TS1 ordered sets. ~~sets on all Lanes receiving a TS1 ordered set.~~

In Section 4.2.6.3.1.2, replace this text:

- ❑ Immediately ... Configuration.Linkwidth.Accept.
- ❑ Optionally, ...
 - The ...
 - Immediately ...

with:

- ❑ If any Lanes receive two consecutive TS1 ordered sets with Link numbers that are different than PAD (K23.7) and Lane number set to PAD (K23.7), a single Link number is selected and transmitted on all Lanes that both detected a receiver and also received two consecutive TS1 ordered sets with Link numbers that are different than PAD (K23.7) and Lane number set to PAD (K23.7). Any left over Lanes that detected a receiver during Detect must transmit TS1 ordered sets with the Link and Lane number set to PAD (K23.7). The next state is Configuration.Linkwidth.Accept.
 - Note: It is recommended that any possible multi-Lane Link that received an error in a TS1 ordered sets on a subset of the received Lanes; delay the evaluation listed above by an additional 2 TS1 ordered sets so as not to pre-maturely configure a smaller Link than possible.
- ❑ Optionally, if crosslinks are supported and any Upstream Lanes first receive two consecutive TS1 ordered sets with Link and Lane numbers set to PAD (K23.7), then:
 - The transmitter continues to send out TS1 ordered sets with Link numbers and Lane numbers set to PAD (K23.7).
 - If any Lanes receive two consecutive TS1 ordered sets with Link numbers that are different than PAD (K23.7) and Lane number set to PAD (K23.7), a single Link number is selected and transmitted on all Lanes that both detected a receiver and also received two consecutive TS1 ordered set with Link numbers that are different than PAD (K23.7) and Lane number set to PAD (K23.7). Any left over Lanes that detected a receiver during Detect must transmit TS1 ordered sets with the Link and Lane number set to PAD (K23.7). The next state is Configuration.Linkwidth.Accept.
 - Note: It is recommended that any possible multi-Lane Link that received an error in a TS1 ordered sets on a subset of the received Lanes; delay the evaluation listed above by an

additional 2 TS1 ordered sets so as not to pre-maturely configure a smaller Link than possible.

In Section 4.2.4.3, edit as shown:

A single FTS training sequence is an ordered set composed of one K28.5 (COM) symbol and three K28.1 symbols. The maximum number of FTS ordered sets (N_FTS) that a component can request is 255, providing a bit time lock of $4 * 255 * 10 * UI$. 4096 FTS ordered sets must be sent when the Extended Synch bit is set in order to provide external Link monitoring tools with enough time to achieve bit and framing synchronization. SKP ordered sets must be scheduled and transmitted between FTS ordered sets as necessary to meet the definitions in <Section 4.2.7> with the exception that no SKP ordered sets can be scheduled during the first ~~N_FTS~~²⁵ FTS ordered sets. A single SKP ordered set is always sent after the last FTS is transmitted.

In Section 4.2.6.6.1.1, edit as shown:

- ❑ Next state is Rx_L0s.Idle after a $T_{TX-IDLE-MIN}$ ~~SET-TO-IDLE~~ (<see Table 4-5>) timeout
 - Note: This ~~is the minimum time the transmitter must be in a~~ **guarantees that the transmitter has established the** Electrical Idle condition.

In Section 4.2.6.8.1, edit as shown:

- ❑ All configured transmitters ~~are~~**must remain** in Electrical Idle **for a minimum time of $T_{TX-IDLE-MIN}$** (<see Table 4-5>).
 - Note: The DC common mode voltage does not have to be within specification.²⁵
 - Note: The receiver needs to wait a minimum of $T_{TX-IDLE-MIN}$ to start looking for Electrical Idle Exit

In section 4.2.6.11, edit as shown:

- ❑ Lanes that were directed by a higher Layer to initiate Hot Reset:
 - All Lanes in the configured Link transmit TS1 ordered sets with the Hot Reset bit (Bit 0) asserted and the configured Link and Lane numbers.
 - If two consecutive TS1 ordered sets **are received on any Lane** with the Hot Reset bit (Bit 0) asserted **and configured Link and Lane numbers are received**, then:
 - ? LinkUp = 0 (False)
 - ? ~~Next state is Detect~~**if no higher Layer is directing the Physical Layer to remain in Hot Reset the next state is Detect.**

²⁵ The common mode being driven ~~must~~**does not need to** meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see <Table 4-5>).

- ? Otherwise, all Lanes in the configured Link continue to transmit TS1 ordered sets with the Hot Reset bit (Bit 0) asserted and the configured Link and Lane numbers.

Otherwise, after a 2 ms timeout next state is Detect.

- ❑ Lanes that were not directed by a higher Layer to initiate Hot Reset (i.e., received two consecutive TS1 ordered sets with the Hot Reset bit (Bit 0) asserted on any configured Lanes):

LinkUp = 0 (False)

If any Lane of an Upstream Port of a Switch receives a training sequence with the Hot Reset bit asserted, all configured Downstream Ports must transition to **Hot Reset** as soon as possible.

- ? **Note: Any optional crosslinks on the switch are an exception to this rule and the behavior is system specific.**

All Lanes in the configured Link transmit TS1 ordered sets with the Hot Reset bit (Bit 0) asserted and the configured Link and Lane numbers.

~~After a 2 ms timeout:~~ If two consecutive TS1 ordered sets were received with the Hot Reset bit (Bit 0) asserted and the configured Link and Lane numbers, the next state is **Hot Reset**.

Otherwise, the next state is **Detect** after a 2 ms timeout.

In Section 4.2.6.10.1, edit as shown:

- ❑ The Loopback Master device transmits TS1 ordered sets with the Loopback bit (Bit 2) asserted until the Loopback Master receives identical TS1 ordered sets with the Loopback bit asserted **on an implementation specific number of Lanes**. The next state is **Loopback.Active**.
- Note: This indicates to the Loopback Master that the Loopback Slave has successfully entered **Loopback**.
 - Note: The Loopback Master timeout is implementation specific **but must be less than 100mS**. The exit is to **Loopback.Exit**.
 - **Note: A boundary condition exists when the Loopback Slave transitions to Loopback.Active that can cause the Loopback Slave to discard a scheduled SKP ordered set. If this occurs, the Loopback Master may not see a SKP ordered set for twice the normal SKP ordered-set scheduling interval.**
- ❑ The next state for the Loopback Slave is **Loopback.Active**.
- Note: The **Loopback Slave must transition on a transmit symbol boundary, and may truncate any ordered set in progress. Loopback Slave will immediately transition to Loopback.Active.**

In Section 4.2.6.10.2, edit as shown:

- ❑ Next state of the Loopback Slave is Loopback.Exit when an Electrical Idle ordered set is received or Electrical Idle is detected **on any Lane**.
- Note: A Loopback Slave must be able to detect **an Electrical Idle condition on any Lane within 1 mS of the Electrical Idle ordered set being received by the Loopback Slave. -was entered within 2 ms in case the Electrical Idle ordered set is not properly detected.**

- Note: During the time after an Electrical Idle ordered-set is received and before Electrical Idle is actually detected by the Loopback Slave the Loopback Slave may receive and transmit undefined 10-bit data.
- The $T_{TX-IDLE-SET-TO-IDLE}$ parameter does not apply in this case since the Loopback Slave may not even detect Electrical Idle until as much as 1 ms after the Electrical Idle Ordered Set.

In Section 4.2.6.10.3, edit as shown:

- ❑ The Loopback Master sends an Electrical Idle ordered set and ~~enters Electrical Idle on all Lanes for 2 ms goes to Electrical Idle for a minimum of 2 ms (<=>).~~
 - The Loopback Master must transition to a valid Electrical Idle condition²⁶ on all Lanes within $T_{TX-IDLE-SET-TO-IDLE}$ after sending the Electrical Idle ordered-set.
 - Note: The Electrical Idle ordered set can be useful in signifying the logical end of transmit and compare operations that occurred by the Loopback Master. Any data received by the Loopback Master after the Electrical Idle ordered set should be ignored since it is undefined.
- ❑ The Loopback Slave ~~must enter~~~~echoes the~~ Electrical Idle on all Lanes for ~~ordered-set and goes to Electrical Idle for a minimum of 2 ms (<=>).~~
 - Note: Before entering Electrical Idle the Loopback Slave must Loopback all symbols that were received prior to detecting Electrical Idle. This ensures that the Loopback Master may see the Electrical Idle ordered-set to signify the logical end of any Loopback send and compare operations.

In Section 4.2.6.9, edit footnote as shown:

- ❑ All Lanes transmit 16 TS1 ordered sets with the Disable Link bit (bit 1) asserted and then transition to Electrical Idle.
 - Note: The Electrical Idle ordered set must be sent prior to entering Electrical Idle.
 - Note: The DC common mode voltage does not have to be within specification.²⁷

In Section 4.2.6.2.3, edit as shown:

- ❑ The next state is Polling.Speed after eight consecutive TS2 ordered sets are received on ~~any~~~~all~~ Lanes that detected a receiver during Detect ...

²⁶ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see <Table 4-5>).

²⁷ The common mode being driven ~~must~~~~does not need to~~ meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see <Table 4-5>).

In Section 4.2.7.1, add bullet as shown:

- ❑ SKIP ordered -sets do not count as an interruption when monitoring for consecutive characters or ordered set (i.e., eight consecutive TS1 ordered sets in **Polling.Active**)
- ❑ SKIP ordered -sets must not be transmitted while the Compliance Pattern (see section <4.2.8>) is in progress during **Polling.Compliance**.

Add text as shown preceeding the text under figure 4-24:

The TX eye diagram in figure 4-24 is specified using the passive compliance/test measurement load (see figure 4-24) in place of any real PCI Express interconnect + RX component.

There are two eye diagrams that must be met for the transmitter. ...

Add text as shown preceeding the text under figure 4-26:

The RX eye diagram in figure 4-26 is specified using the passive compliance/test measurement load (see figure 4-24) in place of any real PCI Express RX component.

Note: In general, the minimum receiver eye diagram measured with the compliance/test measurement load (see figure 4-24) will be larger than the minimum receiver eye diagram measured over a range of systems at the input receiver of any real PCI Express component. The degraded eye diagram at the input receiver is due to traces internal to the package as well as silicon parasitic characteristics which cause the real PCI Express component to vary in impedance from the compliance/test measurement load. The input receiver eye diagram is implementation specific and is not specified. RX component designer should provide additional margin to adequately compensate for the degraded minimum receiver eye diagram (shown in figure 4-26) expected at the input receiver based on some adequate combination of system simulations and the Return Loss measured looking into the RX package and silicon²⁸.

The RX eye diagram must be aligned in time using the jitter median to locate the center of the eye diagram.

In Section 4.2.6.5, edit as shown:

- ❑ Next state is Recovery if directed to this state or if Electrical Idle is detected **on all Lanes** without receiving an Electrical Idle ordered set **on any Lane**.

Note: "if directed" applies to a Port that is instructed by a higher Layer to transition to Recovery.

Note: The transmitter may complete any TLP or DLLP in progress.

- ❑ Next state ~~of transmitter~~ is L0s **for only the transmitter** if directed to this state.

Note: "if directed" applies to a Port that is instructed by a higher Layer to initiate L0s.

Note: This is a point where the TX and RX may diverge into different LTSSM states

²⁸ The reference impedance for return loss measurements is 50 ohms to ground for both the D+ and D- line (i.e., as measured by a Vector Network Analyzer with 50 ohm probes - see <Figure 4-25>). Note: that the series capacitors CTX is optional for the return loss measurement.

- ❑ Next state ~~of receiver~~ is L0s ~~for only the receiver if an receiver detects~~ Electrical Idle ~~ordered set is received on any Lanes and the Port is not directed to L1 or L2 states by any higher Layers~~

Note: This is a point where the TX and RX may diverge into different LTSSM states

- ❑ Next state is L1

- vii. If directed

- and

- viii. an Electrical Idle ordered set is received **on any Lane**

- and

- ix. an Electrical Idle ordered set is transmitted **on all Lanes**.

Note: “if directed” is defined as both ends of the Link having agreed to enter L1 immediately after the condition of both the receipt and transmission of an Electrical Idle ordered set is met (see section <4.3.2.1>).

Note: When directed by a higher Layer one side of the Link always initiates and exits to L1 by transmitting an Electrical Idle ordered set **on all Lanes**, followed by a transition to Electrical Idle.²⁹ The same Port then waits for the receipt of an Electrical Idle ordered set **on any Lane**, and then immediately transitions to L1. Conversely, the side of the Link that first receives an Electrical Idle ordered set **on any Lane** must send an Electrical Idle ordered set **on all Lanes** and immediately transition to L1.

- ❑ Next state is L2:

- x. If directed

- and

- vii. an Electrical Idle ordered set is received **on any Lane**

- and

- viii. an Electrical Idle ordered set is transmitted **on all Lanes**.

Note: “if directed” is defined as both ends of the Link having agreed to enter L2 immediately after the condition of both the receipt and transmission of an Electrical Idle ordered set is met (see section <4.3.2.3>).

Note: When directed by a higher Layer, one side of the Link always initiates and exits to L2 by transmitting an Electrical Idle ordered set **on all Lanes** followed by a transition to Electrical Idle.³⁰ The same Port then waits for the receipt of an Electrical Idle ordered set **on any Lane**, and then immediately transitions to L2. Conversely, the side of the Link that first receives an Electrical Idle ordered set **on any Lane** must send an Electrical Idle ordered set **on all Lane** and immediately transition to L2.

²⁹ The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see <>).

³⁰ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see <>).

In Section 4.2.6.3.2.1, edit as shown:

- ❑ If a **configured** Link can be formed with at least one group of Lanes that received two consecutive TS1 ordered sets with the same received Link number (non PAD and matching one that was transmitted by the Downstream Lanes), then TS1 ordered sets are transmitted with the same Link number and unique **non PAD** Lane numbers are assigned to all these same Lanes. ~~(Lane numbers must range sequentially from 0 to n-1). Any left over Lanes must transmit TS1 ordered sets with the Link and Lane number set to PAD (K23.7).~~ The next state is Configuration.Lanenum.Wait.
 - Note: The assigned non PAD Lane numbers must range from 0 to n-1, be assigned sequentially to the same grouping of Lanes that are receiving the same Link number Lane numbers, and Downstream Lanes which aren't receiving TS1 ordered -sets must not disrupt the initial sequential numbering of the widest possible Link. Any left over Lanes must transmit TS1 ordered sets with the Link and Lane number set to PAD (K23.7).
 - Note: It is recommended that any possible multi-Lane Link that received an error in a TS1 ordered sets on a subset of the received Lanes; delay the evaluation listed above by an additional 2 TS1 ordered sets so as not to pre-maturely configure a smaller Link than possible.
 - Note: A couple of interesting cases to consider here are the following:
 1. A x8 Downstream Port, ...
 2. A x16 Downstream Port, ...
 3. A x8 Downstream Port where only 7 Lanes are receiving TS1s with the same received Link number (non PAD and matching one that was transmitted by the Downstream Lanes) and an eighth Lane, which is in the middle or adjacent to those same Lanes, is not receiving a TS1 ordered-set. In this case, the eighth Lane is treated the same as the other 7 Lanes and Lane numbering for a x8 Lane should occur as described above.

In Section 4.2.6.3.2.2, edit as shown:

- ❑ If a **configured** Link can be formed using Lanes that transmitted a **non PAD** Link number ~~which are receiving-receive~~ two consecutive TS1 ordered sets with the same ~~(non PAD)~~ Link number and any non-PAD Lane number then TS1 Lane numbers are transmitted that if possible match the received Lane numbers or are different if necessary (i.e., Lane reversed). ~~-include at least a Lane number 0, TS1 Lane numbers are transmitted that if possible match the received Lane numbers or are different if necessary (i.e., Lane reversed). Remaining Lanes must transmit TS1 with Link and Lane numbers set to PAD (K23.7).~~ The next state is Configuration.Lanenum.Wait.

Note: The transmitted Lane numbers must range from 0 to m-1, be assigned sequentially **only** to some continuous~~the same~~ grouping of Lanes that are receiving **non PAD** Lane numbers (i.e. Lanes which aren't receiving any TS1 ordered sets always disrupt a continuous grouping and must not be included in this grouping)~~Lane numbers 0 to n-1, include the received Lane 0,~~ and m-1 must be equal to or smaller than the largest received Lane number (n-1). Remaining Lanes must transmit TS1 with Link and Lane numbers set to PAD (K23.7).

Note: It is recommended that any possible multi-Lane Link that received an error in a TS1 ordered sets on a subset of the received Lanes; delay the evaluation listed above by an

additional 2 TS1 ordered sets so as not to pre-maturely configure a smaller Link than possible.

Note: A few interesting cases to consider here are the following:

1. ~~An x8 Upstream Port is attached to a x8 Downstream Port that~~ is presented with Lane numbers that are backward from the preferred numbering. If the optional ... exit Configuration.
 - Note: Optional Lane reversal behavior is required to configure a Link where the Lane numbers are reversed and the Downstream Port does not support Lane reversal. Specifically, the Upstream Port Lane reversal will accommodate the scenario where the default Upstream sequential Lane numbering (0 to n-1) is receiving a reversed Downstream sequential Lane number (n-1 to 0).
2. A x8 Upstream Port is not receiving TS1 ordered-sets on the Upstream Port Lane 0.
 - i. In the case where the Upstream Port can only support a x8 or x1 Link and the Upstream port can support Lane reversal. The Upstream Port will assign a Lane 0 to only the received Lane 7 (received Lane number n-1) and the remaining 7 Lanes must transmit TS1 with Link and Lane numbers set to PAD (K23.7)
 - ii. In the case where the Upstream Port can only support a x8 or x1 Link and the Upstream port can't support Lane reversal. No Link can be formed and the Upstream Port will eventually timeout after 2mS and exit to Detect.

~~Optional Lane reversal behavior is required to configure a Link where the Lane numbers are reversed and the Downstream Port does not support Lane reversal. Specifically, the Upstream Port Lane reversal will accommodate the scenario where the default Upstream sequential Lane numbering (0 to n-1) is receiving a reversed Downstream sequential Lane number (n-1 to 0).~~

3. An optional x8 Upstream crosslink Port, ...

In Section 4.2.6.3.3.1, edit as shown:

- ~~If two consecutive TS1 ordered sets are received with non PAD Link and non PAD Lane numbers that match all the non PAD Link and non PAD Lane numbers (or reversed Lane numbers if Lane reversal is optionally supported) that are being transmitted in Downstream Lane TS1 ordered sets; If a configured Link can be formed with all the Lanes that receive two consecutive TS1 ordered sets with the same transmitted Link numbers and non PAD Lane numbers,~~ the next state is Configuration.Complete.

Note: Reversed Lane numbers is defined strictly as the downstream Lane 0 receiving a TS1 ordered set with a Lane number equal to n-1 and the downstream Lane n-1 receiving a TS1 ordered set with a Lane number equal to 0.

Note: It is recommended that any possible multi-Lane Link that received an error in a TS1 ordered sets on a subset of the received Lanes; delay the evaluation listed above by an additional 2 TS1 ordered sets so as not to pre-maturely configure a smaller Link than possible.

Note: ~~Two possible scenarios exist that can result in the next state being Configuration.Complete:~~

1. ~~The received Link and Lane numbers are the same as transmitted, which is always required to be supported.~~
 2. ~~Optional Lane reversal behavior can occur for either scenario 1. An example of where Lane reversal is required is when the received Lane numbers are the reverse of the Lane numbers being transmitted (case 1). This optional feature would be required to configure a Lane reversed Link where the Upstream Port does not support Lane reversal.~~
- If ~~the Lanes of a configured Link can be formed with any subset of configured by using the Lanes that transmitted a Lane number, which received~~ two consecutive TS1 ordered sets with the same transmitted non PAD Link numbers and any (non-Pad) Lane numbers, then ~~and include a Lane number 0, TS1 ordered sets Lane numbers are transmitted with new Lane numbers assigned and the which must be assigned sequentially either in order or Lane reversed to the same Lanes and no Lane numbers can be repeated. Left over Lanes must transmit TS1 Link and Lane numbers set to PAD (K23.7). The next state is Configuration.Lanenum.Wait.~~

Note: The newly assigned transmitted Lane numbers must range from 0 to m-1, be assigned sequentially only to some continuous grouping of the Lanes that are receiving non PAD Lane numbers (i.e. Lanes which aren't receiving any TS1 ordered sets always disrupt a continuous grouping and must not be included in this grouping), must include either Lane 0 or n-1 (largest received Lane number), and m-1 must be smaller than the largest received Lane number (n-1). Any left over Lanes must transmit TS1 ordered sets with the Link and Lane number set to PAD (K23.7).

Note: It is recommended that any possible multi-Lane Link that received an error in a TS1 ordered sets on a subset of the received Lanes; delay the evaluation listed above by an additional 2 TS1 ordered sets so as not to pre-maturely configure a smaller Link than possible.

In Section 4.2.6.3.3.2, edit as shown:

- If two consecutive TS2 ordered sets are received with non PAD Link and non PAD Lane numbers ~~(non PAD values)~~ that match all non PAD Link and non PAD Lane numbers that are being transmitted in Upstream Lane ~~what is being transmitted in the~~ TS1 ordered sets, the next state is Configuration.Complete.
- If ~~a configured the Lanes of a Link can be formed with any subset of configured by using the Lanes that transmitted a Lane number, which received~~ two consecutive TS1 ordered sets with the same transmitted non PAD Link numbers ~~(non-Pad)~~ and any non-PAD Lane numbers ~~include a Lane number 0;~~ then TS1 ordered sets Lane numbers are transmitted with new Lane numbers assigned and the which must be assigned sequentially either in order or Lane reversed to the same Lanes and no Lane numbers can be repeated. Left over Lanes must transmit TS1 Link and Lane numbers set to PAD (K23.7). The next state is Configuration.Lanenum.Wait.

Note: The newly assigned transmitted Lane numbers must range from 0 to m-1, be assigned sequentially only to some continuous ~~in order or Lane reversed to the same~~ grouping of Lanes that are receiving non PAD Lane numbers (i.e. Lanes which aren't receiving any TS1 ordered sets always disrupt a continuous grouping and must not be included in this grouping) ~~0 to n-1, include the received Lane 0,~~ must include either Lane 0 or n-1 (largest received Lane number), and m-1 must be smaller than the largest received Lane number (n-1). Any left over Lanes must transmit TS1 ordered sets with the Link and Lane number set to PAD (K23.7).

Note: It is recommended that any possible multi-Lane Link that received an error in a TS1 ordered sets on a subset of the received Lanes; delay the evaluation listed above by an additional 2 TS1 ordered sets so as not to pre-maturely configure a smaller Link than possible.

C67 Add Item to Hwlnit Definition

Release Date: 4/15/03

In Section 7.4, Table 7-2, add as shown:

Table 7-2: Register (and Register Bit-Field) Types

Register Attribute	Description
Hwlnit	Hardware Initialized: Register bits are initialized by firmware or hardware mechanisms such as pin strapping or serial EEPROM. (System firmware hardware initialization is allowed is only allowed for system integrated devices.) Bits are read-only after initialization and can only be reset (for write-once by firmware) with Fundamental Reset (see Section 0).
...	

Editorial Errata – Chaps 1-3, 5-7

This section corrects typographical, grammatical, and minor wording errors that do not affect the technical meaning of the specification. These corrections will be included in the next printing of this revision of the specification.

E1. Section Cross References

Release Date: 2/11/03

The following section cross references are incorrect.

Page	Location	Correction
246	"configuration status register defined in Section 7.6."	" configuration status Link Capabilities register defined in Section 7.8.6."
33	"Locked Requests must be supported as specified in Section 6.2."	"Locked Requests must be supported as specified in Section 6.52."

E2. Typographical Errors.

Release Date (1-20): 2/11/03

Release Date (21-23): 3/27/03

Release Date (24): 4/15/03

1. Section "Document Organization", page 17, add text as shown:

The PCI Express Base Specification is applicable to all **variants of PCI Express**.

2. Section 2.6.1.1, p.111, Remove extra paren:

☐ The Transmitter gating function test is performed as follows:

- For each required type of credit, the number of credits required is calculated as:

CUMULATIVE_CREDITS_REQUIRED =

(CREDITS_CONSUMED +
<credit units required for pending TLP>) mod 2^[Field Size]

3. Section 5.3.3.3, p.239, edit as shown:

PM_PME messages are posted Transaction Layer Packets (TLPs) that inform the power management software which agent within the PCI Express Hierarchy requests a PM state change. PM_PME messages, like all other Power Management system messages, must use the general purpose **Transfer-Traffic** Class, TC #0.

4. Section 5.2, p.224, edit as shown:

... the acknowledgment of a PME_~~TURN~~~~Turn_OFF~~~~Off~~ message, (i.e., the injection of a ...

5. Section 5.3.2, p.230, edit as shown:

2. When all functions within a downstream component are programmed to D3hot the downstream component must request the transition of its Link to the L1 state using the PM_~~ENTER~~~~Enter~~_L1 DLLP. Once in D3hot, following the execution of a PM_~~TURN~~~~Turn_OFF~~~~Off~~/PME_TO_Ack handshake sequence, the downstream component must then request a Link transition to L2/3 Ready using the PM_~~ENTER~~~~Enter~~_L23 DLLP. Following the L2/L3 Ready entry transition protocol the downstream component must be ready for loss of main power and reference clock.

6. Section 5.4.1.2.1, p.249, edit as shown:

... immediately issue a TLP after it exists the L1 state.

7. Terms and Acronyms, p.22, edit as shown:

Physical Layer	The Layer of the that directly interacts with the communication medium between two components.
----------------	---

8. Section 2.2.8.1.1, p.69, edit as shown:

☐ INTx Interrupt Signaling is disabled when the Interrupt Disable bit of the Command Register (see Section <!>) **is set to 1b.**

9. Section 2.2.8.1.7, p.77-78, edit as shown:

Table 2-19: Hot Plug Signaling Messages

Name	Code[7:0]	Routing r[2:0]	Support				Req ID	Description/Comments
			R C	E p	S w	B r		
Attention_Indicator_On	0100 0001	100	t	r	tr	r	BDF	This message is issued by the Switch/Root Port when the Attention Indicator Control is set to 01b. The end device receiving the message will terminate the message and initiate appropriate action for toto cause the Attention Indicator located on the card to turn on. See note.
Attention_Indicator_Blink	0100 0011	100	t	r	tr	r	BDF	This message is issued by the Switch/Root Port when the Attention Indicator Control is set to 10b. The end device receiving the message will terminate the message and initiate appropriate action for toto cause the Attention Indicator located on the card to blink. See note.
Attention_Indicator_Off	0100 0000	100	t	r	tr	r	BDF	This message is issued by the Switch/Root Port when the Attention Indicator Control is set to 11b. The end device receiving the message will terminate the message and initiate appropriate action for toto cause the Attention Indicator located on the card to turn off. See note.

Name	Code[7:0]	Routing r[2:0]	Support				Req ID	Description/Comments
			R C	E p	S w	B r		
Power_Indicator_On	0100 0101	100	t	r	tr	r	BDF	This message is issued by the Switch/Root Port when the Power Indicator Command is set to 01b. The end device receiving the message will terminate the message and initiate appropriate action for toto cause the Power Indicator located on the card to turn on. See note.
Power_Indicator_Blink	0100 0111	100	t	r	tr	r	BDF	This message is issued by the Switch/Root Port when the Power Indicator Control is set to 10b. The end device receiving the message will terminate the message and initiate appropriate action for toto cause the Power Indicator located on the card to blink. See note.
Power_Indicator_Off	0100 0100	100	t	r	tr	r	BDF	This message is issued by the Switch/Root Port when the Power Indicator Command is set to 11b. The end device receiving the message will terminate the message and initiate appropriate action for toto cause the Power Indicator located on the card to turn off. See note.
Attention_Button_Pressed	0100 1000	100	t	r	tr	r	BDF	This message is issued by a device in a slot that implements an Attention Button on the card to signal the Switch/Root Port to generate the Attention Button Pressed Event. The Switch Switch/Root Port terminates the message and sets the Attention Button Pressed register to 1b which may result in an interrupt being generated.

10. Section 2.3, p.84-85, edit as shown:

❑ Request according to the routing mechanism indicated in the r[2:0] sub-field of the Type field

- If the value in r[2:0] indicates the Msg/MsgD is routed to the Root Complex (000b), the Switch must route the Msg/MsgD to the Upstream Port of the Switch
 - ? It is an error to receive a Msg/MsgD Request specifying 000b routing at the Upstream Port of a Switch. Switches may check for violations of this rule – TLPs in violation are Malformed TLPs. If checked, this is a reported error associated with the Receiving Port (see Section <>)
- If the value in r[2:0] indicates the Msg/MsgD is routed by address (001b), the Switch must route the Msg/MsgD in the same way it would route a Memory Request by address
- If the value in r[2:0] indicates the Msg/MsgD is routed by ID (010b), the Switch must route the Msg/MsgD in the same way it would route a Completion by ID
- If the value in r[2:0] indicates the Msg/MsgD is a broadcast from the Root Complex (011b), the Switch must route the Msg/MsgD to all Downstream Ports of the Switch
 - ? It is an error to receive a Msg/MsgD Request specifying 011b routing at the Downstream Port of a Switch. Switches may check for violations of this rule – TLPs in violation are Malformed TLPs. If checked, this is a reported error associated with the Receiving Port (see Section <>)
- If the value in r[2:0] indicates the Msg/MsgD terminates at the Receiver (100b or a reserved value), or if the Message Code field value is defined and corresponds to a Message which must be comprehended by the Switch, the Switch must process the message according to the Message processing rules

11. Section 2.3.1, p.89 (in Impl. Note), edit as shown:

Some devices require a lengthy self-initialization sequence **to** complete before they are able to service Configuration Requests... Note that **it** is only legal to ...

12. Section 2.6.1.2, p.113 (in Impl. Note), edit as shown:

For example, since Non-Posted Writes are only allowed on Virtual Channel 0, ~~so~~ there is no need...

13. Section 6.7.1.1, p.292, edit as shown:

...causes the solution to be non-PCI Express ~~complaint-compliant~~ and...

14. Section 1.3.3, p.33, edit as shown:

All Switches are governed by the following base rules (~~advanced~~ **Advanced** Switch components will support additional capabilities beyond those described below).

15. Section 1.4, p.34, edit as shown:

... a PCI Express ~~Hierarchy~~ hierarchy domain from a ...

16. Section 2.2.2, p.51, edit as shown:

... access to cross a 4KB boundary.

17. Section 2.8, p.121, edit as shown:

... granularity ~~than~~ the minimum ...

18. Section 5.3.2, p.230, edit as shown:

... a Link transition to L2/L3 Ready ...

19. Section 5.3.2.1, p.232, edit as shown:

... L1 state."

20. Section 5.4.1.2.1, p.247, edit as shown:

☐ PM_Request_Ack~~CK~~ (DLLP)

21. Throughout, for consistency: Check throughout and correct all cases - The message is called "PME_Turn_Off" – all other variations (e.g. "PM_Turn_Off", "PM_TURN_OFF", etc.) must be corrected.

22. In terms:

UR A Request Packet that specifies some action or access to some space that is not supported by the ~~Completer~~Target.

23. Chapter 3 pg 152, replace "Acknowledgement DLLP" with "ACK DLLP" , as shown:

In addition to the other requirements for sending Ack DLLPs, an Ack or Nak DLLP must be transmitted when all of the following conditions are true:

The Data Link Control and Management State Machine is in the DL_Active state

TLPs have been accepted, but not yet acknowledged by sending an ~~Acknowledgement~~Ack DLLP

"The AckNak_LATENCY_TIMER reaches or exceeds the value specified in Table 3-5

Data Link Layer ~~Acknowledgement~~Ack DLLPs may be Transmitted more frequently than required"

24. A number of capitalization, grammatical, spelling and formatting editorial errata were fixed at the discretion of the technical writer.

E3. Formatting Errors.

Release Date (1): 2/11/03

Release Date (2): 4/15/03

1. Section 3.5.3.1, p.150, clauses are incorrectly shown with separate bullets – remove the bullets (leave text):

❑ If the TLP Sequence Number is not equal to the expected value, stored in NEXT_RCV_SEQ:

- discard the TLP and free any storage allocated for the TLP
- If the TLP Sequence Number satisfies the following equation:

$$2 - (\text{NEXT_RCV_SEQ} - \text{TLP Sequence Number}) \bmod 4096 \leq 2048$$
~~2~~ the TLP is a duplicate, and an Ack DLLP is scheduled for transmission (per transmission priority rules)
- Otherwise, the TLP is out of sequence (indicating one or more lost TLPs):
 - ? if the NAK_SCHEDULED flag is clear,
 - | schedule a Nak DLLP for transmission
 - | set the NAK_SCHEDULED flag
 - | report TLP missing

This is a reported error associated with the Port (see Section 6.2).

2. A number of capitalization, grammatical, spelling and formatting editorial errata were fixed at the discretion of the technical writer.

E4. Minor Wording Errors.

Release Date (1-15): 2/11/03

Release Date (16): 3/27/03

Release Date (17): 4/15/03

1. Section 2.6.1, p.108, edit as shown:

❑ An InitFC1, InitFC2, or UpdateFC FCP that specifies a Virtual Channel ~~as that is~~ disabled is discarded without effect


2. Section 2.2.2, p. 52, edit as shown:

❑ The value in the Length field applies only to data – the ~~Transaction-TLP~~ Digest is not included in the Length

3. Section 2.3, p. 83, edit as shown:

❑ All Received TLPs which fail the required (and implemented optional) checks of TLP formation rules described in this ~~section~~chapter, ...

4. Section 7.10.11, p. 371, Figure 7-37 “Error Source Identification Register”, edit as shown:



5. Section 1.3.3, p. 33, edit as shown:

- ❑ A Switch is not allowed to split a packet into smaller packets, e.g., a single packet with a 256-byte payload must not be divided into two packets ~~each~~ of 128 bytes payload ~~each~~.

6. Section 2.2.2, p.51, edit as shown:

- ❑ For TLPs, that include data, the value in the Length field and the actual amount of data included in the TLP must ~~be equal~~ match.

7. Section 2.7.2, p.119, edit as shown:

... The rules for doing this are specified ~~below~~ in Section <2.7.2.2>.

8. Section 3.5.2.1, p.144, edit as shown:

... number of maximum size TLPs which ~~that~~ can be received ...

9. Section 6.2.3.1, p.265, edit as shown:

... between the ~~first initial~~ Request and the reissued Request.

10. Section 6.8, p.305, edit as shown:

... actual power ~~management-budgeting~~ capabilities ...

11. Section 7.11.9, p.384, edit Figure caption as shown:

Figure 7-47: ~~Structure of an Example VC Arbitration Table with 32- Phases.~~

12. Section 2.1.1.4, p.46, edit as shown:

... support for vendor-defined messages using ~~specific-specified reserved-~~ message codes ~~given in this document~~.

13. Section 2.2.6.2, Implementation Note “Increasing the Number of Outstanding Requests”, p.60, edit as shown:

To increase the maximum possible number of outstanding Requests requiring Completion beyond 256, a ~~single function~~ device may, if the Phantom Function Number Enable bit is set (see Section <!!7.8.3>7.8.4), use Function Numbers not assigned to implemented functions to logically extend the Tag identifier. ~~For a single function device, this can allow~~ing up to an 8-fold increase in the maximum number of outstanding Requests

14. Section 2.2.8.1.1, p.70, edit as shown:

- ~~For a~~Any INTx virtual wires that are active when the Interrupt Disable bit is set, ~~the active wires must~~ must be deasserted by transmitting the appropriate Deassert_INTx Message(s)

15. Section 2.3, p.84, edit as shown:

- ❑ If the value in the Type field indicates the TLP is not a Msg or MsgD Request, the TLP must be routed according to the ~~Switch routing rules~~ routing mechanism used (see Sections <2.2.4.1> and <2.2.4.2>)

16. Section 7.11.10, p.384, edit as shown:

The Port Arbitration Table Entry Size field in the ~~Port VC Capability Register 1~~ ~~VC Resource Capability Register 1~~ determines the table entry size. ...

17. A number of capitalization, grammatical, spelling and formatting editorial errata were fixed at the discretion of the technical writer.

Content and Editorial Errata – Chapter 4

This section includes the revised text for Chapter 4 highlighting changes made since the 11th of December and approved by the EWG on 20 December 2002. Changes made from 22 Jul to 11 December are shown in Appendix A.

Note that page numbering in the later part of this section reflects the page numbers of the modified spec document, and does not correspond to the page numbering of this document itself. Document page numbering resumes in the following sections.

Description of modifications to “PCI Express Base Spec Chapter 4

Per 12/19/02 version.

Note: Simple type-Os and obvious editorial changes are not listed here.

4.2.1: Removed paragraph about optional checking of encoding rules. The required behavior is better described in a later section.

Table 4-1: Fixed type-Os in table and cleaned up wording around Electrical Idle symbol.

4.2.1.3: Clarified required decoding checks and that the Physical layer must notify the data link layer when an error occurs.

4.2.2: Added clarifications that ordered-sets are always transmitted on all lanes of a multi-lane link. Fixed Type-Os.

4.2.2.1: Fixed Type-O. Fixed mistake in note regarding DLLPs and TLP on a multi-lane link at the same time. Fixed a format error. Added statement on what the physical layer does if it detects a framing error.

4.2.3: Fixed Type-O. Inserted clarification that the mechanism the physical layer uses to notify the data link layer of an error is beyond the scope of this spec. Added clarifications of when scrambling is or is not disabled as well as stating it does not apply when in loopback.

4.2.3 Implementation note: Cleaned up clumsy wording.

4.2.4.1: Inserted clarifications regarding the uses of the training ordered-sets. Also corrected the name of the Hot Reset, and disable link and loopback states and their associated bits. Language referring to sending

“at least” 16 TS2s (or Idles) after receiving a TS2 (or Idle) was removed to specify an exact and deterministic TS2 (Idle) count instead of making it open ended.

Tables 4-2 and 4-3: Clarified names of Hot Reset, Disabled, loopback and scrambling bits to be consistent with the rest of the document.

4.2.4.2: Added word independently as clarification.

4.2.4.3: Added clarifications around sending FTS ordered-set; schedule SKPs, don't send until after the first 255 FTS are set, as well as making it clear one SKP ordered set is always sent after the last FTS. Clarification added regarding what to do if you have not realigned (bit and symbol) after the FTS and SKP.

4.2.4.4: Clarify what the physical layer must do when an 8b/10b decode error, symbol lock or lane-lane de-skew error occurs.

4.2.4.5: Refer to link reset using bit in TS1/TS2 as Hot Reset (to be consistent with the rest of the spec).

4.2.4.5.1: Fix poor wording and describe link state during physical layer reset.

4.2.4.5.2: Use term Hot Reset for consistency through out document.

4.2.4.6: Improved wording about all link initialization starts with generation 1 data rate.

4.2.4.7: This section (Link Width and Lane Sequence Negotiation) was made considerably simpler. The required and optional behavior remains here, with some cleaned up wording for clarity. Previously, this section contained a description of the steps required to negotiate the width of a link and ordering of the lanes within the link. The same steps are required, but they have been moved to be part of the LTSSM. Specifically in the Configuration state.

4.2.4.8: Add clarifying words and a reminder that lane-lane de-skew is done in Configuration, Recovery and L0s.

4.2.4.9: Add clarifications around initializing lanes and then combining lanes into configured links.

4.2.5: Added line referencing figure 4-11.

4.2.5.1: Improved (more concise) wording for Detect description.

4.2.5.2: Changed word character to symbol and added cross reference to other section. Added clarification to compliance implementation note (compliance can not be disabled).

4.2.5.3: Inserted wording around entering and exit the configuration state – restricting the number of corner cases.

4.2.5.4: Add Clarification that recovery state is used for bringing a previously configured link back the L0 state.

4.2.5.5: Add line the states all lower power link states are entered from L0.

4.2.5.6: Add clarification L1 is entered only after prior agreement and the receipt of an electrical idle ordered-set. Added reminder to readers the lane-lane de-skew is checked/done in a L0s->L0 transition.

4.2.5.8: Add note stating the ports only need to support beacon if the system and form factors support it and that beacon propagate upstream.

4.2.5.9: Insert an improved description of the Disabled state. State the 2 consecutive IDL symbols must be detected in the electrical idle ordered set.

4.2.5.10: Add clarifications to the description of the loopback state.

4.2.5.11: Use term Hot Reset when resetting a link with the bit in TS1/TS2 (for consistency through out spec) and improve the description. . State the 2 consecutive IDL symbols must be detected in the electrical idle ordered set.

4.2.6: Add cross references to additional sections for clarity. Insert table 4-4, describing the state of link related status bits when links are in various states (mapped into the LTSSM). Added note defining the term in-band. Update figure 4-11 to represent clarifications in LTSSM.

4.2.6.1.1: Remove all references to high impedance detect sequences. The need for any high impedance mode detect sequences were eliminated by specifying the (max) transmitter short circuit current as well as the (max) DC common mode voltage at all transmitters (table 4-5).

4.2.6.1.1 - 4.2.6.1.2 (Detect): The detect state was made simpler when the need for high impedance detects were removed. Clarifications were added stating that all lanes that can become part of one (multi-lane) link must detect powered on receiver terminations to take part in future steps with the LTSSM.

4.2.6.2.1 – 4.2.6.2.4 (Polling): Add clear words that all lanes of a possible link work together here to try to move into configuration together. Add clear words that if any lane of a possible link detect terminations but never see the electrical idle condition broken, all the lanes output the compliance pattern. Insert a TS1/TS2 handshake (link and lane numbers set to PAD) here to establish that both ends of the perspective link have successfully achieved bit and symbol lock and therefore can exchange symbols.

4.2.6.3.1 – 4.2.6.2.6 (Configuration): The cleanup here is to match the intended behavior that was not well described previously. The most noticeable change here is the steps to establish a link number as well as lane numbers are now included in this section. Note that the steps have not changed, just moved into the state machine where they were implemented previously. The one change here is a method to break the potential race condition in some crosslink scenarios was inserted. Improved descriptions of how to enter Disabled and Loopback are here.

4.2.6.4.1 – 4.2.6.4.3 (Recovery): Clarifications and restrictions were added to make it clear exactly what can and can not be done in this state. Clarifications were added regarding entering Disabled, Hot Reset and Loopback from this state. Also, extended synch bit behavior was inserted (it was left out previously). Of note, we define a 24mS timeout in Recovery.RcvrLock since this is also used in Polling.Active and the two states are functionally equivalent for achieving bit and character lock. In the future, we may consider adding a note in the Electrical Receiver Requirement Table (Table4-6) that clearly specifies this time since it is a critical receiver design parameter.

4.2.6.5 (L0): Add clearer words regarding conditions that must be met to enter Recovery, L0s, L1 and L2 from this state. The key observation = detecting Electrical Idle in L0 without an Electrical Idle ordered set should result in Recovery, which should result in the case of a surprise removal transitioning from L0 to Recovery to Detect. This was needed instead of the scenario where detecting Electrical Idle in L0 without an ordered set results in transitioning to L0s. The transition to L0s can cause unintended consequences when only idle data is being transmitted and the link is severed. The edit makes sure that under these conditions the link will try Recovery and then end up in Detect if the link was severed by a hot removal.

4.2.6.6 (L0s): Clarification inserted regarding transmitters and receivers entering L0s and transitioning from L0s to L0 or Recovery. Also point out the when entering low power states, one end initiates the transitions; an improved description of that behavior was described.

4.2.6.7 (L1): Clarification around requirements for DC common mode voltage and conditions to exit to Recovery.

4.2.6.8 (L2): Clarification around requirements termination resistor behavior (must be enabled), for DC common mode voltage and conditions to exit to Detect.

4.2.6.9 (Disabled): Clearer wording and add note to point out when links are considered disabled.

4.2.6.9 (Loopback): Add clear statement that if a link is in loopback the linkup status bit is set to 0. Remove notes on what choosing disparity of EDB, SKP if inserted (it is not desirable to let subsequent symbols have a different disparity than what was transmitted).

4.2.6.9 (Hot Reset): This section previously was short and content-free. The section was lengthened to describe the behavior that must be followed to allow for inter-operability between vendors when using hot reset.

4.2.7: Clarify that SKP do not count as an interruption when looking for consecutive symbols. Add Clarification that the number of SKP symbols in a SKP ordered-set will never vary within lanes in a multi-lane link at receivers. Fixed formatting.

4.2.8: Used clearer words. Fixed mistake regarding the number of lanes the compliance pattern repeats over. Added clarification that the compliance pattern is “close to” worst case crosstalk pattern. Added clarification on exit condition.

4.3.1.4: Changed words to better describe the rules for terminations at receivers.

4.3.1.5: Made DC Common Mode Voltage its own section with clearer words and add maximum the voltage can be.

4.3.1.7: Added text regarding limiting the short circuit current a transmitter can source.

4.3.1.8: Changed words to clearly state a transmitter must detect ZRX-DC. Removed clumsy text and replaced it with simple exit condition description.

4.3.1.9: Concisely stated must receive at least 2 of the 3 IDL symbols in the electrical idle ordered set. State the transmitter may be in either high or low impedance. Add reference to table defining the electrical idle condition.

4.3.2.4: Inserted decision that the use of Beacon or Wake# is form factor and usage dependent. Added clarifications regarding electrical definition of beacon.

Table 4-5: Added the Tx short current limit and DC Common Mode Voltage limit. Also improved the descriptions of some of the parameters.

4.3.3.1: Added clarification the de-emphasized voltage is relative to the transition bit.

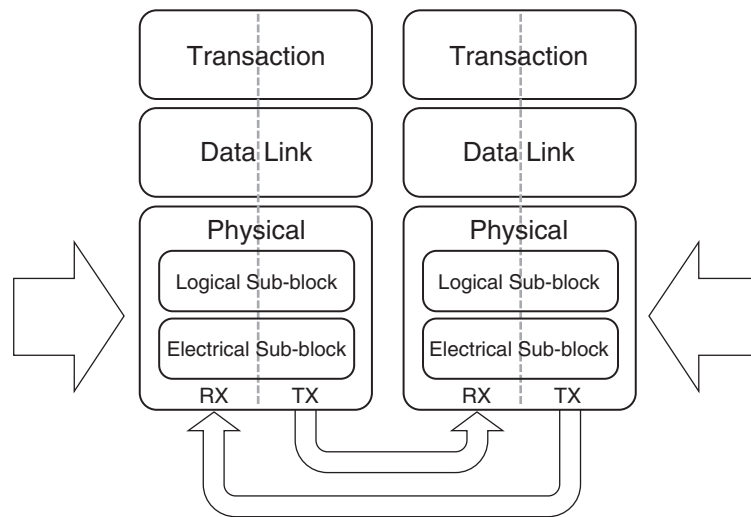
Table 4-6: Improved the descriptions of some of the parameters. Added clarifications in notes.

Text follows.

4. Physical Layer Specification

4.1. Introduction

The Physical Layer isolates the Transaction and Data Link Layers from the signaling technology used for Link data interchange. The Physical Layer is divided into the Logical and Electrical functional sub-blocks (see Figure 4-1).



OM13792

Figure 4-1: Layering Diagram Highlighting Physical Layer

4.2. Logical Sub-block

The Logical sub-block has two main sections: a Transmit section that prepares outgoing information passed from the Data Link Layer for transmission by the Electrical sub-block, and a Receiver section that identifies and prepares received information before passing it to the Data Link Layer.

The Logical sub-block and Electrical sub-block coordinate the state of each transceiver through a status and control register interface or functional equivalent. The Logical sub-block directs control and management functions of the Physical Layer.

4.2.1. Symbol Encoding

PCI Express uses an 8b/10b transmission code. The definition of this transmission code is identical to that specified in ANSI X3.230-1994, clause 11 (and also IEEE 802.3z, 36.2.4). Using this scheme, eight bit data characters are treated as three bits and five bits mapped onto a four-bit code group and a six bit code group, respectively. The control bit in conjunction with the data character is used to identify when to encode one of the 12 special symbols included in the 8b/10b transmission code. These code groups are concatenated to form a ten-bit Symbol. As shown in Figure 4-2, ABCDE maps to abcdei and FGH maps to fghj.

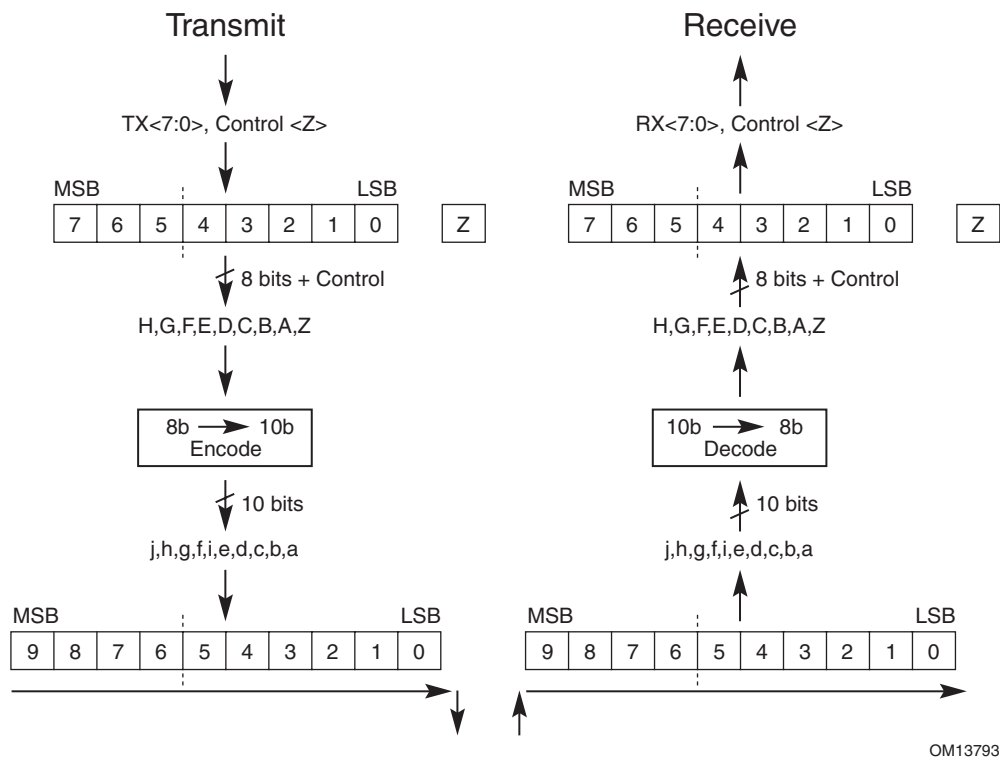
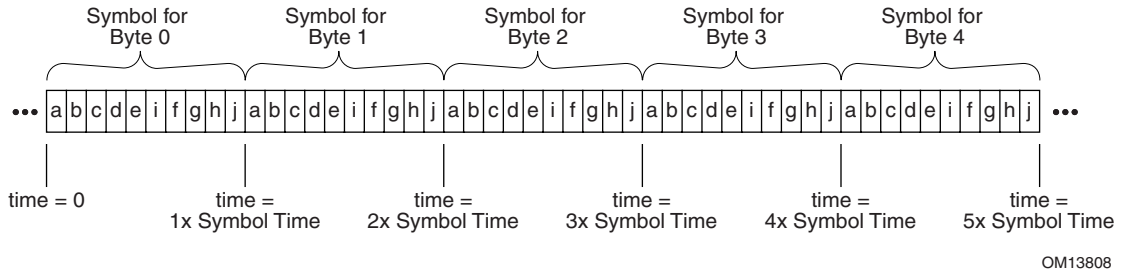


Figure 4-2: Character to Symbol Mapping

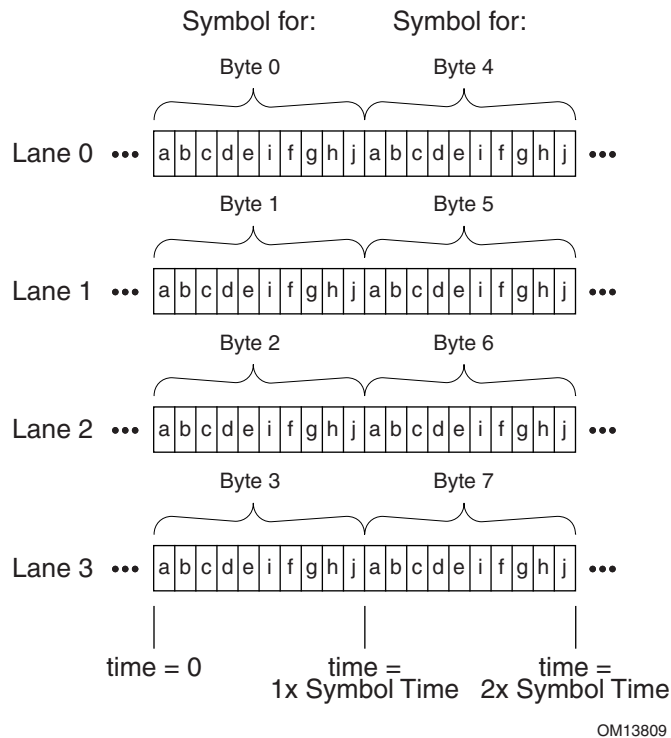
4.2.1.1. Serialization and De-serialization of Data

The bits of a Symbol are placed on a Lane starting with bit 'a' and ending with bit 'j'. Examples are shown in Figure 4-3 and Figure 4-4.



OM13808

Figure 4-3: Bit Transmission Order on Physical Lanes - x1 Example



OM13809

Figure 4-4: Bit Transmission Order on Physical Lanes - x4 Example

4.2.1.2. Special Symbols for Framing and Link Management (K Codes)

The 8b/10b encoding scheme used by PCI Express provides Special Symbols that are distinct from the Data Symbols used to represent Characters. These Special Symbols are used for various Link Management mechanisms described later in this chapter. Special Symbols are also used to frame DLLPs and TLPs, using distinct Special Symbols to allow these two types of Packets to be quickly and easily distinguished.

Table 4-1 shows the Special Symbols used for PCI Express and provides a brief description for the use of each. The use of these Symbols will be discussed in greater detail in following sections.

Table 4-1: Special Symbols

Encoding	Symbol	Name	Description
K28.5	COM	Comma	Used for Lane and Link initialization and management
K27.7	STP	Start TLP	Marks the start of a Transaction Layer Packet
K28.2	SDP	Start DLLP	Marks the start of a Data Link Layer Packet
K29.7	END	End	Marks the end of a Transaction Layer Packet or a Data Link Layer Packet
K30.7	EDB	EnD Bad	Marks the end of a nullified TLP
K23.7	PAD	Pad	Used in Framing and Link Width and Lane ordering negotiations
K28.0	SKP	Skip	Used for compensating for different bit rates for two communicating Ports
K28.1	FTS	Fast Training Sequence	Used within an ordered set to exit from L0s to L0
K28.3	IDL	Idle	Symbol used in the Electrical Idle ordered set
K28.4			Reserved
K28.6			Reserved
K28.7			Reserved

4.2.1.3. 8b/10b Decode Rules

The symbol tables for the valid 8b/10b codes are given in Appendix B. These tables have one column for the positive disparity and one column for the negative disparity.

If a received symbol is found in the proper column corresponding to the current running disparity, then that symbol is valid.

If a received symbol is found in the column corresponding to the incorrect running disparity or if the symbol does not correspond to either column, the Physical Layer must notify the Data Link Layer that the received symbol is invalid. This is a Receiver Error, and is a reported error associated with the Port (see Section 6.2).

4.2.2. Framing and Application of Symbols to Lanes

There are two classes of framing and application of symbols to Lanes. The first class is the ordered sets and the second is TLP and DLLP packet. Ordered sets are always transmitted serially on each Lane, such that a full ordered set appears simultaneously on all Lanes of a multi-Lane Link.

The Framing mechanism uses Special Symbol K28.2 “SDP” to start a DLLP and Special Symbol K27.7 “STP” to start a TLP. The Special Symbol K29.7 “END” is used to mark the end of either a TLP or a DLLP.

The conceptual stream of Symbols must be mapped from its internal representation, which is implementation dependent, onto the external Lanes. The Symbols are mapped onto the Lanes such that the first Symbol (representing Character 0) is placed onto Lane 0; the second is placed onto Lane 1; etc. The x1 Link represents a degenerate case and the mapping is trivial, with all Symbols placed onto the single Lane in order.

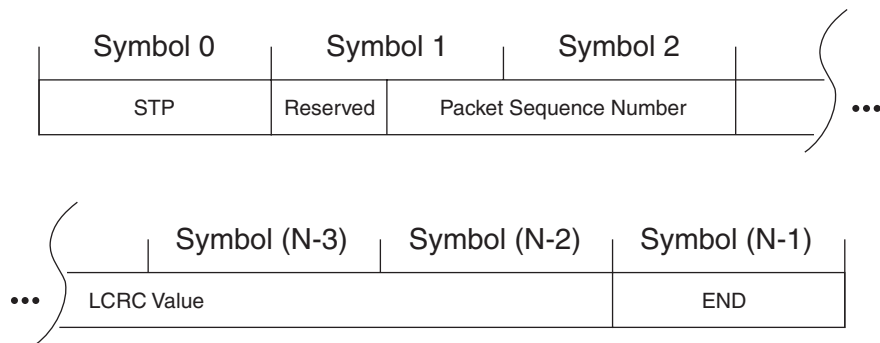
When no packet information or special ordered sets are being transmitted, the Transmitter is in the Logical Idle state. During this time idle data must be transmitted. The idle data must consist of the data byte 0 (00 Hexadecimal), scrambled according to the rules of Section 4.2.3 and 8b/10b encoded according to the rules of Section 4.2.1, in the same way that TLP and DLLP data characters are scrambled and encoded. Likewise, when the Receiver is not receiving any packet information or special ordered sets, the Receiver is in Logical Idle and shall receive idle data as described above. During transmission of the idle data, the SKIP ordered set must continue to be transmitted as specified in Section 4.2.7.

4.2.2.1. Framing and Application of Symbols to Lanes – Rules

In this section, “placed” is defined to mean a requirement on the transmitter to put the symbol into the proper Lane of a Link.

- ❑ TLPs must be framed by placing an STP Symbol at the start of the TLP and an END Symbol or EDB Symbol at the end of the TLP (see Figure 4-5).
- ❑ DLLPs must be framed by placing an SDP Symbol at the start of the DLLP and an END Symbol at the end of the DLLP.
- ❑ Logical Idle is defined to be a period of one or more Symbol times when no information: TLPs, DLLPs or any type of Special Symbol is being Transmitted/Received. Unlike Electrical Idle, during Logical Idle the Idle character (00h) is being transmitted and received.
 - When the Transmitter is in Logical Idle, the Idle data character (00h) shall be transmitted on all Lanes. This is scrambled according to the rules in Section 4.2.3.
 - Receivers must ignore incoming Logical Idle data, and must not have any dependency other than scramble sequencing on any specific data patterns.

- ❑ For Links wider than x1, the STP Symbol (representing the start of a TLP) must be placed in Lane 0 when starting Transmission of a TLP from a Logical Idle Link condition.
- ❑ For Links wider than x1, the SDP Symbol (representing the start of a DLLP) must be placed in Lane 0 when starting Transmission of a DLLP from a Logical Idle Link condition.
- ❑ The STP Symbol must not be placed on the Link more frequently than once per Symbol Time.
- ❑ The SDP Symbol must not be placed on the Link more frequently than once per Symbol Time.
- ❑ As long as the above rules are satisfied, TLP and DLLP Transmissions are permitted to follow each other successively.
- ❑ One STP symbol and one SDP symbol may be placed on the Link in the same symbol time.
 - Note: Links wider than x4 can have STP and SDP Symbols placed in Lane $4 \times N$, where N is a positive integer. For example, for x8, STP and SDP Symbols can be placed in Lanes 0 and 4; and for x16, STP and SDP Symbols can be placed in Lanes 0, 4, 8, or 12.
- ❑ For xN Links where N is 8 or more, if an END or EDB Symbol is placed in a Lane K , where K does not equal $N-1$, and is not followed by a STP or SDP Symbol in Lane $K+1$ (i.e., there is no TLP or DLLP immediately following), then PAD Symbols must be placed in Lanes $K+1$ to Lane $N-1$.
 - Note: For example, on a x8 Link, if END or EDB is placed in Lane 3, PAD must be placed in Lanes 4 to 7, when not followed by STP or SDP.
- ❑ The EDB symbol is used to mark the end of a nullified TLP. Refer to Section 3.5.2.1 for information on the usage of EDB.
- ❑ Receivers may optionally check for violations of the rules of this section. Any such violation is a Receiver Error, and is a reported error associated with the Port (see Section 6.2).



OM13794

Figure 4-5: TLP with Framing Symbols Applied




Figure 4-6: DLLP with Framing Symbols Applied


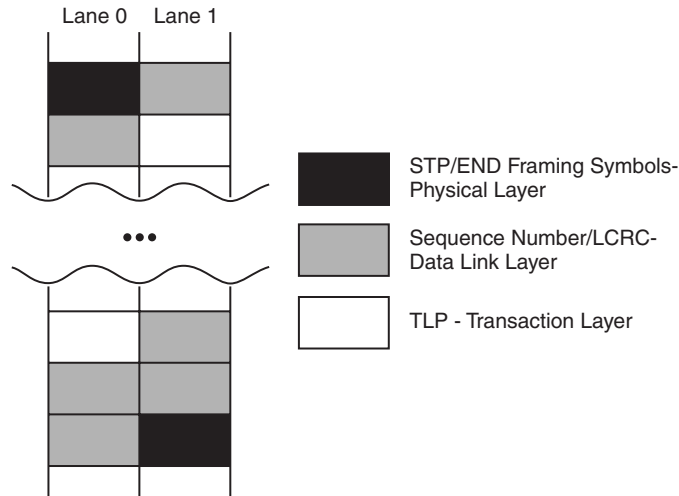
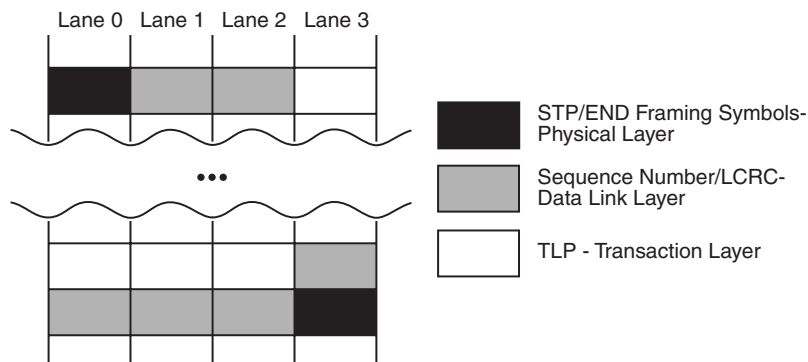


Figure 4-7: Framed TLP on a x1 Link



OM13797

Figure 4-8: Framed TLP on a x2 Link

OM13798

Figure 4-9: Framed TLP on a x4 Link

4.2.3. Data Scrambling

The scrambling function can be implemented with one or many Linear Feedback Shift Registers (LFSRs) on a multi-Lane Link. When there is more than one transmit LFSR per Link, these must operate in concert, maintaining the same simultaneous (see Table 4-5, Lane-to-Lane Output Skew) value in each LFSR. When there is more than one receive LFSR per Link, these must operate in concert, maintaining the same simultaneous (see Table 4-6, Total Skew) value in each LFSR. Regardless of how it's implemented, the LFSRs must interact with data on a Lane-by-Lane basis as if there was a separate LFSR as described here for each Lane within that Link. On the transmit side, scrambling is applied to characters prior to the 8b/10b encoding. On the receive side de-scrambling is applied to characters after 8b/10b decoding.

The LFSR is graphically represented in Figure 4-10. Scrambling or unscrambling is performed by serially XORing the 8-bit (D0-D7) character with the 16-bit (S0-S15) output of the LFSR. An output of the LFSR, S15, is XORed with D0 of the data to be processed.

The LFSR and data register are then serially advanced and the output processing is repeated for D1 through D7. The LFSR is advanced after the data is XORed. The LFSR implements the polynomial:

$$G(X)=X^{16}+X^{15}+X^{13}+X^4+1$$

The mechanism(s) and/or interface(s) utilized by the Data Link Layer to notify the Physical Layer to disable scrambling is implementation specific and beyond the scope of this specification.

The data scrambling rules are the following:

- ☐ The COM character initializes the LFSR.
- ☐ The LFSR value is advanced eight serial shifts for each character except the SKP.
- ☐ All data characters (D codes) except those within a Training Sequence Ordered sets (TS1, TS2) and the Compliance Pattern (see Section 4.2.8) are scrambled.
- ☐ All special characters (K codes) are not scrambled.
- ☐ The initialized value of an LFSR seed (S0-S15) is FFFFh. Immediately after a COM exits the transmit LFSR, the LFSR on the transmit side is initialized. Every time a COM enters the receive LFSR on any Lane of that Link, the LFSR on the receive side is initialized.
- ☐ Scrambling can only be disabled at the end of Configuration (see Section 4.2.6.3.5).
- ☐ Scrambling does not apply to a Loopback Slave.
- ☐ Scrambling is always enabled in Detect by default.



IMPLEMENTATION NOTE

Disabling Scrambling

Disabling Scrambling is intended to help simplify test and debug equipment. Control of the exact data patterns is useful in a test and debug environment. Since scrambling is reset at the Physical Layer there is no reasonable way to reliably control the state of the data transitions through software. Thus, the disable scrambling bit is provided for these purposes.

The mechanism(s) and/or interface(s) utilized by the data Link Layer to notify the physical Layer to disable scrambling is component implementation specific and beyond the scope of this specification.

For more information on scrambling, see Appendix C.

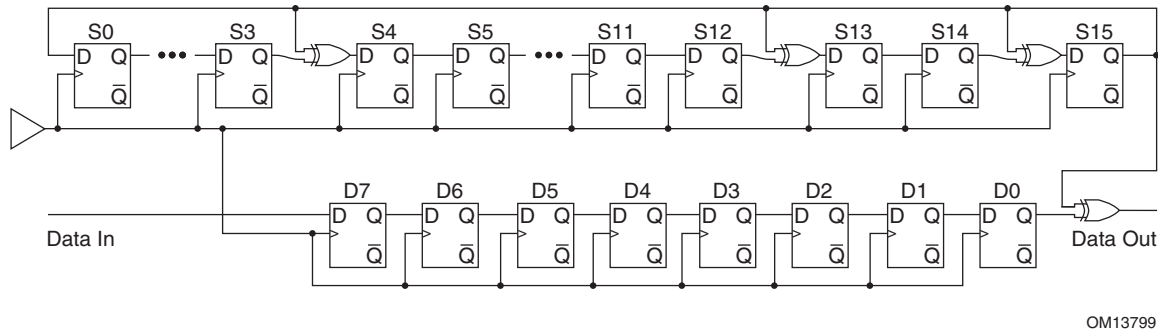


Figure 4-10: LFSR with Scrambling Polynomial

4.2.4. Link Initialization and Training

This section defines the Physical Layer control process that configures and initializes each Link for normal operation. This section covers the following functions:

- ☐ Configuring and initializing the Link.
- ☐ Supporting normal packet transfers.
- ☐ Supported state transitions when recovering from Link errors.
- ☐ Restarting a Port from low power states.

The following are discovered and determined during the training process:

- ☐ Link width.
- ☐ Link data rate.¹⁹
- ☐ Lane reversal.
- ☐ Polarity inversion.

Training does:

- ☐ Link data rate²⁰ negotiation.
- ☐ Bit lock per Lane.
- ☐ Lane polarity.
- ☐ Symbol lock per Lane.
- ☐ Lane ordering within a Link.
- ☐ Link width negotiation.
- ☐ Lane-to-Lane de-skew within a multi-Lane Link.

¹⁹ This specification only defines one data rate. Future revisions will define additional rates.

²⁰ This specification defines the mechanism for negotiating the Link operational bit rate to the highest supported operational data rate.

Receivers may optionally check for violations of the Link Initialization and Training Protocols. If such checking is implemented, any violation is a Training Error. A Training Error is a reported error associated with the Port (see Section 6.2). A Training Error is considered fatal to the Link.

4.2.4.1. Training Sequence Ordered Sets

Training sequences are composed of ordered sets used for initializing bit alignment, symbol alignment and to exchange Physical Layer parameters. Training sequence ordered sets are never scrambled but are always 8b/10b encoded.

SKP ordered sets (see Section 4.2.7) must be transmitted during but never interrupt a TS1 or TS2 ordered set.

The Training control bits for Hot Reset, Disable Link, and Enable Loopback are mutually exclusive, only one of these bits may be set at a time as well as transmitted on all Lanes in a configured (all Lanes in L0) or possible (all Lanes in Configuration) Link. If more than one of the Hot Reset, Disable Link, or Enable Loopback bits are set at the same time, the Link behavior is undefined.

Table 4-2: TS1 Ordered Set

Symbol Number	Allowed Values	Encoded Values	Description
0		K28.5	COMMA code group for symbol alignment
1	0 - 255	D0.0 - D31.7, K23.7	Link Number within component
2	0 - 31	D0.0 - D31.0, K23.7	Lane Number within Port
3	0 - 255	D0.0 - D31.7	N_FTS. This is the number of fast training ordered sets required by the receiver to obtain reliable bit and symbol lock.
4	2	D2.0	Data Rate Identifier Bit 0 - Reserved, set to 0 Bit 1 = 1, generation 1 (2.5 Gb/s) data rate supported Bit 2:7 - Reserved, set to 0

Symbol Number	Allowed Values	Encoded Values	Description
5	Bit 0 = 0, 1 Bit 1 = 0, 1 Bit 2 = 0, 1 Bit 3 = 0, 1 Bit 4:7 = 0	D0.0, D1.0, D2.0, D4.0, D8.0	Training Control Bit 0 – Hot Reset Bit 0 = 0, De-assert Bit 0 = 1, Assert Bit 1 – Disable Link Bit 1 = 0, De-assert Bit 1 = 1, Assert Bit 2 – Loopback Bit 2 = 0, De-assert Bit 2 = 1, Assert Bit 3 – Disable Scrambling Bit 3 = 0, De-assert Bit 3 = 1, Assert Bit 4:7, Reserved, set to 0
6-15		D10.2	TS1 Identifier

Table 4-3: TS2 Ordered Set

Symbol Number	Allowed Values	Encoded Values	Description
0		K28.5	COMMA code group for symbol alignment
1	0 - 255	D0.0 - D31.7, K23.7	Link Number within component
2	0 - 31	D0.0 - D31.0, K23.7	Lane Number within Port
3	0 - 255	D0.0 - D31.7	N_FTS. This is the number of fast training ordered sets required by the receiver to obtain reliable bit and symbol lock.
4	2	D2.0	Data Rate Identifier Bit 0 – Reserved, set to 0 Bit 1 = 1, generation 1 (2.5 Gb/s) data rate supported Bit 2:7 – Reserved, set to 0

Symbol Number	Allowed Values	Encoded Values	Description
5	Bit 0 = 0, 1 Bit 1 = 0, 1 Bit 2 = 0, 1 Bit 3 = 0, 1 Bit 4:7 = 0	D0.0, D1.0, D2.0, D4.0, D8.0	Training Control Bit 0 – Hot Reset Bit 0 = 0, De-assert Bit 0 = 1, Assert Bit 1 – Disable Link Bit 1 = 0, De-assert Bit 1 = 1, Assert Bit 2 – Loopback Bit 2 = 0, De-assert Bit 2 = 1, Assert Bit 3 – Disable Scrambling Bit 3 = 0, De-assert Bit 3 = 1, Assert Bit 4:7, Reserved, set to 0
6-15		D5.2	TS2 Identifier

4.2.4.2. Lane Polarity Inversion

During the training sequence, the receiver looks at symbols 6-15 of TS1 and TS2 as the indicator of Lane polarity inversion (D+ and D- are swapped). If Lane polarity inversion occurs, the TS1 symbols 6-15 received will be D21.5 as opposed to the expected D10.2. Similarly, if Lane polarity inversion occurs, symbols 6-15 of the TS2 ordered set will be D26.5 as opposed to the expected D5.2. This provides the clear indication of Lane polarity inversion.

If polarity inversion is detected the receiver must invert the received data. The transmitter must never invert the transmitted data. Support for Lane Polarity Inversion is required on all PCI Express Receivers across all Lanes independently.

4.2.4.3. Fast Training Sequence (FTS)

FTS is the mechanism that is used for bit and symbol lock when transitioning from L0s to L0. The FTS is used by the receiver to detect the exit from Electrical Idle and align the receiver's bit/symbol receive circuitry to the incoming data. See Section 4.2.5 for a description of L0 and L0s.

A single FTS training sequence is an ordered set composed of one K28.5 (COM) symbol and three K28.1 symbols. The maximum number of FTS ordered sets (N_FTS) that a component can request is 255, providing a bit time lock of $4 * 255 * 10 * UI$. 4096 FTS ordered sets must be sent when the Extended Synch bit is set in order to provide external Link monitoring tools with enough time to achieve bit and framing synchronization. SKP ordered sets must be scheduled and transmitted between FTS ordered sets as necessary to meet the definitions in Section 4.2.7 with the exception that no SKP ordered sets can be

scheduled during the first 255 FTS ordered sets. A single SKP ordered set is always sent after the last FTS is transmitted.

N_FTS defines the number of FTS ordered sets that must be transmitted when transitioning from L0s to L0. At generation 1 data rate, the value that can be requested by a component corresponds to a symbol lock time of 16 ns (N_FTS set to 0 and one SKP ordered set) to ~4 μ s (N_FTS set to 255), except when the Extended Synch bit is set, which requires the transmission of 4096 FTS ordered sets resulting in a bit lock time of 64 μ s. Note that the N_FTS value reported by a component may change; for example, due to software modifying the value in the Common Clock Configuration bit (Section 7.8.7).

If the N_FTS period of time expires before the receiver obtains bit, symbol, and Lane-to-Lane de-skew on all Lanes of the configured Link, the receiver must transition to the Recovery. This sequence is detailed in the LTSSM in Section 4.2.5.

4.2.4.4. Link Error Recovery

- ❑ Link Errors are defined as 8b/10b decode errors, loss of symbol lock, or loss of Lane-to-Lane de-skew.
 - Note: 8b/10b decode errors trigger a Receiver Error (see Table 4-4), which is a reported error associated with the Port (see Section 6.2).
- ❑ On a configured Link, which is in L0, error recovery will at a minimum be managed in a Layer above the Physical Layer (as described in Section 3.5) by directing the Link to transition to Recovery.
 - Note: Link Errors may also result in the Physical Layer initiating a LTSSM state transition from L0 to Recovery.
- ❑ All LTSSM states other than L0 make progress²¹ when Link Errors occur.
 - Note: Link errors that occur, while in LTSSM states other than L0, must not result in the Physical Layer initiating a LTSSM state transition.
- ❑ If a Lane detects an implementation specific number of 8b/10b errors, symbol lock must be verified or re-established as soon as possible.²²

²¹ In this context, progress is defined as the LTSSM not remaining indefinitely in one state with the possible exception of Detect.

²² The method to verify and re-establish symbol lock is implementation specific.

4.2.4.5. Reset

Reset is described from a system point of view in Section 6.6.

4.2.4.5.1. Fundamental Reset

When Fundamental Reset is asserted:

- ☐ The receiver terminations are required to only meet $Z_{RX-HIGH-IMP-DC}$ (see Table 4-6).
- ☐ The transmitter terminations is required to only meet Z_{TX-DC} (see Table 4-6).
- ☐ The transmitter holds a constant DC common mode voltage.²³

When Fundamental Reset is de-asserted:

- ☐ The Port LTSSM (see Section 4.2.5) is initialized (i.e., Detect is immediately entered).

4.2.4.5.2. Hot Reset

Hot Reset is a protocol reset defined in Section 4.2.5.11.

4.2.4.6. Link Data Rate Negotiation

All devices are required to start Link initialization using a generation 1 data rate on each Lane. A field in the training sequence ordered set (see Section 4.2.4) is used to advertise all supported data rates, and any higher speed supported by both sides of the Link will be initiated during Polling.Speed.

4.2.4.7. Link Width and Lane Sequence Negotiation

PCI Express Links must consist of 1, 2, 4, 8, 12, 16, or 32 Lanes in parallel, referred to as x1, x2, x4, x8, x12, x16, and x32 Links, respectively. All Lanes within a Link shall simultaneously (as defined by $L_{TX-SKEW}$ in Table 4-5) transmit data based on the exact same frequency. The negotiation process is described as a sequence of steps.

The negotiation establishes values for Link number and Lane number for each Lane that is part of a valid Link; each Lane that is not part of a valid Link exits the negotiation to become a separate Link or remain in Electrical Idle.

During Link width and Lane number negotiation, the two communicating Ports must accommodate the maximum allowed Lane-Lane skew as specified by $L_{RX-SKEW}$ in Table 4-6.

²³ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode during L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 4-5).

Optional Link negotiation behaviors include Lane reversal, variable width Links, splitting of Ports into multiple Links and the configuration of a crosslink.

Annex specifications to this specification may impose other rules and restrictions that must be comprehended by components compliant to those annex specifications; it is the intent of this specification to comprehend interoperability for a broad range of component capabilities.

4.2.4.7.1. Required and Optional Port Behavior

1. The ability for a xN Port to form a xN Link as well as a x1 Link (where N can be 32, 16, 12, 8, 4, 2, and 1) is required.
 - Note: Designers must connect Ports between two different components in a way that allows those components to meet the above requirement. If the Ports between components are connected in ways that are not consistent with intended usage as defined by the component's Port descriptions/data sheets behavior is undefined.
2. The ability for a xN Port to form any Link width between N and 1 is optional.
 - Note: An example of this behavior includes a x16 Port which can only configure into only one Link, but the width of the Link can be configured to be x12, x8, x4, x2 as well as the required widths of x16 and x1.
3. The ability to split a Port into two or more Links is optional.
 - Note: An example of this behavior would be a x16 Port that may be able to configure 2 x8 Links, 4 x4 Links, or even 16 x1 Links.
4. Support for Lane reversal is optional.
 - Note: Lane reversal must be done for both the transmitter and receiver of a given Port for a multi-Lane Link.
 - Note: An example of Lane reversal consists of Lane 0 of an Upstream Port attached to Lane N-1 of a Downstream Port where either the Downstream or Upstream device may reverse the Lane order to configure a xN Link.

Support for formation of a crosslink is optional. In this context, a Downstream Port connected to a Downstream Port or an Upstream Port connected to an Upstream Port is a crosslink.

Current and future electromechanical and/or form factor specifications may require the implementation of some optional features listed above. Component designers must read the specifications that the component(s) they are designing will used in to ensure compliance to those specifications.

4.2.4.8. Lane-to-Lane De-skew

The receiver must compensate for the allowable skew between Lanes within a multi-Lane Link (see Table 4-5 and Table 4-6) before delivering the data and control to the Data Link Layer.

Lane-to-Lane de-skew shall be done across all Lanes within multi-Lane Links. An unambiguous de-skew mechanism is the COM symbol transmitted during training sequence or SKIP ordered sets across all Lanes within a configured Link. Other de-skew mechanisms may also be employed. Lane-to-Lane de-skew must be performed during Configuration, Recovery, and L0s in the LTSSM.

4.2.4.9. Lane vs. Link Training

The Link initialization process builds unassociated Lanes of a Port into associated Lanes that form a Link. For Lanes to configure properly into a desired Link, the TS1 and TS2 ordered sets must have the appropriate fields (symbol 3, 4, and 5) set to the same value on all Lanes.

Links are formed at the conclusion of Configuration.

- ❑ Note: If the optional behavior of a Port being able to configure multiple Links is employed, the following observations can be made:
- A separate LTSSM is needed for the maximum number of Links that are desired to be configured by any given Port.
 - The LTSSM Rules are written for configuring one Link. The decision to configure Links in a serial fashion or parallel is implementation specific.

4.2.5. Link Training and Status State Machine (LTSSM) Descriptions

The LTSSM states are illustrated in Figure 4-11. These states are described in following sections.

All timeout values specified in the Link training and status state machine (LTSSM) timeout values are minus 0 seconds and plus 50% unless explicitly stated otherwise. All timeout values must be set to the specified values after power-on/reset. All counter values must be set to the specified values after power-on/reset.

4.2.5.1. Detect

The purpose of this state is to detect when a far end termination is present. This state can be entered at any time if directed.

4.2.5.2. Polling

The Port transmits training ordered sets and responds to the received training ordered sets. In this state, bit lock and symbol lock are established, Lane polarity is configured, and Lane data rate is established.

The polling state includes Polling.Compliance (see Section 4.2.6.2.2). This state is intended for use with test equipment used to assess if the transmitter and the interconnect present in the device under test setup is compliant with many voltage and timing specifications in Table 4-5 and Table 4-6.



IMPLEMENTATION NOTE

Use of Polling.Compliance

Polling.Compliance is intended for a compliance test environment and not entered during normal operation and cannot be disabled for any reason. Polling.Compliance is entered based on the physical system environment as described in Section 4.2.6.2.1. Any other mechanism that causes a transmitter to output the compliance pattern is implementation specific and is beyond the scope of this specification.

4.2.5.3. Configuration

In Configuration, both the transmitter and receiver are sending and receiving data at the negotiated data rate. The Lanes of a Port configure into a Link through a width and Lane negotiation sequence. Also, Lane-to-Lane de-skew must occur, scrambling can be disabled, the N_FTS is set, and the Disable or Loopback states can be entered.

4.2.5.4. Recovery

In Recovery, both the transmitter and receiver are sending and receiving data using the configured Link and Lane number as well as the previously negotiated data rate. Recovery allows a configured Link to re-establish bit lock, symbol lock, and Lane-to-Lane de-skew. Recovery is also used to set a new N_FTS and enter the Loopback, Disable, Hot Reset, and Configuration states.

4.2.5.5. L0

L0 is the normal operational state where data and control packets can be transmitted and received. All power management states are entered from this state.

4.2.5.6. L0s

L0s is intended as a power savings state.

L0s allows a Link to quickly enter and recover from a power conservation state without going through Recovery.

The entry to L0s occurs after receiving an Electrical Idle ordered set.

The exit from L0s to L0 must re-establish bit lock, symbol lock, and Lane-to-Lane de-skew.

A transmitter and receiver Lane pair on a Port are not required to both be in L0s simultaneously.

4.2.5.7. L1

L1 is intended as a power savings state.

The L1 state allows an additional power savings over L0s at the cost of additional resume latency.

The entry to L1 occurs after being directed by the Data Link Layer and receiving an Electrical Idle ordered-set.

4.2.5.8. L2

Power can be aggressively conserved in L2. Most of the Transmitter and Receiver may be shut off.²⁴ Main power and clocks are not guaranteed, but aux²⁵ power is available.

When Beacon support is required by the associated system or form-factor specification, an Upstream Port that supports the wakeup capability must be able to send; and a Downstream Port must be able to receive; a wakeup signal referred to as a Beacon.²⁶

The entry to L2 occurs after being directed by the Data Link Layer and receiving an Electrical Idle ordered set.

4.2.5.9. Disabled

The intent of the Disabled state is to allow a group of Lanes to be disabled until directed or Electrical Idle is exited (i.e., due to a hot removal and insertion) after entering Disabled.

Disabled uses Bit 1 (Disable Link) in the Training Control Register (see Table 4-2 and Table 4-3) which is sent within the TS1 and TS2 training ordered set.

²⁴ The exception is the receiver termination, which must remain in a low impedance state.

²⁵ In this context, “aux” power means a power source which can be used to drive the Beacon circuitry.

²⁶ Certain form factor specifications require the use of a Beacon for a device to request main power reactivation, for example to wake a system that is in D3_{cold}. See Section 4.3.2.4 for information on the electrical requirements of the Beacon. Refer to Chapter 5 for more information on how a device may use the Beacon as the wakeup mechanism.

A Link can enter Disabled if directed by a higher Layer. A Link can also reach the Disable state by receiving two consecutive TS1 ordered sets with the Disable Bit asserted (see Section 4.2.6.9).

4.2.5.10. Loopback

Loopback is intended for test and fault isolation use. Only the entry and exit behavior is specified, all other details are implementation specific. Loopback can operate on either a per Lane or configured Link basis.

A Loopback Master is the component requesting Loopback.

A Loopback Slave is the component looping back the data.

Loopback uses Bit 2 (Loopback) in the Training Control Register (see Table 4-2 and Table 4-3) which is sent within the TS1 and TS2 training ordered set.

The entry mechanism for Loopback Master is device specific.

The Loopback Slave device enters Loopback whenever two consecutive TS1 ordered sets are received with the Loopback bit set.



IMPLEMENTATION NOTE

Use of Loopback

Once in the Loopback state, the master can send any pattern of symbols as long as the rules of 8b/10b encoding (including disparity) are followed. Once in Loopback, the concept of data scrambling is no longer relevant; what is sent out is looped back. The mechanism(s) and/or interface(s) utilized by the Data Link Layer to notify the Physical Layer to enter the Loopback state is component implementation specific and beyond the scope of this specification.

4.2.5.11. Hot Reset

Hot Reset uses Bit 0 (Hot Reset) in the Training Control Register (see Table 4-2 and Table 4-3) which is sent within the TS1 and TS2 training ordered set.

A Link can enter Hot Reset if directed by a higher Layer. A Link can also reach the Hot Reset state by receiving two consecutive TS1 ordered sets with the Hot Reset Bit asserted (see Section 4.2.6.11).

4.2.6. Link Training and Status State Rules

Various Link status bits are monitored through software with the exception of LinkUp which is monitored by the Data Link Layer. Table 4-4 describes how the Link status bits

must be handled throughout the LTSSM (for more information, see Section 3.1 for LinkUp; Section 7.8.8 for Link Speed, LinkWidth, and Link Training; Section 6.2 for Receiver Error; and Section 6.7 for In-Band Presence).

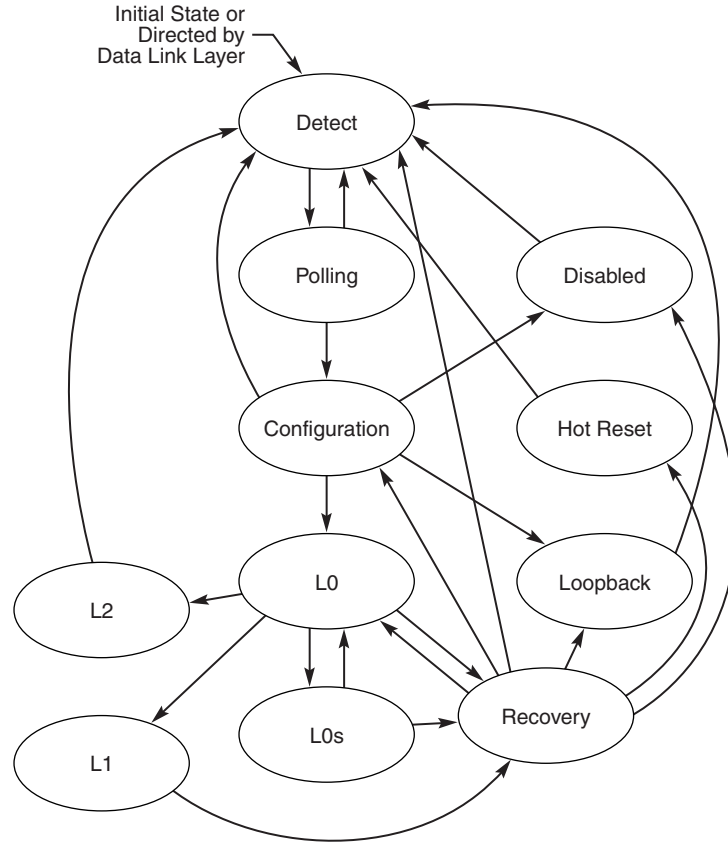
Table 4-4: Table of Link Status Mapped to the LTSSM

LTSSM State	Link Width	Link Speed	LinkUp	Link Training	Receiver Error	In-Band Presence²⁷
Detect	Undefined	Undefined	0	0	No action	0
Polling	Undefined	Set	0	0	No action	1
Configuration	Set	Not action	0/1 ²⁸	1	Set on 8b/10b Error	1
Recovery	No action	No action	1	1	No action	1
L0	No action	No action	1	0	Set on 8b/10b Error or optionally on Framing Violation	1
L0s	No action	No action	1	0	No action	1
L1	No action	No action	1	0	No action	1
L2	No action	No action	1	0	No action	1
Disabled	Undefined	Undefined	0	0	Set on 8b/10b Error	1
Loopback	No action	No action	0	0	No action	1
Hot Reset	No action	No action	0	0	Set on 8b/10b Error	1

The state machine rules for configuring and operating a PCI Express Link are defined in the following sections.

²⁷ In-band refers to the fact that no sideband signals are used to calculate the presence of a powered up device on the other end of a Link.

²⁸ LinkUp will always be 0 if coming into Configuration via Detect -> Polling -> Configuration and LinkUp will always be 1 if coming into Configuration from any other state.



OM13800A

Figure 4-11: Main State Diagram for Link Training and Status State Machine

4.2.6.1. Detect


4.2.6.1.1. Detect.Quiet

- ☐ Transmitter is in an Electrical Idle state.
 - Note: The DC common mode voltage does not have to be within specification.²⁹
- ☐ Generation 1 data rate is selected.
 - Note: This does not affect the advertised data rate in TS1 and TS2.
- ☐ LinkUp = 0 (status is cleared).
- ☐ The next state is Detect.Active after a 12 ms timeout or if Electrical Idle is broken on any Lane.

²⁹ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 4-5).

4.2.6.1.2. Detect.Active

- ❑ The transmitter performs a Receiver Detection sequence on all un-configured Lanes that can form one or more Links (see Section 4.3.1.8 for more information).
- ❑ Next state is Polling if a receiver is detected on all unconfigured Lanes
- ❑ Next state is Detect.Quiet if a receiver is not detected on any Lanes.
- ❑ If at least one but not all un-configured Lanes detect a receiver, then:
 1. Wait for 12 ms.
 2. The transmitter performs a Receiver Detection sequence on all un-configured Lanes that can form one or more Links (see Section 4.3.1.8 for more information),
 - i) The next state is Polling if exactly the same Lanes detect a receiver as the first Receiver Detection sequence.
 - Note: Lanes that did not detect a receiver must:
 - i) Be associated with a new LTSSM if this optional feature is supported.
 - or
 - ii) All Lanes that cannot be associated with an optional new LTSSM must transition to Electrical Idle.³⁰
 - Note: These Lanes must be re-associated with the LTSSM immediately after the LTSSM in progress transitions back to Detect.
 - ii) Otherwise, the next state is Detect.Quiet.



OM14313A

Figure 4-12: Detect Sub-State Machine

³⁰ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 4-5).

4.2.6.2. Polling

4.2.6.2.1. Polling.Active

- ❑ Transmitter sends TS1 ordered sets with Lane and Link numbers set to PAD (K23.7) on all Lanes that detected a receiver during Detect.
- ❑ Next state is Polling.Configuration after eight consecutive TS1 or TS2 ordered sets or their complement is received with the Lane and Link numbers set to PAD (K23.7) on all Lanes that detected a receiver during Detect and at least 1024 TS1 ordered sets were transmitted.
- ❑ Otherwise, after a 24 ms timeout the next state is:
 1. Polling.Configuration if,
 - a) Any Lane received eight consecutive TS1 or TS2 ordered sets (or their complement) with the Lane and Link numbers set to PAD (K23.7), and a minimum of 1024 TS1s are transmitted after receiving one TS1.
And
 - b) All Lanes that detected a receiver during Detect have detected an exit from Electrical Idle at least once since entering Polling.Active.
 - Note: This prevents one or more bad receivers or transmitters from holding up a valid Link from being configured, and allows for additional training in Polling.Configuration.
 2. Polling.Compliance if at least one Lane's receiver in Polling has never detected an exit from Electrical Idle since entering Polling.Active.
 - Note: This indicates the presence of a passive test load on at least one Lane, which will force all Lanes to enter Polling.Compliance.
 3. Detect if no TS1 or TS2 ordered set is received with Link and Lane number set to Pad on any Lane. The highest advertised speed in TS1 and TS2 is lowered (unless generation 1 is the highest advertised speed).

4.2.6.2.2. Polling.Compliance

- ❑ Transmitter sends out the compliance pattern on all Lanes that detected a receiver during Detect at the data rate which was employed upon entry to Polling.Compliance (see Section 4.2.8).
- ❑ Next state is Polling.Active if Electrical Idle exit has been detected at the receiver of all Lanes that detected a receiver during Detect.

4.2.6.2.3. Polling.Configuration

- ☐ Receiver must invert polarity if necessary (see Section 4.2.4.2).
- ☐ Transmitter sends TS2 ordered sets with Link and Lane numbers set to PAD (K23.7) on all Lanes that detected a receiver during Detect .
- ☐ The next state is Configuration after eight consecutive TS2 ordered sets are received on any Lanes that detected a receiver during Detect, 16 TS2 ordered sets are transmitted after receiving one TS2 ordered set, and none of those same Lanes is transmitting and receiving a higher Data Rate Identifier.³¹
- ☐ The next state is Polling.Speed after eight consecutive TS2 ordered sets are received on all Lanes that detected a receiver during Detect, 16 TS2 ordered sets are transmitted after receiving one TS2 ordered set, and at least one of those same Lanes is transmitting and receiving a higher Data Rate Identifier.³²
- ☐ Otherwise, next state is Detect after a 48 ms timeout.

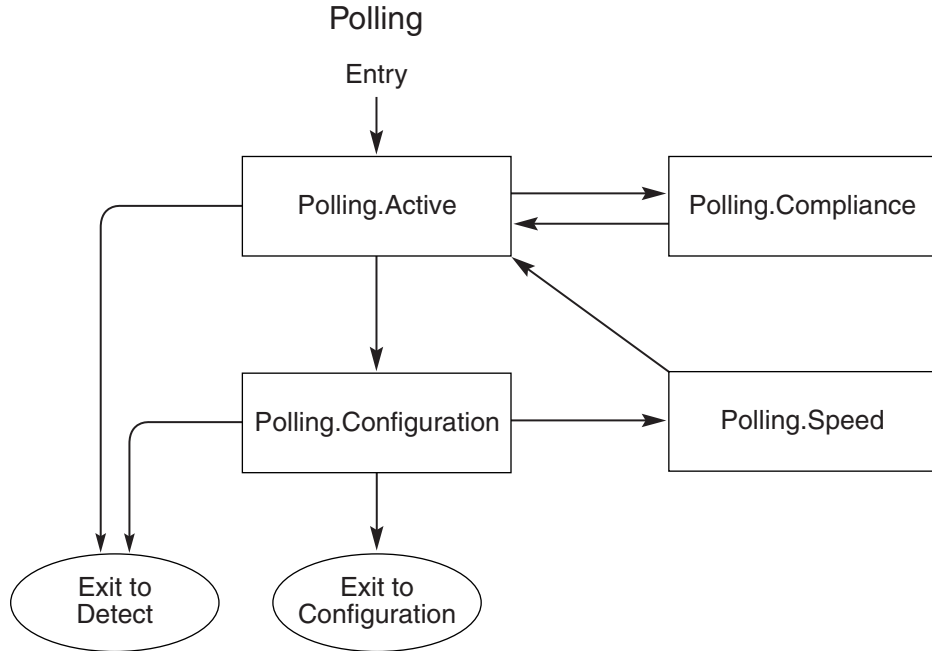
4.2.6.2.4. Polling.Speed

- ☐ The transmitter enters Electrical Idle for a minimum of $T_{TX-IDLE-MIN}$ (see Table 4-5) and no longer than 2 ms..
 - Note: Electrical Idle ordered set is sent prior to entering Electrical Idle.
 - Note: The DC common mode voltage does not have to be within specification.³³
- ☐ Data rate is changed on all Lanes to highest common data rate supported on both sides of the Link indicated by the training sequence (see Section 4.2.4.1).
- ☐ Next state is Polling.Active.

³¹ Higher data rate than what is currently being executed. Data Rate support on a Link is determined by the highest common speed being transmitted and received in TS1 and TS2 ordered sets.

³² Higher data rate than what is currently being executed. Data Rate support on a Link is determined by the highest common speed being transmitted and received in TS1 and TS2 ordered sets.

³³ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 4-5).



OM13801A

Figure 4-13: Polling Sub-State Machine

4.2.6.3. Configuration

4.2.6.3.1. Configuration.Linkwidth.Start

4.2.6.3.1.1. Downstream Lanes

- ❑ Next state is Disable if directed.
 - Note: “if directed” applies to a Downstream Port that is instructed by a higher Layer to assert the Disable Link bit (TS1 and TS2) on all Lanes that detected a receiver during Detect.
- ❑ Next state is Loopback if directed to this state, and the transmitter is capable of being a Loopback Master.
 - Note: “if directed” applies to a Port that is instructed by a higher Layer to assert the Loopback bit (TS1 and TS2) on all Lanes that detected a receiver during Detect.
- ❑ In the optional case where a crosslink is supported, the next state is Disable after all Lanes that detected a receiver during Detect and are receiving TS1 ordered sets with the Disable Link bit asserted in two consecutive TS1 ordered sets.

- ❑ Next state is Loopback if all Lanes that detected a receiver during Detect receive the Loopback bit asserted in two consecutive TS1 ordered sets on all Lanes receiving a TS1 ordered set.
 - Note that the device receiving the ordered set with the Loopback bit set becomes the Loopback Slave.
- ❑ The transmitter sends TS1 ordered sets with selected Link numbers and sets Lane numbers to PAD (K23.7) on Downstream Lanes that are in Configuration.
 - Note: Link numbers can only be different for groups of Lanes capable of being a unique Link.
 - Note: An example of Link number assignments includes a set of eight Lanes on an Upstream component (Downstream Lanes) capable of negotiating to become one x8 Port when connected to one Downstream component (Upstream Lanes) or two x4 Ports when connected to two different Downstream components: The Upstream component (Downstream Lanes) sends out TS1 ordered sets with the Link number set to N on four Lanes and Link number set to N+1 on the other four Lanes. The Lane numbers are all set to PAD (K23.7).
- ❑ If any Lanes first received at least one or more TS1 ordered sets with a Link and Lane number set to PAD (K23.7), the next state is Configuration.Linkwidth.Accept immediately after all those same Downstream Lanes receive two consecutive TS1 ordered sets with any Link number that was transmitted and different than PAD (K23.7).
 - Note: If the crosslink configuration is not supported the condition of first receiving a Link and Lane number set to PAD is always true.
- ❑ Optionally, if crosslinks are supported and all Downstream Lanes initially receive two consecutive TS1 ordered sets with a Link number different than PAD (K23.7), then the Downstream Lanes are now designated as Upstream Lanes and a new random cross Link timeout is chosen (see $T_{\text{crosslink}}$ in Table 4-5). The next state is Configuration.Linkwidth.Start as Upstream Lanes.
 - Note: This supports the optional crosslink where both sides may try to act as a Downstream Port. This is resolved by making both Ports become Upstream and assigning a random timeout until one side of the Link becomes a Downstream and the other side remains an Upstream Port. This timeout must be random even when hooking up two of the exact same devices so as to eventually break any possible deadlock.
- ❑ The next state is Detect after a 24 ms timeout.

4.2.6.3.1.2. Upstream Lanes

- ❑ In the optional case where crosslinks are supported the next state is Disable if directed.
 - Note: “if directed” only applies to an optional crosslink Port that is instructed by a higher Layer to assert the Disable Link bit (TS1 and TS2) on all Lanes that detected a receiver during Detect.
- ❑ Next state is Loopback if directed to this state, and the transmitter is capable of being a Loopback Master.
 - Note: “if directed” applies to a Port that is instructed by a higher Layer to assert the Loopback bit (TS1 and TS2) on all Lanes that detected a receiver during Detect.
- ❑ Next state is Disable after any Lanes that detected a receiver during Detect and are receiving TS1 ordered sets with the Disable Link bit asserted in two consecutive TS1 ordered sets.
 - Note: In the optional case where a crosslink is supported, the next state is Disable only after all Lanes that detected a receiver during Detect and are receiving TS1 ordered sets with the Disable Link bit asserted in two consecutive TS1 ordered sets.
- ❑ Next state is Loopback if all Lanes that detected a receiver during Detect receive the Loopback bit asserted in two consecutive TS1 ordered sets on all Lanes receiving a TS1 ordered set.
 - Note: The device receiving the ordered set with the Loopback bit set becomes the Loopback Slave.
- ❑ The transmitter sends out TS1 ordered sets with Link numbers and Lane numbers set to PAD (K23.7) on Upstream Lanes that detected a receiver during Detect.
- ❑ Immediately after all Upstream Lanes receive two consecutive TS1 ordered sets with Link numbers that are different than PAD (K23.7), a single Link number is selected and transmitted on all Lanes that can form a Link. Any left over Lanes that detected a receiver during Detect must transmit TS1 ordered sets with the Link and Lane number set to PAD (K23.7). The next state is Configuration.Linkwidth.Accept.
- ❑ Optionally, if crosslinks are supported and all Upstream Lanes first receive two consecutive TS1 ordered sets with Link and Lane numbers set to PAD (K23.7), then:
 - The transmitter continues to send out TS1 ordered sets with Link numbers and Lane numbers set to PAD (K23.7).
 - Immediately after all Upstream Lanes receive two consecutive TS1 ordered sets with Link numbers that are different than PAD (K23.7), a single Link number is selected and transmitted on all Lanes that can form a Link. Any left over Lanes that detected a receiver during Detect must transmit TS1 ordered sets with the Link and Lane number set to PAD (K23.7). The next state is Configuration.Linkwidth.Accept.

- Otherwise, after a $T_{\text{crosslink}}$ timeout the Upstream Lanes become Downstream Lanes and the next state is Configuration.Linkwidth.Start as Downstream Lanes.
 - ♦ Note: This optional behavior is required for crosslink behavior where two Ports may start off with Upstream Ports, and one will eventually take the lead as a Downstream Port.

□ The next state is Detect after a 24 ms timeout.

4.2.6.3.2. Configuration.Linkwidth.Accept

4.2.6.3.2.1. Downstream Lanes

- If a Link can be formed with at least one group of Lanes that received two consecutive TS1 ordered sets with the same received Link number (non PAD and matching one that was transmitted by the Downstream Lanes), then TS1 ordered sets are transmitted with the same Link number and unique Lane numbers are assigned to all these same Lanes (Lane numbers must range sequentially from 0 to n-1). Any left over Lanes must transmit TS1 ordered sets with the Link and Lane number set to PAD (K23.7). The next state is Configuration.Lanenum.Wait.
 - Note: A couple of interesting cases to consider here are the following:
 1. A x8 Downstream Port, which can be divided into two x4 Links, sends two different Link numbers on to two x4 Upstream Ports. The Upstream Ports respond simultaneously by picking the two Link numbers. The Downstream Port will have to choose one of these sets of Link numbers to configure as a Link, and leave the other for a secondary LTSSM to configure (which will ultimately happen in Configuration.Complete).
 2. A x16 Downstream Port, which can be divided into two x8 Links, is hooked up to a x12 Upstream Port that can be configured as a x12 Link or a x8 and a x4 Link. During Configuration.Linkwidth.Start the Upstream Port returned the same Link number on all 12 Lanes. The Downstream Port would then return the same received Link number and assign Lane numbers on the eight Lanes that can form a x8 Link with the remaining four Lanes transmitting a Lane number and a Link number set to PAD (K23.7).
- The next state is Detect after a 2 ms timeout or if no Link can be configured or if all Lanes receive two consecutive TS1 ordered sets with Link and Lane numbers set to Pad (K23.7).

4.2.6.3.2.2. Upstream Lanes

- If Lanes that transmitted a Link number receive two consecutive TS1 ordered sets with the same (non PAD) Link number and include at least a Lane number 0, TS1 Lane numbers are transmitted that if possible match the received Lane numbers or are different if necessary (i.e., Lane reversed). Remaining Lanes must transmit TS1 with Link and Lane numbers set to PAD (K23.7). The next state is Configuration.Lanenum.Wait.
 - Note: The transmitted Lane numbers must range from 0 to m-1, be assigned sequentially to the same grouping of Lanes that are receiving Lane numbers 0 to n-1, include the received Lane 0, and m-1 must be equal to or smaller than the largest received Lane number (n-1).
 - Note: A few interesting cases to consider here are the following:
 1. An x8 Upstream Port is attached to a x8 Downstream Port that is presented with Lane numbers that are backward from the preferred numbering. If the optional behavior of Lane reversal is supported by the Upstream Port, the Upstream Port transmits the same Lane numbers back to the Downstream Port. Otherwise the opposite Lane numbers are transmitted back to the Downstream Port, and it will be up to the Downstream Port to optionally fix the Lane ordering or exit Configuration.
 2. Optional Lane reversal behavior is required to configure a Link where the Lane numbers are reversed and the Downstream Port does not support Lane reversal. Specifically, the Upstream Port Lane reversal will accommodate the scenario where the default Upstream sequential Lane numbering (0 to n-1) is receiving a reversed Downstream sequential Lane number (n-1 to 0).
 3. An optional x8 Upstream crosslink Port, which can be divided into two x4 Links, is attached to two x4 Downstream Ports that present the same Link number, and each x4 Downstream Port presents Lane numbers simultaneously that were each numbered 0 to 3. The Upstream Port will have to choose one of these sets of Lane numbers to configure as a Link, and leave the other for a second pass through Configuration.
- The next state is Detect after a 2 ms timeout or if no Link can be configured or if all Lanes receive two consecutive TS1 ordered sets with Link and Lane numbers set to Pad (K23.7).

4.2.6.3.3. Configuration.Lanenum.Accept

4.2.6.3.3.1 Downstream Lanes

- ❑ If a configured Link can be formed with all the Lanes that receive two consecutive TS1 ordered sets with the same transmitted Link numbers and non-PAD Lane numbers, the next state is Configuration.Complete.
 - Note: Two possible scenarios exist that can result in the next state being Configuration.Complete:
 1. The received Link and Lane numbers are the same as transmitted, which is always required to be supported.
 2. Optional Lane reversal behavior can occur for either scenario 1. An example of where Lane reversal is required is when the received Lane numbers are the reverse of the Lane numbers being transmitted (case 1). This optional feature would be required to configure a Lane reversed Link where the Upstream Port does not support Lane reversal.
- ❑ If the Lanes of a Link can be configured by using the Lanes that transmitted a Lane number, which received two consecutive TS1 ordered sets with the same transmitted Link number (non-Pad) and include a Lane number 0, TS1 Lane numbers are transmitted which must be assigned sequentially either in order or Lane reversed to the same Lanes and no Lane numbers can be repeated. Left over Lanes must transmit TS1 Link and Lane numbers set to PAD (K23.7). The next state is Configuration.Lanenum.Wait.
 - Note: The transmitted Lane numbers must range from 0 to m-1, be assigned sequentially in order or Lane reversed to the same grouping of Lanes that are receiving Lane numbers 0 to n-1, include the received Lane 0, and m-1 must be smaller than the largest received Lane number (n-1).
- ❑ The next state is Detect if no Link can be configured or if all Lanes receive two consecutive TS1 ordered sets with Link and Lane numbers set to Pad (K23.7).

4.2.6.3.3.2 Upstream Lanes

- ❑ If two consecutive TS2 ordered sets are received with Link and Lane numbers (non PAD values) that match what is being transmitted in the TS1 ordered sets, the next state is Configuration.Complete.
- ❑ If the Lanes of a Link can be configured by using the Lanes that transmitted a Lane number, which received two consecutive TS1 ordered sets with the same transmitted Link number (non-Pad) and include a Lane number 0, TS1 Lane numbers are transmitted which must be assigned sequentially either in order or Lane reversed to the same Lanes and no Lane numbers can be repeated. Left over Lanes must transmit TS1 Link and Lane numbers set to PAD (K23.7). The next state is Configuration.Lanenum.Wait.
 - Note: The transmitted Lane numbers must range from 0 to m-1, be assigned sequentially in order or Lane reversed to the same grouping of Lanes that are receiving Lane numbers 0 to n-1, include the received Lane 0, and m-1 must be smaller than the largest received Lane number (n-1).
- ❑ The next state is Detect if no Link can be configured or if all Lanes receive two consecutive TS1 ordered sets with Link and Lane numbers set to Pad (K23.7).

4.2.6.3.4. Configuration.Lanenum.Wait**4.2.6.3.4.1. Downstream Lanes**

- ❑ The next state is Configuration.Lanenum.Accept if any of the Lanes receive two consecutive TS1 which have a Lane number different from when the Lane first entered Configuration.Lanenum.Wait, and not all the Lanes Link numbers are set to Pad (K23.7).
- ❑ The next state is Detect after a 2 ms timeout or if all Lanes receive two consecutive TS1 ordered sets with Link and Lane numbers set to Pad (K23.7).

4.2.6.3.4.2. Upstream Lanes

- ❑ The next state is Configuration.Lanenum.Accept
 1. If any of the Lanes receive two consecutive TS1s that have a Lane number different from when the Lane first entered Configuration.Lanenum.Wait, and not all the Lanes' Link numbers are set to Pad (K23.7)
 - or
 2. If any Lane receives two consecutive TS2 ordered sets
- ❑ The next state is Detect after a 2 ms timeout or if all Lanes receive two consecutive TS1 ordered sets with Link and Lane numbers set to Pad (K23.7).

4.2.6.3.5. Configuration.Complete

4.2.6.3.5.1. Downstream Lanes

- ❑ TS2 ordered sets are transmitted using Link and Lane numbers that match the received TS1 Link and Lane numbers.
- ❑ N_FTS must be noted for use in L0s when leaving this state.
- ❑ Lane-to-Lane de-skew must be completed when leaving this state.
- ❑ Scrambling is disabled if all configured Lanes have the Disable Scrambling bit asserted in two consecutively received TS2 ordered sets.
 - Note: It is required that the Port that is sending the Disable Scrambling bit on all of the configured Lanes will also disable scrambling.
- ❑ The next state is Configuration.Idle immediately after all Lanes that are transmitting TS2 ordered sets receive eight consecutive TS2 ordered sets with matching Lane and Link numbers (non-Pad) and 16 TS2 ordered sets are sent after receiving one TS2 ordered set.
 - Note: All remaining Lanes that are not part of the configured Link are no longer associated with the LTSSM in progress and must:
 - i. Be associated with a new LTSSM if this optional feature is supported.
 or
 - ii. All Lanes that cannot be associated with an optional new LTSSM must transition to Electrical Idle.³⁴
 - Note: In the case of an optional crosslink, the receiver terminations are required to meet $Z_{RX-HIGH-IMP-DC}$ (see Table 4-6).
 - Note: These Lanes must be re-associated with the LTSSM immediately after the LTSSM in progress transitions back to Detect.
- ❑ The next state is Detect after a 2 ms timeout.

4.2.6.3.5.2. Upstream Lanes

- ❑ TS2 ordered sets are transmitted using Link and Lane numbers that match the received TS2 Link and Lane numbers.
- ❑ N_FTS must be noted for use in L0s when leaving this state.
- ❑ Lane-to-Lane de-skew must be completed when leaving this state.

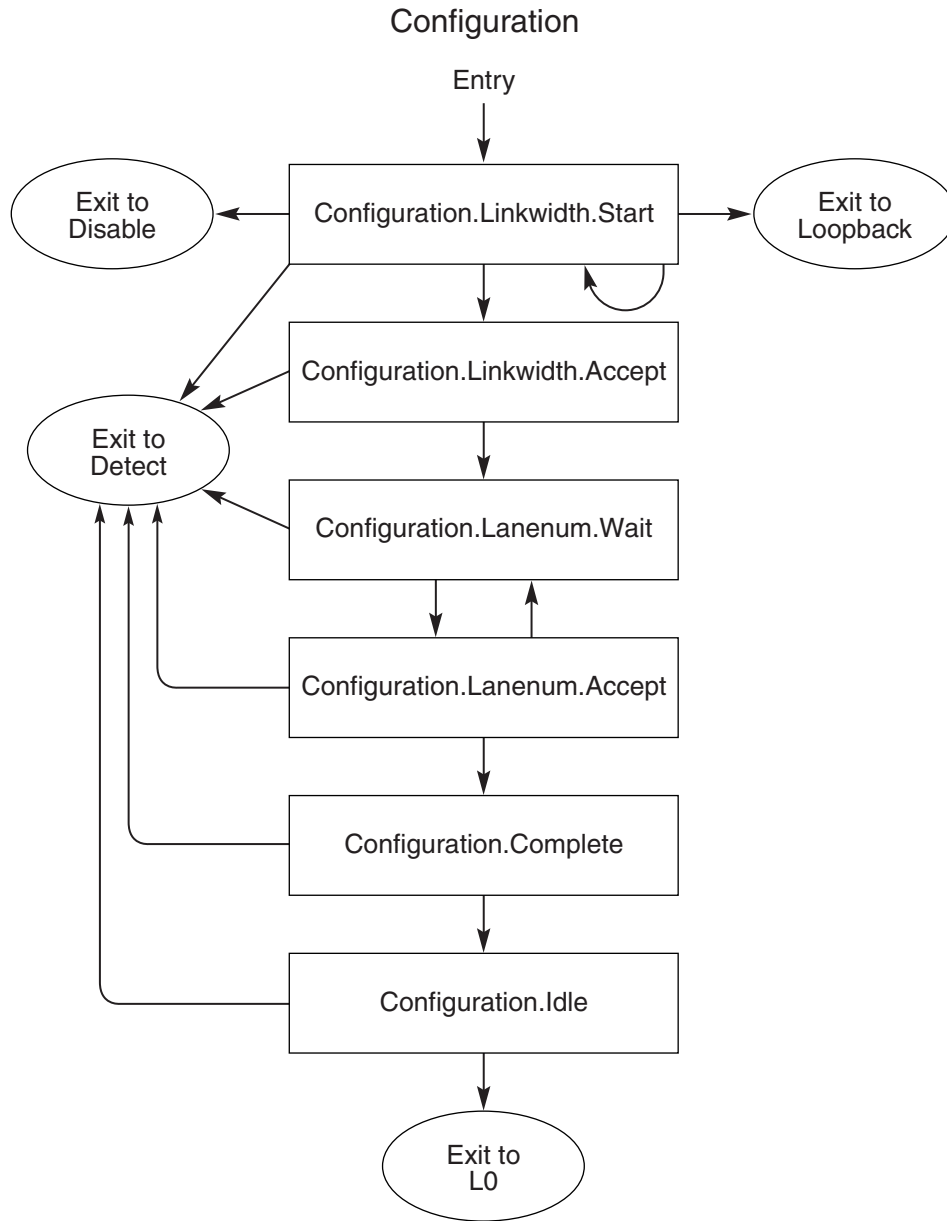
³⁴ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 4-5).

- ❑ Scrambling is disabled if all configured Lanes have the Disable Scrambling bit asserted in two consecutively received TS2 ordered sets.
 - Note: It is required that the Port that is sending the Disable Scrambling bit on all of the configured Lanes will also disable scrambling.
- ❑ The next state is Configuration.Idle immediately after all Lanes that are transmitting TS2 ordered sets receive eight consecutive TS2 ordered sets with matching Lane and Link numbers (non-Pad) and 16 consecutive TS2 ordered sets are sent after receiving one TS2 ordered set.
 - Note: All remaining Lanes that are not part of the configured Link are no longer associated with the LTSSM in progress and must:
 - i. Optionally be associated with a new crosslink LTSSM if this feature is supported.
 - or
 - ii. All remaining Lanes that are not associated with a new crosslink LTSSM must transition to Electrical Idle³⁵, and receiver terminations are required to meet $Z_{RX-HIGH-IMP-DC}$ (see Table 4-6).
 - Note: These Lanes must be re-associated with the LTSSM immediately after the LTSSM in progress transitions back to Detect.
- ❑ The next state is Detect after a 2 ms timeout.

4.2.6.3.6. Configuration.Idle

- ❑ Transmitter sends Idle data symbols on all configured Lanes.
- ❑ Receiver waits for Idle data.
- ❑ LinkUp = 1 (status is set true).
- ❑ Next state is L0 if eight consecutive symbol times of Idle data received on all configured Lanes and 16 Idle data symbols are sent after receiving one Idle data symbol.
- ❑ Otherwise, the next state is Detect after a minimum 2 ms timeout.

³⁵ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 4-5).



OM13802A

Figure 4-14: Configuration Sub-State Machine

4.2.6.4. Recovery

4.2.6.4.1. Recovery.RcvrLock

- ❑ Transmitter sends TS1 ordered sets on all configured Lanes using the same Link and Lane numbers that were set after leaving Configuration.
- ❑ Next state is Recovery.RcvrCfg if eight consecutive TS1 or TS2 ordered sets are received on all configured Lanes with the same Link and Lane numbers that match what is being transmitted on those same Lanes.
 - Note: If the Extended Synch bit is set, the transmitter must send a minimum of 1024 consecutive TS1 ordered sets before transitioning to Recovery.RcvrCfg.
 - ◆ Note: The Extended Synch allows external Link monitoring tools (e.g., logic analyzers) enough time to achieve bit and symbol lock.
- ❑ Otherwise, after a 24 ms timeout:
 1. The next state is Configuration if all the configured Lanes that are receiving a TS1 or TS2 ordered set have received at least one TS1 or TS2 with Link and Lane numbers that match what is being transmitted on those same Lanes.
 2. Otherwise, the next state is Detect.

4.2.6.4.2. Recovery.RcvrCfg

- ❑ Transmitter sends TS2 ordered sets on all configured Lanes using the same Link and Lane numbers that were set after leaving Configuration.
- ❑ Next state is Recover.Idle if eight consecutive TS2 ordered sets are received on all configured Lanes with the same Link and Lane number that match what is being transmitted on those same Lanes and 16 TS2 ordered sets are sent after receiving one TS2 ordered set.
 - Note: If the N_FTS value was changed, the new value must be used for future L0s states.
 - Note: Lane-to-Lane de-skew must be completed before leaving Recovery.RcvrCfg.
- ❑ Next state is Configuration if eight consecutive TS1 ordered sets are received on any configured Lanes with Link or Lane numbers that do not match what is being transmitted on those same Lanes and 16 TS2 ordered sets are sent after receiving one TS2 ordered set.
 - Note: If the N_FTS value was changed, the new value must be used for future L0s states.
- ❑ Otherwise, after a 48 ms timeout the next state is Detect.

4.2.6.4.3. Recovery.Idle

- ☐ Next state is Disabled if directed.
 - Note: “if directed” applies to a Downstream or optional crosslink Port that is instructed by a higher Layer to assert the Disable Link bit (TS1 and TS2) on the Link.
- ☐ Next state is Hot Reset if directed.
 - Note: “if directed” applies to a Downstream or optional crosslink Port that is instructed by a higher Layer to assert the Hot Reset bit (TS1 and TS2) on the Link.
- ☐ Next state is Configuration if directed.
 - Note: “if directed” applies to a Port that is instructed by a higher Layer to optionally re-configure the Link (i.e., different width Link).
- ☐ Next state is Loopback if directed to this state and the transmitter is capable of being a Loopback Master.
 - Note: “if directed” applies to a Port that is instructed by a higher Layer to assert the Loopback bit (TS1 and TS2) on the Link.
- ☐ Next state is Disabled immediately after any configured Lane has the Disable Link bit asserted in two consecutively received TS1 ordered sets.
 - Note: This is behavior only applicable to Upstream and optional crosslink Ports.
- ☐ Next state is Hot Reset immediately after any configured Lane has the Hot Reset bit asserted in two consecutive TS1 ordered sets.
 - Note: This is behavior only applicable to Upstream and optional crosslink Ports.
- ☐ Next state is Configuration if two consecutive TS1 ordered sets are received on any configured Lane with a Lane number set to PAD.
 - Note: A Port that optionally transitions to Configuration to change the Link configuration is guaranteed to send Lane numbers set to PAD on all Lanes.
- ☐ Next state is Loopback if any configured Lane has the Loopback bit asserted in two consecutive TS1 ordered sets.
 - Note: The device receiving the ordered set with the Loopback bit set becomes the Loopback Slave.
- ☐ Transmitter sends Idle data on all configured Lanes.
 - Note: If directed to other states, Idle Characters do not have to be sent before transitioning to the other states (i.e., Disable, Hot Reset, Configuration, or Loopback)
- ☐ Next state is L0 if eight consecutive symbol times of Idle data received on all configured Lanes and 16 Idle data symbols are sent after receiving one Idle data symbol.
- ☐ Otherwise, after a 2 ms timeout the next state is Detect.

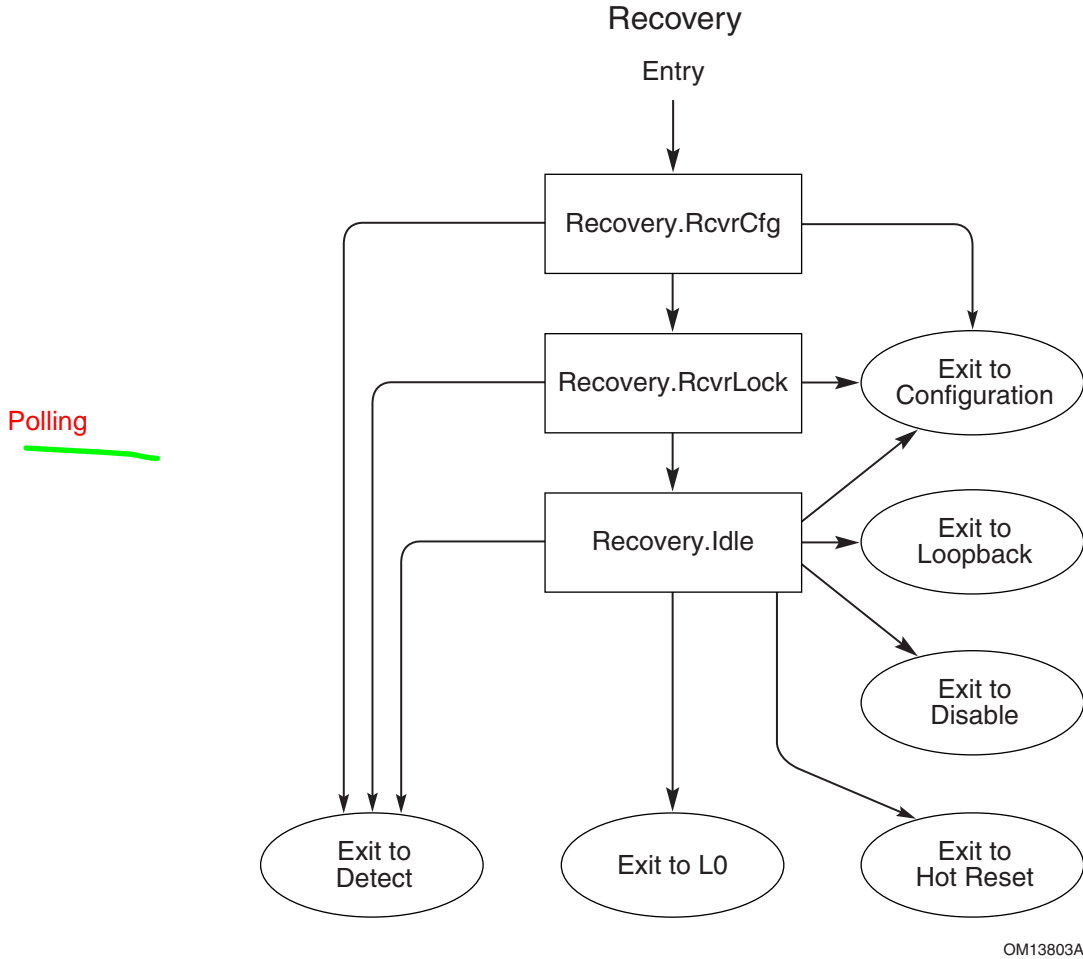


Figure 4-15: Recovery Sub-State Machine

4.2.6.5. L0

This is the normal operational state.

- ☐ LinkUp = 1 (status is set true).
- ☐ Next state is Recovery if a TS1 or TS2 ordered set is received on any configured Lane.
- ☐ Next state is Recovery if directed to this state or if Electrical Idle is detected without receiving an Electrical Idle ordered set.
 - Note: “if directed” applies to a Port that is instructed by a higher Layer to transition to Recovery.
 - Note: The transmitter may complete any TLP or DLLP in progress.
- ☐ Next state of transmitter is L0s if directed to this state.
 - Note: “if directed” applies to a Port that is instructed by a higher Layer to initiate L0s.

- ❑ Next state of receiver is L0s if receiver detects Electrical Idle ordered set and is not directed to L1 or L2 states.
- ❑ Next state is L1:
 - i. If directed
 - and
 - ii. an Electrical Idle ordered set is received
 - and
 - iii. an Electrical Idle ordered set is transmitted.
 - Note: “if directed” is defined as both ends of the Link having agreed to enter L1 immediately after the condition of both the receipt and transmission of an Electrical Idle ordered set is met.
 - Note: When directed by a higher Layer one side of the Link always initiates and exits to L1 by transmitting an Electrical Idle ordered set, followed by a transition to Electrical Idle.³⁶ The same Port then waits for the receipt of an Electrical Idle ordered set, and then immediately transitions to L1. Conversely, the side of the Link that first receives an Electrical Idle ordered set must send an Electrical Idle ordered set and immediately transition to L1.
- ❑ Next state is L2:
 - i. If directed
 - and
 - ii. an Electrical Idle ordered set is received
 - and
 - iii. an Electrical Idle ordered set is transmitted.
 - Note: “if directed” is defined as both ends of the Link having agreed to enter L2 immediately after the condition of both the receipt and transmission of an Electrical Idle ordered set is met.
 - Note: When directed by a higher Layer, one side of the Link always initiates and exits to L2 by transmitting an Electrical Idle ordered set followed by a transition to Electrical Idle.³⁷ The same Port then waits for the receipt of an Electrical Idle ordered set, and then immediately transitions to L2. Conversely, the side of the Link that first receives an Electrical Idle ordered set must send an Electrical Idle ordered set and immediately transition to L2.

³⁶ The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 4-5).

³⁷ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 4-5).

4.2.6.6. L0s

4.2.6.6.1. Receiver L0s

4.2.6.6.1.1. Rx_L0s.Entry

- Next state is Rx_L0s.Idle after a $T_{\text{TX-IDLE-SET-TO-IDLE}}$ (Table 4-5) timeout
 - Note: This guarantees that the transmitter has established the Electrical Idle condition.

4.2.6.6.1.2. Rx_L0s.Idle

- Next state is Rx_L0s.FTS if the receiver detects an exit from Electrical Idle on any Lane of the configured Link.

4.2.6.6.1.3. Rx_L0s.FTS

- The next state is L0 if a SKP ordered set is received on all configured Lanes of the Link.
 - Note: The receiver must be able to accept valid data immediately after the SKP ordered set.
 - Note: Lane-to-Lane de-skew must be completed before leaving Rx_L0s.FTS.
- Otherwise, next state is Recovery after the N_FTS timeout.
 - Note: The N_FTS timeout is approximately $40 \times [N_FTS + 1 \text{ (SKP)}] \times \text{UI}$. This time would be exact except for that SKPs may be longer than a four symbol ordered set.
 - Note: The transmitter must also transition to Recovery, but may complete any TLP or DLLP in progress.
 - Note: It is recommended that the N_FTS field be increased when transitioning to Recovery to prevent future transitions to Recovery from Rx_L0s.FTS.

4.2.6.6.2. Transmitter L0s

4.2.6.6.2.1. Tx_L0s.Entry

- ❑ Transmitter sends the Electrical Idle ordered set and enters Electrical Idle.
 - Note: The DC common mode voltage must be within specification by $T_{TX-IDLE-SET-TO-IDLE}$ ³⁸
- ❑ Next state is Tx_L0s.Idle after a $T_{TX-IDLE-MIN}$ (Table 4-5) timeout.

4.2.6.6.2.2. Tx_L0s.Idle

- ❑ Next state is Tx_L0s.FTS if directed.

4.2.6.6.2.3. Tx_L0s.FTS

- ❑ Transmitter sends N_FTS Fast Training Sequences on all configured Lanes.
 - Note: No SKP ordered sets can be inserted before all FTS ordered sets as defined by the agreed upon N_FTS parameter are transmitted.
 - Note: If the Extended Synch bit is set, the Transmitter sends 4096 Fast Training Sequences.
- ❑ Transmitter sends a single SKP ordered set on all configured Lanes.
- ❑ Next state is L0.

³⁸ The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 4-5).

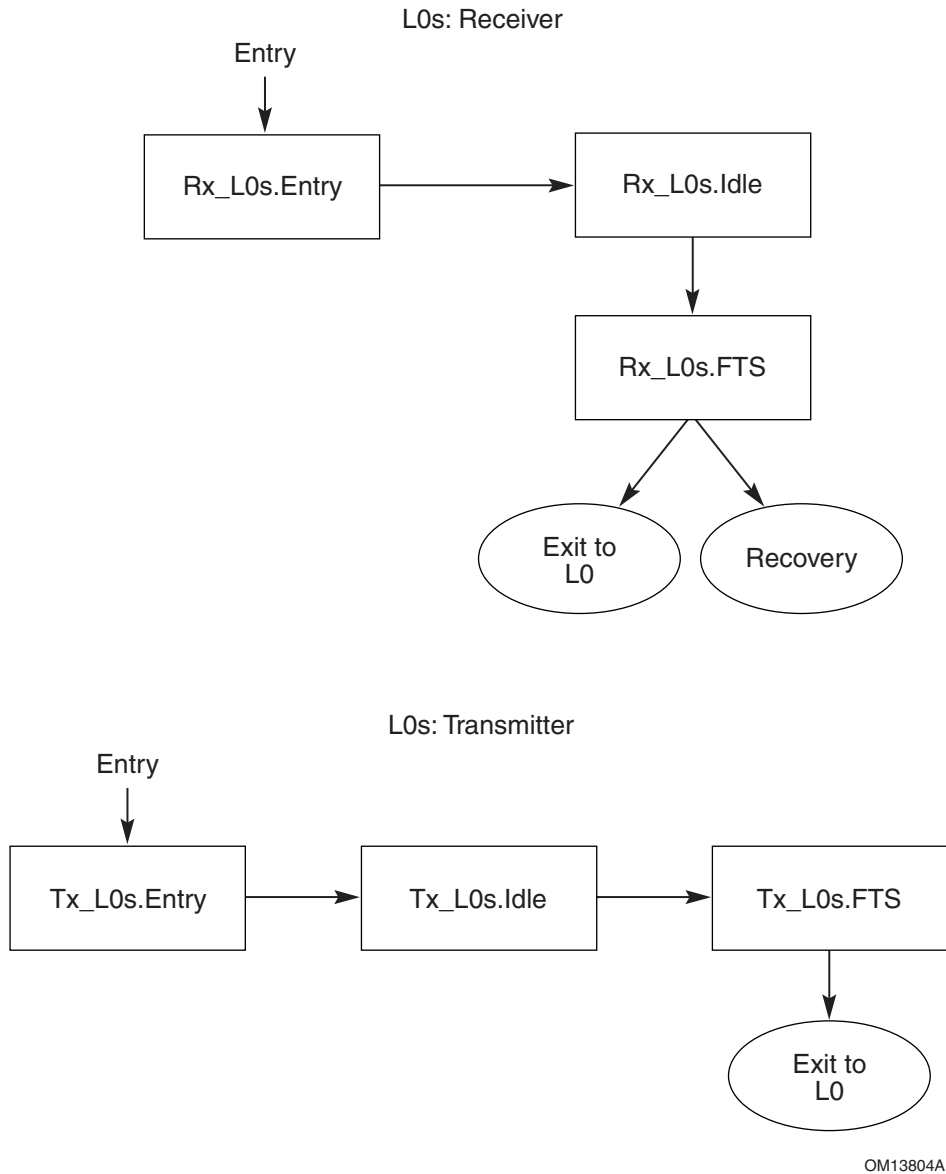


Figure 4-16: L0s Sub-State Machine


4.2.6.7. L1

4.2.6.7.1. L1.Entry

- ❑ All configured transmitters are in Electrical Idle.
 - Note: The DC common mode voltage must be within specification by $T_{TX-IDLE-SET-TO-IDLE}$.³⁹
- ❑ The next state is L1.Idle after a $T_{TX-IDLE-MIN}$ (Table 4-5) timeout.
 - Note: This guarantees that the transmitter has established the Electrical Idle condition.

4.2.6.7.2. L1.Idle

- ❑ Transmitter remains in Electrical Idle.
 - Note: The DC common mode voltage must be within specification.⁴⁰
- ❑ Next state is Recovery if any receiver detects exit from Electrical Idle or directed.



OM13805A

Figure 4-17: L1 Sub-State Machine

³⁹ The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 4-5).

⁴⁰ The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 4-5).

4.2.6.8. L2

4.2.6.8.1. L2.Idle

- ❑ All RX Termination must remain enabled in low impedance.
- ❑ All configured transmitters are in Electrical Idle.
 - Note: The DC common mode voltage does not have to be within specification.⁴¹
- ❑ For Downstream Lanes:
 - For a Root Port, the next state is Detect if a Beacon is received on at least Lane 0.
 - ♦ Note: Main power must be restored before entering Detect.
 - For a Switch, if a Beacon is received on at least Lane 0, the Upstream Port must transition to L2.TransmitWake.
- ❑ For Upstream Lanes:
 - The next state is Detect if Electrical Idle Exit is detected on any Lane.
 - ♦ Note: A Switch must transition any Downstream Lanes to Detect.
 - Next state is L2.TransmitWake for an Upstream Port if directed to transmit a Beacon.
 - ♦ Note: Beacons may only be transmitted on Upstream Ports in the direction of the Root Complex.

4.2.6.8.2. L2.TransmitWake

Note: This state only applies to Upstream Ports.

- ❑ Transmit the Beacon on at least Lane 0 (Refer to Section 4.3.2.4).
- ❑ Next state is Detect if Electrical Idle exit is detected on any Upstream Port's receiver that is in the direction of the Root Complex.
 - Note: Power is guaranteed to be restored when Upstream receivers see Electrical Idle exited, but it may also be restored prior to Electrical Idle being exited.

⁴¹ The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle (V_{TX-CM-DC-ACTIVE-IDLE-DELTA}) specification (see Table 4-5).




Figure 4-18: L2 Sub-State Machine

4.2.6.9. Disabled

- ❑ All Lanes transmit 16 TS1 ordered sets with the Disable Link bit (bit 1) asserted and then transition to Electrical Idle.
 - Note: The Electrical Idle ordered set must be sent prior to entering Electrical Idle.
 - Note: The DC common mode voltage does not have to be within specification.⁴²
- ❑ If an Electrical Idle ordered set was transmitted and received (even while transmitting TS1 with the Disable Link bit asserted), then:
 - LinkUp = 0 (False)
 - ◆ Note: At this point, the Lanes are considered Disabled.
 - The next state is Detect when directed or if Electrical Idle is exited.
- ❑ Otherwise, if no Electrical Idle ordered set is received after a 2 ms timeout, the next state is Detect.

⁴² The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle (V_{TX-CM-DC-ACTIVE-IDLE-DELTA}) specification (see Table 4-5).

4.2.6.10. Loopback

4.2.6.10.1. Loopback.Entry

- ❑ LinkUp = 0 (False)
- ❑ The Loopback Master device transmits TS1 ordered sets with the Loopback bit (Bit 2) asserted until the Loopback Master receives identical TS1 ordered sets with the Loopback bit asserted. The next state is Loopback.Active.
 - Note: This indicates to the Loopback Master that the Loopback Slave has successfully entered Loopback.
 - Note: The Loopback Master timeout is implementation specific. The exit is to Loopback.Exit.
- ❑ The next state for the Loopback Slave is Loopback.Active.
 - Note: The Loopback Slave will immediately transition to Loopback.Active.

4.2.6.10.2. Loopback.Active

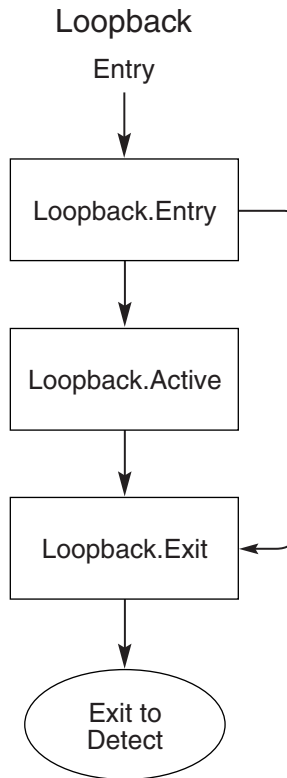
- ❑ The Loopback Master must send valid 8b/10b data. The next state of the Loopback Master is Loopback.Exit if directed
- ❑ A Loopback Slave is required to retransmit each 10b data and control symbol exactly as received, without applying scrambling/descrambling or disparity corrections, with three important exceptions:
 - If a received 10b symbol is determined to be an invalid 10b code (i.e., no legal translation to a control or data value possible) then the slave must instead transmit the EDB symbol in the corresponding time slot of the invalid symbol. Either a positive or negative disparity can be chosen for the EDB symbol.
 - If a SKP ordered set retransmission requires adding a SKP symbol to accommodate timing tolerance correction, the SKP symbol is inserted in the retransmitted symbol stream anywhere in the SKP ordered set following the COM symbol. Either a positive or negative disparity can be chosen for the inserted SKP symbol.
 - If a SKP ordered set retransmission requires dropping a SKP symbol to accommodate timing tolerance correction, the SKP symbol is simply not retransmitted and transmission continues with the next received symbol or an EDB, as defined above.

As result of these rules, received valid 10b codes are retransmitted even if they fail to match expected disparity in the receiver and result in retransmission violating normal disparity rules.

- ❑ Next state of the Loopback Slave is Loopback.Exit when an Electrical Idle ordered set is received or Electrical Idle is detected.
 - Note: A Loopback Slave must be able to detect Electrical Idle was entered within 2 ms in case the Electrical Idle ordered set is not properly detected.
- ❑ The next state of the Loopback Master is Loopback.Exit if directed.

4.2.6.10.3. Loopback.Exit

- ❑ The Loopback Master sends an Electrical Idle ordered set and goes to Electrical Idle for a minimum of 2 ms (Table 4-5).
- ❑ The Loopback Slave echoes the Electrical Idle ordered set and goes to Electrical Idle for a minimum of 2 ms (Table 4-5).
- ❑ The next state of the Loopback Master and Loopback Slave is Detect.



OM13807A

Figure 4-19: Loopback State Machine

4.2.6.11. Hot Reset

- ❑ Lanes that were directed by a higher Layer to initiate Hot Reset:
 - All Lanes in the configured Link transmit TS1 ordered sets with the Hot Reset bit (Bit 0) asserted and the configured Link and Lane numbers.
 - If two consecutive TS1 ordered sets with the Hot Reset bit (Bit 0) asserted are received, then:
 - ◆ LinkUp = 0 (False)
 - ◆ Next state is Detect if no higher Layer is directing the Physical Layer to remain in Hot Reset.
 - Otherwise, after a 2 ms timeout next state is Detect.
- ❑ Lanes that were not directed by a higher Layer to initiate Hot Reset (i.e., received two consecutive TS1 ordered sets with the Hot Reset bit (Bit 0) asserted on any configured Lanes):
 - LinkUp = 0 (False)
 - If any Lane of an Upstream Port of a Switch receives a training sequence with the Hot Reset bit asserted, all configured Downstream Ports must transition to Hot Reset as soon as possible.
 - All Lanes in the configured Link transmit TS1 ordered sets with the Hot Reset bit (Bit 0) asserted and the configured Link and Lane numbers.
 - After a 2 ms timeout:
 - ◆ If two consecutive TS1 ordered sets were received with the Hot Reset bit (Bit 0) asserted and the configured Link and Lane numbers, the next state is Hot Reset.
 - ◆ Otherwise, the next state is Detect.

4.2.7. Clock Tolerance Compensation

Skip ordered sets (defined below) are used to compensate for differences in frequencies between bit rates at two ends of a Link. The Receiver Physical Layer Logical sub-block must include elastic buffering which performs this compensation. The interval between SKIP ordered set transmissions is derived from the absolute value of the Transmit and Receive clock frequency difference specified in Table 4-5. Having worse case clock frequencies at the limits of the tolerance specified will result in a 600 ppm difference between the transmit and receive clocks of a Link. As a result, the transmit and receive clocks can shift one clock every 1666 clocks.

4.2.7.1. Rules for Transmitters

- ❑ All Lanes shall transmit Symbols at the same frequency (the difference between bit rates is 0 ppm within all multi-Lane Links).
- ❑ When transmitted, the SKIP ordered set shall be transmitted simultaneously on all Lanes of a multi-Lane Link (See Section 4.2.4.8 and Table 4-5 for the definition of simultaneous in this context).
- ❑ The transmitted SKIP ordered set is: one COM Symbol followed by three consecutive SKP Symbols
- ❑ The SKIP ordered set shall be scheduled for insertion at an interval between 1180 and 1538 Symbol Times.
- ❑ Scheduled SKIP ordered sets shall be transmitted if a packet or ordered set is not already in progress, otherwise they are accumulated and then inserted consecutively at the next packet or ordered set boundary.
- ❑ SKIP ordered sets do not count as an interruption when monitoring for consecutive characters or ordered set (i.e., eight consecutive TS1 ordered sets in Polling.Active)

4.2.7.2. Rules for Receivers

- ❑ Receivers shall recognize received SKIP ordered set consisting of one COM Symbol followed consecutively by one to five SKP Symbols.
 - Note: The number of received SKP symbols in an ordered set shall not vary from Lane-to-Lane in a multi-Lane Link.
- ❑ Receivers shall be tolerant to receive and process SKIP ordered sets at an average interval between 1180 to 1538 symbol times.
- ❑ Receivers shall be tolerant to receive and process consecutive SKIP ordered sets.
 - Note: Receivers shall be tolerant to receive and process SKIP ordered sets separated from each other at most 5664 symbol times – measured as the distance between the leading COM symbols.

4.2.8. Compliance Pattern

During Polling, the Polling.Compliance sub-state must be entered based on the presence of test equipment being attached to one Lane of a possible Link (see Section 4.2.5.3). The compliance pattern consists of the sequence of 8b/10b symbols K28.5, D21.5, K28.5, and D10.2 repeating. Current running disparity must be set to negative before sending the first symbol.

The compliance sequence is:

Symbol	K28.5	D21.5	K28.5	D10.2
Current Disparity	0	1	1	0
Pattern	0011111010	1010101010	1100000101	0101010101

For any given device that has multiple Lanes, every eighth Lane is delayed by a total of four symbols. A two symbol delay occurs at both the beginning and end of the four symbol sequence.

This delay sequence on every eight Lane is illustrated below:

Symbol:	D	D	K28.5	D21.5	K28.5	D10.2	D	D
----------------	---	---	-------	-------	-------	-------	---	---

D a K28.5 symbol.

After the eight symbols are sent, the delay symbols are advanced to the next Lane and the process is repeated. An illustration of this process is shown below:

Lane 0	D	D	K28.5-	D21.5	K28.5+	D10.2	D	D	K28.5-	D21.5	K28.5+	D10.2
Lane 1	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2	D	D	K28.5-	D21.5
Lane 2	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2
Lane 3	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2
Lane 4	K28.5-	D21.5	K28.5-	D10.2	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2
Lane 5	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2
Lane 6	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2
Lane 7	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2
Lane 8	D	D	K28.5-	D21.5	K28.5+	D10.2	D	D	K28.5-	D21.5	K28.5+	D10.2
Lane 9	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2	D	D	K28.5-	D21.5

Key:

K28.5- Comma when disparity is negative, specifically: "0011111010"

K28.5+ Comma when disparity is positive, specifically: "1100000101"

D21.5 Out of phase data character, specifically: "1010101010"

D10.2 Out of phase data character, specifically: "0101010101"

D Delay Character – K28.5

This sequence of delays ensures interference between adjacent Lanes, enabling measurement of the compliance pattern under close to worst-case Inter-symbol Interference and cross-talk conditions.

The compliance pattern can be exited only when an Electrical Idle Exit is detected at all the Lanes that detected a receiver during Detect.

4.3. Electrical Sub-Block

The Electrical sub-block contains a Transmitter and a Receiver. The Transmitter is supplied by the Logical sub-block with Symbols which it serializes and transmits onto a Lane. The Receiver is supplied with serialized Symbols from the Lane. It transforms the electrical signals into a bit stream which is de-serialized and supplied to the Logical sub-block along with a Link clock recovered from the incoming serial stream.

4.3.1. Electrical Sub-Block Requirements

4.3.1.1. Clocking Dependencies

The Ports on the two ends of a Link must transmit data at a rate that is within 600 parts per million (ppm) of each other at all times. This is specified to allow bit rate clock sources with a +/- 300 ppm tolerance.

4.3.1.1.1. Spread Spectrum Clock (SSC) Sources

The data rate can be modulated from +0% to -0.5% of the nominal data rate frequency, at a modulation rate in the range not exceeding 30 kHz – 33 kHz. The +/- 300 ppm requirement still holds, which requires the two communicating Ports be modulated such that they never exceed a total of 600 ppm difference. For most implementations this places the requirement that both Ports require the same bit rate clock source when the data is modulated with an SSC.

4.3.1.2. AC Coupling

Each Lane of a Link must be AC coupled. The minimum and maximum value for the capacitance is given in Table 4-5. The requirement for the inclusion of AC coupling capacitors on the interconnect media is specified by the transmitter.

4.3.1.3. Interconnect

In the context of this specification, the interconnect comprises everything between the pins at of a transmitter package and the pins of a receiver package. Often, this will consist of traces on a printed circuit board or other suitable medium, AC coupling capacitors and perhaps connectors. The interconnect total capacitance to ground seen by the receiver detection circuit (see Section 4.3.1.8) must not exceed 3 nF, including capacitance added by attached test instrumentation. Note that this capacitance is separate and distinct from the AC Coupling capacitance value (see Section 4.3.1.2).

4.3.1.4. Termination

- ❑ The transmitter is required to meet RL_{TX-DM} , RL_{TX-CM} , $Z_{TX-DIFF-DC}$, Z_{TX-DC} (see Table 4-5) any time functional differential signals are being transmitted.
- ❑ The transmitter is required to only meet Z_{TX-DC} (see Table 4-5) anytime functional differential signals are not being transmitted.
 - Note: The differential impedance during this same time is not defined.
- ❑ The receiver is required to meet RL_{RX-DM} , RL_{RX-CM} , $Z_{RX-DIFF-DC}$, Z_{RX-DC} (see Table 4-6) during all LTSSM states excluding only times during when the device is powered down, Fundamental Reset is asserted, or when explicitly specified.
- ❑ The receiver is required to meet $Z_{RX-HIGH-IMP-DC}$ (see Table 4-6) anytime adequate power is not provided to the receiver, Fundamental Reset is asserted (see Section 4.2.4.5.1), or when explicitly specified.
 - Note: The receiver differential impedance is not defined.

4.3.1.5. DC Common Mode Voltage

The receiver DC common mode voltage is always 0 V during all states.

The transmitter DC common mode voltage is held at the same value during all states unless otherwise specified. The range of allowed transmitter DC common mode values is specified in Table 4-5 ($V_{TX-DC-CM}$).

4.3.1.6. ESD

All signal and power pins must withstand 2000 V of ESD using the human body model and 500 V using the charged device model without damage. Class 2 per JEDEC JESE22-A114-A.

This ESD protection mechanism also helps protect the powered down receiver from potential common mode transients during certain possible reset or surprise insertion situations.

4.3.1.7. Short Circuit Requirements

All Transmitters and Receivers must support surprise hot insertion/removal without damage to the component. The transmitter and receiver must be capable of withstanding sustained short circuit to ground of D+ and D-. The transmitter short circuit current limit $I_{TX-SHORT}$ is provided in Table 4-5.

4.3.1.8. Receiver Detection

The Receiver Detection circuit is performed by a transmitter and must correctly detect whether a load impedance of Z_{RX-DC} or lower is present.

The behavior of the receiver detection sequence is described below.

Step 1. Transmitter is in a stable Electrical Idle state.⁴³

Step 2. A transmitter changes the common mode voltage on D+ and D- to a different value.

- a. A receiver is detected based on the rate that the lines change to the new voltage.
 - i. The receiver is not present if the voltage at the transmitter charges at a rate dictated only by the transmitter impedance, and capacitance of the interconnect, and series capacitor.
 - ii. The receiver is present if the voltage at the transmitter charges at a rate dictated by the transmitter impedance, the series capacitor, the interconnect capacitance, and the receiver termination.

Anytime Electrical Idle is exited the detect sequence does not have to execute or can be aborted on that Lane.

4.3.1.9. Electrical Idle

Electrical Idle is a steady state condition where the Transmitter D+ and D- voltages are held constant at the same value. Electrical Idle is primarily used in power saving and inactive states (i.e., Disable).

Before a transmitter enters Electrical Idle, it must always send the Electrical Idle ordered set, a K28.5 (COM) followed by three K28.3 (IDL) (see Table 4-5), unless otherwise specified. After sending the last symbol of the Electrical Idle ordered set, the transmitter must be in a valid Electrical Idle state as specified by $T_{TX-IDLE-SET-TO-IDLE}$ (see Table 4-5).

The successful reception of an Electrical Idle ordered set occurs upon the receipt of two out of the three K28.3 (IDL) characters in the transmitted Electrical Idle ordered set.

The low impedance common mode and differential receiver terminations values (see Section 4.3.1.4) must be met in Electrical Idle. The transmitter can be in either a low or high impedance mode during Electrical Idle.

Any time a transmitter enters Electrical Idle it must remain in Electrical Idle for a minimum of $T_{TX-IDLE-MIN}$ (see Table 4-5). The receiver should expect the Electrical Idle ordered set followed by a minimum amount of time in Electrical Idle ($T_{TX-IDLE-SET-TO-IDLE}$) to arm its Electrical Idle Exit detector.

⁴³ The common mode being driven does not have to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 4-5).

Electrical Idle Exit occurs when a signal larger than the minimum $V_{RX-IDLE-DET-DIFFP-P}$ is detected at a receiver.

4.3.2. Electrical Signal Specifications


A Differential Signal is defined by taking the voltage difference between two conductors. In this specification, a differential signal or differential pair is comprised of a voltage on a positive conductor, V_{D+} , and a negative conductor, V_{D-} . The differential voltage (V_{DIFF}) is defined as the difference of the positive conductor voltage and the negative conductor voltage ($V_{DIFF} = V_{D+} - V_{D-}$). The Common Mode Voltage (V_{CM}) is defined as the average or mean voltage present on the same differential pair ($V_{CM} = [V_{D+} + V_{D-}]/2$). This document's electrical specifications often refer to peak-to-peak measurements or peak measurements, which are defined by the following equations.

- $V_{DIFFP-P} = (2 * \max |V_{D+} - V_{D-}|)$ (This applies to a symmetric differential swing.)
- $V_{DIFFP-P} = (\max |V_{D+} - V_{D-}| \{V_{D+} > V_{D-}\} + \max |V_{D+} - V_{D-}| \{V_{D+} < V_{D-}\})$ (This applies to an asymmetric differential swing.)
- $V_{DIFFP} = (\max |V_{D+} - V_{D-}|)$ (This applies to a symmetric differential swing.)
- $V_{DIFFP} = (\max |V_{D+} - V_{D-}| \{V_{D+} > V_{D-}\})$ or $(\max |V_{D+} - V_{D-}| \{V_{D+} < V_{D-}\})$ whichever is greater (This applies to an asymmetric differential swing.)
- $V_{CMP} = (\max |V_{D+} + V_{D-}| / 2)$

Note: The maximum value is calculated on a per unit interval evaluation. The maximum function as described is implicit for all peak-to-peak and peak equations throughout the rest of this chapter, and thus a maximum function will not appear in any following subsequent representations of these equations.

In this section DC is defined as all frequency components below $F_{dc} = 30$ kHz. AC is defined as all frequency components at or above $F_{dc} = 30$ kHz. These definitions pertain to all voltage and current specifications.

An example waveform is shown in Figure 4-20. In this waveform the differential peak-peak signal is approximately 0.6 V, the differential peak signal is approximately 0.3 V and the common mode voltage is approximately 0.25 V.



OM13810

Figure 4-20: Sample Differential Signal

4.3.2.1. Loss

Loss (attenuation of the differential voltage swing) in this system is a critical parameter that must be properly considered and managed in order to ensure proper system functionality. Failure to properly consider the loss may result in a differential signal swing arriving at the Receiver that does not meet specifications. The interconnect loss is specified in terms of the amount of attenuation or loss that can be tolerated between the Transmitter (Tx) and Receiver (Rx). The Tx is responsible for producing the specified differential eye height at the pins of its package. Together, the Tx and the interconnect are responsible for producing the specified differential eye height at the Rx pins (see Figure 4-26).

The worst-case operational loss budget is calculated by taking the minimum output voltage ($V_{TX-DIFFP-P} = 800 \text{ mV}$) divided by the minimum input voltage to the receiver ($V_{RX-DIFFP-P} = 175 \text{ mV}$), which results in 13.2 dB. Additional headroom in the loss budget can be achieved by driving a larger differential output voltage (up to the maximum specified in Table 4-5) at the transmitter.

4.3.2.2. Jitter and BER

Jitter is categorized into random sources (R_j) and deterministic sources (D_j). The total jitter (T_j) is the convolution of the probability density functions for all the independent jitter sources, R_j and D_j . The nature of R_j can be approximated as Gaussian and is used to establish the bit error rate (BER) of the Link.

The UI allocation is given as the allowable T_j at the target BER. The UI allocation must meet a maximum BER of 10^{-12} for the T_j . The allocation to R_j and D_j is not specified.

The methods for measuring the BER compliance are beyond the scope of this specification.

4.3.2.3. De-emphasis


De-emphasis is included (i.e., Generation 1 fundamental band = 250 MHz to 1.25 GHz).

De-emphasis must be implemented when multiple bits of the same polarity are output in succession. Subsequent bits are driven at a differential voltage level 3.5 dB (+/- .5 dB) below the first bit. Note that individual bits, and the first bit from a sequence in which all bits have the same polarity, must always be driven between the Min and Max values as specified by $V_{TX-DIFFP-P}$ in Table 4-5.

The only exception pertains to transmitting the Beacon (see Section 4.3.2.4).

Note: The specified amount of de-emphasis allows for PCI Express designs to optimize maximum interoperability while minimizing complexity of managing configurable de-emphasis values. Thus, the primary benefits of de-emphasis are targeted for worst-case loss budgets of 11-13.2 dB, while being slightly less optimal for lower loss systems. However, this tradeoff is more than offset by the fact that there is inherently more voltage margin in lower loss systems.

An example waveform illustrating de-emphasis and representing the bit sequence (from left to right) of “1001000011” is shown in Figure 4-21.



OM14497

Figure 4-21: Sample Transmitted Waveform Showing -3.5 dB De-emphasis Around a 0.5 V Common Mode Voltage

4.3.2.4. Beacon

Support for Beacon is required for all “universal” PCI Express components that support a wakeup mechanism in order to function in form factors that require the use of Beacon. However, not all systems and form factor specifications require the use of Beacon, and for components that are restricted to use in such environments it is not necessary to support Beacon. This section applies to all components that support Beacon.

The Beacon is a signal sent by a Downstream Component to start the exit from an L2 state.


All transmitter electrical specifications (Table 4-5) must be met while sending a Beacon with the following exceptions and clarifications.

- ☐ The Beacon is a DC Balanced signal of periodic arbitrary data, which is required to contain some pulse widths ≥ 2 ns but no larger than 16 μ s.
- ☐ The maximum time between qualifying pulses ($2 \text{ ns} \leq x \leq 16 \mu\text{s}$) can be no larger than 16 μ s.
- ☐ DC Balance must be always be restored within a maximum time of 32 μ s.
- ☐ Beacon is transmitted in a low impedance mode.
- ☐ All Beacons must be transmitted and received on at least Lane 0 of multi-Lane Links.⁴⁴
- ☐ The output Beacon voltage level must be at a -6 dB de-emphasis level for Beacon pulses with a width greater than 500 ns.
- ☐ The output Beacon voltage level can range between the pre-emphasized and corresponding -3.5 dB de-emphasized voltage levels for Beacon pulses smaller than 500 ns.
- ☐ The Lane-to-Lane Output Skew (see Table 4-5) and SKP ordered set output (see Section 4.2.7) specifications do not apply.
- ☐ When any bridge and/or a Switch receives a Beacon at a Downstream Port, that component the Switch must propagate a Beacon wakeup indication upstream. This wakeup indication must use the appropriate wakeup mechanism required by the system or form factor associated with the Upstream Port of the Switch (see Section 5.3.3.2).

4.3.2.4.1. Beacon Example


An example receiver waveform driven at the -6 dB level for a 30 kHz Beacon is shown in Figure 4-22. An example receiver waveform using the COM character at full speed signaling is shown in Figure 4-23. It should be noted that other waveforms and signaling are possible other than the examples shown in Figure 4-22 and Figure 4-23 (i.e., Polling is another valid Beacon signal).

⁴⁴ Lane 0 as defined after Link Width and Lane reversal negotiations are complete.



OM13812

Figure 4-22: A 30 kHz BEACON Signaling Through a 75 nF Capacitor



OM14314

Figure 4-23: BEACON, Which Includes a 2 ns Pulse Through a 75 nF Capacitor

4.3.3. Differential Transmitter (TX) Output Specifications

The following table defines the specification of parameters for the differential output at all Transmitters (TXs). The parameters are specified at the component pins.

Table 4-5: Differential Transmitter (TX) Output Specifications

Symbol	Parameter	Min	Nom	Max	Units	Comments
UI	Unit Interval	399.88	400	400.12	ps	Each UI is 400 ps +/-300 ppm. UI does not account for SSC dictated variations. See Note 1.
$V_{TX-DIFFp-p}$	Differential Peak to Peak Output Voltage	0.800		1.2	V	$V_{TX-DIFFp-p} = 2 \cdot IV_{TX-D+} - V_{TX-D-}$ See Note 2.
$V_{TX-DE-RATIO}$	De-Emphasized Differential Output Voltage (Ratio)	-3.0	-3.5	-4.0	dB	This is the ratio of the $V_{TX-DIFFp-p}$ of the second and following bits after a transition divided by the $V_{TX-DIFFp-p}$ of the first bit after a transition. See Note 2.
T_{TX-EYE}	Minimum TX Eye Width	0.70			UI	The maximum transmitter jitter can be derived as $T_{TX-MAX-JITTER} = 1 - T_{TX-EYE} = .3$ UI See Notes 2 and 3.
$T_{TX-EYE-MEDIAN-to-MAX-JITTER}$	Maximum time between the jitter median and maximum deviation from the median.			0.15	UI	Jitter is defined as the measurement variation of the crossing points ($V_{TX-DIFFp-p} = 0$ V) in relation to an appropriate average TX UI. See Note 2 and 3.
$T_{TX-RISE},$ $T_{TX-FALL}$	D+/D- TX Output Rise/Fall Time	0.125			UI	See Notes 2 and 5.
$V_{TX-CM-ACp}$	AC Peak Common Mode Output Voltage			20	mV	$V_{TX-CM-ACp} = IV_{TX-D+} + V_{TX-D-}/2 - V_{TX-CM-DC}$ $V_{TX-CM-DC} = DC_{(avg)}$ of $IV_{TX-D+} + V_{TX-D-}/2$ during L0 See Note 2.
$V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$	Absolute Delta of DC Common Mode Voltage During L0 and Electrical Idle.	0		100	mV	$IV_{TX-CM-DC} [during L0] - V_{TX-CM-Idle-DC} [During Electrical Idle.] \leq 100$ mV $V_{TX-CM-DC} = DC_{(avg)}$ of $IV_{TX-D+} + V_{TX-D-}/2$ [L0] $V_{TX-CM-Idle-DC} = DC_{(avg)}$ of $IV_{TX-D+} + V_{TX-D-}/2$ [Electrical Idle] See Note 2.


Symbol	Parameter	Min	Nom	Max	Units	Comments
$V_{TX-CM-DC-LINE-DELTA}$	Absolute Delta of DC Common Mode Voltage between D+ and D-	0		25	mV	$ V_{TX-CM-DC-D+} [\text{during L0}] - V_{TX-CM-DC-D-} [\text{During L0}] \leq 25 \text{ mV}$ $V_{TX-CM-DC-D+} = DC_{(avg)} \text{ of } IV_{TX-D+} [\text{during L0}]$ $V_{TX-CM-DC-D-} = DC_{(avg)} \text{ of } IV_{TX-D-} [\text{during L0}]$ See Note 2.
$V_{TX-IDLE-DIFFP}$	Electrical Idle Differential Peak Output Voltage	0		20	mV	$V_{TX-IDLE-DIFFP} = IV_{TX-Idle-D+} - V_{TX-Idle-D-} \leq 20 \text{ mV}$ See Note 2.
$V_{TX-RCV-DETECT}$	The amount of voltage change allowed during Receiver Detection			600	mV	The total amount of voltage change that a transmitter can apply to sense whether a low impedance receiver is present. See Section 4.3.1.8.
$V_{TX-DC-CM}$	The TX DC Common Mode Voltage	0		3.6	V	The allowed DC Common Mode voltage under any conditions. See Section 4.3.1.8.
$I_{TX-SHORT}$	TX Short Circuit Current Limit			90	mA	The total current the transmitter can provide when shorted to its ground
$T_{TX-IDLE-MIN}$	Minimum time spent in Electrical Idle	50			UI	Minimum time a transmitter must be in Electrical Idle
$T_{TX-IDLE-SET-TO-IDLE}$	Maximum time to transition to a valid Electrical Idle after sending an Electrical Idle ordered set			20	UI	After sending an Electrical Idle ordered set, the transmitter must meet all Electrical Idle specifications within this time.
$RL_{TX-DIFF}$	Differential Return Loss	12			dB	Measured over 50 MHz to 1.25 GHz. See Note 4.
RL_{TX-CM}	Common Mode Return Loss	6			dB	Measured over 50 MHz to 1.25 GHz. See Note 4.
$Z_{TX-DIFF-DC}$	DC Differential TX Impedance	80	100	120	Ω	TX DC Differential Mode Low impedance
Z_{TX-DC}	Transmitter DC Impedance	40 k			Ω	Required TX D+ as well as D- DC impedance during all states
$L_{TX-SKEW}$	Lane-to-Lane Output Skew			$500 + 2 \text{ UI}$	ps	Static skew between any two Transmitter Lanes within a single Link
C_{TX}	AC Coupling Capacitor	75		200	nF	All transmitters shall be AC coupled. The AC coupling is required either within the media or within the transmitting component itself.

Symbol	Parameter	Min	Nom	Max	Units	Comments
$T_{\text{crosslink}}$	Crosslink Random Timeout	0		1	ms	This random timeout helps resolve conflicts in crosslink configuration by eventually resulting in only one Downstream and one Upstream Port (see 4.2.6.3).

Notes:

1. No test load is necessarily associated with this value.
2. Specified at the measurement point into a timing and voltage compliance test load as shown in Figure 4-25 and measured over any 250 consecutive TX UIs. (Also refer to the Transmitter Compliance Eye Diagram as shown in Figure 4-24.)
3. A $T_{\text{TX-EYE}} = 0.70$ UI provides for a total sum of deterministic and random jitter budget of $T_{\text{TX-JITTER-MAX}} = 0.30$ UI for the transmitter collected over any 250 consecutive TX UIs. The $T_{\text{TX-EYE-MEDIAN-to-MAX-JITTER}}$ specification ensures a jitter distribution in which the median and the maximum deviation from the median is less than half of the total TX jitter budget collected over any 250 consecutive TX UIs. It should be noted that the median is not the same as the mean. The jitter median describes the point in time where the number of jitter points on either side is approximately equal as opposed to the averaged time value.
4. The transmitter input impedance shall result in a differential return loss greater than or equal to 12 dB and a common mode return loss greater than or equal to 6 dB over a frequency range of 50 MHz to 1.25 GHz. This input impedance requirement applies to all valid input levels. The reference impedance for return loss measurements is 50 ohms to ground for both the D+ and D- line (i.e., as measured by a Vector Network Analyzer with 50 ohm probes - see Figure 4-25). Note that the series capacitors C_{TX} is optional for the return loss measurement.
5. Measured between 20-80% at Transmitter package pins into a test load as shown in Figure 4-25 for both $V_{\text{TX-D+}}$ and $V_{\text{TX-D-}}$.

4.3.3.1. Transmitter Compliance Eye Diagram



OM13816

Figure 4-24: Minimum Transmitter Timing and Voltage Output Compliance Specification

There are two eye diagrams that must be met for the transmitter. Both eye diagrams must be aligned in time using the jitter median to locate the center of the eye diagram. The different eye diagrams will differ in voltage depending whether it is a transition bit or a de-emphasized bit. The exact reduced voltage level of the de-emphasized bit will always be relative to the transition bit.

The eye diagram must be valid for any 250 consecutive UIs.

An appropriate average TX UI must be used as the interval for measuring the eye diagram.

4.3.3.2. Compliance Test and Measurement Load

The AC timing and voltage parameters must be verified at the measurement point, as specified by the device vendor within 0.2 inches of the package pins, into a test/measurement load shown in Figure 4-25.

Note: The allowance of the measurement point to be within 0.2 inches of the package pins is meant to acknowledge that package/board routing may benefit from D+ and D- not being exactly matched in length at the package pin boundary. If the vendor does not explicitly

state where the measurement point is located, the measurement point is assumed to be the D+ and D-package pins.

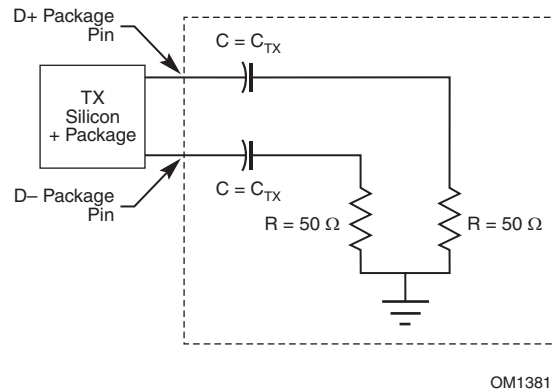


Figure 4-25: Compliance Test/Measurement Load

The test load is shown at the transmitter package reference plane, but the same Test/Measurement load is applicable to the receiver package reference plane.

C_{TX} is an optional portion of the measurement test load. The measurement should be taken on the opposite side of the capacitor from the package, and the value of the C_{TX} must be in the range of 75 nF to 200 nF.

4.3.4. Differential Receiver (RX) Input Specifications

The following table defines the specification of parameters for all differential Receivers (RXs). The parameters are specified at the component pins.

Table 4-6: Differential Receiver (RX) Input Specifications

Symbol	Parameter	Min	Nom	Max	Units	Comments
UI	Unit Interval	399.88	400	400.12	ps	The UI is 400 ps +/-300 ppm. UI does not account for SSC dictated variations. See Note 6.
$V_{RX-DIFFP-P}$	Differential Input Peak to Peak Voltage	0.175		1.200	V	$V_{RX-DIFFP-P} = 2 * V_{RX-D+} - V_{RX-D-} $ See Note 7.
T_{RX-EYE}	Minimum Receiver Eye Width	0.4			UI	The maximum interconnect media and transmitter jitter that can be tolerated by the receiver can be derived as $T_{RX-MAX-JITTER} = 1 - T_{RX-EYE} = 0.6$ UI See Notes 7 and 8.

Symbol	Parameter	Min	Nom	Max	Units	Comments
$T_{RX-EYE-MEDIAN-to-MAX-JITTER}$	Maximum time between the jitter median and maximum deviation from the median.			0.3	UI	Jitter is defined as the measurement variation of the crossing points ($V_{RX-DIFFp-p} = 0$ V) in relation to an appropriate average TX UI. See Notes 7 and 8.
$V_{RX-CM-ACp}$	AC Peak Common Mode Input Voltage			150	mV	$V_{RX-CM-AC} = V_{RX-D+} + V_{RX-D-} /2 - V_{RX-CM-DC}$ $V_{RX-CM-DC} = DC_{(avg)}$ of $ V_{RX-D+} + V_{RX-D-} /2$ during L0 See Note 7.
$RL_{RX-DIFF}$	Differential Return Loss	15			dB	Measured over 50 MHz to 1.25 GHz with the D+ and D- lines biased at +300 mV and -300 mV, respectively. See Note 9.
RL_{RX-CM}	Common Mode Return Loss	6			dB	Measured over 50 MHz to 1.25 GHz with the D+ and D- lines biased at 0 V. See Note 9.
$Z_{RX-DIFF-DC}$	DC Differential Input Impedance	80	100	120	Ω	RX DC Differential Mode impedance. See Note 10.
Z_{RX-DC}	DC Input Impedance	40	50	60	Ω	Required RX D+ as well as D- DC impedance (50 Ω +/- 20% tolerance). See Notes 7 and 10.
$Z_{RX-HIGH-IMP-DC}$	Powered Down DC Input Impedance	200 k			Ω	Required RX D+ as well as D- DC impedance when the receiver terminations do not have power. See Note 11.
$V_{RX-IDLE-DET-DIFFp-p}$	Electrical Idle Detect Threshold	65		175	mV	$V_{RX-IDLE-DET-DIFFp-p} = 2 * V_{RX-D+} - V_{RX-D-} $ Measured at the package pins of the Receiver
$T_{RX-IDLE-DET-DIFF-ENTERTIME}$	Unexpected Electrical Idle Enter Detect Threshold Integration Time			10	ms	An unexpected Electrical Idle ($V_{RX-DIFFp-p} < V_{RX-IDLE-DET-DIFFp-p}$) must be recognized no longer than $T_{RX-IDLE-DET-DIFF-ENTERTIME}$ to signal an unexpected idle condition.
$L_{RX-SKEW}$	Total Skew			20	ns	Skew across all Lanes on a Link. This includes variation in the length of a SKIP ordered set (e.g., COM and one to five SKP symbols) at the RX as well as any delay differences arising from the interconnect itself.

Notes:

6. No test load is necessarily associated with this value.

7. Specified at the measurement point and measured over any 250 consecutive UIs. The test load in Figure 4-25 should be used as the RX device when taking measurements (also refer to the Receiver Compliance Eye Diagram as shown in Figure 4-26). If the clocks to the RX and TX are not derived from the same clock chip the TX UI must be used as a reference for the eye diagram.
8. A $T_{RX-EYE} = 0.40$ UI provides for a total sum of 0.60 UI deterministic and random jitter budget for the transmitter and interconnect collected any 250 consecutive UIs. The $T_{RX-EYE-MEDIAN-to-MAX-JITTER}$ specification ensures a jitter distribution in which the median and the maximum deviation from the median is less than half of the total .6 UI jitter budget collected over any 250 consecutive TX UIs. It should be noted that the median is not the same as the mean. The jitter median describes the point in time where the number of jitter points on either side is approximately equal as opposed to the averaged time value. If the clocks to the RX and TX are not derived from the same clock chip, the appropriate average TX UI must be used as the reference for the eye diagram.
9. The receiver input impedance shall result in a differential return loss greater than or equal to 15 dB with the D+ line biased to 300 mV and the D- line biased to -300 mV and a common mode return loss greater than or equal to 6 dB (no bias required) over a frequency range of 50 MHz to 1.25 GHz. This input impedance requirement applies to all valid input levels. The reference impedance for return loss measurements for is 50 ohms to ground for both the D+ and D- line (i.e., as measured by a Vector Network Analyzer with 50 ohm probes - see Figure 4-25). Note: that the series capacitors C_{TX} is optional for the return loss measurement.
10. Impedance during all LTSSM states. When transitioning from a Fundamental Reset to Detect (the initial state of the LTSSM) there is a 5 ms transition time before receiver termination values must be met on all un-configured Lanes of a Port.
11. The Rx DC Common Mode Impedance that exists when no power is present or Fundamental Reset is asserted. This helps ensure that the Receiver Detect circuit will not falsely assume a receiver is powered on when it is not. This term must be measured at 300 mV above the RX ground.




Figure 4-26: Minimum Receiver Eye Timing and Voltage Compliance Specification

The RX eye diagram must be aligned in time using the jitter median to locate the center of the eye diagram.

The eye diagram must be valid for any 250 consecutive UIs.

An appropriate average TX UI must be used as the interval for measuring the eye diagram.

Appendix A – Chapter 4 w/Changebars 22Jul to 11 Dec (Released 11 Feb 2003 – Does Not Include C54 onward)

This appendix includes the revised text for Chapter 4 highlighting changes made up to the 11th of December. Final changes are shown in the version above. Note that, due to Word quirks and the editors inability to work around those quirks, all section numbers are shown as deleted. They are not in fact deleted, and these changes should be ignored.

4

4. Physical Layer Specification (*Released 11 Feb 2003 – Does Not Include C54 onward*)

4.1. Introduction

The Physical Layer isolates the Transaction and Data Link Layers from the signaling technology used for Link data interchange. The Physical Layer is divided into the Logical and Electrical functional sub-blocks (see <>).

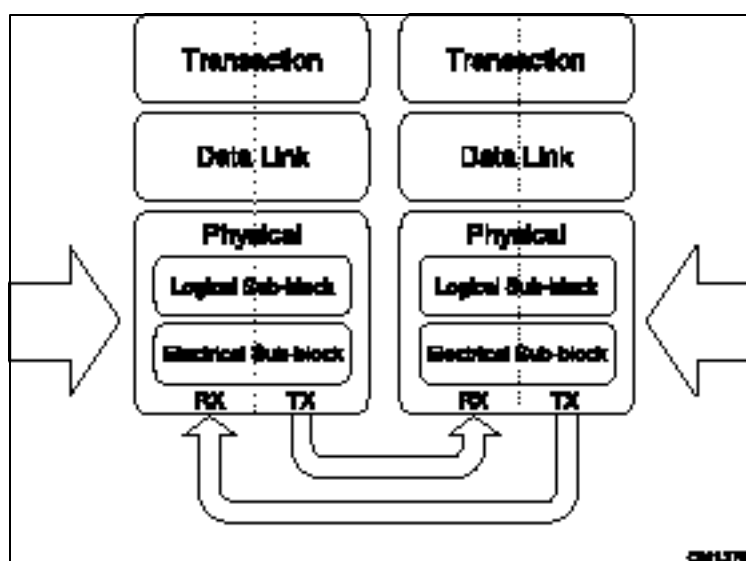


Figure 4-1: Layering Diagram Highlighting Physical Layer

4.2. Logical Sub-block

The Logical sub-block has two main sections: a Transmit section that prepares outgoing information passed from the Data Link Layer for transmission by the Electrical sub-block, and a Receiver section that identifies and prepares received information before passing it to the Data Link Layer.

The Logical sub-block and Electrical sub-block coordinate the state of each transceiver through a status and control register interface or functional equivalent. The Logical sub-block directs control and management functions of the Physical Layer.

Receivers may optionally check for violations of the rules associated with Receiver functions such as Symbol decoding and the like. If such checking is implemented, violations cause the indication of a Receiver Error to the Data Link Layer. A Receiver Error is a reported error associated with the Port (see Section 6.2).

4.2.1 Symbol Encoding

PCI Express uses an 8b/10b transmission code. The definition of this transmission code is identical to that specified in ANSI X3.230-1994, clause 11 (and also IEEE 802.3z, 36.2.4). Using this scheme, eight bit data characters are treated as three bits and five bits mapped onto a four-bit code group and a six bit code group, respectively. The control bit in conjunction with the data character is used to identify when to encode one of the 12 special symbols included in the 8b/10b transmission code. These code groups are concatenated to form a ten-bit Symbol. As shown in Figure 4-2, ABCDE maps to abcdei and FGH maps to fghj.

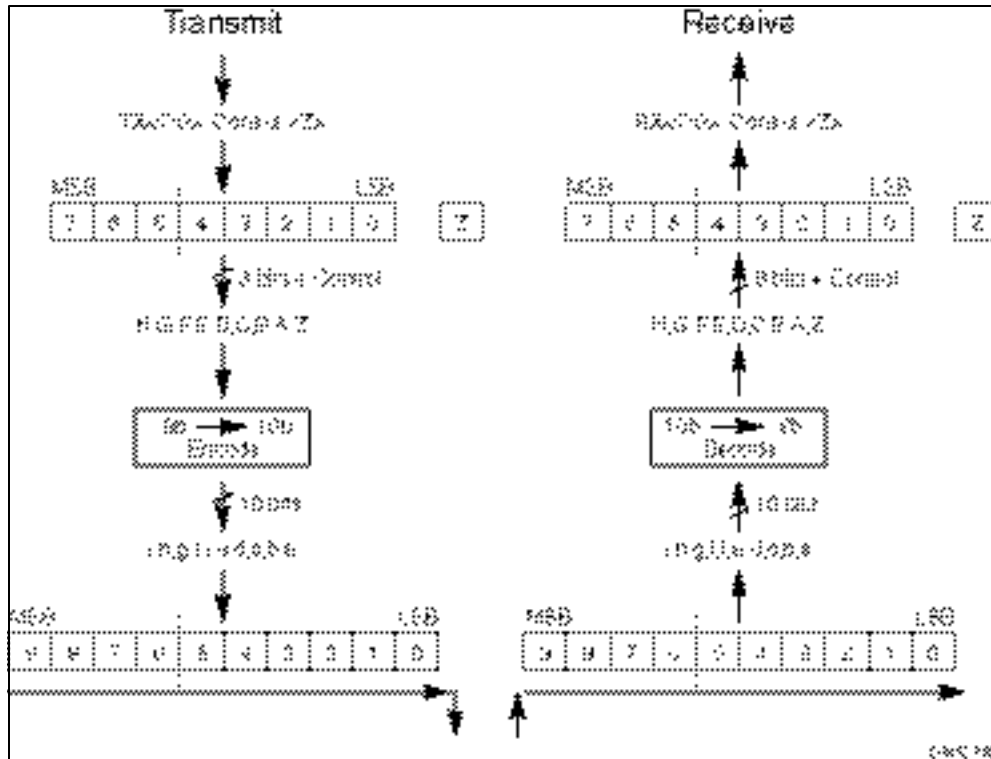


Figure 4-2: Character to Symbol Mapping

4.2.1.1 Serialization and De-serialization of Data

The bits of a Symbol are placed on a Lane starting with bit 'a' and ending with bit 'j'. Examples are shown in < and >.

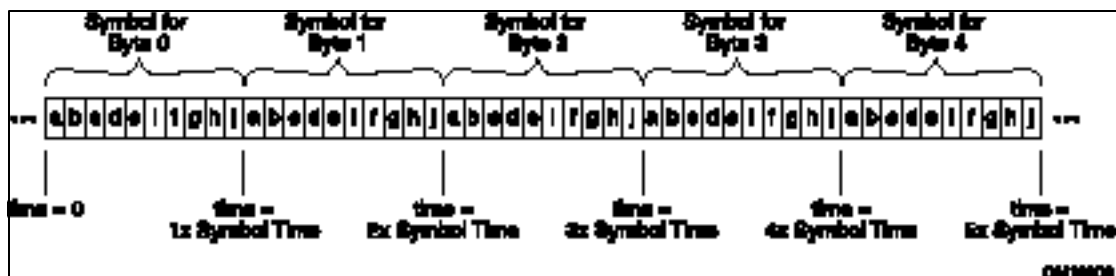


Figure 0-1: Bit Transmission Order on Physical Lanes - x1 Example




Figure 0-2: Bit Transmission Order on Physical Lanes - x4 Example

4.2.1.2. Special Symbols for Framing and Link Management (K Codes)

The 8b/10b encoding scheme used by PCI Express provides Special Symbols that are distinct from the Data Symbols used to represent Characters. These Special Symbols are used for various Link Management mechanisms described later in this chapter. Special Symbols are also used to frame DLLPs and TLPs, using distinct Special Symbols to allow these two types of Packets to be quickly and easily distinguished.

Table 4-1 shows the Special Symbols used for PCI Express and provides a brief description for the use of each. The use of these Symbols will be discussed in greater detail in following sections.

Table 4-1: Special Symbols

Encoding	Symbol	Name	Description
K28.5	COM	Comma	Used for Lane and Link initialization and management
K27.7	STP	Start TLP	Marks the start of a Transaction Layer Packet
K28.2	SDP	Start DLLP	Marks the start of a Data Link Layer Packet
K29.7	END	End	Marks the end of a Transaction Layer Packet or a Data Link Layer Packet
K30.7	EDB	EnD Bad	Marks the end of a nullified TLP

Encoding	Symbol	Name	Description
K23.7	PAD	Pad	Used in Framing and Link Width and Lane ordering negotiations
K28.0	SKP	Skip	Used for compensating for different bit rates for two communicating port Ports
K28.1	FTS	Fast Training Sequence	Used within an ordered-set to exit from L0s to L0
K28.3	IDL	Idle	Symbol Electrical Idle symbol used in the E lectrical Idle ordered-set
K28.4			Reserved
K28.6			Reserved
K28.7			Reserved

4.2.1.3.8b/10b Decode Rules

The symbol tables for the valid 8b/10b codes are given in Appendix B. These tables have one column for the positive disparity and one column for the negative disparity.

If a received symbol is found in the proper column corresponding to the current running disparity, then that symbol is valid.

If a received symbol is found in the column corresponding to the incorrect running disparity, ~~then a symbol has been received with a disparity error. In this case the current running disparity shall be set to the disparity value of the received symbol. The symbol is considered valid.~~ or if the symbol does not correspond to either column, then the Physical Layer must notify the Data Link Layer that the received symbol is ~~considered~~ invalid. This is a Receiver Error, and is a reported error associated with the Port (see Section <6.2>).

4.2.2. Framing and Application of Symbols to Lanes

There are two classes of framing and application of symbols to lanes. The first class is the ordered sets and the second is TLP and DLLP packet. Ordered sets are always transmitted serially on each lane, such that a full ordered set appears simultaneously on all lanes of a multi-lane link.

The Framing mechanism uses Special Symbol K28.2 “SDP” to start a DLLP and Special Symbol K27.7 “STP” to start a TLP. The Special Symbol K29.7 “END” is used to mark the end of either a TLP or a DLLP.

The conceptual stream of Symbols must be mapped from its internal representation, which is implementation dependent, onto the external Lanes. The Symbols are mapped onto the Lanes such that the first Symbol (representing Character 0) is placed onto Lane 0; the second is placed onto Lane 1, etc. The x1 Link represents a degenerate case, and the mapping is trivial, with all Symbols placed onto the single Lane in order.

When no packet information or special ordered-sets are being transmitted, the Transmitter is in the Logical Idle state. During this time Idle data must be transmitted. The Idle data must consist of the data byte 0 (00 Hexadecimal), scrambled according to the rules of Section 0 and 8b/10b encoded according to the rules of Section 0, in the same way that TLP and DLLP data characters are scrambled and encoded. Likewise, when the Receiver is not receiving any packet information or special ordered-sets, the Receiver is in Logical Idle and shall receive idle data as described above. During transmission of the idle data, the skip ordered-set must continue to be transmitted as specified in Section 0.

4.2.2.1. Framing and Application of Symbols to Lanes – Rules

In this section, “placed” is defined to mean a requirement on the transmitter to put the symbol into the proper Lane of a Link.

- 2-TLPs must be framed by placing an STP Symbol at the start of the TLP and an END Symbol or EDB Symbol at the end of the TLP (see Figure 0-3).
- 2-DLLPs must be framed by placing an SDP Symbol at the start of the DLLP and an END Symbol at the end of the DLLP.
- 2-Logical Idle is defined to be a period of one or more Symbol times when no information: TLPs, DLLPs or any type of Special Symbol is being Transmitted/Received. Unlike Electrical Idle, during Logical Idle the Idle character (00h) is being transmitted and received.
 - When the Transmitter is in Logical Idle, the Idle data character (00h) shall be transmitted on all Lanes. This is scrambled according to the rules in Section 0.
 - Receivers must ignore incoming Logical Idle data, and must not have any dependency other than scramble sequencing on any specific data patterns.
- 2-For Links wider than x1, the STP Symbol (representing the start of a TLP) must be placed in Lane 0 when starting Transmission of a TLP from a Logical Idle Link condition.
- 2-For Links wider than x1, the SDP Symbol (representing the start of a DLLP) must be placed in Lane 0 when starting Transmission of a DLLP from a Logical Idle Link condition.

- 2-The STP Symbol must not be placed on the Link more frequently than once per Symbol Time.
- 2-The SDP Symbol must not be placed on the Link more frequently than once per Symbol Time.
- 2-As long as the above rules are satisfied, TLP and DLLP Transmissions are permitted to follow each other successively.
- 2-One STP symbol and one SDP symbol may be placed on the Link in the same symbol time.
 - o Note: ~~For Links that are wider than x48 and wider Links, this means that can have~~ STP and SDP Symbols ~~can be~~ placed in Lane $4*N$, where N is a positive integer. For example, for x8, STP and SDP Symbols can be placed in Lanes 0 and 4; and for x16, STP and SDP Symbols can be placed in Lanes 0, 4, 8, or 12.
- 2-For xN Links where N is 8 or more, if an END or EDB Symbol is placed in a Lane K , where K does not equal $N-1$, and is not followed by a STP or SDP Symbol in Lane $K+1$ (i.e., there is no TLP or DLLP immediately following), then PAD Symbols must be placed in Lanes $K+1$ to Lane $N-1$.
 - Note: For example, on a x8 Link, if END or EDB is placed in Lane 3, PAD must be placed in Lanes 4 to 7, when not followed by STP or SDP.
- 2-The EDB symbol is used to mark the end of a nullified TLP. Refer to Section <3.5.2.1> for information on the usage of EDB.
- 2-~~Receivers may optionally check for violations of the rules of this section. Any such violation is a Receiver Error, and is a reported error associated with the Port (see Section <6.2>). If such checking is implemented, violations cause the indication of a Receiver Error to the Data Link Layer. A Receiver Error is a reported error associated with the Port (see Section 6.2).~~

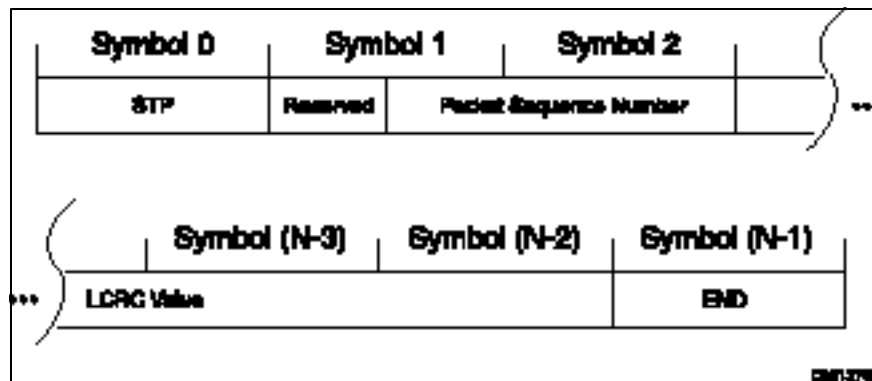


Figure 0-3: TLP with Framing Symbols Applied




Figure 0-4: DLLP with Framing Symbols Applied




Figure 0-5: Framed TLP on a x1 Link




Figure 0-6: Framed TLP on a x2 Link




Figure 0-7: Framed TLP on a x4 Link

4.2.3 Data Scrambling

The scrambling function can be implemented with one or many Linear Feedback Shift Register's (LFSR's) on a multi-Lane Link. When there is more than one transmit LFSR per Link, these must operate in concert, maintaining the same simultaneous (see Table 0-4, Lane-to-Lane Out Skew) value in each LFSR. When there is more than one receive LFSR per Link, these must operate in concert, maintaining the same simultaneous (see Table 0-5, Total Skew) value in each LFSR. Regardless of how it's implemented, the LFSRs must interact with data on a Lane-by-Lane basis as if there was a separate LFSR as described here for each Lane within that Link. On the transmit side, scrambling is applied to characters prior to the 8b/10b encoding. On the receive side de-scrambling is applied to characters after 8b/10b decoding.

The LFSR is graphically represented in <>. Scrambling or unscrambling is performed by serially XORing the 8-bit (D0-D7) character with the 16-bit (S0-S15) output of the LFSR. An output of the LFSR, S15, is XORed with D0 of the data to be processed. The LFSR and data register are then serially advanced and the output processing is repeated for D1 through D7. The LFSR is advanced after the data is XORed. The LFSR implements the polynomial:

$$G(X)=X^{16}+X^{15}+X^{13}+X^4+1$$

The mechanism(s) and/or interface(s) utilized by the data Link layer to notify the physical layer to disable scrambling is implementation specific and beyond the scope of this specification.

Data scrambling rules are the following:

2- The COM character initializes the LFSR.

2- The LFSR value is advanced eight serial shifts for each character except the SKP.

2- All data characters (D codes) except those within a Training Sequence Ordered-sets (TS1, TS2) and the Compliance Pattern (see section 4.2.8) are scrambled.

2- All special characters (K codes) are not scrambled.

The initialized value of an LFSR seed (S0-S15) is FFFFh. Immediately after a COM exits the transmit LFSR, the LFSR on the transmit side is initialized. Every time a COM enters the receive LFSR on any Lane of that Link, the LFSR on the receive side is initialized.

~~Scrambling is disabled by setting the scrambling disable bit in the TS1 ordered set and sending the TS1 ordered set until the other device responds with the scrambling disable bit set in the TS1 ordered set.~~

~~Scrambling must be disabled on all lanes of a link. Failure to disable scrambling on all lanes of a link results in undefined behavior.~~

Scrambling can only be disabled at the end of **Configuration** (see section 0) ~~in Polling, never in Configuration or Recovery.~~

Scrambling does not apply to a Loopback Slave.

Scrambling is always enabled in **Polling-Detect** by default.



IMPLEMENTATION NOTE

Disabling Scrambling

Disabling Scrambling is intended to help simplify test and debug equipment. Control of the exact data patterns is useful in a test and debug environment. Since scrambling is reset at the physical layer there is no reasonable way to reliably control the state of the data transitions through software. Thus, the disable scrambling bit is provided for these purposes for use in a test and/or debug environment. The minimal requirements to disable scrambling are; the successful detection that the lane(s) is terminated to a corresponding transmitter receiver pair and the physical layer was directed by the data link layer to set the Disable Scrambling bit (see Tables 4-2 and 4-3, Symbol 5, bit 3). As stated in Section 4.2.6.2, if the lane is terminated by nominal 50 single-ended (nominal 100 ohms differential) termination and TS1/TS2 ordered-sets are not returned, the transmitter will enter the Polling Compliance state and scrambling or disabling scrambling does not apply. Also note that when the physical layer enters polling and is directed to set the disable scrambling bit, if there are no state transition restrictions such as remain in the Polling state, the LTSSM will continue to attempt to progress through the Polling State to the Configuration state to the L0 state. In that case and barring design flaws or system faults, the lanes will become configured as part of a link. An alternate but similar scenario is that the data link layer directs the physical layer to enter the Polling state and disable scrambling (set bit 3 in symbol 5 in the TS1/TS2 ordered-sets) after the links have been configured (i.e. from L0). In those scenarios, all lanes within a link must disable scrambling. This choice is made to simplify test and debug equipment requirements as well as remain consistent with other sections of this specification by requiring all lanes within a link behave in a similar fashion.

The mechanism(s) and/or interface(s) utilized by the Data Link Layer to notify the physical layer to disable scrambling is component implementation specific and beyond the scope of this specification.

For more information on scrambling, see Appendix C.

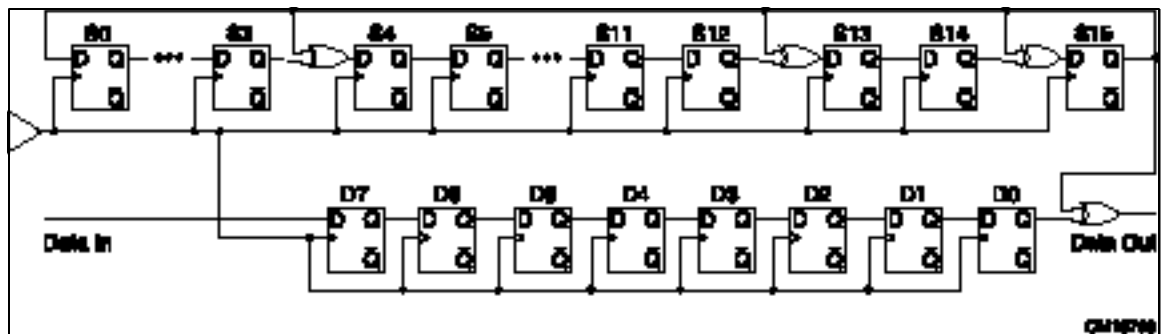


Figure 0-8: LFSR with Scrambling Polynomial

4.2.4 Link Initialization and Training

This section defines the Physical Layer control process that configures and initializes each Link for normal operation. This section covers the following functions:

- 2-Configuring and initializing the Link.

- 2-Supporting normal packet transfers.
- 2-Supported state transitions when recovering from Link errors.
- 2-Restarting a Port from low power states.

The following are discovered and determined during the training process:

- 2-Link width.
- 2-Link data rate³¹.
- 2-Lane reversal.
- 2-Polarity inversion.

Training does:

- 2-Link data rate³² negotiation.
- 2-Bit synchronization per Lane.
- 2-Lane polarity.
- 2-Symbol synchronization per Lane.
- 2-Lane ordering within a Link.
- 2-Link width negotiation.
- 2-Lane-to-Lane de-skew within a multi-Lane Link.

Receivers may optionally check for violations of the Link Initialization and Training Protocols. If such checking is implemented, any violation is a Training Error. A Training Error is a reported error associated with the Port (see Section **Error! Reference source not found.**). A Training Error is considered fatal to the Link.

4.2.4.1. Training Sequence Ordered-sets

Training sequences are composed of ordered-sets used for **initializing** bit alignment, symbol alignment and to exchange physical layer parameters. Training sequence ordered-sets are never scrambled but are always 8b/10b encoded.

SKP ordered-sets (see section <4.2.7>) **must** be transmitted during ~~training sequences~~ but never interrupt a TS1 or TS2 ordered-set.

~~Any reference in the state machine section indicating that 16 ordered-sets are to be transmitted after receiving “n” number of ordered-sets means to send at least 16 additional ordered-sets after the reception of at least “n” ordered-sets. This is in addition to the ordered-sets sent while waiting for “n” ordered-sets to be received.~~

~~In order for N_FTS to be valid two or more TSx ordered-sets must be received with the same value. The value of N_FTS must be the same for all lanes of a configured link or the behavior is undefined.~~

³¹ This specification only defines one data rate. Future revisions will define additional rates.

³² This specification defines the mechanism for negotiating the Link operational bit rate to the highest supported operational data rate.

~~Anytime two consecutive TS1 or TS2 ordered sets are received in the L0, Polling, Configuration, or Recovery states with the Reset Bit set the Link Control Reset state must be entered directly. The Reset Bit must be set in all TS1 or TS2 ordered sets on all lanes of a configured link or the behavior is undefined.~~

~~Anytime two consecutive TS1 or TS2 ordered sets are received in the L0, Polling, Configuration, or Recovery states with the Loopback Bit set, the Loopback state must be entered directly. The Loopback Bit must be set in all TS1 or TS2 ordered sets on all lanes of a configured link or the behavior on that configured link is undefined.~~

~~Anytime two consecutive TS1 or TS2 ordered sets are received in the L0, Polling, Configuration, or Recovery states with the Disable Bit set, the Disable state must be entered directly. The Disable Bit must be set in all TS1 or TS2 ordered sets on all lanes of a configured link or the behavior is undefined.~~

~~Scrambling is disabled during Polling by setting the scrambling disable bit in the TS1 ordered set and sending the TS1 ordered set until the other device responds with the scrambling disable bit set in the TS1 ordered set. A minimum of two TS1 ordered sets must be received with the scrambling disable bit set before responding with the scrambling disable bit set.~~

~~Skip ordered sets may be sent between consecutive TS1 or TS2 ordered sets. Idle data is not allowed between consecutive TS1 or TS2 ordered sets.~~

The Training control bits for ~~Hot Scrambling Disable~~, Reset, ~~Disable Link Disable~~, and Enable Loopback ~~Enable~~ are mutually exclusive, only one of these bits may be set at a time as well as transmitted on all Lanes in a configured (all Lanes in **L0**) or possible (all Lanes in **Configuration**) Link. If more than one of the ~~Hot Scrambling Disable~~, Reset, Link Disable, or Loopback Enable bits are set at the same time then the Link behavior is undefined.

Table 0-1: TS1 Ordered-Set

Symbol Number	Allowed Values	Encoded Values	Description
0		K28.5	COMMA code group for symbol alignment
1	0-255	D0.0 - D31.7, K23.7	Link Number within component
2	0-31	D0.0 - D31.0, K23.7	Lane Number within Port
3	0 – 255	D0.0 - D31.7	N_FTS. This is the number of fast training ordered-sets required by the receiver to obtain reliable bit and symbol lock.
4	2	D2.0	Data Rate Identifier Bit 0 – Reserved, set to 0 Bit 1 = 1, generation 1 (2.5 Gb/s) data rate supported Bit 2:7 – Reserved, set to 0
5	Bit 0 = 0, 1 Bit 1 = 0, 1 Bit 2 = 0, 1 Bit 3 = 0, 1 Bit 4:7 = 0	D0.0, D1.0, D2.0, D4.0, D8.0	Training Control Bit 0 - Hot Reset Bit 0 = 0, De-assert Reset Bit 0 = 1, Assert Reset Bit 1 - Disable Link Bit 1 = 0, De-assert Enable Link Bit 1 = 1, Assert Disable Link Bit 2 - Loopback Bit 2 = 0, De-assert No Loopback Bit 2 = 1, Assert Enable Loopback Bit 3 – Disable Scrambling Bit 3 = 0, De-Assert Enable Scrambling Bit 3 = 1, Assert Disable Scrambling Bit 4:7, Reserved, set to 0
6-15		D10.2	TS1 Identifier

Table 0-2: TS2 Ordered-Set

Symbol Number	Allowed Values	Encoded Values	Description
0		K28.5	COMMA code group for symbol alignment
1	0-255	D0.0 - D31.7, K23.7	Link Number within component
2	0-31	D0.0 - D31.0, K23.7	Lane Number within Port
3	0 – 255	D0.0 - D31.7	N_FTS. This is the number of fast training ordered-sets required by the receiver to obtain reliable bit and symbol lock.
4	2	D2.0	Data Rate Identifier Bit 0 – Reserved, set to 0 Bit 1 = 1, generation 1 (2.5 Gb/s) data rate supported Bit 2:7 – Reserved, set to 0
5	Bit 0 = 0, 1 Bit 1 = 0, 1 Bit 2 = 0, 1 Bit 3 = 0, 1 Bit 4:7 = 0	D0.0, D1.0, D2.0, D4.0, D8.0	Training Control Bit 0 - Hot Reset Bit 0 = 0, De-assert Bit 0 = 1, Assert Bit 1 - Disable Link Bit 1 = 0, De-assert Bit 1 = 1, Assert Bit 2 - Loopback Bit 2 = 0, De-assert Bit 2 = 1, Assert Bit 3 - Disable Scrambling Bit 3 = 0, De-Assert Bit 3 = 1, Assert Training Control Bit 0 = 0, De-assert Reset Bit 0 = 1, Assert Reset Bit 1 = 0, Enable Link Bit 1 = 1, Disable Link Bit 2 = 0, No Loopback Bit 2 = 1, Enable Loopback Bit 3 = 0, Enable Scrambling Bit 3 = 1, Disable Scrambling Bit 4:7, Reserved, set to 0
6-15		D5.2	TS2 Identifier

4.2.4.2. Lane Polarity Inversion

During the training sequence, the receiver looks at symbols 6-15 of TS1 and TS2 as the indicator of Lane polarity inversion (D+ and D- are swapped). If Lane polarity inversion occurs, the TS1 symbols 6-15 received will be D21.5 as opposed to the expected D10.2. Similarly, if Lane polarity inversion occurs, symbols 6-15 of the TS2 ordered-set will be D26.5 as opposed to the expected D5.2. This provides the clear indication of Lane polarity inversion.

If polarity inversion is detected the receiver must invert the received data. The transmitter must never invert the transmitted data. Support for Lane Polarity Inversion is required on all PCI Express Receivers across all Lanes **independently**.

4.2.4.3. Fast Training Sequence (FTS)

FTS is the mechanism that is used for bit and symbol synchronization when transitioning from **L0s** to **L0**. The FTS is used by the receiver to detect the exit from Electrical Idle and align the receiver's bit/symbol receive circuitry to the incoming data. See Section 0 for a description of **L0** and **L0s**.

A single FTS training sequence is an ordered-set composed of one K28.5 (COM) symbol and three K28.1 symbols. The maximum number of FTS ordered-sets (N_FTS) that a component can request is 255, providing a bit time synchronization of $4 * 255 * 10 * UI$. 4096 FTS ordered sets must be sent when the Extended Synch bit is set in order to provide external Link monitoring tools with enough time to achieve bit and framing synchronization. **SKP ordered sets must be scheduled and transmitted between FTS ordered sets as necessary to meet the definitions in section <4.2.7> with exception that no SKP ordered sets can be scheduled during the first 255 FTS ordered-sets. SKPs may not be sent during the FTS.** A single SKP ordered set is always sent after the last FTS is transmitted complete.

~~After initial power up, the N_FTS value is exchanged in the TS1/TS2 training sequence.~~ "N_FTS defines the number of FTS ordered-sets that must be transmitted when transitioning from L0s to L0. At a generation 1 ~~For the data rate in this specification,~~ the value that can be requested by a component corresponds to a character bit lock time of 16 ns (N_FTS set to 0 and one SKP ordered set) to ~4 μ s (N_FTS set to 255), except when the Extended Synch bit is set, which requires the transmission of 4096 FTS ordered sets resulting in a bit lock time of 64 μ s.

~~When transitioning from L0s to L0, the receiver shall observe the period of time from Electrical Idle Exit to the time that the receiver obtains bit and symbol alignment. If the N_FTS period of time expires prior to the receiver obtaining character alignment on all lane Lanes of the configured Link, the receiver must transition to the Recovery state in order to recover the Link alignment.~~ This sequence is detailed in the LTSSM in Section 0.

4.2.4.4. Link Error Recovery

~~At any time~~ After configuring a Link the Physical Layer can be directed to enter the **Recovery** state from **L0**, as described in Section <3.5>. Refer to Section <6.2> for more information on behavior when the physical layer reports errors.

4.2.4.5. Link Reset

There are two types of reset, one at the physical layer that is platform specific ("Power Good Reset" or cold/warm reset) and one that is passed in the Training Control Register (bit number 0 of symbol 5) in the TS1 and TS2 ordered-sets. This reset is called the ~~Training Control Reset~~ Hot Reset or hot reset.

4.2.4.5.1. Physical Layer Reset (“Power Good”)

~~A~~ Physical Layer Reset is provided by the system to the logical sub-block and is used to properly initialize the Port. ~~This~~ Physical Layer Reset must be asserted when the power to the device does not meet the device specifications. ~~This~~ Physical Layer Reset may also be asserted by other control agents in the device (for instance the Link Layer, the Transaction Layer or a software mechanism) to assert reset to the Physical Layer. The following ~~should~~~~must~~ be met ~~when possible if~~~~when this Physical Layer Reset~~~~reset~~ is asserted:

- ~~?~~ The receiver terminations are enabled (see $Z_{RX-COM-INITIAL-DC}$ in Table 4-5).
- ~~?~~ The transmitter ~~must~~ holds a constant DC common mode voltage ~~on the differential pair using a high impedance driver. For the definition of high impedance in this context, see Table 4-4.~~

4.2.4.5.2. Training Control Reset Hot Reset (Hot Reset)

~~Hot Reset~~ is ~~In addition to Physical Layer Reset, a protocol reset is defined in section 0. This Training Control Reset~~ uses a reset indicator bit defined in the Training Control Register (Table 4-2, Table 4-3) that is sent during the training sequence. ~~An upstream device sets this bit to force a reset of all of the downstream devices and links. Optionally a downstream device may use this reset to reset other logic within the device. The method and mechanisms to do this are implementation specific.~~

~~When a bridge receives a training sequence with the reset bit asserted, it must propagate that reset onto all downstream links by transmitting the TS1 ordered sets with the reset bit asserted. Training Control Reset shall not propagate upstream. All other physical layer information exchanged in those ordered sets must be accurate and correct.~~

~~All Lanes within a multi-Lane Link transmit the TS1, TS2 ordered sets during Training Control Reset. When Training Control Reset is removed, each transmitter and receiver must enter Detect.~~

~~Unless otherwise specified the terms “reset” and “power on/reset” in this chapter refer to the Physical Layer Reset.~~

4.2.4.6. Link Disable

~~A Link can be disabled if directed. When directed to this state the following behavior occurs:~~

- ~~?~~ The Port drives its transmitters to high impedance.
- ~~?~~ The receiver terminations must remain enabled.
- ~~?~~ There should be no response to any received data.

~~When directed to disable a Link, all Lanes within a multi-Lane Link transmit a minimum of 4 and a maximum of 16 TS1 ordered sets with the Disable Link bit set. The Link remains disabled until directed or a physical reset occurs.~~

~~After a physical reset or after being directed out of Link disable, the next state is Detect.~~

4.2.4.7. Link Data Rate Negotiation

All devices are required to ~~initialize~~~~start and configure~~ Link initialization ~~with~~ using a generation 1 data rate on each Lane. A field in the training sequence ordered set (see Section 0) is used to advertise~~indicate~~ all the supported~~maximum capable~~ data rates, and any higher speed supported by both sides of the Link will be initiated during ~~Polling~~~~Speed for the Lane~~.

~~This document (Generation 1) specifies a data rate of 2.5 Gb/s in each direction on each Lane.~~

Link Width and Lane Sequence Negotiation

PCI Express Links must consist of 1, 2, 4, 8, 12, 16, or 32 Lanes in parallel, referred to as x1, x2, x4, x8, x12, x16, and x32 links respectively. All Lanes within a Link shall simultaneously (as defined by $L_{TX-SKEW}$ in <insert x-ref to table 4.5>) transmit data based on the exact same frequency. The negotiation process is described as a sequence of steps.

The negotiation establishes values for Link number and Lane number for each Lane that is part of a valid link; each Lane that is not part of a valid Link exits the negotiation to become a separate Link or remain in Electrical Idle.

During Link width and Lane number negotiation, the two communicating Ports must accommodate the maximum allowed Lane-Lane skew as specified by $L_{RX-SKEW}$ in <insert x-ref to table 4.6>.

Optional Link negotiation behaviors include Lane reversal, variable width links, splitting of Ports into multiple links and the configuration of cross-link.

Annex specifications to this specification may impose other rules and restrictions that must be comprehended by components compliant to those annex specifications; it is the intent of this specification to comprehend interoperability for a broad range of component capabilities.

Required and Optional Port Behavior

The ability for a xN Port to form a xN Link as well as a x1 Link (where N can be 32, 16, 12, 8, 4, 2, and 1) is required

- Note: Designers must connect Ports between two different components in a way that allows those components to meet the above requirement. If the Ports between components are connected in ways that are not consistent with intended usage as defined by the component's Port descriptions/data sheets behavior is undefined.

The ability for a xN Port to form any Link width between N and 1 is optional.

- Note: An example of this behavior includes a x16 Port which can only configure into only one link, but the width of the Link can be configured to be x12, x8, x4, x2 as well as the required widths of x16 and x1.

The ability to split a Port into two or more links is optional.

- Note: An example of this behavior would be a x16 Port that may be able to configure 2 x8 links, 4 x4 links, or even 16 x1 links.

Support for Lane reversal is optional.

- Note: Lane reversal is accomplished by reassigning Lane numbers to transmitter and receiver pairs (i.e. a specific transmitter is always associated with the same receiver).
- Note: An example of Lane reversal consists of Lane 0 of an upstream Port attached to Lane N-1 of a downstream Port where either the downstream or upstream device may reverse the Lane order to configure a xN link.

Support for formation of a crosslink is optional. In this context a downstream Port connected to a downstream Port or an upstream Port connected to an upstream Port is a crosslink.

Current and future electromechanical and/or form factor specifications may require the implementation of some optional features listed above. Component designers must read the specifications that the component(s) they are designing will used in to ensure compliance to those specifications.

4.2.4.8. Link Width and Lane Sequence Negotiation

PCI Express Links must consist of 1, 2, 4, 8, 12, 16, or 32 Lanes in parallel, referred to as x1, x2, x4, x8, x12, x16, and x32 links respectively. All Lanes within a Link shall transmit data based on the exact same frequency.

The negotiation process is described as a sequence of steps. The negotiation establishes values for Link Number and Lane Number for each Lane that is part of a valid Link; each Lane that is not part of a valid Link exits the negotiation with values of K23.7 (PAD-out of range) for Link Number and Lane number.

During Link width and Lane sequence negotiation, the two communicating ports must accommodate the maximum allowed Lane-Lane skew as specified by L_{RX_SKEW} in Table 4-5.

Optional behaviors are described to comprehend fixed configuration components and components to be used in the implementation of advanced switching cross-links (Section 1.6). Annex specifications to this specification may impose other rules and restrictions that must be comprehended by components compliant to those annex specifications; it is the intent of this specification to comprehend interoperability for a broad range of component capabilities.

4.2.4.8.1. Required/Optional Port Behavior

A xN port must be capable of forming a xN Link as well as a x1 Link (where N can be 32, 16, 12, 8, 4, 2, and 1). All other widths between N and 1 are optional. Note: Current and future electromechanical and/or form factor specifications may state that negotiation of additional widths is required. For example, the *PCI Express Card Electromechanical Specification* may specify additional requirements for downstream ports.

The ability for a set of transceivers to become one port and form one Link or become multiple ports and form multiple links is optional.

Support for Lane reversal (see step 2, section 4.2.4.8.2) at any and all ports is optional.

Support for crosslink connections is optional.

4.2.4.8.2.Steps to Negotiate the Width and Lane Ordering of Links

While in the configuration state, for Lanes that have successfully completed the bit synchronization, polarity inversion and symbol synchronization training, components negotiate the Link width and sequence of the Lanes within each Link via the steps of:

Step 1:

The upstream component initializes Link numbers:

An upstream component (downstream port) assigns unique Link numbers to groups of Lanes capable of being unique links³³. Indivisible groups of Lanes (those that can only be configured as a Lane within one Link) must connect to at most one downstream component (upstream port). The initial Link numbers are presented on each Lane to the downstream component(s). Until indicated (step 3), Lane numbers are presented as K23.7 (PAD-out of range). Upstream ports present their Link numbers and Lane numbers as K23.7 (PAD-out of range).

Mechanism: The downstream port (upstream component) shall send out the TS1 ordered-sets with the assigned Link numbers inserted into the Link number field (symbol 1) on the groups of Lanes capable of being unique Links and the Lane number field (symbol 2) set to K23.7.

Example of a set of eight lanes on an upstream component capable of negotiating to become one x8 port when connected to one downstream component or two x4 ports when connected to two different downstream components: The upstream component (downstream port(s)) sends out TS1 ordered-sets with the Link number set to N on four lanes and Link number set to N+1 on the other four lanes. The Lane numbers are all set to K23.7. The resultant number of links that are formed as well as their width(s) is dependant upon the system configuration as well as the capabilities of the downstream components.

Note: From this point on the rules are written to describe how each individual Link is configured. Regardless of the number of links a component supports, each Link is negotiated with the same rules that follow. Lanes within unique (only capable of being configured into one Link) and aggregated (capable of being configured into more than one) links must comply with timing rules in (Section 4.2.4.9). Independent, unique links have independent timing and control of negotiation. Unique and aggregated links are mapped with one and only one PCI-to-PCI bridge structure (Section 1.4).

³³ The most flexible case being all Lanes could be separate x1 Links. The most restrictive case being all Lanes as part of one Link.

Step 2:

~~The upstream port (downstream component) responds with Link number assignments:~~

~~A downstream component (upstream port) assigns the Link number (label) by assigning a common Link number to each of its lanes connected to the upstream component (downstream port), where the assigned Link number is selected from one of the Link numbers the downstream component received from the upstream component. If the downstream component is restricted as to its placement of Lane number 0³⁴, it must select the Link number received on that Lane. If the downstream component is restricted to Link widths other than what is presented, it must only transition the Link number of the subset of lanes that it can support within the Link.~~

~~Mechanism: The upstream port (downstream component) shall send out the TS1 ordered set with the assigned common Link number inserted into the Link number field (symbol 1) and the Lane number field (symbol 2) set to K23.7 for all lanes within the widths supported by the port. Lanes which cannot be included due to supported width restrictions shall continue sending TS1 ordered sets with the Link number and Lane number fields both set to K23.7.~~

~~Example 1: a x8 port: The upstream port (downstream component) sends out TS1 ordered sets with the Link number set to one of the Link numbers presented from the upstream component and the Lane number set to K23.7 on all 8 lanes. Per the example under step 1 above, it must choose between N and N+1. If the upstream port (downstream component) did not support Lane reversal, it must choose the Lane number presented on its Lane 0.~~

~~Example 2: a x16 port which is not capable of becoming a x8 Link, but is capable of being a x4 Link: The upstream port (downstream component) sends out TS1 ordered sets with the Link number set to the Link number presented from the downstream port (upstream component) on the four lanes it can support within the Link; the Lane numbers remain set to K23.7 on those four lanes. It shall send out TS1 ordered sets with the Link numbers and Lane numbers set to K23.7 on the twelve remaining lanes.~~

~~Note for Step 2: There may be times when an upstream port (downstream component) may be connected to another upstream port (downstream component)(cross-link). The rule below defines the behavior in this situation. Support for this behavior is optional.~~

~~Upstream port (downstream component) connected to upstream port (downstream component):~~

~~If after a minimum of 2 ms, no TS1 ordered sets with a Data Symbol for its Link number have been received on any Lane within the perspective Link, the upstream port (downstream component) may optionally assume the role of a downstream port and transition this port to Step 1. If this feature is not supported, it must maintain the values of K23.7 for Link and Lane number fields and exit the negotiation. Note Figure 4-12 does not illustrate this optional behavior.~~

Step 3:

~~The downstream port (upstream component) initializes Lane numbers:~~

~~The downstream port (upstream component) must acknowledge the assigned Link number received from the upstream port (downstream component) by transitioning the Link number to the~~

³⁴ A simple example of this is when the port does not support Lane reversal.

assigned Link number on each Lane to the upstream port (downstream component) as well as transitioning the Lane number fields to its preferred Lane numbers, while maintaining Link widths consistent with the width restrictions above. The preferred Lane numbers must be consecutive and one Lane number must be assigned to 0. If the downstream port (upstream component) has not received a Data Symbol for its Link number after 2 ms, all lanes of the Link must maintain values of K23.7 for Link and Lane number fields and exit the negotiation. If the assigned Link number does not match any of its initial Link numbers, see note below.

~~Mechanism: The downstream port (upstream component) shall send out the TS1 ordered set with Link number field (symbol 1) set to the assigned Link number and the Lane number field (symbol 2) set to its preferred number on all lanes which received a Link number from the upstream port (downstream component) that can be accommodated within the Link widths that port can support. It must transition the Link numbers and Lane numbers to K23.7 on the lanes that were previously rejected by the upstream port (downstream component) and any additional lanes the downstream port (upstream component) cannot accommodate within its supported Link widths.~~

~~Note for Step 3: There may be times when a downstream port may be connected to a downstream port. The rules below define the behavior in this situation. Support for this behavior is optional.~~

~~Downstream port (upstream component) connected to downstream port (upstream component):~~

~~If the downstream port (upstream component) receives an assigned Link number that does not match any of its initial Link numbers, it may optionally compare the received Link number to its initial Link number. If the received Link number is less, it must transition these lanes to Step 2, assuming the role of an upstream port of a downstream component. A downstream port (upstream component) must remain in Step 3 if its assigned Link number was greater than the received Link number or it does not support this feature. If after a minimum of 16 TS1 ordered sets have been received on each Lane within the perspective Link, the downstream port (upstream component) has not received a Link number that matches any of its initial Link numbers, it must maintain the values of K23.7 for Link and Lane number fields and exit the negotiation.~~

Step 4:

The upstream port (downstream component) responds with Lane number assignments:

The upstream port (downstream component) must accommodate Link width and Lane numbers presented by the downstream port (upstream component) if it is possible to do so (see system designer rules below in this section). If the upstream port (downstream component) can accept the Link width presented but not the Lane numbers, it must acknowledge with its preferred ordering of Lane numbers at this time. The preferred Lane numbers must be consecutive and one Lane number must be assigned to 0. If the upstream port (downstream component) is restricted to Link widths other than what is presented, it must only transition the Lane numbers on the subset of lanes that can be accommodated within the Link widths that port can support. It must transition the Link numbers and Lane numbers to K23.7 on the lanes that cannot be accommodated within the widths that port supports.

Mechanism: If the upstream port (downstream component) has not received Data Symbols for its Lane numbers after receiving an additional 16 or more TS1 ordered sets, all respective lanes of the Link must maintain values of K23.7 for Link and Lane number fields and exit the negotiation. If the upstream port (downstream component) can accommodate the Lane numbers received from the downstream port (upstream component), it shall send out the TS1 ordered sets with Link number field (symbol 1) set to the assigned Link number and the Lane number field (symbol 2) set to the Lane numbers assigned by the downstream port (upstream component). Otherwise, it shall insert its preferred numbers on all lanes with Lane numbers currently assigned by the downstream

port (upstream component) that it can accommodate within the Link widths that port can support. It must transition the Link numbers and Lane numbers to K23.7 on the lanes that were previously rejected by the downstream port (upstream component) and any additional lanes the upstream port (downstream component) cannot accommodate in the Link width.

Step 5:

~~The downstream port (upstream component) confirms Link number and Lane assignments:~~

~~The downstream port (upstream component) must accommodate any Lane numbers which are consistent with all system and component rules that do not match its preferred ordering, completing the Link width and lane ordering negotiation (see system designer rules below in this section). If the upstream port (downstream component) has assigned Lane numbers to a number of lanes resulting in Link width that port cannot support, the downstream port (upstream component) must accommodate upstream port's (downstream component's) Lane 0, establishing a Link of width 1.~~

~~Mechanism: If the downstream port (upstream component) has not received Data Symbols for its Lane numbers after receiving an additional 16 or more TS1 ordered sets, all lanes of the Link must maintain values of K23.7 for Link and Lane number fields and exit the negotiation. If the downstream port (upstream component) is to further reduce the width requested by the upstream port (downstream component) to a width greater than x1, the downstream port (upstream component) must return to step 3³⁵. Otherwise, the downstream port (upstream component) shall transition to sending the TS2 ordered sets with the Link number fields (symbol 1) and Lane number fields (symbol 2) set to negotiated values. It must transition the Link numbers and Lane numbers to K23.7 on the lanes that were previously rejected by the upstream port (downstream component) and any additional lanes the downstream port (upstream component) cannot accommodate in the Link width.~~

Step 6:

~~The upstream port (downstream component) confirms Link number and Lane assignments:~~

~~Mechanism: After receiving at least one TS2 ordered set, the upstream port (downstream component) shall transition to sending the TS2 ordered sets with the Link number fields (symbol 1) and Lane number fields (symbol 2) set to negotiated values. It must transition the Link numbers and Lane numbers to K23.7 on the lanes that were previously rejected by the downstream port (upstream component).~~

Step 7:

~~The downstream and upstream ports (upstream and downstream components, respectively) settle on Link number and Lane assignments:~~

³⁵ It is only possible to return to step 3 one time due to the limited number of Link widths allowed (x1, x2, x4, x8, x12, x16, x32). The longest sequence of width negotiation consists of an upstream component (downstream port), which supports x16, x8, x2, x1 connected to a downstream component (upstream port), which supports x32, x12, x4, x1. Step 1 would attempt to create a x16 Link, step 2 a x12 Link, step 3 (first pass) a x8 Link, step 4 (first pass) a x4 Link, step 3 (second pass) a x2 Link, step 4 (second pass) a x1 Link.

Both ports continue sending TS2 ordered sets. If after completing the negotiation (steps 1—6), either port again transitions the Lane numbers (should only occur if this Link is an advanced switching cross-link (refer to Section 1.6) where both ports have been implemented as downstream ports of upstream components), the port may optionally silently accept the received Lane numbers as the correct labeling of the other port's transmitters and therefore its own receivers; otherwise, all lanes of the Link must maintain values of K23.7 for Link and Lane number fields and exit the negotiation. No further changes to Link number and Lane number are allowed at this point without a complete re-training and re-configuration of the ports and associated Link(s). Label numbers are to be retained, skipping the Link width and Lane sequence negotiation steps unless transitioned to Link state Polling.Quiet. When these steps are skipped, the previously negotiated Link and Lane numbers are retained and inserted into the appropriate fields of the TS1 and TS2 ordered sets.

Mechanism: At least 16 TS2 ordered sets are sent after receiving one TS2 ordered set. If the received TS2 ordered sets have Lane numbers that do not match those transmitted in the transmitted TS2 ordered set, the port may internally disassociate the transmitter and receiver from the same Lane number label, associating the receiver with the Lane number received. The transmitter association to Lane number must not change once TS2 ordered sets have been sent. If the receiver cannot be associated with the Lane number received, all lanes of the Link must maintain values of K23.7 for Link and Lane number fields and exit the negotiation. The port returns to Config.Idle after at least 8 TS2 ordered sets which contain the previously negotiated link and lane numbers are received consecutively.

All lanes that are connected to the other port but not included in the negotiated Link must maintain values of K23.7 for Link and Lane number fields assigned by upstream and downstream ports.

Any lanes that fail to establish Data Symbol values for Link and Lane number fields are inactive, and will not exchange information with the Data Link Layer. All data and control to the Data Link Layer from active lanes shall be consistent with the agreed Lane numbering. When negotiating Link width and Lane sequence; each downstream port (upstream component) must transition between steps in unison across all of its lanes and each upstream port (downstream component) must transition between steps in unison across all of its lanes.

The following graphical flow diagrams demonstrate interoperability of components with simplified negotiation machines. Components/ports with complex negotiation machines are added to facilitate clarity.

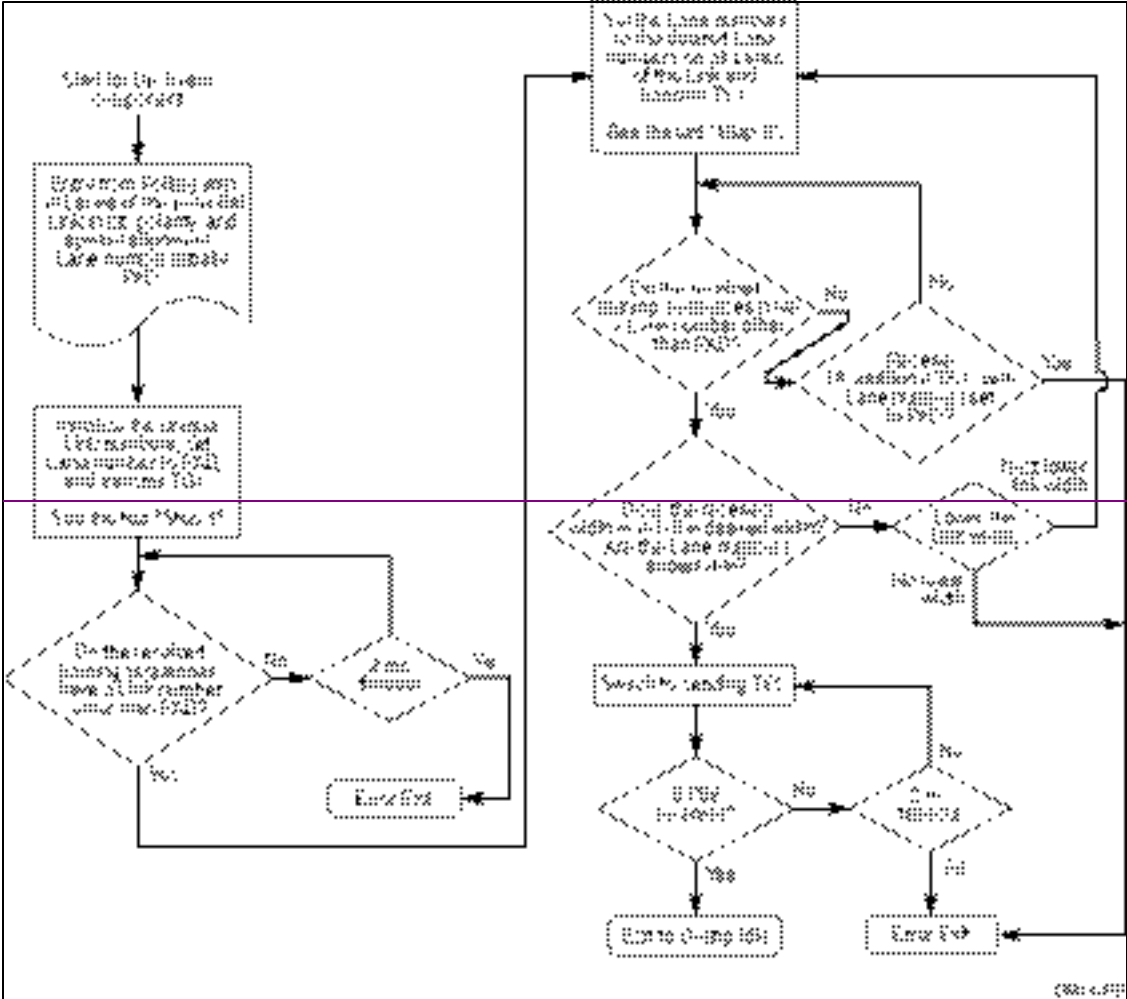


Figure 4-11: Width Negotiation, Simplified State Machine, Upstream Component

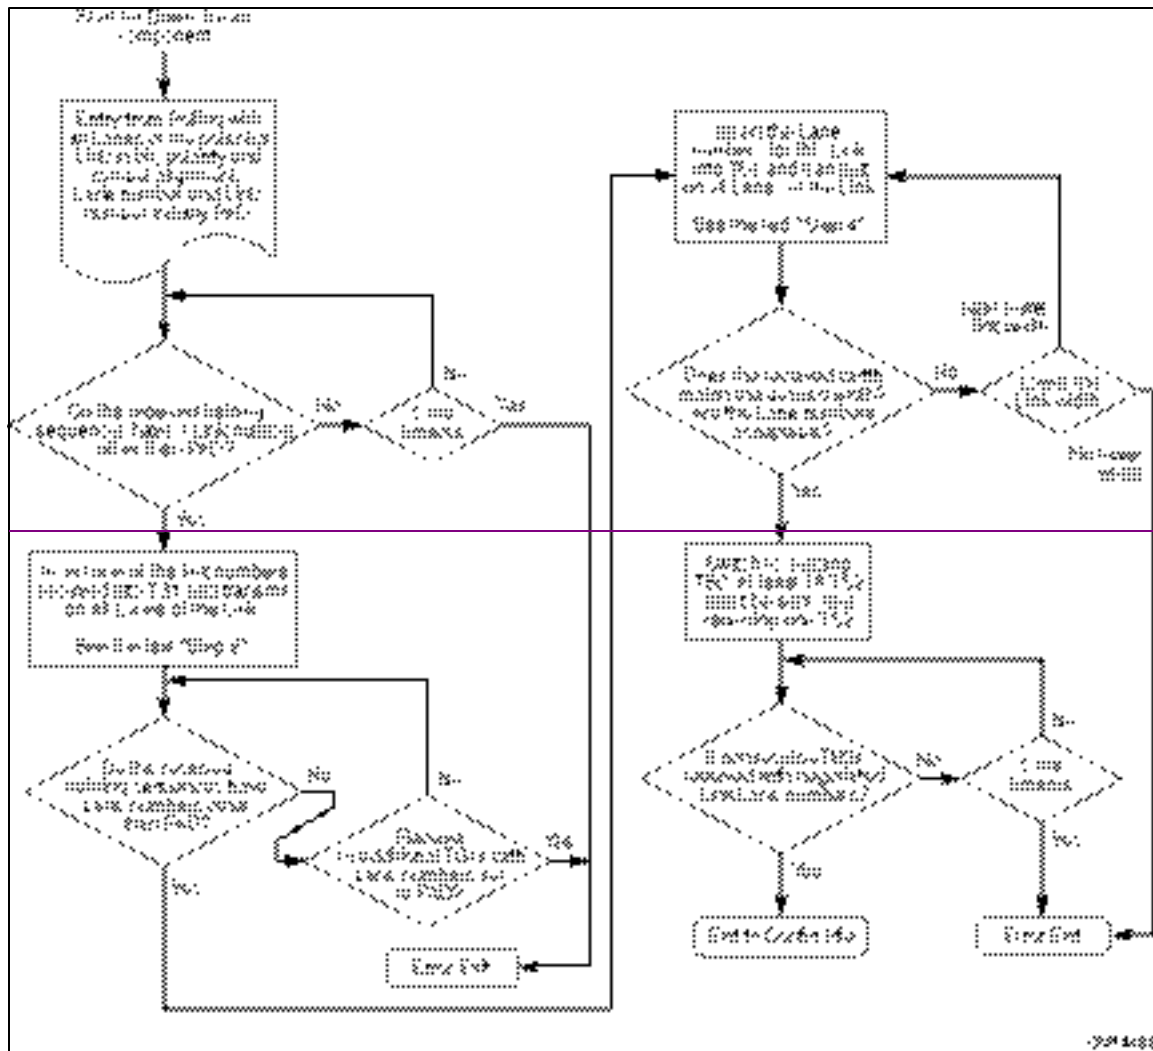


Figure 4-12: Width Negotiation, Simplified State Machine, Downstream Component

System Designer Rules:

System designers must connect Lanes within a Link interconnected through a connector or other suitable reference such that components are capable of labeling Lanes consistent with consecutive Lanes of the reference, inclusive of reference Lane 0. Components with fixed Lane ordering will interoperate with other compliant components flexible enough to also support other labelings. A simple example of increased flexibility would be to accommodate Lanes connected in reverse order. It is straightforward for a component to reverse Lane order upon receiving a Lane number of 0 on its most significant Lane or failure to detect an active in-bound Lane on its own Lane 0.

Example of Simple, Compliant Components:

Illustrated in Figure 4-13 are the transitions in the training exchange of Link and Lane numbers between an upstream component's fifth downstream Port that supports Link widths of x32, x12, x4, or x1 only and a downstream component's upstream Port that supports Link widths of x16, x8, x2, or x1 only. This is a scenario, with negotiation occurring at each step and the only common width is x1; at each step, the next largest Link width supported is implicit in the next exchange in the sequence. Link number is first (above), and Lane number second (below), with K representing the K23.7 symbol; transitions trigger the next step in

the negotiation. The upstream component's (downstream port's) Lane transitions are shown in white and downstream component's (upstream port's) Lane transitions are shown in gray. Only the 16 Lanes that have successfully completed the bit synchronization, polarity reversal and symbol synchronization training are shown.

t	Port Lanes															
	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	K	K	K	K	K	K	K	K	K	K	K	K	K	K	K	K
1					5	5	5	5	5	5	5	5	5	5	5	5
2	K	K	K	K	K	K	K	K	K	K	K	K	K	K	K	K
3					K	K	K	K	K	K	K	K	5	5	5	5
4	K	K	K	K	K	K	K	K	K	K	K	K	K	K	5	5
5					K	K	K	K	K	K	K	K	K	K	K	5

Figure 4-13: Width Negotiation Example

4.2.4.8.3. Port to Port Width Negotiation Example

The following text and diagrams show additional examples of two ports negotiating the width and Lane ordering of a Link.

Configuring groups of Lanes to form single logical links is done via a negotiation process that modifies the values of TS1 and TS2 symbols representing Link number and Lane number for each Lane. The process can be viewed as presentation of proposals and counter-proposals in the form of transitioning symbol values in discrete steps. Ordered-sets are symbol-synchronized across Lanes that potentially make up a single logical Link, allowing the transitions of symbol values across Lanes to be examined together. The ordered-sets containing each proposal are repeated until a new counter-proposal is detected via the receipt of appropriate symbol transitions. The association of a Lane to a particular logical Link is indicated by its final Link number symbol and its position (ordering) within the Link is indicated by its final Lane number symbol. Lanes that have not been included in any logical Link will have final symbol values identical to the pre-negotiation value of PAD.

Steps 1 and 2 establish the number of links an upstream component (downstream port(s)) is attached to. Steps 1 and 2 also begin (but not necessarily) complete establishing the width of the resultant Link(s).

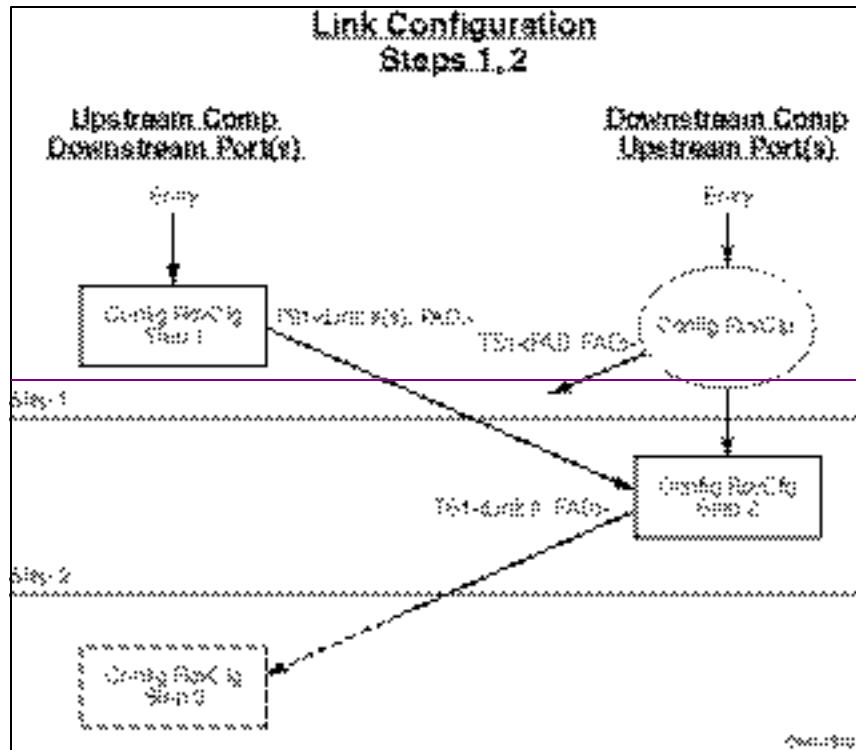


Figure 4-14: Link Width Negotiation; Steps 1 and 2

In order to enter the configuration state, Lanes within a particular Link have already exchanged TS1 ordered-sets and completed the bit synchronization, polarity inversion (if needed) and symbol synchronization functions. Prior to entering the configuration state, the Link number and Lane number fields have been set to PAD (K23.7) and TS1 ordered-sets are sent repeatedly.

Step 1:

Upon entering Config.RcvrCfg, the downstream port(s) starts the Link width and Lane ordering negotiations by sending out the TS1 ordered-set with a unique Link number on sets of Lanes, which that component could support as unique links; the Lane numbers continue to be set to PAD.

Step 2:

Upon receipt of the TS1 ordered-set with Link numbers (non-PADs) present in the Link number field, the upstream port shall respond by choosing one of the Link numbers it received. This step of returning the one Link number determines (for the downstream port(s)) the number of Links that are to be negotiated.

The upstream port responds with a Link number only on the Lanes in which it received a Link number and Lanes that it can support in one Link. A simple example: a port may be designed to support a x32 Link. Only 16 of those Lanes may have been attached, and therefore TS1s received only on 16 Lanes. The port may not support a x16 Link, but may support a x12 Link. In that case, the upstream port returns TS1 ordered-sets with a Link number only on the 12 Lanes that it is capable of supporting in a x12, and with the Link number set to PAD on the 4 remaining Lanes. This is the first counter proposal towards establishing the final Link width.

Additional notes on steps 1 and 2:

One method to create a cross-link Section 1.6 is to connect a downstream port to another downstream port. One of two scenarios can occur; a.) The two ports choose different Link numbers to begin negotiations, or b.) The two ports choose the same Link number to begin negotiation. If a.) occurs, the rules are described as part of step 3. If b.) occurs, the two ports will not yet be able to differentiate the Step 1 behavior of a cross-link condition from the Step 2 behavior of a normal upstream port.

Note: If a system designer connects two (or more) downstream ports on one upstream component that is capable of being aggregated into one Link (Link aggregation) in a cross-link to two downstream ports on a different upstream component, the configuration results are undefined.

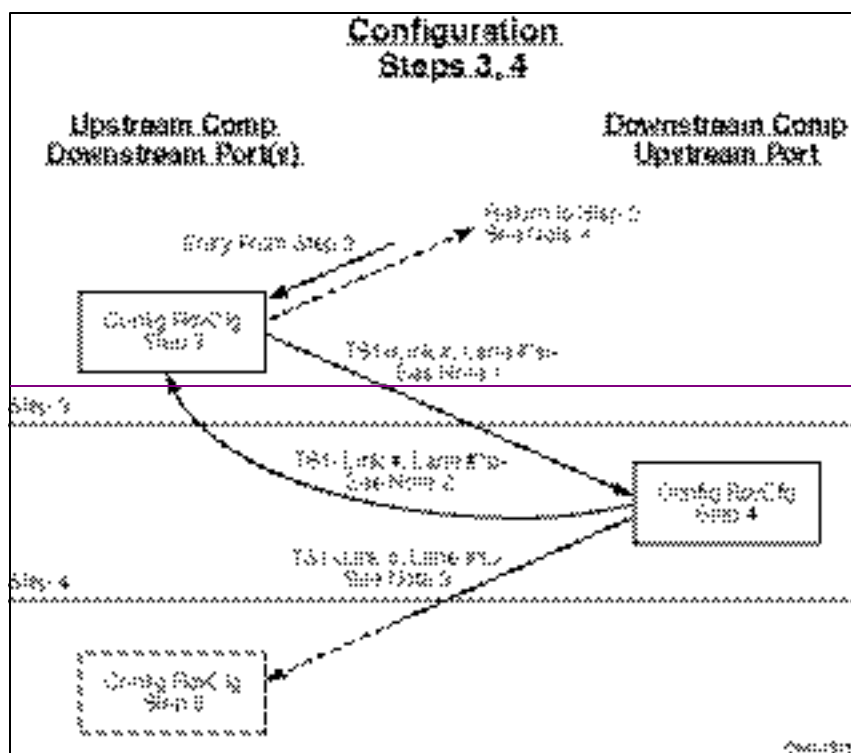


Figure 4-15: Link Width Negotiation; Steps 3 and 4

Steps 3 and 4 establish Lane ordering within each Link established in Steps 1 and 2. To find a supported Link width common to both components, Steps 3 and 4 continue to reduce the Link width by removing select Lanes from the negotiation; Lanes never join a Link negotiation through these steps. Returning a Lane's Link number to the value of PAD indicates removal; otherwise Link numbers persist with the value assigned in Step 2.

Step 3:

~~Upon receipt of TS1 ordered sets with a Link number inserted in that field on each Lane, the downstream port transitions to Step 3, making its first proposal for Lane numbers within each group of Lanes with common received Link numbers.~~

~~Note 1 (Figure 4-15): In the event that a set of ports were connected to a single upstream port, those ports would all see the same Link number returned. This is the mechanism that allows those ports to be aggregated into one Link. If those ports are not capable of being aggregated into one Link, the upstream component must continue negotiation with only one of those ports and transition the Link number to PAD on Lanes of the remaining ports, removing them from the negotiation process.~~

~~Note 2 (Figure 4-15): Components in a cross-link as described in scenario a.) above, start negotiations by presenting different Link numbers to each other. Components designed to comprehend the cross-link condition implement the optional compare of the two Link numbers. The component that receives a Link number smaller than the Link number it presented on its port assumes the role of an upstream port and transitions to Step 2. The other component receives a Link number greater than the Link number presented on its port remains in Step 3 to await a further Link number transition matching its own.~~

Step 4:

~~Upon receipt of the TS1 ordered set with Lane numbers presented in the Lane number fields and a common Link number present in the Link number fields, the upstream port transitions to Step 4, asserting an appropriate set of Lane numbers. The upstream port should only counter-propose Lane numbers if it has a fixed ordering of its Lanes and the downstream port is connected in a reversed Lane fashion; otherwise, Step 4 acknowledges the downstream port proposal.~~

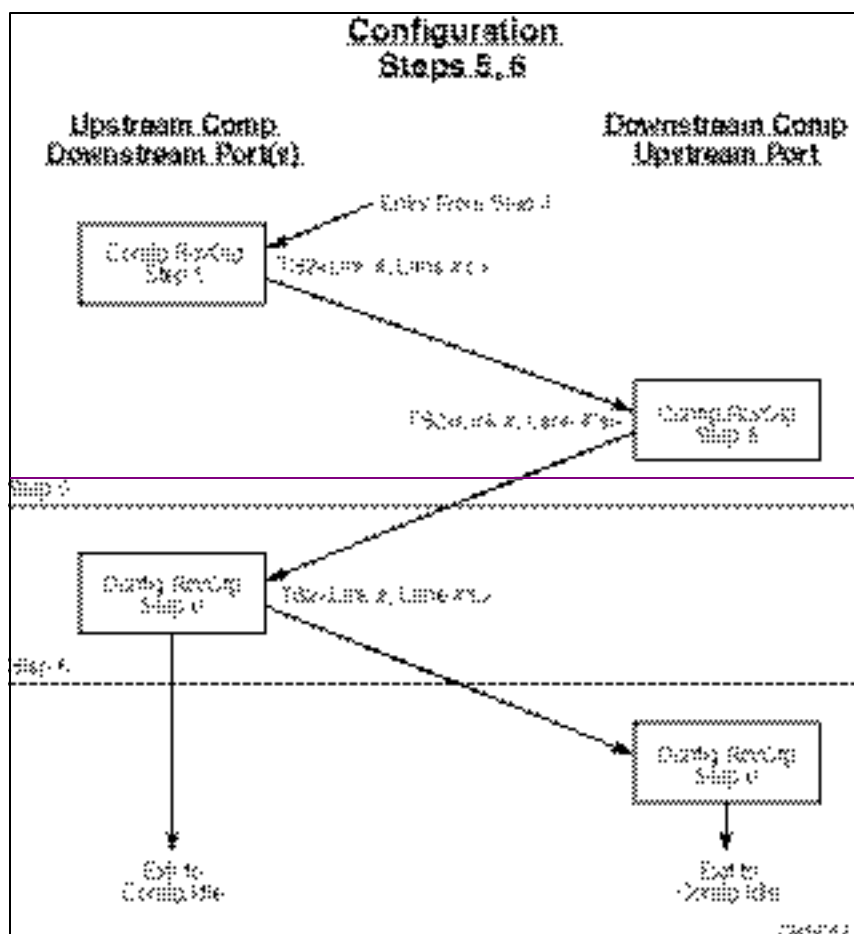


Figure 4-16: Link Width Negotiation; Steps 5 and 6

Steps 5 and 6 acknowledge the completed Link width and Lane sequence negotiation.

Step 5:

The downstream port transitions to sending TS2 ordered sets with the Link number and the Lane numbers inserted in the defined fields. At this time, the downstream port also removes Lanes that have been removed from the negotiation by the upstream port.

Note 3 (Figure 4-16): If additional Lanes are removed from the negotiation as may occur in the extreme mismatch of supported Link widths described in Step 4 of the rules, the downstream port can only transition to TS2 if a x1 Link is to be formed. Fewer additional Lanes removed resulting in a x2 Link require the downstream port to return to Step 3, remaining in TS1.

Step 6:

The upstream port transitions to sending TS2 ordered sets when it receives TS2 ordered sets with the agreed upon Link number and the Lane ordering it last presented to the downstream port. At least Lane 0 is retained in this step; the upstream port removes Lanes that have been removed from the negotiation by the downstream port.

Step 7: (not shown in Figure 4-16)

As noted earlier, there is a case where two downstream ports have been negotiating with each other and not realizing it until this step. This can occur in a cross-link when the two downstream ports both choose the same Link number to begin negotiation; both ports present their Link number in their Step 1, receiving that same Link number in Step 3. In Step 3, they both presented their preferred Lane numbers; if connected such that these align, Step 5 will not modify Lane numbers and simply transition to sending TS2 ordered-sets. However, if one component is connected in reverse Lane fashion and both ports support Lane reversal, Step 5 will cause both to accommodate the other. The mechanism to observe this is when a TS2 ordered-set arrives with Lane numbers that do not match the Lane numbers being transmitted. The crosslink optional behavior to resolve this conflict is to continue sending TS2 ordered-set with the agreed-upon Link number and conflicting Lane numbers. However, the Lane numbers now represent the transmitter Lane number. The port must then disassociate its transmitter with a receiver and reverse the ordering of the receivers to match the Lane ordering of the other port.

4.2.4.9. Lane-to-Lane De-skew

The receiver must compensate for the allowable skew between Lanes within a multi-Lane Link before delivering the data and control to the Data Link Layer.

Lane-to-Lane de-skew shall be done across all Lanes within multi-Lane links. An unambiguous de-skew mechanism is the COM symbol transmitted during training sequence or skip ordered-sets across all Lanes within a configured the Link. ~~(at what the transmitter believes is) simultaneously.~~ Other de-skew mechanisms may also be employed. ~~The receiver must compensate for the allowable skew between Lanes within a multi-Lane Link before delivering the data and control to the Data Link Layer.~~

4.2.4.10. Lane vs. Link Training

The ~~initialization~~ Link initialization ~~training~~ process builds unassociated Lanes of a Port ~~device~~ into associated Lanes that form a Link. For Lanes to configure properly into a desired link, the TS1 and TS2 ordered sets must have the appropriate fields (symbol 3, 4, and 5) set to the same value on all Lanes.

Links are formed at the conclusion of **Configuration**.

- Note: If the optional behavior of a Port being able to configure multiple links is employed the following observations can be made.
 - A separate LTSSM is needed for the maximum number of links that are desired to be configured by any given Port.
 - The LTSSM Rules are written for configuring one Link. The decision to configure Links in a serial fashion or parallel is implementation specific.

~~This occurs during the first state of the configuration state machine where the links are configured (e.g. width negotiation and optional Lane reversal). State machines prior to the operate on a per Lane basis, after the operations are on a Link basis.~~

~~For example, transmitted data prior to sends the specified data on all Lanes of the device; after the state the transmitter sends the specified data on all Lanes of the configured Link.~~

4.2.5. Link Training and Status State Machine (LTSSM) Descriptions

The LTSSM states are illustrated in Figure 0-9. These states are described in following sections.

All timeout values specified in the Link training and status state machine (LTSSM) timeout values are minus 0 seconds and plus 50% unless explicitly stated otherwise. All timeout values must be set to the specified values after power-on/reset. All counter values must be set to the specified values after power-on/reset.

~~The LTSSM states are illustrated in Figure 4-17. These states are described in following sections.~~

4.2.5.1. Detect

The purpose of this state is to detect when a far end termination is present. This state can be entered at anytime if directed. ~~receiver is powered on in order to avoid transferring common mode between the transmitter and receiver.~~

4.2.5.2. Polling

The Port transmits training ordered-sets and responds to the received training ordered-sets. In this state bit lock and ~~character~~symbol lock are established, Lane polarity is configured, and Lane data rate is established.

The polling state includes ~~a substate called~~ **Polling.Compliance** (see Section 0). This state is intended for use with test equipment used to assess if the transmitter and the interconnect present in the device under test setup is compliant with many voltage and timing specifications in Table 0-4 and Table 0-5.



IMPLEMENTATION NOTE

Use of the Polling.Compliance Substate

Polling.Compliance is intended for a compliance test environment and not entered during normal operation and ~~can not be disabled for any reason.~~ ~~The Polling.Compliance substate~~ is entered based on the physical system environment as described in Section 0. Any other mechanism(s) that causes a transmitter to put the compliance pattern are implementation specific and are beyond the scope of this specification.

4.2.5.3. Configuration

In **Configuration** both the transmitter and receiver are sending and receiving data at the negotiated data rate. The Lanes of a Port ~~configure into a Link through a width and lane negotiation sequence configures width, and Lane reversal and manages Lane-to-Lane skew is addressed,~~ Scrambling can be disabled, the N_FTS is set, and the Disable or Loopback states can be entered ~~within the Link.~~

4.2.5.4. Recovery

In **Recovery** both the transmitter and receiver are sending and receiving data ~~using the configured Link and Lane number as well as at~~ the previously negotiated data rate. This state is mainly used to re-establish bit and byte lock, but **Recovery** is also used to set a new N_FTS, enter the following states; Loopback, Disable, Hot Reset, and Configuration.

~~The Port transmits training ordered-sets and responds to the received training ordered-sets. In this state bit lock and symbol lock are re-established.~~

~~4.2.5.5.L0~~

~~L0~~ is the normal operational state where data and control packets can be transmitted and received. ~~All power management states are entered from this state.~~

~~4.2.5.6.L0s~~

~~L0s~~ is intended as a power savings state.

~~L0s~~ allows a Link to quickly enter and recover from a power conservation state without going through ~~the Configuration or Recovery states.~~

The entry to ~~L0s~~ occurs after receiving an Electrical Idle ordered-set.

A transmitter and receiver Lane pair on a Port are not required to both be in ~~L0s~~ simultaneously.

~~4.2.5.7.L1~~

~~L1~~ is intended as a power savings state.

The ~~L1~~ state allows an additional power savings over ~~L0s~~ at the cost of additional resume latency.

~~The receiver must be able to recover from this state within 64 μ s, including reacquiring bit and symbol synchronization.~~

The entry to ~~L1~~ occurs after being directed by the Data Link Layer and receiving an Electrical Idle ordered-set.

4.2.5.8.L2

Power can be aggressively conserved in **L2**. Most of the transmitter and receiver may be disabled³⁶. Main power and clocks are not guaranteed, but aux³⁷ power is available.

When Beacon support is required by the associated system or form-factor specification, an upstream port that supports the wakeup capability must be able to send; and a downstream port must be able to receive; a wakeup signal referred to as a Beacon.³⁸

~~An upstream Port must be able to send and a downstream Port must be able to receive a wakeup signal referred to as a Beacon.~~³⁹

The entry to **L2** occurs after being directed by the Data Link Layer and receiving an Electrical Idle ordered-set.

Disabled

The intent of **Disabled** state is to allow a group of Lanes to be disabled until directed or Electrical Idle is exited after entering **Disabled** (i.e. due to a hot removal and insertion).

Disabled uses Bit 1 (Disable Link) in the Training Control Register (Table 0-1, Table 0-2) which is sent within the TS1 and TS2 training ordered set.

A Link can enter **Disabled** if directed by a higher Layer. A Link can also reach the **Disable** state by receiving a TS1 ordered set with the Disable Bit asserted (see section 0).

4.2.5.9.External Loopback

Loopback is intended for test and fault isolation use. Only the entry and exit behavior is specified, all other details are implementation specific. **Loopback** can operate on ~~either a per Lane or configured-lane or a link~~ Link basis.

A Loopback Master is the component requesting **L**oopback.

A Loopback Slave is the component looping back the data.

Loopback uses Bit 2 (Loopback) in the Training Control Register (Table 0-1, Table 0-2) which is sent within the TS1 and TS2 training ordered set.

The entry mechanism for Loopback Master is device specific. ~~The system designer is responsible to ensure that two connected devices do not initiate loopback master requests at the same time.~~

³⁶ The exception is the receiver termination, which must remain in a low impedance state.

³⁷ In this context, “aux” power means a power source which can be used to drive the Beacon ~~and Receiver Detection~~ circuitry.

³⁸ ~~Certain form factor specifications require the use of A device generates a B~~beacons for a device to request main power reactivation, for example in order to wake a system that is in D3_{cold}. See Section <4.3.2.4> for information on the electrical requirements of the ~~beacon~~Beacon. Refer to Chapter <5> for more information on how a device may use the ~~beacon~~Beacon as the wake^{up} mechanism.

³⁹ A device generates ~~B~~beacons in order to wake a system that is in D3_{cold}. See Section 0 for information on the electrical requirements of the beacon. Refer to Chapter Error! Reference source not found. for more information on how a device may use the ~~B~~beacon as the wake mechanism.

The **Loopback** Slave device enters **Loopback** whenever two or more consecutive TS1 ordered-sets are received with the **L**oopback bit set.



IMPLEMENTATION NOTE

Use of the **Polling-Compliance Substate Loopback**

Once in the Loopback state, the master can send any pattern of symbols as long as the rules of 8b/10b encoding (including disparity) are followed. Once in Loopback, the concept of data scrambling is no longer relevant; what is sent out is looped back. The mechanism(s) and/or interface(s) utilized by the Data Link Layer to notify the physical layer to enter the Loopback state is component implementation specific and beyond the scope of this specification.

4.2.5.11.Training Control ResetHot Reset

Hot Reset uses Bit 0 (Hot Reset) in the Training Control Register (Table 0-1, Table 0-2) which is sent within the TS1 and TS2 training ordered set.

Hot Reset must be propagated to all downstream components.

•**Training Control Reset** is entered when directed or when two consecutive TS1 or TS2 ordered-sets are received with the **R**eset bit set.

4.2.6.Link Training and Status State Descriptions Rules

Various Link status bits are monitored through software. Table 0-3 describes how the Link status bits must be handled throughout the LTSSM (for more information see section <3.1> for LinkUp, <7.8????> for Link Speed, LinkWidth, Link Training, <6.2> for Receiver Error, and <6.7> for In-Band Presence)

Table 0-3: Table of Link Status mapped to the LTSSM

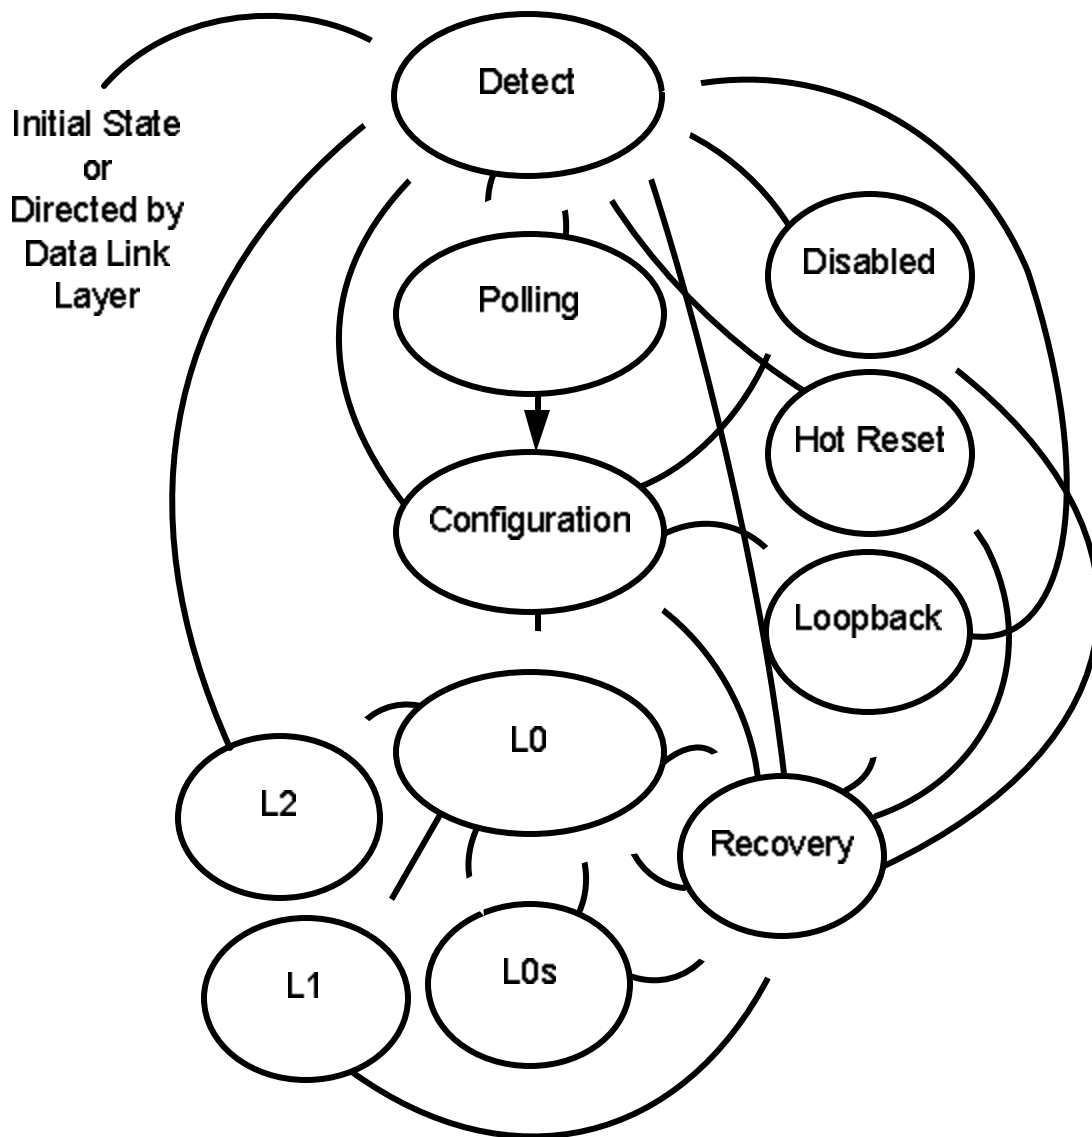
LTSSM State	Link Width	Link Speed	LinkUp	Link Training	Receiver Error	In-Band Presence
Detect	Undefined	Undefined	0	0	No action	0
Polling	Undefined	Undefined	0	0	No action	1
Configuration	Undefined	Undefined	0/1 ⁴⁰	1	Set on 8b/10b Error	1
Recovery	No action	No action	1	1	No action	1
L0	Set	Set	1	0	Set on 8b/10b Error ⁴¹	1
L0s	No action	No action	1	0	No action	1
L1	No action	No action	1	0	No action	1
L2	No action	No action	1	0	No action	1
Disabled	Undefined	Undefined	0	0	Set on 8b/10b Error	1
Loopback	No action	No action	1	0	No action	1
Hot Reset	No action	No action	1 ⁴²	0	Set on 8b/10b Error	1

The state machine rules for configuring and operating a PCI Express link are defined in the following sections.

⁴⁰ LinkUp will always be 0 if coming into **ConfigurationPolling** from **Detect** to **Polling** and is LinkUp will always be 1 if coming into **ConfigurationPolling** from any other state

⁴¹ Receiver Error must be set on a 8b/10b violation and optionally set on a Framing Violation.

⁴² Hot Reset will exit to **Detect** and will immediately result in Linkup=0



4.2.6.1.2. Detect.Active

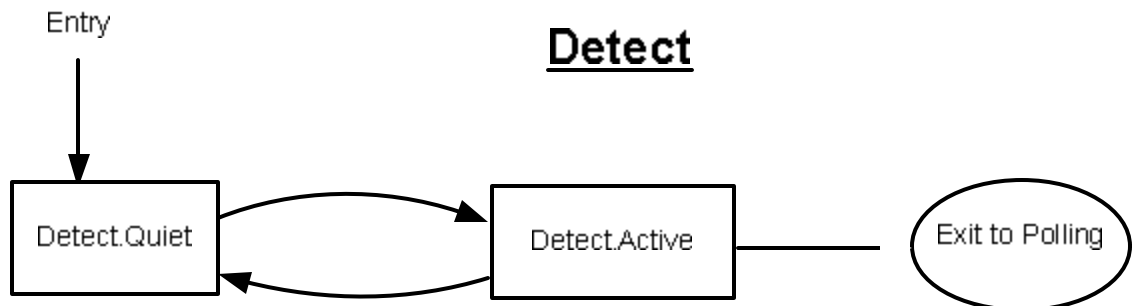
~~2~~ The transmitter performs a ~~high impedance~~ Receiver Detection sequence on all un-configured Lanes that can form one or more Links (see Section 0 for more information).

Next state is ~~Polling~~**Detect.Charge** if a receiver is detected on all un-configured Lanes

~~2~~ Next state is **Detect.Quiet** if a receiver is not detected on any Lanes.

If at least one but not all un-configured Lanes detect a receiver then:

5. Wait for 12ms.
6. The transmitter performs a Receiver Detection sequence on all un-configured Lanes that can form one or more Links (see Section 0 for more information),
 - i) The next state is **Polling** if exactly the same Lanes detect a receiver as the first Receiver Detection sequence.
 - Note: All Lanes that did not detect a receiver must transition to Electrical Idle⁴⁴, and are no longer associated with the Link training status state machine in progress or these same Lanes may then be associated with a new LTSSM if this optional feature is supported.
 - ii) Otherwise, the next state is **Detect.Quiet**



4.2.6.1.3. ~~Detect.Charge~~

~~2~~ Transmitter is in a high impedance Electrical Idle state.

⁴⁴ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).

? Next state is **Polling** after 12 ms timeout or when the operating DC common mode voltage is stable and within specification.⁴⁵

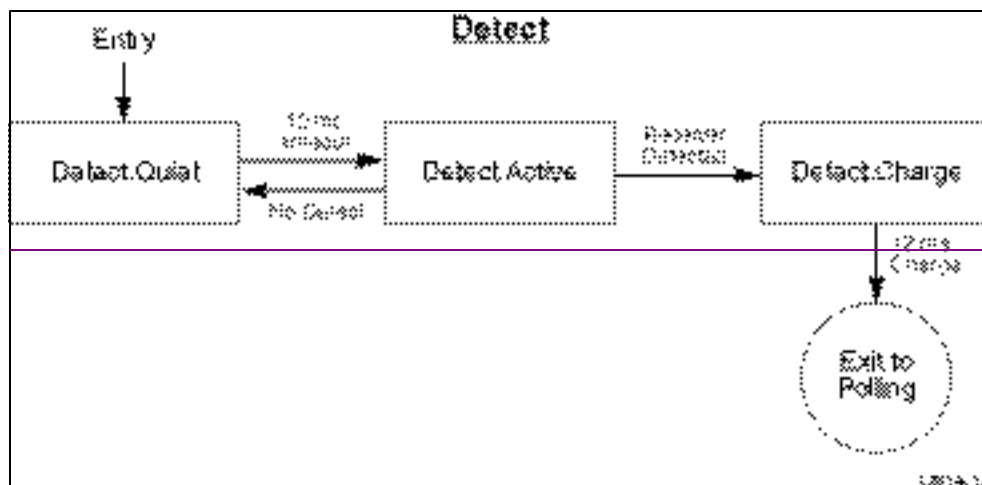


Figure 0-10: Detect Sub-State Machine

4.2.6.2. Polling

4.2.6.2.1. Polling.Quiet

? Transmitter is in Electrical Idle.

? A Receiver Detection sequence (see Section 4.3.1.8 for more information) is performed. If no receiver is present, next state is Detect

•Note: Detect must be done here in order to detect a device removed during polling.

? LinkUp = 0 (status is cleared).

? Next state is **Polling.Active** if a single TS1 or TS2 ordered set or their complement is received.

? Next state is **Polling.Active** after a minimum of 12 ms.

4.2.6.2.2. Polling.Active

? Transmitter sends a minimum of 1024 consecutive TS1 ordered sets on with Lane and Link numbers set to PAD (K23.7) on all Lanes that detected a receiver during **Detect**.

•Note: If the Extended Synch bit is set then the transmitter must send a minimum of 1024 consecutive TS1 ordered sets to allow external Link

⁴⁵ The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).

monitoring tools (i.e. logic analyzers) enough time to achieve bit and byte lock.

- ~~Note: This guarantees a minimum of 64 μ s for the bit lock time at generation 1 data rates.~~

~~2~~ Next state is **Polling.Configuration** after if a 8 consecutive single TS1 or TS2 ordered-sets or their complement is received with the Lane and Link numbers set to PAD (K23.7) on all Lanes that detected a receiver during **Detect** and at least 1024 TS1 Ordered sets were transmitted.

Otherwise, after a 24 ms timeout the next state is:

3. **Polling.Configuration** if

- Any Lane received 8 consecutive TS1 or TS2 ordered-sets (or their complement) with the Lane and Link numbers set to PAD (K23.7), and a minimum of 1024 TS1s are transmitted after receiving one TS1.

And

- If all lanes that detected a receiver during **Detect** have detected an exit from Electrical Idle at least once since entering **Polling**.

- Note: This prevents one or more bad receivers or transmitters from holding up a valid link from being configured, and allows for additional training in **Polling.Configuration**.

4. **Polling.Compliance** if at least one Lane's receiver in **Polling** has never detected an exit from Electrical Idle since entering **Polling**.

- Note: This indicates the presence of a passive test load on at least one Lane, which will force all Lanes to enter **Polling.Compliance**.

5. **Detect** if no TS1 or TS2 ordered set is received with Link and Lane number set to Pad on any Lane. The highest advertised speed in TS1 and TS2 is lowered (unless generation 1 is the highest advertised speed).

~~2~~ Next state is **Polling.Compliance** if ~~the transmitter has entered Polling.Active 16 consecutive times without~~ receiving a single TS1 or TS2 ordered set and the receiver has never detected an exit from Electrical Idle after the first time entering **Polling**.

~~Note: The compliance mode is entered only if electrical idle was never exited on the Lane since the time of reset.~~

~~2~~ ext state is **Polling.Quiet** if the transmitter sends 1024 TS1 ordered-sets without receiving a single TS1 or TS2 ordered set.

~~4.2.6.2.2~~ Polling.Compliance

~~2~~ Transmitter sends out the compliance pattern on all Lanes that detected a receiver during **Detect** at the data rate which was employed upon entry to **Polling.Compliance** (see Section 0).

? Next state is **Polling.Active** if Electrical Idle exit has been detected ~~is no longer detected~~ at the receiver ~~of all Lanes that detected a receiver during Detect~~

~~4.2.6.2.4.~~ Polling.Configuration

Receiver inverts polarity if necessary (see Section 0).

? Transmitter sends TS2~~1~~ ordered-sets with link and lane numbers set to PAD (K23.7) on all Lanes that detected a receiver during **Detect** ~~the Port. At least 16 TS1 ordered-sets are sent after receiving one TS1 or TS2 ordered-set.~~

The next state is **Configuration** after 8 consecutive TS2 ordered-sets are received on any Lanes that detected a receiver during **Detect**, 16 TS2 ordered-sets are transmitted after receiving one TS2 ordered-set, and none of those same Lanes is transmitting and receiving a higher Data Rate Identifier⁴⁶.

The next state is **Polling.Speed** after 8 consecutive TS2 ordered-sets are received on all Lanes that detected a receiver during **Detect**, 16 TS2 ordered-sets are transmitted after receiving one TS2 ordered-set, and at least one of those same Lanes is transmitting and receiving a higher Data Rate Identifier⁴⁷.

? Next state is ~~Configuration~~ if eight consecutive TS1 or TS2 ordered-sets are received and no higher data rate is supported

? Otherwise, next state is **Polling.Speed** if eight consecutive TS1 or TS2 ordered-sets are received.

? Otherwise, next state is ~~Detect~~**Polling.Active** after a ~~482~~ ms timeout.

~~4.2.6.2.5.~~ Polling.Speed

? The transmitter enters Electrical Idle for a minimum of $T_{TX-IDLE-MIN}$ (see Table 0-4) ~~and no longer than 2ms.~~

- Note : Electrical Idle ordered set is sent prior to entering ~~E~~lectrical Idle
 - Note: The DC common mode voltage does not have to be within specification.⁴⁸

? Data rate is changed on all Lanes to highest common data rate supported on both sides of the Link indicated by the training sequence (see Section 0).

⁴⁶ Higher data rate than what is currently being executed. Data Rate support on a link is determined by the highest common speed being transmitted and received in TS1 and TS2 ordered sets.

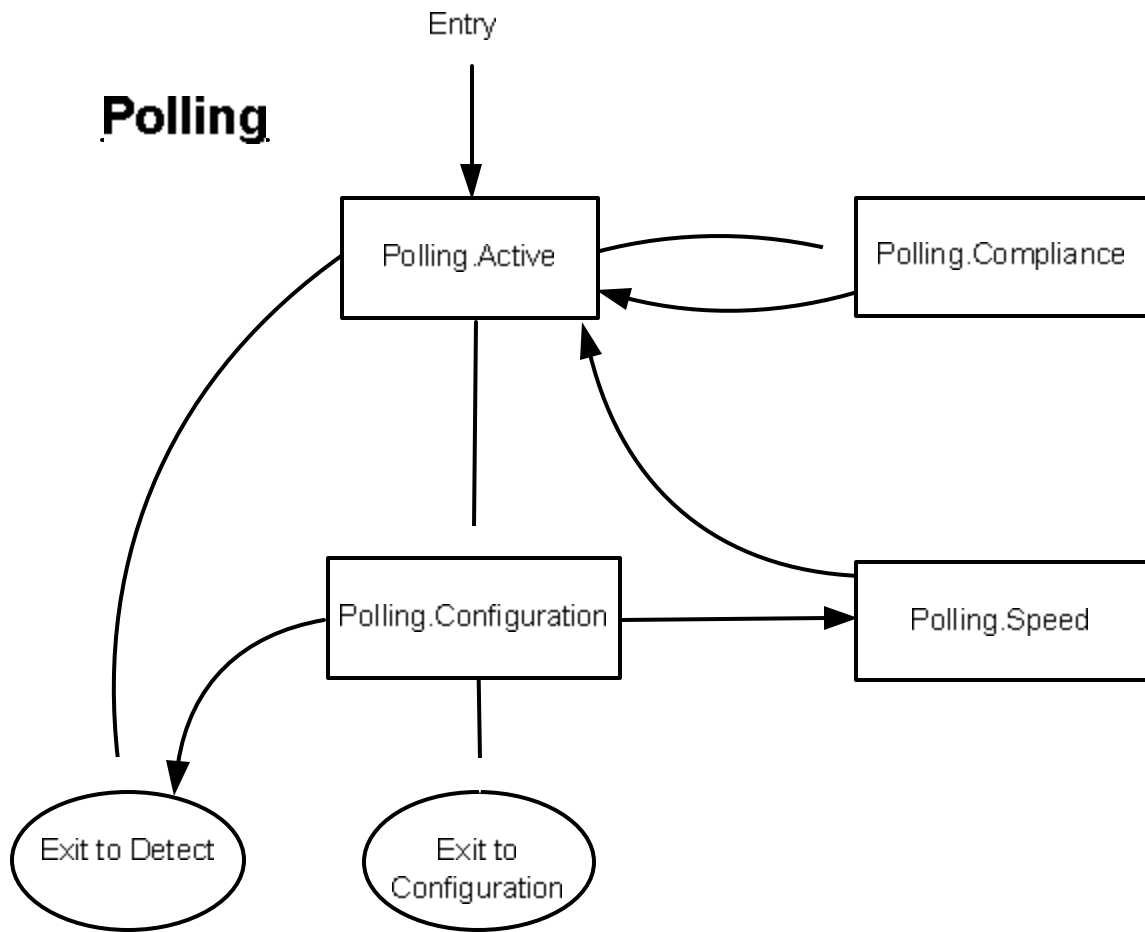
⁴⁷ Higher data rate than what is currently being executed. Data Rate support on a link is determined by the highest common speed being transmitted and received in TS1 and TS2 ordered sets.

⁴⁸ The common mode being driven does not need to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).

? ~~Transmitter sends a minimum of 1024 consecutive TS1 ordered sets on all lanes.~~

• ~~Note: This guarantees a minimum bit lock time.~~

? Next state is **Polling.ActiveConfiguration**.



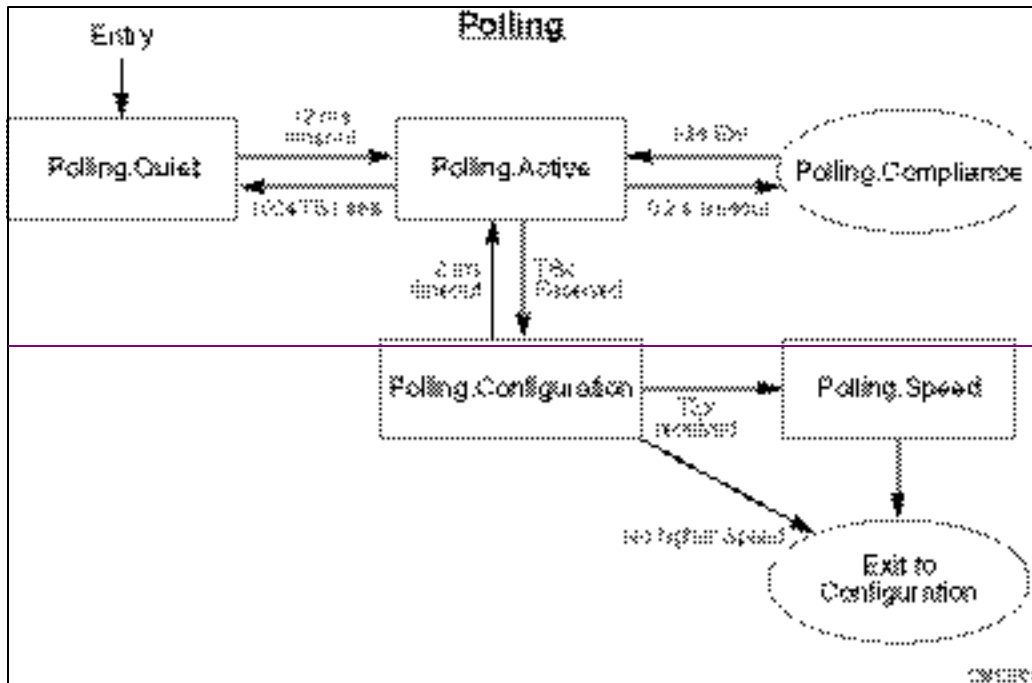


Figure 0-11: Polling Sub-State Machine

Configuration

Configuration.Linkwidth.Start

Downstream Lanes

↳ Next state is **Disable** if directed.

Note: “if directed” applies to a downstream Port that is instructed by a higher Layer to assert the Disable Link bit (TS1 and TS2) on all Lanes that detected a receiver during **Detect**.

Next state is **Loopback** if directed to this state, and the transmitter is capable of being a Loopback Master.

Note: “if directed” applies to a Port that is instructed by a higher Layer to assert the Loopback bit (TS1 and TS2) on all Lanes that detected a receiver during **Detect**.

↳ Next state is **Disable** after any Lanes that detected a receiver during **Detect** and are receiving TS1 ordered sets with the disable bit set in 2 consecutive TS1 ordered sets.

- Note: In the optional case where a cross-link is being configured the next state is **Disable** only after all Lanes that detected a receiver during **Detect** and are receiving TS1 ordered sets with the disable bit set in 2 consecutive TS1 ordered sets

Next state is **Loopback** if all Lanes that detected a receiver during **Detect** receive the enable loopback bit set in 2 consecutive TS1 ordered sets on all Lanes receiving a TS1 ordered-set.

- Note the device receiving the ordered set with the loopback bit set becomes the loopback slave.

The transmitter sends TS1 ordered sets with selected Link numbers and sets Lane numbers to PAD (K23.7) on downstream Lanes that are in **Configuration**.

- Note: Link numbers can only be different for groups of Lanes capable of being a unique link.
- Note: An example of Link number assignments includes a set of eight Lanes on an upstream component (downstream Lanes) capable of negotiating to become one x8 Port when connected to one downstream component (upstream Lanes) or two x4 Ports when connected to two different downstream components: The upstream component (Downstream Lanes) sends out TS1 ordered-sets with the Link number set to N on four Lanes and Link number set to N+1 on the other four Lanes. The Lane numbers are all set to PAD (K23.7).

If any Lanes first received at least one or more TS1 ordered sets with a Link and Lane number set to PAD (K23.7), the next state is **Configuration.Linkwidth.Accept** immediately after all those same downstream Lanes receive 2 consecutive TS1 ordered sets with any Link number that was transmitted and different than PAD (K23.7).

- Note: If the cross-link configuration is not supported the condition of first receiving a Link and Lane number set to PAD is always true.

Optionally, if all downstream Lanes initially receive a TS1 ordered set with a Link number different than PAD (K23.7), then the downstream Lanes are now designated as upstream Lanes and a new random cross Link timeout is chosen (see $T_{\text{crosslink}}$ in Table 4-4). The next state is **Configuration.Linkwidth.Start**

- Note: This supports the optional cross-link where both sides may try to act as a downstream Port. This is resolved by making both Ports become upstream and assigning a random timeout until one side of the Link becomes a downstream and the other side remains an upstream Port. This timeout must be random even when hooking up two of the exact same devices so as to eventually break any possible deadlock.

The next state is **Detect** after a 24 ms timeout.

Upstream Lanes

The transmitter sends out TS1 ordered sets with Link numbers and Lane numbers set to PAD (K23.7) on upstream Lanes that detected a receiver during **Detect**.

Immediately after all upstream Lanes receive 2 consecutive TS1 ordered sets with Link numbers that are different than PAD (K23.7), a single Link number is selected and transmitted on all Lanes that can form a Link. Any left over Lanes that detected a

receiver during **Detect** must transmit TS1 ordered sets with the Link and Lane number set to PAD (K23.7). The next state is **Configuration.Linkwidth.Accept**.

Optionally, if all upstream Lanes first receive 2 consecutive TS1 ordered sets with Link and Lane numbers set to PAD (K23.7), the upstream Lanes are changed to downstream Lanes after a $T_{\text{crosslink}}$ timeout and the next state is **Configuration.Linkwidth.Start**.

- Note: This optional behavior is required for cross-link behavior where two Ports may start off with upstream Ports, and one will eventually take the lead as a downstream Port.

□ The next state is **Detect** after a 24 ms timeout.

Configuration.Linkwidth.Accept

Downstream Lanes

If a Link can be formed with at least one group of Lanes that received 2 consecutive TS1 ordered sets with the same received Link number (non PAD and matching one that was transmitted by the Downstream Lanes), then TS1 ordered sets are transmitted with the same Link number and unique Lane numbers are assigned to all these same Lanes (Lane numbers must range sequentially from 0 to n-1). Any left over Lanes must transmit TS1 ordered sets with the Link and Lane number set to PAD (K23.7). The next state is **Configuration.Lanenum.Wait**.

Note: A couple interesting cases to consider here are the following.

4. A x8 downstream Port, which can be divided into two x4 Links, sends two different Link numbers on to two x4 upstream Ports. The upstream Ports respond simultaneously by picking the two Link numbers. The downstream Port will have to choose one of these sets of Link numbers to configure as a Link, and leave the other for a second pass through **Configuration**.
5. A x16 downstream Port, which can be divided into two x8 links, is hooked up to an upstream Port that can be configured as a x12 Link or x8 and x4 link. During **Configuration.Linkwidth.Start** the upstream Port returned the same Link number on all 12 Lanes. The downstream Port would then return the same received Link number and assign Lane numbers on the 8 Lanes that can form a x8 Link with the remaining 4 Lanes transmitting a Lane number and Link number set to PAD (K23.7).

□ The next state is **Detect** after a 2ms timeout or if no Link can be configured or if all lanes receive 2 consecutive TS1 ordered sets with Link and Lane numbers set to Pad (K23.7).

Upstream Lanes

If Lanes that transmitted a Link number receive 2 consecutive TS1 ordered sets with the same (non PAD) Link number and include at least a Lane number 0, then TS1 Lane numbers are transmitted that if possible match the received Lane numbers or are

different if necessary (i.e. lane reversed). Remaining Lanes must transmit TS1 with Link and Lane numbers set to PAD (K23.7). The next state is

Configuration.Lanenum.Wait.

- Note: The transmitted Lane numbers must range from 0 to m-1, be assigned sequentially to the same grouping of Lanes that are receiving Lane numbers 0 to n-1, include the received Lane 0, and m-1 must be equal to or smaller than the largest received Lane number (n-1).
- Note: A couple interesting cases to consider here are the following.

2.i. An x8 upstream Port is attached to a x8 downstream Port that is presented with Lane numbers that are backward from the preferred numbering. If the optional behavior of Lane reversal is supported by the upstream Port, the upstream Port transmits the same Lane numbers back to the downstream Port. Otherwise the opposite Lane numbers are transmitted back to the downstream Port, and it will be up to the downstream Port to optionally fix the Lane ordering or exit

Configuration.

- ii. Optional Lane reversal behavior is required to configure a Link where the Lane numbers are reversed and the downstream Port doesn't support Lane reversal. Specifically, the upstream Port Lane reversal will accommodate the scenario where the default upstream sequential Lane numbering (0 to n-1) is receiving a reversed downstream sequential Lane number (n-1 to 0).
- iii. An optional x8 upstream cross-link Port, which can be divided into two x4 links, is attached to two x4 downstream Ports that present the same Link number, and each x4 downstream Port presents Lane numbers simultaneously that were each numbered 0 to 3. The upstream Port will have to choose one of these sets of Lane numbers to configure as a link, and leave the other for a second pass through **Configuration**.

The next state is **Detect** after a 2ms timeout or if no Link can be configured or if all lanes receive 2 consecutive TS1 ordered sets with Link and Lane numbers set to Pad (K23.7).

Configuration.Lanenum.Accept

Downstream Lanes

If a configured Link can be formed with all the Lanes that receive 2 consecutive TS1 ordered sets with the same transmitted Link numbers and different Lane numbers (non PAD) the next state is **Configuration.Complete**.

- Note: Two possible scenarios exist that can result in the next state being **Configuration.Complete**.
- 3. The received Link and Lane numbers are the same as received, which is always required to be supported.

4. Optional Lane reversal behavior can occur for either scenario 1. An example of where Lane reversal is required is when the received Lane numbers are the reverse of the Lane numbers being transmitted (case 1). This optional feature would be required to configure a Lane reversed Link where the Upstream Port doesn't support Lane reversal.

If the Lanes of a Link can be configured by using the Lanes that transmitted a Lane number, which received 2 consecutive TS1 ordered sets with the same transmitted Link number (non-Pad) and include a Lane number 0, then TS1 Lane numbers are transmitted which must be assigned sequentially either in order or lane reversed to the same Lanes and no Lane numbers can be repeated. Left over Lanes must transmit TS1 Link and Lane numbers set to PAD (K23.7). The next state is **Configuration.Lanenum.Wait**.

- Note: The transmitted Lane numbers must range from 0 to m-1, be assigned sequentially in order or lane reversed to the same grouping of Lanes that are receiving Lane numbers 0 to n-1, include the received Lane 0, and m-1 must be smaller than the largest received Lane number (n-1).

The next state is **Detect** if no Link can be configured or if all lanes receive 2 consecutive TS1 ordered sets with Link and Lane numbers set to Pad (K23.7).

Upstream Lanes

If 2 consecutive TS2 ordered sets are received with Link and Lane numbers (non PAD values) that match what is being transmitted in the TS1 ordered sets, the next state is **Configuration.Complete**.

If the Lanes of a Link can be configured by using the Lanes that transmitted a Lane number, which received 2 consecutive TS1 ordered sets with the same transmitted Link number (non-Pad) and include a Lane number 0, then TS1 Lane numbers are transmitted which must be assigned sequentially either in order or lane reversed to the same Lanes and no Lane numbers can be repeated. Left over Lanes must transmit TS1 Link and Lane numbers set to PAD (K23.7). The next state is **Configuration.Lanenum.Wait**.

- Note: The transmitted Lane numbers must range from 0 to m-1, be assigned sequentially in order or lane reversed to the same grouping of Lanes that are receiving Lane numbers 0 to n-1, include the received Lane 0, and m-1 must be smaller than the largest received Lane number (n-1).

The next state is **Detect** if no Link can be configured or if all lanes receive 2 consecutive TS1 ordered sets with Link and Lane numbers set to Pad (K23.7).

Configuration.Lanenum.Wait

Downstream Lanes

The next state is **Configuration.Lanenum.Accept** if any of the Lanes receive 2 consecutive TS1 has which a Lane number different from when it first entered

Configuration.Lanenum.Wait, and not all the Lanes Link numbers are set to Pad (K23.7).

The next state is **Detect** after a 2ms timeout or if all lanes receive 2 consecutive TS1 ordered sets with Link and Lane numbers set to Pad (K23.7).

Upstream Lanes

The next state is **Configuration.Lanenum.Accept**

- 1) If any of the Lanes receive 2 consecutive TS1 which has a Lane number different from when it first entered

Configuration.Lanenum.Wait, and not all the Lanes Link numbers are set to Pad (K23.7).

or

- 2) If any lane receives 2 consecutive TS2 ordered-set

The next state is **Detect** after a 2ms timeout or if all lanes receive 2 consecutive TS1 ordered sets with Link and Lane numbers set to Pad (K23.7).

Configuration.Complete

Downstream Lanes

TS2 ordered sets are transmitted using Link and Lane numbers that match the received TS1 Link and Lane numbers.

N_FTS must be noted for use in L0s when leaving this state.

Lane to Lane deskew must be completed when leaving this state.

Scrambling is disabled if all configured Lanes have the Disable Scrambling bit asserted in 2 consecutively received TS2 ordered sets.

- Note: It is required that the Port that is sending the disable scrambling bit on all of the configured Lanes will also disable scrambling.

The next state is **Configuration.Idle** immediately after all Lanes that are transmitting TS2 ordered sets receive 8 consecutive TS2 ordered sets with matching Lane and Link numbers (non-Pad).

- Note: All remaining Lanes that aren't part of the configured Link transition to Electrical Idle, and are no longer associated with the Link training status state machine in progress. These same Lanes may then be associated with a new LTSSM if this optional feature is supported.

The next state is **Detect** after a 2ms timeout.

Upstream Lanes

TS2 ordered sets are transmitted using Link and Lane numbers that match the received TS2 Link and Lane numbers.

N_FTS must be noted for use in L0s when leaving this state.

Lane to Lane deskew must be completed when leaving this state.

Scrambling is disabled if all configured Lanes have the Disable Scrambling bit asserted in 2 consecutively received TS2 ordered sets.

- Note: It is required that the Port that is sending the disable scrambling bit on all of the configured Lanes will also disable scrambling.

The next state is **Configuration.Idle** immediately after all Lanes that are transmitting TS2 ordered sets receive 8 consecutive TS2 ordered sets with matching Lane and Link numbers (non-Pad).

- Note: All remaining Lanes that aren't part of the configured Link transition to Electrical Idle, and are no longer associated with the Link training status state machine in progress. These same Lanes may then be associated with a new LTSSM if this optional feature is supported.

The next state is **Detect** after a 2ms timeout.

4.2.6.4.Configuration

4.2.6.4.1.Config.RevrCfg

? Transmitter sends TS1 ordered sets on the lanes, starting the link width and lane ordering process. If configuration of the link is successful, state will complete by sending TS2 ordered sets. At least 16 TS1 ordered sets are sent after receiving one TS1 or TS2 ordered set.

- Note: All lanes must have achieved bit and symbol lock by this state as ensured by Polling.

? Link width and Lane reversal is performed as described in Section 4.2.4.8.

? Lanes that do not configure successfully and are not capable of becoming part of a new link must be disabled by putting the transmitter into high impedance. Lanes that do not configure successfully and are capable of becoming part of a new link must enter the detect state.

- Note: Disabled Lanes should be re-enabled if any active Lanes within the same Link enter Detect.

- Note: Disabled lanes are not required to send the electrical idle ordered set.

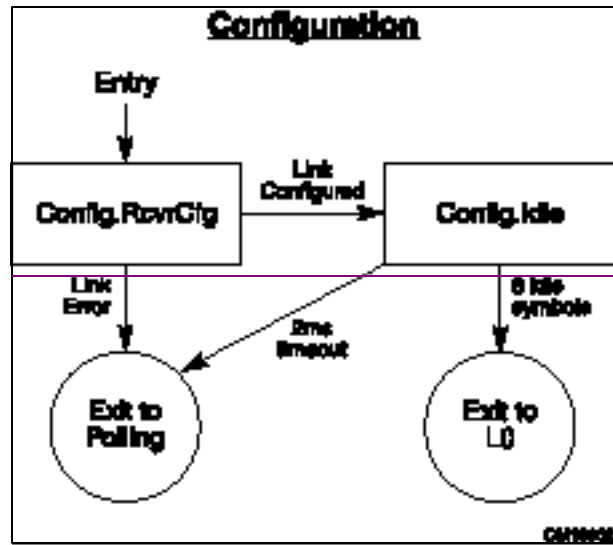
- Note: All Lanes on a configured Link must operate at the same data rate.

? If a link negotiates a valid configuration then the link width and lane reversal procedure will return directly to Config.Idle as described in Section 4.2.4.8.

? Otherwise, the data rate that the Port indicates it supports is dropped down to the next lower data rate and the next state is Polling. See Section 4.2.4.7 for information on data rate negotiation.

4.2.6.4.2. Configuration.Idle

- ? Transmitter sends Idle data symbols on all configured Lanes. ~~At least 16 idle data symbols are sent after receiving one Idle data symbol.~~
- ? Receiver waits for Idle data.
- ? LinkUp = 1 (status is set true).
- ? Next state is **L0** if ~~eight~~ consecutive symbol times of Idle data received on all configured Lanes ~~and 16 idle data symbols are sent after receiving one Idle data symbol.~~
- ? Otherwise, ~~after a minimum 2 ms timeout the data rate that the Port indicates it supports is dropped down to the next lower data rate and the~~ next state is **PollingDetect** after a minimum 2 ms timeout. ~~See Section 4.2.4.7 for information on data rate negotiation.~~





4.2.6.4.Recovery

4.2.6.4.1.Recovery.RcvrLock

Transmitter sends TS1 ordered-sets on all configured Lanes using the same Link and Lane numbers that were set after leaving **Configuration**.

Next state is **Recovery.RcvrCfg** if 8 consecutive TS1 or TS2 ordered-sets are received on all configured Lanes with the same Link and Lane numbers that match what is being transmitted on those same Lanes.

Note: If the Extended Synch bit is set then the transmitter must send a minimum of 1024 consecutive TS1 ordered sets before transitioning to **Recovery.RcvrCfg**.

Note: The Extended Synch allows external Link monitoring tools (i.e. logic analyzers) enough time to achieve bit and byte lock.

Otherwise, after a 24 ms timeout:

1. ~~after 2 ms the~~ The next state is **Configuration** if all the configured Lanes that are receiving a TS1 or TS2 ordered-set have received at least one TS1 or TS2 with Link and Lane numbers that match what is being transmitted on those same Lanes.
2. Otherwise, the next state is **Detect**

4.2.6.4.2.Recovery.RcvrCfg

? Transmitter sends TS2 ordered-sets on all configured Lanes using the same Link and Lane numbers that were set after leaving **Configuration**. ~~At least 16 TS2 ordered-sets are sent after receiving one TS2 ordered-set.~~

Next state is **Recover.Idle** if 8 consecutive TS2 ordered-sets are received on all configured Lanes with the same Link and Lane number that match what is being transmitted on those same Lanes and 16 TS2 ordered-sets are sent after receiving one TS2 ordered-set.

- Note: If the N_FTS value was changed this must be used for future L0s states.

Next state is **Configuration** if 8 consecutive TS1 ordered-sets are received on any configured Lanes with Link or Lane number that don't match what is being transmitted on those same Lanes and 16 TS2 ordered-sets are sent after receiving one TS2 ordered-set.

- Note: If the N_FTS value was changed this must be used for future L0s states.

? Next state is **Recovery.Idle** if 8 consecutive TS2 ordered-sets are received on all configured Lanes

? Otherwise, after a 482 ms timeout the next state is ~~Polling~~ **Detect**

~~4.2.6.4.3.Recovery.Idle~~

~~2~~Next state is **Disabled** if directed.

Note: “if directed” applies to a downstream Port that is instructed by a higher Layer to assert the Disable Link bit (TS1 and TS2) on the Link.

~~2~~Next state is **Hot Reset** if directed.

Note: “if directed” applies to a downstream Port that is instructed by a higher Layer to assert the Hot Reset bit (TS1 and TS2) on the Link.

~~2~~Next state is **Configuration** if directed.

Note: “if directed” applies to a Port that is instructed by a higher Layer to re-configure the Link (i.e. different width Link).

Next state is **Loopback** if directed to this state, and the transmitter is capable of being a Loopback Master.

Note: “if directed” applies to a Port that is instructed by a higher Layer to assert the Loopback bit (TS1 and TS2) on the Link.

Next state is **Disabled** immediately after any configured Lane has the Disable Link bit asserted in 2 consecutively received TS1 ordered sets.

Next state is **Hot Reset** immediately after any configured Lane has the Hot Reset bit asserted in 2 consecutive TS1 ordered sets.

Next state is **Configuration** if 2 consecutive TS1 ordered-sets are received on all configured Lanes and any Link and/or Lane numbers have changed since last exiting **Configuration**.

Next state is **Loopback** if any configured Lane has the Loopback bit asserted in 2 consecutive TS1 or TS2 ordered sets.

- Note the device receiving the ordered set with the loopback bit set becomes the loopback slave.

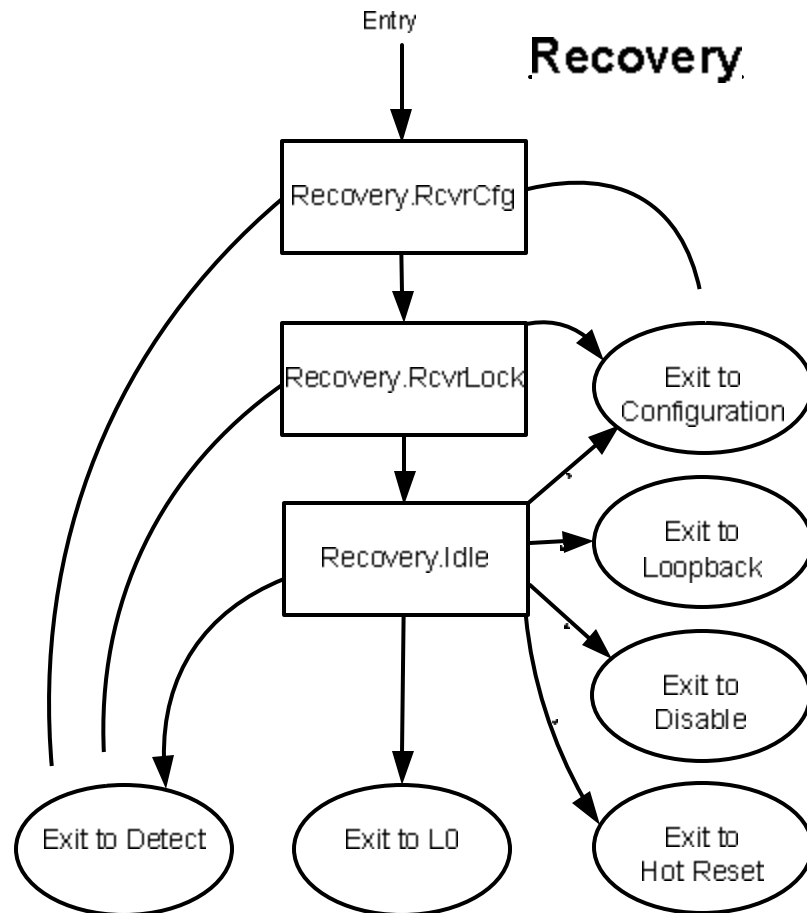
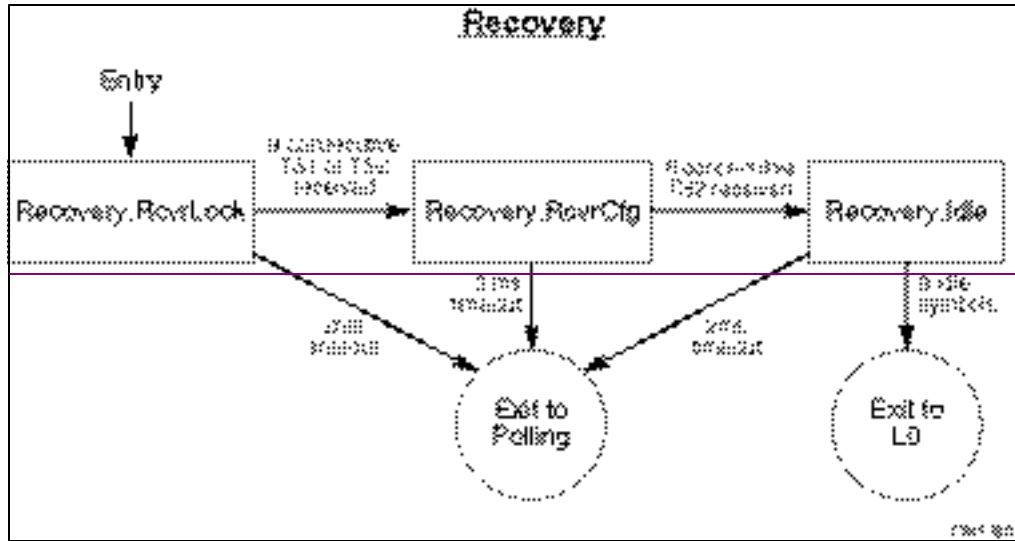
~~2~~Transmitter sends Idle data ~~(minimum of 16 symbol times)~~ on all configured Lanes.

- Note: If directed to other states then no Idle Characters have to be sent before transitioning to the other states (i.e. Disable, Hot Reset, Configuration, or Loopback)

~~2~~Receiver waits for Idle data.

~~2~~Next state is **L0** if ~~eight~~ consecutive symbol times of Idle data received on all configured Lanes and 16 Idle data symbols are sent after receiving one Idle data symbol.

~~2~~Otherwise, after a 2 ms timeout the next state is **PollingDetect**.



4.2.6.5.L0

This is the normal operational state.

LinkUp = 1 (status is set true).

~~? Transmitter and receiver are enabled in a low impedance state.~~

~~? Next state is **Recovery** if a TS1 or TS2 ordered-set is received on any configured Lane.~~

Next state is **Recovery** if directed to this state or if Electrical Idle is detected without receiving an Electrical Idle ordered-set.

Note: “if directed” applies to a Port that is instructed by a higher Layer to transition to **Recovery**.

Note: The transmitter may complete any TLP or DLLP in progress.

Next state of transmitter is **L0s** if directed to this state.

Note: “if directed” applies to a Port that is instructed by a higher Layer to initiate **L0s**.

~~? Next state is Polling if directed to this state.~~

~~? Next state is **Detect** if directed to this state.~~

~~? Next state of receiver is **L0s** if receiver detects Electrical Idle ordered-set and not directed to L1 or L2 states.~~

~~? Next state of transmitter is **L0s** if directed to this state.~~

Next state is **L1**:

xi. If directed.

And

xii. An Electrical Idle ordered-set is received.

And

xiii. An Electrical Idle ordered-set is transmitted.

Note: “if directed” is defined as both ends of the Link having agreed to enter **L2** immediately after the condition of both the receipt and transmission of an Electrical Idle ordered-set is met.

~~? Next state is **L2**:~~

xiv. If directed.

And

xv. An Electrical Idle ordered-set is received.

And

xvi. An Electrical Idle ordered-set is transmitted.

Note: “if directed” is defined as both ends of the Link having agreed to enter **L2** immediately after the condition of both the receipt and transmission of an Electrical Idle ordered-set is met.

4.2.6.6.L0s

4.2.6.6.1.Receiver L0s

4.2.6.6.1.1. Rx_L0s.Entry

2-Next state is **Rx_L0s.Idle** after a $T_{TX-IDLE-SET-TO-IDLE}$ (Table 0-4) timeout.

Note: This guarantees that transmitter has established the Electrical Idle condition.

4.2.6.6.1.2. Rx_L0s.Idle

2-Next state is **Rx_L0s.FTS** if the receiver detects an exit from Electrical Idle on any Lane of the configured Link.

4.2.6.6.1.3. Rx_L0s.FTS

The next state is **L0** if a SKP ordered set is received on all configured Lanes of the link.

- Note: The receiver must be able to accept valid data immediately after the SKP ordered-set. ~~Receiver locks to incoming bit stream and acquires symbol alignment~~

2-Otherwise, nNext state is **Recovery** if the receiver does not detect bit, symbol alignment, and one FTS ordered set wafterithin the- N_FTS timeN_FTSout-duration plus the maximum SKP ordered set duration on all Lanes of the Link.

- Note: The N_FTS timeout is approximately $40 \cdot [N_FTS + 1 \text{ (SKP)}] \cdot UI$. This time would be exact except for that SKPs may be longer than a 4 symbol ordered set.
- Note: The transmitter must also transition to **Recovery**, but may complete any TLP or DLLP in progress.
- Note: It is recommended that the N_FTS field be increased when transitioning to **Recovery** to prevent future transitions to **Recovery** from **Rx_L0s.FTS**.

4.2.6.6.2. Transmitter L0s

4.2.6.6.2.1. *Tx_L0s.Entry*

Transmitter sends the Electrical Idle ordered set and enters Electrical Idle.

- **Note:** The DC common mode voltage must be within specification by $T_{TX-IDLE-SET-TO-IDLE}$.⁴⁹ ~~Note: The L0s electrical idle can be held at low impedance or high impedance.~~

Next state is **Tx_L0s.Idle** after a $T_{TX-IDLE-MIN}$ (Table 0-4) timeout.

4.2.6.6.2.2. *Tx_L0s.Idle*

Next state is **Tx_L0s.FTS** if directed.

4.2.6.6.2.3. *Tx_L0s.FTS*

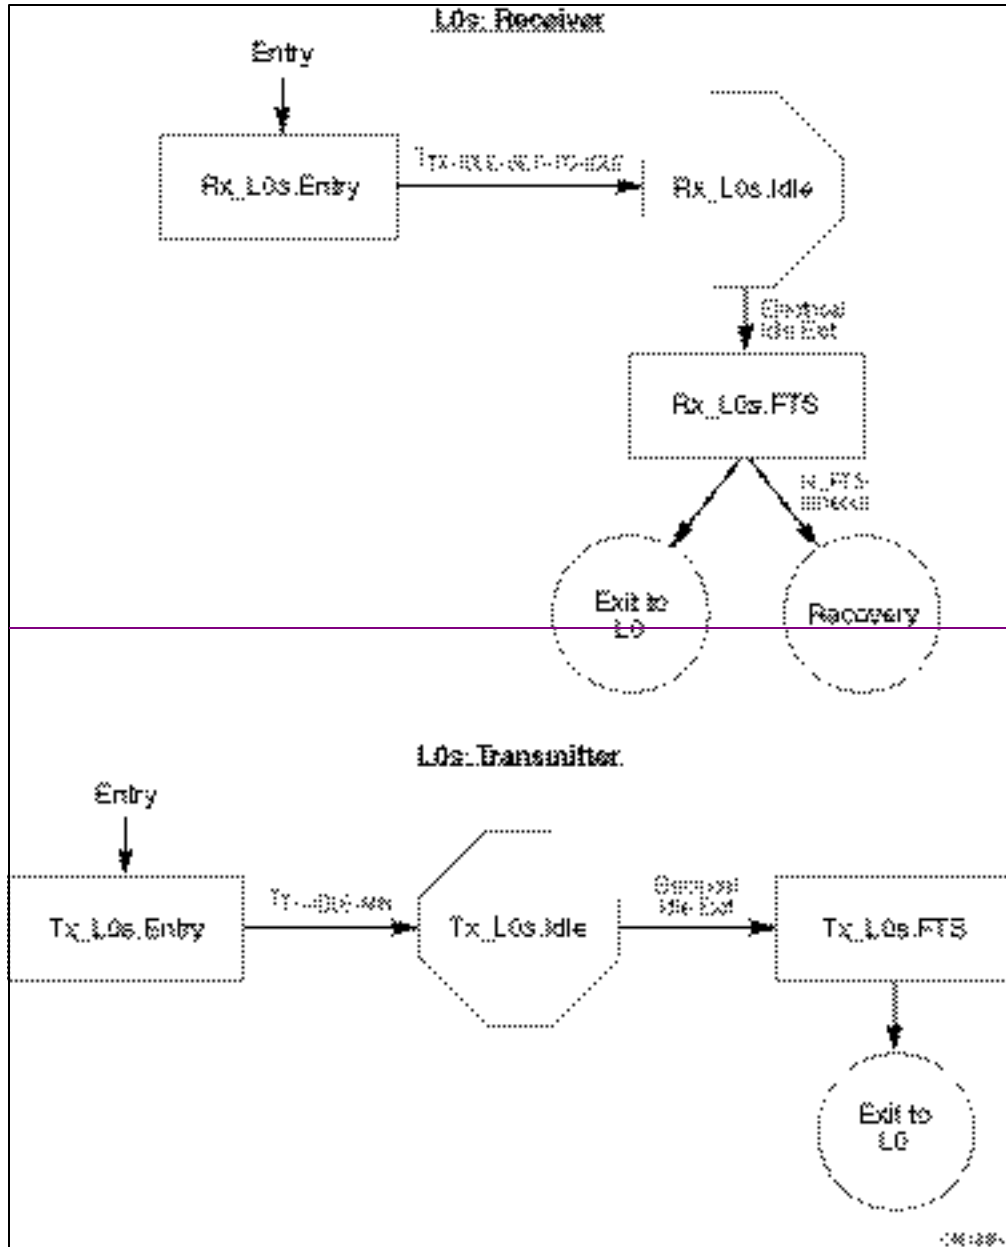
Transmitter sends N_FTS Fast Training Sequences.

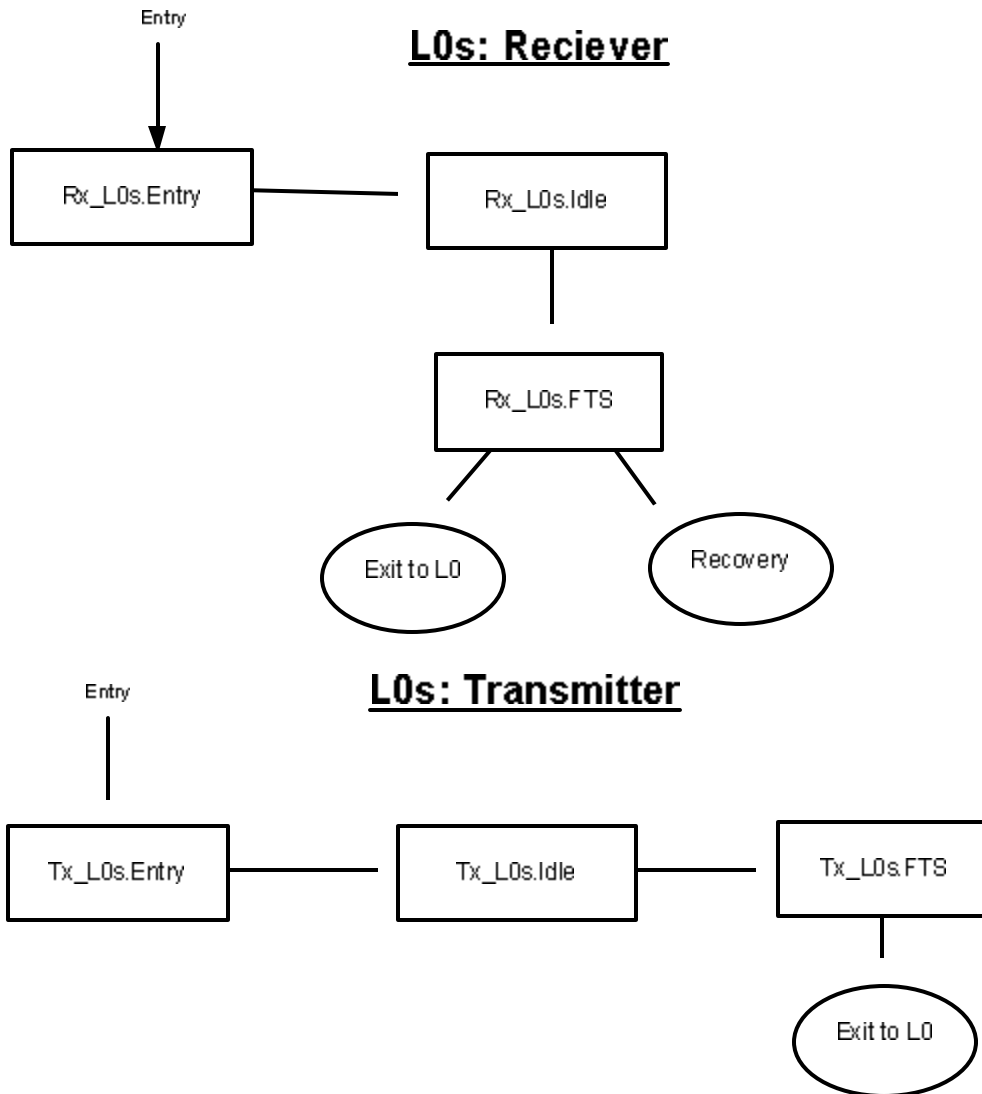
- **Note:** No SKP ordered sets can be inserted before all FTS ordered-sets as defined by the agreed upon N_FTS parameter are transmitted.
- **Note:** If the Extended Synch bit is set then the Transmitter sends 4096 Fast Training Sequences.

Transmitter sends a single SKP ordered set on all configured Lanes.

Next state is **L0**.

⁴⁹ The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).





4.2.6.7.L1

4.2.6.7.1.L1.Entry

Transmitter sends the Electrical Idle ordered set and enters Electrical Idle.

- Note: The DC common mode voltage must be within specification by $T_{TX-IDLE-SET-TO-IDLE}$ ⁵⁰.

~~2 Receiver waits for at least Electrical Idle $T_{TX-IDLE-SET-TO-IDLE}$ time given in Table 4-4.~~

~~2 The next state is **L1.Idle** after a $T_{TX-IDLE-MIN}$ (Table 0-4) timeout.~~

Note: This guarantees that transmitter has established the Electrical Idle condition.

~~4.2.6.7.2.L1.Idle~~

~~2 Transmitter remains~~ in Electrical Idle.

- Note: The DC common mode voltage must be within specification⁵¹.


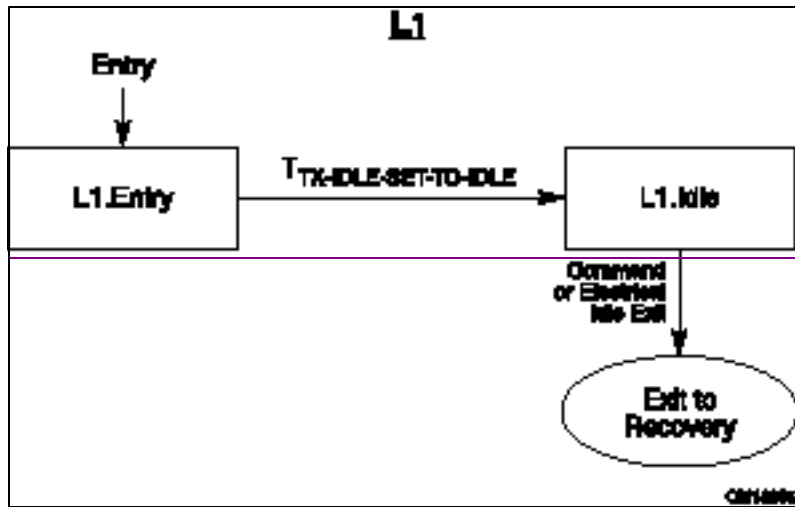
~~•Note: The L1 electrical idle can be held at low impedance or high impedance.~~

~~2 Next state is **Recovery** if Extended Synch bit is not set and either directed or if any receiver detects exit from Electrical Idle or directed.~~

~~○ If Extended Synch bit is set and either directed or if any receiver detects exit from Electrical Idle the transmitter first sends 1024 TS1 and then exits to **Recovery**.~~

⁵⁰ The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).

⁵¹ The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).



4.2.6.8.L2

4.2.6.8.1.L2.Idle

All RX Termination must remain enabled in low impedance.

Transmitter sends the Electrical Idle ordered set and enters Electrical Idle.

- Note: The DC common mode voltage does not have to be within specification.⁵²

For downstream Lanes:

For a Root Port ~~Transmitter is in high impedance Electrical Idle.~~

~~? Note: The transmitter is not required to perform receiver detection during L2 Electrical Idle.~~

~~? Note: Transmitters are not required to hold DC common mode in L2 Electrical Idle~~

the next state is **Detect** if a Beacon is received on at least Lane 0.

Note: Main power must be restored before entering **Detect**

- For a Switch if a Beacon is received on at least Lane 0 the upstream Port must transition to **L2.TransmitWake**.
- ~~If a Beacon is detected on any receiver, then upstream the transmitters in the direction of the Root Complex must transition to L2.Detect. See section 4.3.2.4 for more information~~

For upstream Lanes:

The next state is **Detect** if Electrical Idle Exit is detected on any Lane.

Note: A switch must transition any downstream lanes to **Detect**.

~~— If Electrical Idle Exit is detected on any receiver that is in the direction of the Root Complex then all ports transition to **Detect**~~

☐ Next state is **L2.TransmitWake** for an upstream port ~~L2.Detect~~ if directed to transmit a Beacon.

- Note: Beacons may only be transmitted on ~~upstream~~ Ports in the direction of the Root Complex.

●

⁵² The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle (V_{TX-CM-DC-ACTIVE-IDLE-DELTA}) specification (see Table 0-4).

4.2.6.8.2.L2.Detect

- ? ~~Transmitter is in a high impedance Electrical Idle state for a minimum of 12 ms.~~
- ? ~~A high impedance Receiver Detection sequence is performed (see Section 4.3.1.8 for more information)~~
- ? ~~Next state is L2.TransmitWake if a receiver is detected.~~
- ? ~~Next state is Detect if a receiver is not detected.~~

4.2.6.8.3.L2.TransmitWake

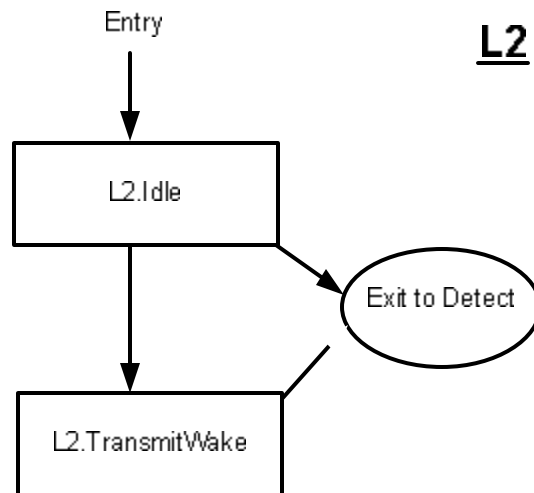
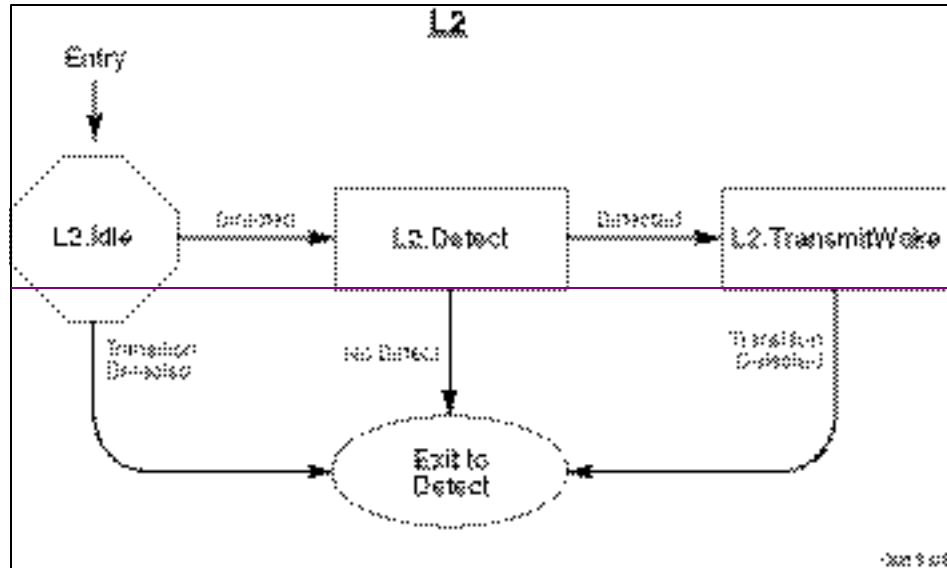
Note: This state only applies to upstream Ports.

~~Transmitter is in a high impedance Electrical Idle state for a minimum of 12 ms.~~

~~The transmitter transmits the Beacon on at least Lane 0 of the Port in the direction of the Root Complex (Refer to Section 0).~~

- ? Next state is **Detect** if Electrical Idle exit is detected on any upstream Port's receiver ~~that is in the direction of the Root Complex.~~

Note: Power is guaranteed to be restored when upstream receivers see Electrical Idle exited, but it may also be restored prior to Electrical Idle being exited.



?All Lanes in the configured Link tEntrance to and exit from this state only when directed.

? Transmitter sends between 4 and 16 TS1 ordered-sets with the Disable Link bit (Bit 1) assertedset and then transition to Electrical Idle.

Note: The Electrical Idle ordered set must be sent prior to entering Electrical Idle.

Note: The DC common mode voltage does not have to be within specification.

53

? If an Electrical Idle ordered-set was received (even while transmitting TS1 with the Disable Link bit asserted) Transmitter then goes into a high impedance Electrical Idle state.

then:

LinkUp = 0 (False)

The nNext state is **Detect** when directed or if Electrical Idle is exited.

Otherwise, after a 2ms timeout the next state is **Detect**.

4.2.6.10. Loopback

~~This mode is intended for test and fault isolation use only, and is not a normal operational mode. Only the entry and exit behavior is specified. All other details are implementation specific.~~

4.2.6.10.1. Loopback.Entry

? ~~The master device enters loopback when directed.~~ The Loopback Master device asserts the Loopback bit (Bit 2) in the TS1 ordered-set while preserving the configured Link and Lane numbers and transmits TS1 ordered sets until the Loopback Master receives TS1 ordered sets with the loopback bit set. The next state is **Loopback.Active**.

? Note: This indicates to the Loopback Master that the Loopback Slave has successfully entered **Loopback**.

? Note: The Loopback Master device timeout from loopback entry is implementation specific. The exit is to **Loopback.Exit**.

⁵³ The common mode being driven must meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle ($V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$) specification (see Table 0-4).

~~? The slave device enters loopback when two or more consecutive TS1 ordered-sets are received with the loopback bit set.~~

~~? Note: Loopback must be entered within the receipt of 1024 consecutive TS1 ordered sets with the loopback bit set.~~

~~? The next state for the Loopback Slave is **Loopback.Active**.~~

~~4.2.6.10.2: Loopback.Active~~

The Loopback Master ~~mustis~~ sending valid 8b/10b characters. The next state of the Loopback Master is **Loopback.Exit** if directed.

~~? The Loopback Slave must receive valid 8b/10b data. If SKP ordered-sets are received they are also looped back to the Loopback Master. SKP symbols must be added or removed by the Loopback Slave as needed.~~

A Loopback Slave is required to retransmit each 10b data and control symbol exactly as received, without applying scrambling/descrambling or disparity corrections, with three important exceptions:

If a received 10b symbol is determined to be an invalid 10b code (i.e. no legal translation to a control or data value possible) then the slave must instead transmit the EDB symbol in the corresponding time slot of the invalid symbol. The disparity version of the EDB symbol is selected using normal transmission rules based on the running transmitter disparity history up to the EDB transmission.

If a SKP ordered set retransmission requires adding a SKP symbol to accommodate timing tolerance correction, the SKP symbol is inserted in the retransmitted symbol stream anywhere in the ordered set following the COM symbol. The disparity of the added SKP symbol is selected using normal transmission rules based on the running transmitter disparity history up to the added SKP transmission.

If a SKP ordered set retransmission requires dropping a SKP symbol to accommodate timing tolerance correction, the SKP symbol is simply not retransmitted, and transmission continues with the next received symbol or an EDB, as defined above.

As result of these rules, received valid 10b codes are retransmitted even if they fail to match expected disparity in the receiver and result in retransmission violating normal disparity rules.

~~? The Loopback Slave's transmitter is sending the 10 bit data as received. If the received data was not 8b/10b valid, the transmitter sends back the special symbol EDB control character in place of the invalid character.~~

~~? The Loopback Slave must transmit the data with the same disparity as was received.~~

~~? Data received with disparity violations are transmitted with the new disparity and are not replaced with EDB.~~

~~? Next state of the Loopback Slave is **Loopback.Exit** when an Electrical Idle ordered-set is received or Electrical Idle is detected.~~

~~? Next state of the Loopback Slave is Loopback.Exit if an Electrical Idle is detected continuously for a minimum of 1 UI at the Loopback Slave receiver.~~

~~2~~ The next state of the Loopback Master is **Loopback.Exit** if directed.

~~4.2.6.10.3~~ Loopback.Exit

~~2~~ The **Loopback** Master sends an Electrical Idle ordered set and goes to ~~high impedance~~ Electrical Idle ~~for a minimum of $T_{TX-IDLE-MIN}$~~ (Table 0-4).

~~2~~ The slave echoes the Electrical Idle ordered set and goes to ~~high impedance~~ Electrical Idle ~~for a minimum of $T_{TX-IDLE-MIN}$~~ (Table 0-4)..

~~2~~ The next state of the ~~Loopback Master and Slave~~ ~~and the master~~ is **Detect.Detect**

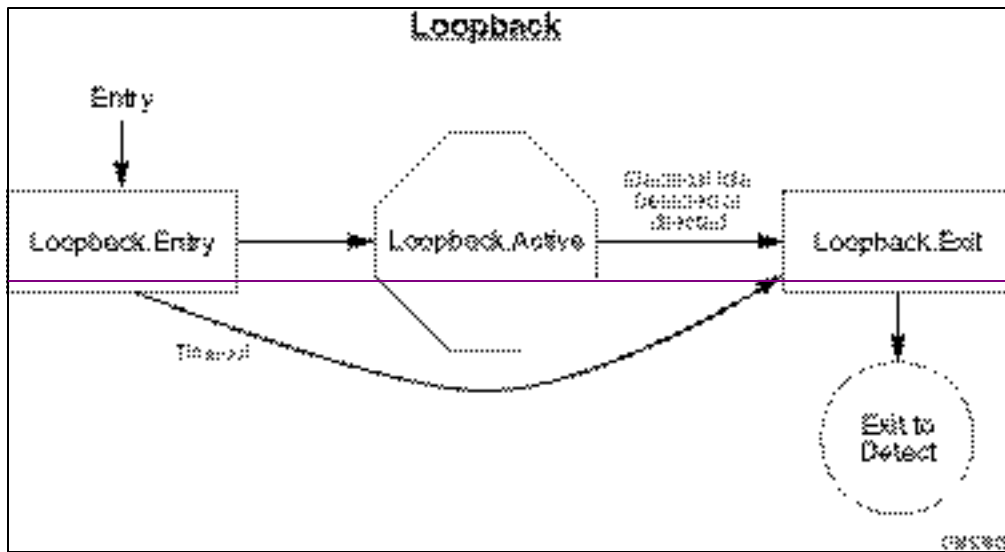



Figure 0-17: Loopback State Machine

~~4.2.6.11. Training Control Reset~~ Hot Reset

~~Training Control Reset Active~~

~~? All Lanes in the configured Link tLink enters reset state and transmit a minimum of 1024 TS1 ordered-sets with the Hot Rreset bit (Bit 0) assertedset on all downstream ports.~~

If any Lane of an upstream Port of a Switch receives a training sequence with the Hot Reset bit asserted, all configured downstream Ports must transition to **Hot Reset** as soon as possible.

~~? All transmitters on upstream ports transmit one Electrical Idle ordered-set, then enter Electrical Idle.~~

~~? Next state is Detect~~

~~4.2.7. Clock Tolerance Compensation~~

Skip ordered-sets (defined below) are used to compensate for differences in frequencies between bit rates at two ends of a Link. The receiver physical layer logical sub-block must include elastic buffering which performs this compensation. The interval between skip ordered-set transmissions is derived from the absolute value of the transmit and receive clock frequency difference specified in Table 0-4. Having worse case clock frequencies at the limits of the tolerance specified will result in a 600 ppm difference between the transmit and receive clocks of a Link. As a result, the transmit and receive clocks can shift one clock every 1666 clocks.

Rules for Transmitters

- 2-All Lanes shall transmit Symbols at the same frequency (the difference between bit rates is 0 ppm within all multi-Lane links).
 - 2-When transmitted, the skip ordered-set shall be transmitted simultaneously on all Lanes of a multi-Lane Link (See Section 0 and Table 0-4 for the definition of simultaneous in this context).
 - 2-The transmitted skip ordered-set is: one COM Symbol followed by three consecutive SKP Symbols
 - 2-The skip ordered-set shall be scheduled for insertion at an interval between 1180 and 1538 Symbol Times.
 - 2-Scheduled SKIP ordered-sets shall be transmitted if a packet or ordered-set is not already in progress, otherwise they are accumulated and then inserted consecutively at the next packet or ordered-set boundary.
- SKIP ordered-sets do not count as an interruption when monitoring for consecutive characters or ordered set (I.E. 8 consecutive TS1 ordered sets in **Polling.Active**)

Rules for Receivers

- 2-Receivers shall recognize received skip ordered-set consisting of one COM Symbol followed consecutively by one to five SKP Symbols.
- Note: The number of received SKP symbols in an ordered-set shall not vary from Lane to Lane in a multi-Lane Link.
- 2-Receivers shall be tolerant to receive and process SKIP ordered-sets at an average interval between 1180 to 1538 symbol times.
 - 2-Receivers shall be tolerant to receive and process consecutive SKIP ordered-sets.
- 2-Note: Receivers shall be tolerant to receive and process SKIP ordered-sets separated from each other at most 5664 symbol times – measured as the distance between the leading COM symbols.

4.2.8.Compliance Pattern

During **Polling** the **Polling.Compliance** sub-state ~~the compliance substate of the polling state machine~~ must be entered based on the presence of a test equipment being attached to one Lane of a possible Link (see Section 0). The compliance pattern consists of the sequence of 8b/10b symbols K28.5, D21.5, K28.5, and D10.2 repeating. Current running disparity must be set to negative before sending the first symbol.

~~The compliance pattern is not entered if the receiver has previously detected an exit from Electrical Idle.~~

The compliance sequence is:

Symbol	K28.5	D21.5	K28.5	D10.2
Current Disparity	0	1	1	0
Pattern	0011111010	1010101010	1100000101	0101010101

For any given device that has multiple Lanes, every eighth Lane is delayed by a total of four symbols. A two symbol delay occurs at both the beginning and end of the four symbol sequence, ~~for a total of eight symbols.~~

The delay sequence on every ~~eighth~~^{fourth} Lane is ~~illustrated below~~^{then}:

Symbol:	D	D	K28.5	D21.5	K28.5	D10.2	D	D
----------------	---	---	-------	-------	-------	-------	---	---

~~Where~~ D is a K28.5 symbol.

After the eight symbols are sent, the delay symbols are advanced to the next Lane and the process is repeated. ~~An illustration of this process is shown below~~^{This looks like:}

Lane 0	D	D	K28.5-	D21.5	K28.5+	D10.2	D	D	K28.5-	D21.5	K28.5+	D10.2
Lane 1	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2	D	D	K28.5-	D21.5
Lane 2	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2
Lane 3	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2
Lane 4	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2
Lane 5	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2
Lane 6	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2
Lane 7	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2
Lane 8	D	D	K28.5-	D21.5	K28.5+	D10.2	D	D	K28.5-	D21.5	K28.5+	D10.2
Lane 9	K28.5-	D21.5	K28.5+	D10.2	K28.5-	D21.5	K28.5+	D10.2	D	D	K28.5-	D21.5

Key:

K28.5- Comma when disparity is negative, specifically: "0011111010"

K28.5+ Comma when disparity is positive, specifically: "1100000101"

D21.5 Out of phase data character, specifically: "1010101010"

D10.2 Out of phase data character, specifically: "0101010101"

D Delay Character – K28.5

This sequence of delays ensures maximum possible interference between adjacent Lanes, enabling measurement of the compliance pattern under ~~close to~~ worst-case ~~Inter-Symbol Interference and cross-talk~~ conditions.

The compliance pattern ~~can~~^{is} only be exited ~~when~~^{if} an Electrical Idle exit-~~condition~~ is detected at ~~all the~~ ~~Lanes~~^{receiver} that detected a receiver during ~~Detect~~ or if a physical reset occurs.

4.3. Electrical Sub-Block

The Electrical sub-block contains a Transmitter and a Receiver. The Transmitter is supplied by the Logical sub-block with Symbols which it serializes and transmits onto a Lane. The Receiver is supplied with serialized Symbols from the Lane. It transforms the electrical signals into a bit stream which is de-serialized and supplied to the Logical sub-block along with a Link clock recovered from the incoming serial stream.

4.3.1. Electrical Sub-Block Requirements

4.3.1.1. Clocking Dependencies

The Ports on the two ends of a Link must transmit data at a rate that is within 600 parts per million (ppm) of each other at all times. This is specified to allow bit rate clock sources with a +/- 300 ppm tolerance.

4.3.1.1.1. Spread Spectrum Clock (SSC) Sources

The data rate can be modulated from +0% to -0.5% of the nominal data rate frequency, at a modulation rate in the range not exceeding 30 kHz – 33 kHz. The +/- 300 ppm requirement still holds, which requires the two communicating Ports be modulated such that they never exceed a total of 600 ppm difference. For most implementations this places the requirement that both Ports require the same bit rate clock source when the data is modulated with an SSC.

4.3.1.2. AC Coupling

Each Lane of a Link must be AC coupled. The minimum and maximum value for the capacitance is given in Table 0-4. The requirement for the inclusion of AC coupling capacitors on the interconnect media is specified by the transmitter.

4.3.1.3. Interconnect

In the context of this specification, the interconnect comprises everything between the pins at of a transmitter package and the pins of a receiver package. Often, this will consist of traces on a printed circuit board or other suitable medium, AC coupling capacitors and perhaps connectors. The interconnect total capacitance to ground seen by the receiver detection circuit (see Section 0) must not ~~by~~ exceed 3 nF, including capacitance added by attached test instrumentation. Note that this capacitance is separate and distinct from the AC Coupling capacitance value (see Section 0).

4.3.1.4. Termination

~~Low and high impedance states are defined for both the transmitter and the receiver and are listed in Table 4-4 and Table 4-5. The transmitter is required to meet RL_{TX-DM} , RL_{TX-CM} , $Z_{TX-DIFF-DC}$, Z_{TX-DC} (see Table 0-4) anytime functional differential signals are being transmitted.~~

~~The transmitter is required to only meet Z_{TX-DC} (see Table 0-4) anytime functional differential signals are not being transmitted.~~

~~Note: The differential impedance during this same time is not defined.~~

The receiver is required to meet RL_{RX-DM} , RL_{RX-CM} , $Z_{RX-DIFF-DC}$, Z_{RX-DC} (see Table 0-5) during all LTSSM states.

The receiver is required to meet $Z_{RX-HIGH-IMP-DC}$ (see Table 0-5) anytime adequate power is not provided to the receiver.

Note: The differential impedance during this same time is not defined.

~~The only time the transmitter high impedance state is required is to initialize and maintain Electrical Idle during times when a hot plug/removal or asynchronous power up event could occur.⁵⁴~~

~~The transmitter low impedance state is required any time differential data is to be sent.~~

~~The only time the receiver must be in a high impedance state is when the receiver does not have power. Otherwise, the receiver must always be in a low impedance state.~~

❑ DC Common Mode Voltage

The receiver DC common mode voltage is always 0 V during all states.

The transmitter DC common mode voltage ~~is initially established during Detect and~~ is held at the same value during all ~~subsequent~~ states unless otherwise specified. The range of allowed transmitter DC common mode values is specified in Table 0-4 ($V_{TX-DC-CM}$).

4.3.1.6.ESD

All signal and power pins must withstand (2000 V) of ESD using the human body model and 500 V using the charged device model without damage. Class 2 per JEDEC JESE22-A114-A.

This ESD protection mechanism also helps protect the powered down receiver from potential common mode transients during certain possible reset or surprise insertion situations.

4.3.1.7.Short Circuit Requirements

All Transmitters and Receivers must support surprise hot insertion/removal without damage to the component. The transmitter and receiver must be capable of withstanding sustained short circuit to ground of D+ and D-. The transmitter short circuit current limit $I_{TX-SHORT}$ is provided in Table 0-4 respectively.

4.3.1.8.Receiver Detection

The Receiver Detection circuit is performed by a transmitter and must correctly detect whether a load impedance of Z_{RX-DC} or lower is present.

~~The receiver detection sequence is used to avoid unwanted common mode transfers between the receiver and transmitter.~~

~~The receiver detection can be performed in either a low or high impedance state unless explicitly specified.~~

⁵⁴ Any time high impedance is required it is explicitly stated in the Link Training and Status State Machine (LTSSM).

The behavior of the receiver detection sequence is described below.

Step 1. Transmitter is in a stable Electrical Idle state.⁵⁵

Step 2. ~~A transmitter doing a high impedance detection changes the common mode or differential voltage on both D+ and D- lines⁵⁶ to a different value. A transmitter doing a low impedance detection changes the common mode voltage on D+ and D- to a different value.~~

- a. A receiver is detected based on the rate that the lines change to the new voltage.
 - i. The receiver is not present if the voltage at the transmitter charges at a rate dictated only by the transmitter impedance, and capacitance of the interconnect, and series capacitor.
 - ii. The receiver is present if the voltage at the transmitter charges at a rate dictated by the transmitter impedance, the series capacitor, the interconnect capacitance, and the receiver termination.

~~The Receiver Detection circuit must correctly detect that receiver terminations are in high impedance (or are disconnected) with a total interconnect capacitance to ground of up to 3 nF, including capacitance of any attached test instrumentation. Note that this capacitance is separate and distinct from the AC Coupling capacitance value (see Section 4.3.1.2). The minimum and maximum AC capacitance for the AC coupling capacitors is given in Table 4-4.~~

~~While in electrical idle the receiver detection must be performed periodically (unless explicitly stated otherwise) to ensure the device is still present. The maximum period between receiver detection sequences is given in Table 4-4: T_{TX IDLE RCV DETECT MAX}.~~

~~Anytime data-Electrical Idle is exited being received by a port this is sufficient to establish the existence of a device. In this case the detect sequence does not have to execute or can be aborted on that Lane.~~

~~If a device is directed to exit Electrical Idle the receiver detection sequence must be immediately aborted.~~

4.3.1.9. ~~Disable/Surprise Removal Detection~~

~~Two separate events may signal that one end of a Link has either been disabled or disconnected.~~

- ~~1. During L0, Recovery, Configuration, or Loopback if Electrical Idle is detected without receiving the Electrical Idle ordered set, the Link immediately enters Detect~~
- ~~2. During Electrical Idle and certain specified times the transmitter must poll for the presence of a powered receiver as described in Section 4.3.1.8. If a receiver is no longer present, the Link immediately enters Detect.~~

4.3.1.10. ~~Electrical Idle~~

Electrical idle is a steady state condition where the Transmitter ~~and Receiver~~ D+ and D- voltages are held constant at the same value. Electrical idle is primarily used in power saving and inactive states (i.e. ~~Disable~~) common mode voltage initialization.

⁵⁵ The common mode being driven ~~must does not~~ to meet the Absolute Delta Between DC Common Mode During L0 and Electrical Idle (V_{TX-CM-DC-ACTIVE-IDLE-DELTA}) specification (see Table 0-4).

⁵⁶ The maximum change in common mode voltage can be no more than V_{TX-RCV-DETECT} in Table 0-4.

Before a transmitter enters Electrical Idle, it must **always** send the Electrical Idle ordered-set, a K28.5 (COM) followed by three K28.3 (IDL) (see Table 0-4), unless otherwise specified. After sending the last symbol of the Electrical Idle ordered-set the transmitter must be in a valid Electrical Idle state as specified by $T_{TX-IDLE-SET-TO-IDLE}$ (see Table 0-4).

The receiver ~~must~~**shall** enter Electrical Idle upon receipt of ~~two~~ **two** K28.3 (IDL) characters ~~this ordered-set to enter electrical idle.~~

The ~~low impedance common mode and differential~~ receiver terminations values (see section 0) must be ~~remain metenabled~~ in Electrical Idle. The transmitter ~~can be in either a low or high impedance mode during~~ **must meet the DC common mode voltage specification while transitioning into and out of** Electrical Idle, ~~which can be done in a low or high impedance state unless explicitly specified.~~

Any time a transmitter enters Electrical Idle it must remain in electrical idle for a minimum of $T_{TX-IDLE-MIN}$ (see Table 0-4). The receiver should expect the Electrical Idle ordered-set followed by a minimum amount of time in Electrical Idle ($T_{TX-IDLE-SET-TO-IDLE}$) to arm its Electrical Idle Exit detector.

Electrical Idle exit occurs when a signal larger than the minimum $V_{RX-IDLE-DET-DIFF_{pp}}$ is detected at a receiver.

~~See Section 4.3.1.9 for additional notes related to Electrical Idle.~~

4.3.2. Electrical Signal Specifications

A Differential Signal is defined by taking the voltage difference between two conductors. In this specification, a differential signal or differential pair is comprised of a voltage on a positive conductor, V_{D+} , and a negative conductor, V_{D-} . The differential voltage (V_{DIFF}) is defined as the difference of the positive conductor voltage and the negative conductor voltage ($V_{DIFF} = V_{D+} - V_{D-}$). The Common Mode Voltage (V_{CM}) is defined as the average or mean voltage present on the same differential pair ($V_{CM} = [V_{D+} + V_{D-}] / 2$). This document's electrical specifications often refer to peak-to-peak measurements or peak measurements, which are defined by the following equations.

$$\frac{2}{3} V_{DIFF_{pp}} = (2 * \max |V_{D+} - V_{D-}|) \text{ (This applies to a symmetric differential swing)}$$

$$\frac{2}{3} V_{DIFF_{pp}} = (\max |V_{D+} - V_{D-}| \{V_{D+} > V_{D-}\} + \max |V_{D+} - V_{D-}| \{V_{D+} < V_{D-}\}) \text{ (This applies to an asymmetric differential swing.)}$$

$$\frac{2}{3} V_{DIFF_p} = (\max |V_{D+} - V_{D-}|) \text{ (This applies to a symmetric differential swing)}$$

$$\frac{2}{3} V_{DIFF_p} = (\max |V_{D+} - V_{D-}| \{V_{D+} > V_{D-}\}) \text{ or } (\max |V_{D+} - V_{D-}| \{V_{D+} < V_{D-}\}) \text{ whichever is greater (This applies to an asymmetric differential swing.)}$$

$$\frac{2}{3} V_{CM_p} = (\max |V_{D+} + V_{D-}| / 2)$$

Note: The maximum value is calculated on a per unit interval evaluation. The maximum function as described is implicit for all peak-to-peak and peak equations throughout the rest of this chapter, and thus a maximum function will not appear in any following subsequent representations of these equations.

In this section (4.3.2), DC is defined as all frequency components below $F_{dc} = 30$ kHz. AC is defined as all frequency components at or above $F_{dc} = 30$ kHz. These definitions pertain to all voltage and current specifications.

An example waveform is shown in Figure 0-18. In this waveform the differential peak-peak signal is approximately 0.6 V, the differential peak signal is approximately 0.3 V and the common mode voltage is approximately 0.25 V.




Figure 0-18: Sample Differential Signal

4.3.2.1 Loss

Loss (attenuation of the differential voltage swing) in this system is a critical parameter that must be properly considered and managed in order to ensure proper system functionality. Failure to properly consider the loss may result in a differential signal swing arriving at the Receiver that does not meet specifications. The interconnect loss is specified in terms of the amount of attenuation or loss that can be tolerated between the Transmitter (Tx) and Receiver (Rx). The Tx is responsible for producing the specified differential eye height at the pins of its package. Together, the Tx and the interconnect are responsible for producing the specified differential eye height at the Rx pins (see Figure 0-24).

The worst-case operational loss budget is calculated by taking the minimum output voltage ($V_{TX-DIFFP-P} = 800 \text{ mV}$) divided by the minimum input voltage to the receiver ($V_{RX-DIFFP-P} = 175 \text{ mV}$), which results in 13.2 dB. Additional headroom in the loss budget can be achieved by driving a larger differential output voltage (up to the maximum specified in Table 0-4) at the transmitter.

4.3.2.2. Jitter and BER

Jitter is categorized into random sources (R_j) and deterministic sources (D_j). The total jitter (T_j) is the convolution of the probability density functions for all the independent jitter sources, R_j and D_j . The nature of R_j can be approximated as Gaussian and is used to establish the bit error rate (BER) of the link.

The UI allocation is given as the allowable T_j at the target BER. The UI allocation must meet a maximum BER of 10^{-12} for the T_j . The allocation to R_j and D_j is not specified.

The methods for measuring the BER compliance are beyond the scope of this specification.

4.3.2.3. De-emphasis

De-emphasis is included (i.e., Generation 1 fundamental band = 250 MHz to 1.25 GHz).

De-emphasis must be implemented when multiple bits of the same polarity are output in succession. Subsequent bits are driven at a differential voltage level 3.5 dB (+/- .5 dB) below the first bit. Note that individual bits, and the first bit from a sequence in which all bits have the same polarity, must always be driven between the Min and Max values as specified by $V_{TX-DIFFp-p}$ in Table 0-4.

The only exception pertains to transmitting the Beacon (see Section 0).

Note: The specified amount of de-emphasis allows for PCI Express designs to optimize maximum interoperability while minimizing complexity of managing configurable de-emphasis values. Thus, the primary benefits of de-emphasis are targeted for worst-case loss budgets of 11-13.2 dB, while being slightly less optimal for lower loss systems. However, this tradeoff is more than offset by the fact that there is inherently more voltage margin in lower loss systems.

An example waveform illustrating de-emphasis and representing the bit sequence (from left to right) of “1001000011” is shown in Figure 0-19.




Figure 0-19: Sample Transmitted Waveform Showing -3.5 dB De-emphasis Around a 0.5 V Common Mode Voltage

Beacon

Support for Beacon is required for all “universal” PCI Express components that support a wakeup mechanism in order to function in form factors that require the use of Beacon. However, not all systems and form factor specifications require the use of Beacon, and for components that are restricted to use in such environments it is not necessary to support Beacon. The following section applies to all components that support Beacon.

The Beacon is a signal sent by a Downstream Component to start the exit from an L2 state.

All transmitter electrical specifications (Table 4-4) must be met while sending a Beacon with the following exceptions and clarifications.

The Beacon is a DC Balanced signal of periodic arbitrary data, which is required to contain some pulse widths $\geq 2\text{ns}$ but no larger than $16\mu\text{s}$.

The maximum time between qualifying pulses ($2\text{ns} \leq x \leq 16\mu\text{s}$) can be no larger than $16\mu\text{s}$.

DC Balance must be always be restored within a maximum time of $32\mu\text{s}$.

Beacon is transmitted in a low impedance mode.

All Beacons must be transmitted and received on at least Lane 0 of multi-Lane Links⁵⁷.

The output Beacon voltage level must be at a -6 dB de-emphasis level for Beacon pulses with a width greater than 500ns .

The output Beacon voltage level can range between the pre-emphasized and corresponding -3.5 dB de-emphasized voltage levels for Beacon pulses smaller than 500ns .

⁵⁷ Lane 0 as defined after Link Width and Lane reversal negotiations are complete.

~~? The Lane-to-Lane Output Skew and SKP ordered set output (see section 0) specifications do not apply.~~

When any bridge and/or a switch receives a Beacon at a downstream Port, that component the switch must propagate a Beacon wakeup indication upstream. This wakeup indication must use the appropriate wakeup mechanism required by the system or form factor associated with the upstream port of the switch (see Section <5.3.3.2>).

~~All transmitter electrical specifications (Table 4-4) must be met while sending a Beacon with the following exceptions and clarifications.~~

~~? The period of the Beacon must be no greater than 33.333 μ s maximum.~~

~~? All Beacons must be transmitted and received on at least Lane 0 of multi-lane links⁵⁸.~~

~~? The Beacon signal must contain minimum width pulses that are greater than or equal to 2 ns.~~

~~? The Beacon must be DC Balanced (i.e., any Beacon must contain an equal number of 1's and 0's).~~

~~? The output Beacon voltage level must be at a -6 dB de-emphasis level for Beacon pulses with a width greater than 500 ns.~~

~~? The output Beacon voltage level can range between the pre-emphasized and corresponding -3.5 dB de-emphasized voltage levels for Beacon pulses smaller than 500 ns.~~

~~? A Receiver Detection sequence (Section 4.3.1.8) must occur every 100 ms, and if no receiver is found then the Link returns to Detect.~~

~~? The Lane-to-Lane Output Skew and Skip Symbol Output specifications do not apply.~~

~~When any bridge and/or switch receives a Beacon, that component must propagate a Beacon upstream.~~

4.3.2.4.1. Beacon Example

An example receiver waveform driven at the -6 dB level for a 30 kHz Beacon is shown in Figure 0-20. An example receiver waveform using the COM character at full speed signaling is shown in Figure 0-21. It should be noted that other waveforms and signaling are possible other than the two examples shown below (i.e., Polling is another valid Beacon signal).

⁵⁸ Lane 0 as defined after Link Width and Lane reversal negotiations are complete.




Figure 0-20: A 30 kHz BEACON Signaling Through a 75 nF Capacitor




Figure 0-21: BEACON, Which Includes a 2 ns Pulse Through a 75 nF Capacitor

4.3.3. Differential Transmitter (TX) Output Specifications

The following table defines the specification of parameters for the differential output at all Transmitters (TXs). The parameters are specified at the component pins.

Table 0-4: Differential Transmitter (TX) Output Specifications

Symbol	Parameter	Min	Nom	Max	Units	Comments
--------	-----------	-----	-----	-----	-------	----------

Symbol	Parameter	Min	Nom	Max	Units	Comments
UI	Unit Interval	399.88	400	400.12	ps	Each UI is 400 ps +/-300 ppm. UI does not account for SSC dictated variations. See Note 1.
$V_{TX-DIFFp-p}$	Differential Peak to Peak Output Voltage	0.800		1.2	V	$V_{TX-DIFFp-p} = 2 * V_{TX-D+} - V_{TX-D-} $ See Note 2.
$V_{TX-DE-RATIO}$	De-Emphasized Differential Output Voltage (Ratio)	-3.0	-3.5	-4.0	dB	This is the ratio of the $V_{TX-DIFFp-p}$ of the second and following bits after a transition divided by the $V_{TX-DIFFp-p}$ of the first bit after a transition. See Note 2.
T_{TX-EYE}	Minimum TX Eye Width	0.70			UI	The maximum transmitter jitter can be derived as $T_{TX-MAX-JITTER} = 1 - T_{TX-EYE} = .3$ UI See Notes 2 and 3.
$T_{TX-EYE-MEDIAN-to-MAX-JITTER}$	Maximum time between the jitter median and maximum deviation from the median.			0.15	UI	Jitter is defined as the measurement variation of the crossing points ($V_{TX-DIFFp-p} = 0$ V) in relation to an appropriate average TX UI. See Note 2 and 3.
$T_{TX-RISE}, T_{TX-FALL}$	D+/D- TX Output Rise/Fall Time	0.125			UI	See Notes 2 and 5.
$V_{TX-CM-ACp}$	AC Peak Common Mode Output Voltage			20	mV	$V_{TX-CM-ACp} = V_{TX-D+} + V_{TX-D-} /2 - V_{TX-CM-DC}$ $V_{TX-CM-DC} = DC_{(avg)}$ of $ V_{TX-D+} + V_{TX-D-} /2$ during L0 See Note 2.
$V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$	Absolute Delta of DC Common Mode Voltage During L0 and Electrical Idle.	0		100	mV	$ V_{TX-CM-DC} [during L0] - V_{TX-CM-Idle-DC} [During Electrical Idle] \leq 100$ mV $V_{TX-CM-DC} = DC_{(avg)}$ of $ V_{TX-D+} + V_{TX-D-} /2$ [L0] $V_{TX-CM-Idle-DC} = DC_{(avg)}$ of $ V_{TX-D+} + V_{TX-D-} /2$ [electrical idle] See Note 2.


Symbol	Parameter	Min	Nom	Max	Units	Comments
$V_{TX-CM-DC-LINE-DELTA}$	Absolute Delta of DC Common Mode Voltage between D+ and D-.	0		25	mV	$ V_{TX-CM-DC-D+}[\text{during L0}] - V_{TX-CM-DC-D-}[\text{During L0.}] \leq 25 \text{ mV}$ $V_{TX-CM-DC-D+} = DC_{(avg)} \text{ of } V_{TX-D+} [\text{during L0}]$ $V_{TX-CM-DC-D-} = DC_{(avg)} \text{ of } V_{TX-D-} [\text{during L0}]$ See Note 2.
$V_{TX-IDLE-DIFFp}$	Electrical Idle Differential Peak Output Voltage	0		20	mV	$V_{TX-IDLE-DIFFp} = V_{TX-Idle-D+} - V_{TX-Idle-D-} \leq 20 \text{ mV}$ See Note 2.
$V_{TX-RCV-DETECT}$	The amount of voltage change allowed during Receiver Detection.			600	mV	The total amount of voltage change that a transmitter can apply to sense whether a low impedance receiver is present. See Section 0.
$V_{TX-DC-CM}$	The TX DC Common Mode Voltage.	0		3.6	V	The allowed DC Common Mode voltage under any conditions. See Section 0.
$I_{TX-SHORT}$	TX Short Circuit Current Limit			90	mA	The total current the transmitter can provide when shorted to its ground.
$T_{TX-IDLE-MIN}$	Minimum time spent in Electrical Idle	50			UI	Minimum time a transmitter must be in electrical idle.
$T_{TX-IDLE-SET-TO-IDLE}$	Maximum time to transition to a valid Electrical Idle after sending an Electrical Idle ordered-set			20	UI	After sending an electrical idle ordered-set, the transmitter must meet all electrical idle specifications within this time.
$T_{TX-IDLE-RCV-DETECT-MAX}$	Maximum time spent in Electrical Idle before initiating a receiver detect sequence.			100	ms	Maximum time spent in Electrical Idle before initiating a receiver detect sequence. See Section 4.3.1.8

Symbol	Parameter	Min	Nom	Max	Units	Comments
$RL_{TX-DIFF}$	Differential Return Loss	12			dB	Measured over 50 MHz to 1.25 GHz See Note 4.
RL_{TX-CM}	Common Mode Return Loss	6			dB	Measured over 50 MHz to 1.25 GHz See Note 4.
$Z_{TX-DIFF-DC}$	DC Differential TX Impedance	80	100	120	Ω	TX DC Differential Mode Low impedance
$Z_{TX-COM-High-IMP-DC}$	Transmitter DC Common Mode High-Impedance State (DC)	405-k		20-k	Ω	Required TX D+ as well as D- DC-DC High-Impedance during all states.
$L_{TX-SKEW}$	Lane-to-Lane Output Skew			500 + 2UI	ps	Static skew between any two Transmitter Lanes within a single Link Transmitter.
C_{TX}	AC Coupling Capacitor	75		200	nF	All transmitters shall be AC coupled. The AC coupling is required either within the media or within the transmitting component itself.
$T_{crosslink}$	Crosslink random timeout	0		1	ms	This random timeout helps resolve conflicts in crosslink configuration by eventually resulting in only one downstream and one upstream Port (See 0).

Notes:

- No test load is necessarily associated with this value.
- Specified at the measurement point into a timing and voltage compliance test load as shown in Figure 0-23 and measured over any 250 consecutive TX UIs. (Also refer to the Transmitter Compliance Eye Diagram as shown in ~~Change this diagram->~~Figure 0-22.)
- A $T_{TX-EYE} = 0.70$ UI provides for a total sum of deterministic and random jitter budget of $T_{TX-JITTER-MAX} = 0.30$ UI for the transmitter collected over any 250 consecutive TX UIs. The $T_{TX-EYE-MEDIAN-to-MAX-JITTER}$ specification ensures a jitter distribution in which the median and the maximum deviation from the median is less than half of the total TX jitter budget collected over any 250 consecutive TX UIs. It should be noted that the median is not the same as the mean. The jitter median describes the point in time where the number of jitter points on either side is approximately equal as opposed to the averaged time value.
- The transmitter input impedance shall result in a differential return loss greater than or equal to 12 dB and a common mode return loss greater than or equal to 6 dB over a frequency range of 50 MHz to 1.25 GHz. This input impedance requirement applies to all valid input levels. The reference impedance for return loss measurements is 50 ohms to ground for both the D+ and D- line (i.e., as measured by a Vector Network Analyzer with 50 ohm probes - see Figure 0-23). Note that the series capacitors C_{TX} is optional for the return loss measurement.
- Measured between 20-80% at Transmitter package pins into a test load as shown in Figure 0-23 for both V_{TX-D+} and V_{TX-D-} .

4.3.3.1. Transmitter Compliance Eye Diagram



Change this diagram->Figure 0-22: Minimum Transmitter Timing and Voltage Output Compliance Specification

There are two eye diagrams that must be met for the transmitter. Both eye diagrams must be aligned in time using the jitter median to locate the center of the eye diagram. The different eye diagrams will differ in voltage depending whether it is a transition bit or a de-emphasized bit. **The exact reduced voltage level of the de-emphasized bit will always be relative to the transition bit.**

The eye diagram must be valid for any 250 consecutive UIs.

An appropriate average TX UI must be used as the interval for measuring the eye diagram.

4.3.3.2. Compliance Test and Measurement Load

The AC timing and voltage parameters must be verified at the measurement point, as specified by the device vendor within 0.2 inches of the package pins, into a test/measurement load shown in Figure 0-23.

Note: The allowance of the measurement point to be within 0.2 inches of the package pins is meant to acknowledge that package/board routing may benefit from D+ and D- not being exactly matched in length at the package pin boundary. If the vendor does not explicitly state where the measurement point is located, the measurement point is assumed to be the D+ and D-package pins.

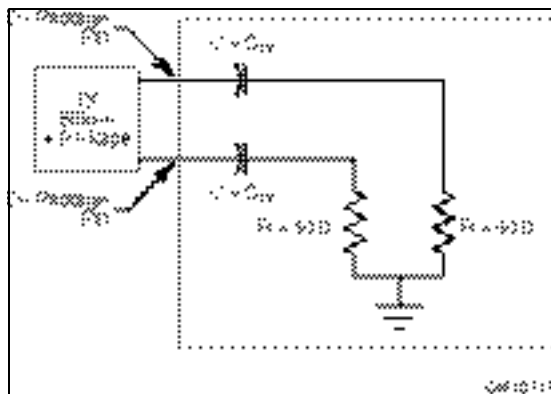


Figure 0-23: Compliance Test/Measurement Load

The test load is shown at the transmitter package reference plane, but the same Test/Measurement load is applicable to the receiver package reference plane.

C_{TX} is an optional portion of the measurement test load. The measurement should be taken on the opposite side of the capacitor from the package, and the value of the C_{TX} must be in the range of 75 nF to 200 nF.

4.3.4. Differential Receiver (RX) Input Specifications

The following table defines the specification of parameters for all differential Receivers (RXs). The parameters are specified at the component pins.

Table 0-5: Differential Receiver (RX) Input Specifications

Symbol	Parameter	Min	Nom	Max	Units	Comments
UI	Unit Interval	399.88	400	400.12	ps	The UI is 400 ps +/-300 ppm. UI does not account for SSC dictated variations. See Note 6.
$V_{RX-DIFFp-p}$	Differential Input Peak to Peak Voltage	0.175		1.200	V	$V_{RX-DIFFp-p} = 2 * V_{RX-D+} - V_{RX-D-} $ See Note 7.
T_{RX-EYE}	Minimum Receiver Eye Width	0.4			UI	The maximum interconnect media and transmitter jitter that can be tolerated by the receiver can be derived as $T_{RX-MAX-JITTER} = 1 - T_{RX-EYE} = 0.6$ UI See Notes 7 and 8.
$T_{RX-EYE-MEDIAN-to-MAX-JITTER}$	Maximum time between the jitter median and maximum deviation from the median.			0.3	UI	Jitter is defined as the measurement variation of the crossing points ($V_{RX-DIFFp-p} = 0$ V) in relation to an appropriate average TX UI. See Notes 7 and 8.

Symbol	Parameter	Min	Nom	Max	Units	Comments
$V_{RX-CM-ACp}$	RMS AC Peak Common Mode Input Voltage			150	mV	$V_{RX-CM-AC} = V_{RX-D+} + V_{RX-D-} /2 - V_{RX-CM-DC}$ $V_{RX-CM-DC} = DC_{(avg)}$ of $ V_{RX-D+} + V_{RX-D-} /2$ during L0 See Note 7.
$RL_{RX-DIFF}$	Differential Return Loss	15			dB	Measured over 50 MHz to 1.25 GHz with the D+ and D- lines biased at +300mV and -300mV respectively. See Note 9.
RL_{RX-CM}	Common Mode Return Loss	6			dB	Measured over 50 MHz to 1.25 GHz with the D+ and D- lines biased at 0V See Note 9
$Z_{RX-DIFF-DC}$	DC Differential Input Impedance	80	100	120	Ω	RX DC Differential Mode impedance. See Note 10.
$Z_{RX-COM-DC}$	DC Input Common Mode Input Impedance	40	50	60	Ω	Required RX D+ as well as D- DC Common Mode impedance. (50 Ω +/-20% tolerance). See Notes 7 and 10.
$Z_{RX-COM-INITIAL-DC}$	Initial DC Input Common Mode Input Impedance	5	50	60	Ω	RX DC Common Mode impedance allowed when the receiver terminations are first enabled power on . The enabling time between the first and the last the terminations must occur a within 5 ms. See Note 11.
$Z_{RX-COM-HIGH-IMP-DC}$	Powered Down DC Input Common Mode Input Impedance	200 k			Ω	Required RX D+ as well as D- DC Impedance RX-DC Common Mode impedance when the receiver terminations do not have are not power. ed (i.e., no power) . See Note 12.
$V_{RX-IDLE-DET-DIFFp-p}$	Electrical Idle Detect Threshold	65		175	mV	$V_{RX-IDLE-DET-DIFFp-p} = 2 * V_{RX-D+} - V_{RX-D-} $ Measured at the package pins of the Receiver.

Symbol	Parameter	Min	Nom	Max	Units	Comments
$T_{RX-IDLE-DET-DIFF-ENTERTIME}$	Unexpected Electrical Idle Enter Detect Threshold Integration Time			10	ms	An unexpected electrical idle ($V_{RX-DIFF-P} < V_{RX-IDLE-DET-DIFF-P}$) must be recognized no longer than $T_{RX-IDLE-DET-DIFF-ENTERTIME}$ to signal an unexpected idle condition.
$L_{RX-SKEW}$	Total Skew			20	ns	Skew across all Lanes on a port Link. This includes variation in the length of a skip ordered-set (e.g., COM and 1 to 5 SKP symbols) at the RX as well as any delay differences arising from the interconnect itself.

Notes:

6. No test load is necessarily associated with this value.
7. Specified at the measurement point and measured over any 250 consecutive UIs. The test load in Figure 0-23 should be used as the RX device when taking measurements (also refer to the Receiver Compliance Eye Diagram as shown in Figure 0-24). If the clocks to the RX and TX are not derived from the same clock chip the TX UI must be used as a reference for the eye diagram.
8. A $T_{RX-EYE} = 0.40$ UI provides for a total sum of 0.60 UI deterministic and random jitter budget for the transmitter and interconnect collected any 250 consecutive UIs. The $T_{RX-EYE-MEDIAN-TO-MAX-JITTER}$ specification ensures a jitter distribution in which the median and the maximum deviation from the median is less than half of the total .6 UI jitter budget collected over any 250 consecutive TX UIs. It should be noted that the median is not the same as the mean. The jitter median describes the point in time where the number of jitter points on either side is approximately equal as opposed to the averaged time value. If the clocks to the RX and TX are not derived from the same clock chip, the appropriate average TX UI must be used as the reference for the eye diagram.
9. The receiver input impedance shall result in a differential return loss greater than or equal to 15 dB with the D+ line biased to 300mV and the D- line biased to -300mV and a common mode return loss greater than or equal to 6 dB (no bias required) over a frequency range of 50 MHz to 1.25 GHz. This input impedance requirement applies to all valid input levels. The reference impedance for return loss measurements for is 50 ohms to ground for both the D+ and D- line (i.e., as measured by a Vector Network Analyzer with 50 ohm probes - see Figure 0-23). Note: that the series capacitors C_{TX} is optional for the return loss measurement.
10. Impedance during all operating conditions.
11. The Rx DC Common Mode Impedance that must be present when the receiver terminations are first enabled to ensure that the Receiver Detect occurs properly. The enabling time between the first and the last the terminations must occur within 5ms. Compensation of this impedance can start immediately and the ($Z_{RX-COM-DC}$) Rx DC Common Mode Impedance must be within the specified range by the time Detect is entered.
12. The Rx DC Common Mode Impedance that exists when the receiver terminations are disabled or when no power is present. This helps ensure that the Receiver Detect circuit will not falsely assume a receiver is powered on when it is not. This term must be measured at 300mV above the RX ground.




Figure 0-24: Minimum Receiver Eye Timing and Voltage Compliance Specification

The RX eye diagram must be aligned in time using the jitter median to locate the center of the eye diagram.

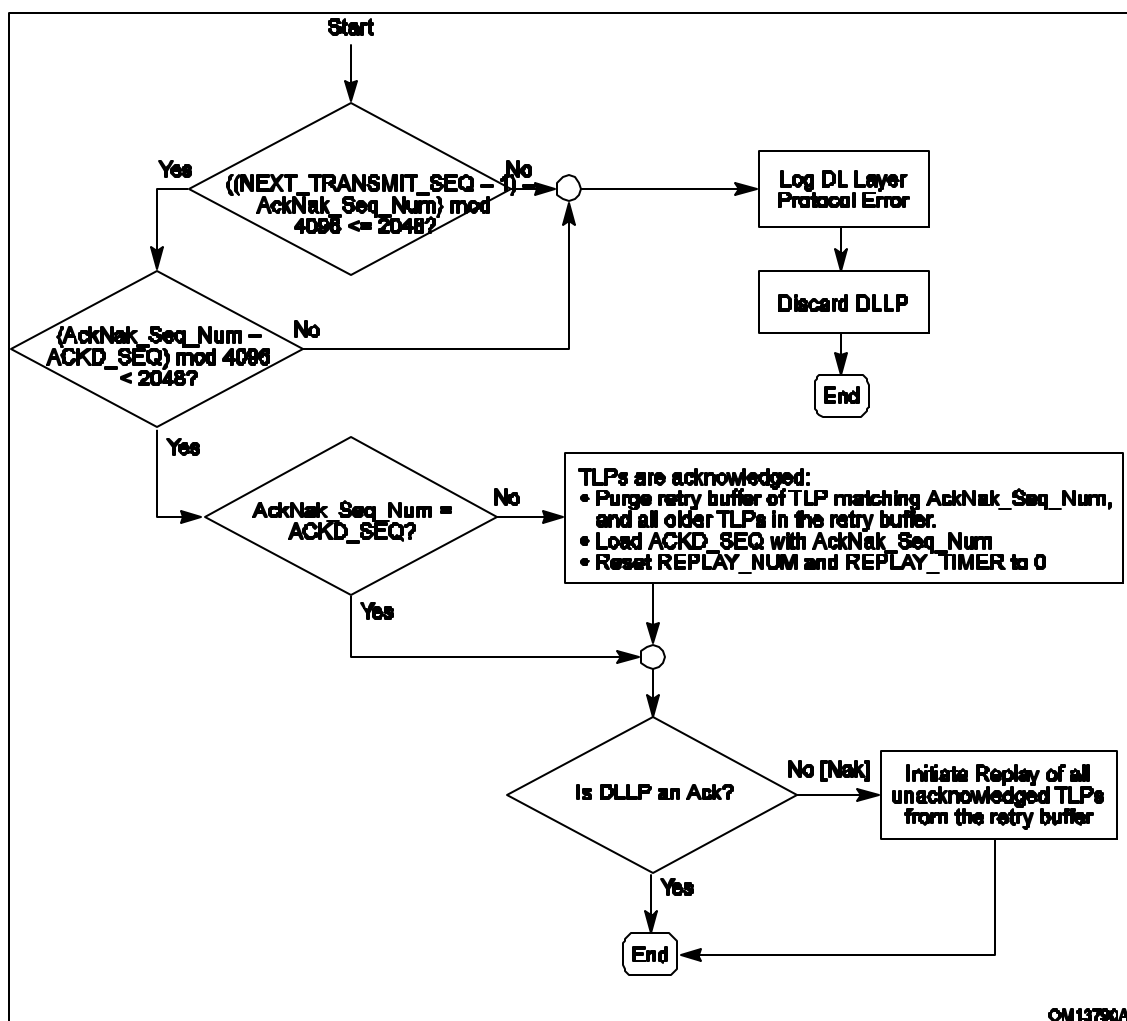
The eye diagram must be valid for any 250 consecutive UIs.

An appropriate average TX UI must be used as the interval for measuring the eye diagram.

Appendix B – New and Redrawn Figures

This appendix includes all available new and redrawn figures, as called for in the body of this document.

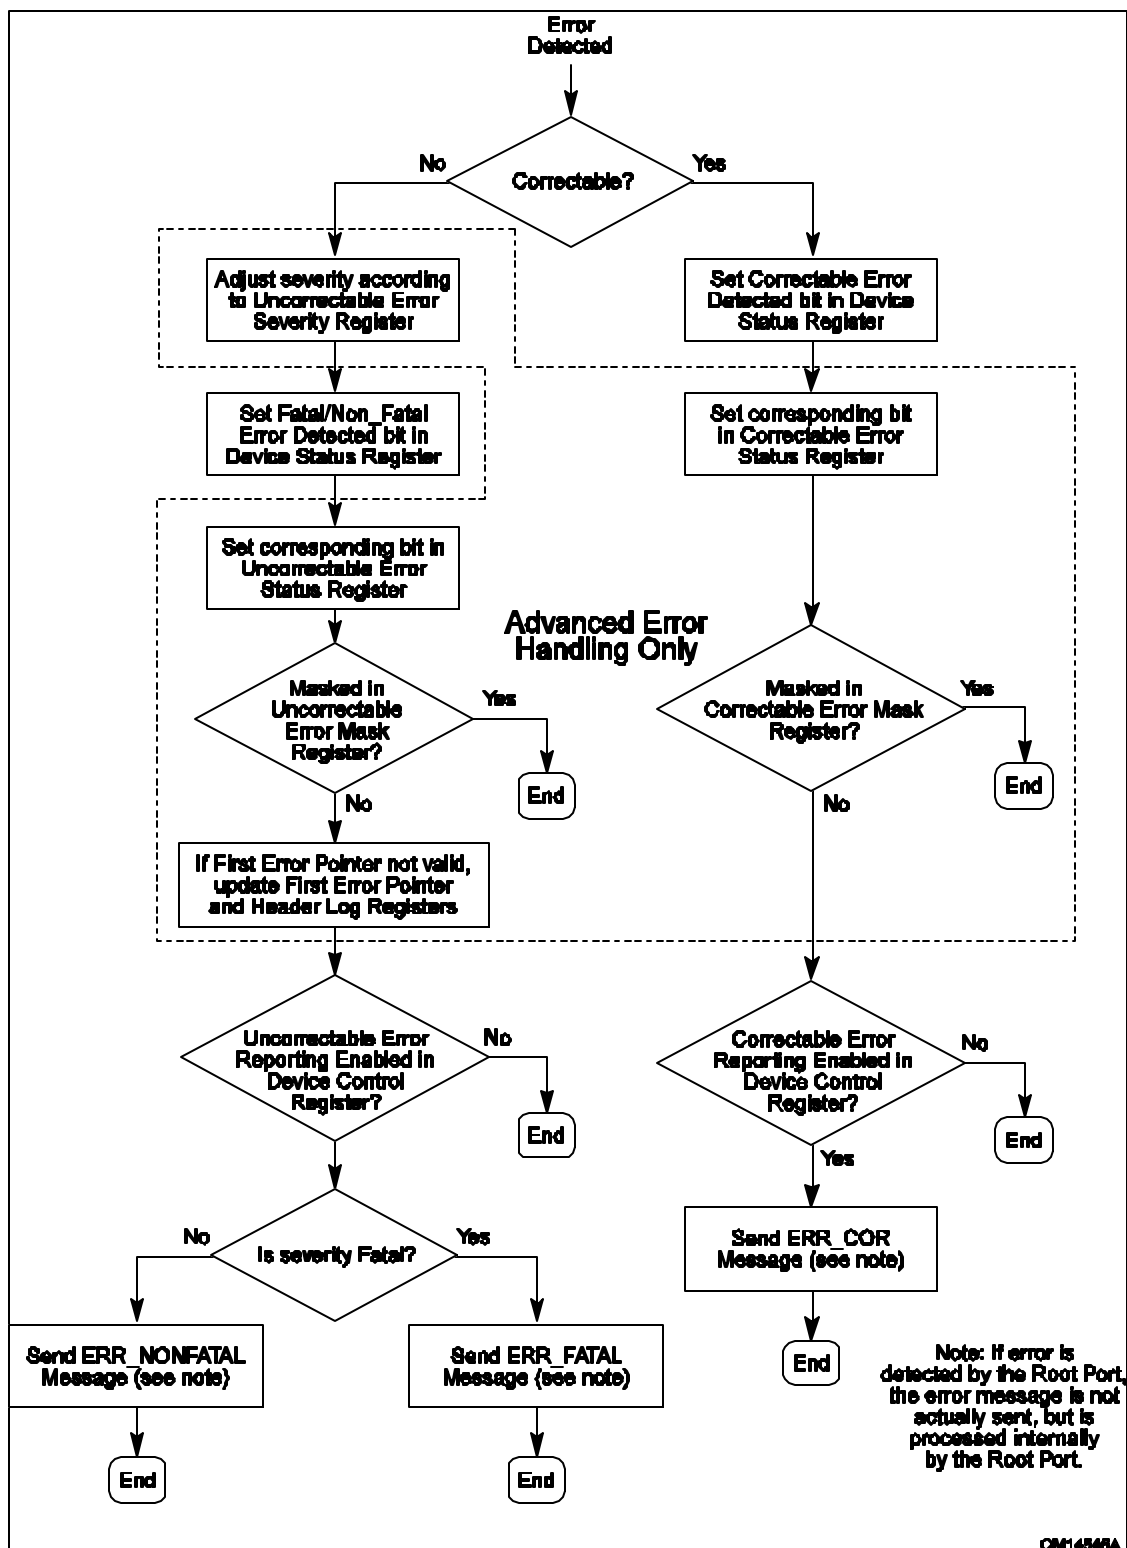
Note: At this time, not all figures have been created/redrawn.




Revised Figure 3-17

31	24	23	16	15	8	7	0
Header Log Register (1st DW)							
Header Byte 0		Header Byte 1		Header Byte 2		Header Byte 3	
Header Log Register (2nd DW)							
Header Byte 4		Header Byte 5		Header Byte 6		Header Byte 7	
Header Log Register (3rd DW)							
Header Byte 8		Header Byte 9		Header Byte 10		Header Byte 11	
Header Log Register (4th DW)							
Header Byte 12		Header Byte 13		Header Byte 14		Header Byte 15	
DM1456							

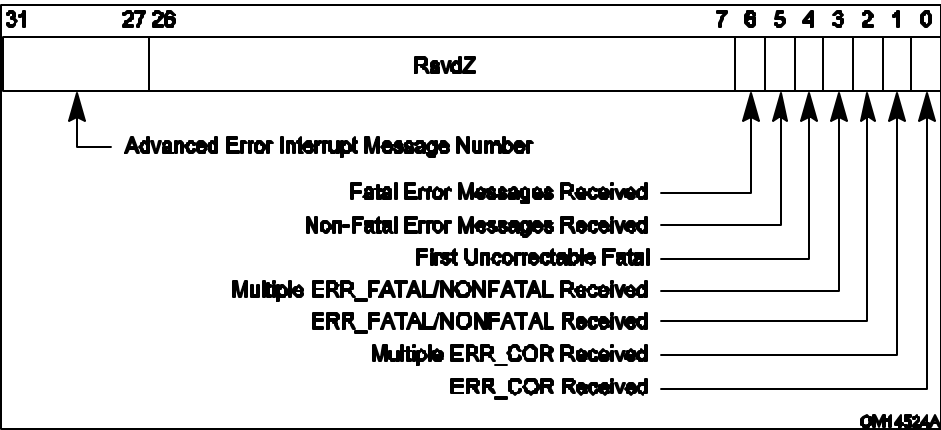
Revised Figure 7-34




Revised Figure 6-2



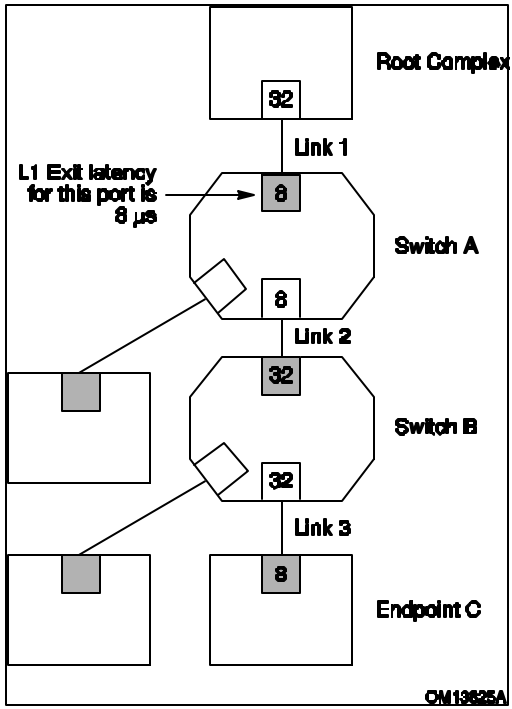
Revised Figure 7-37



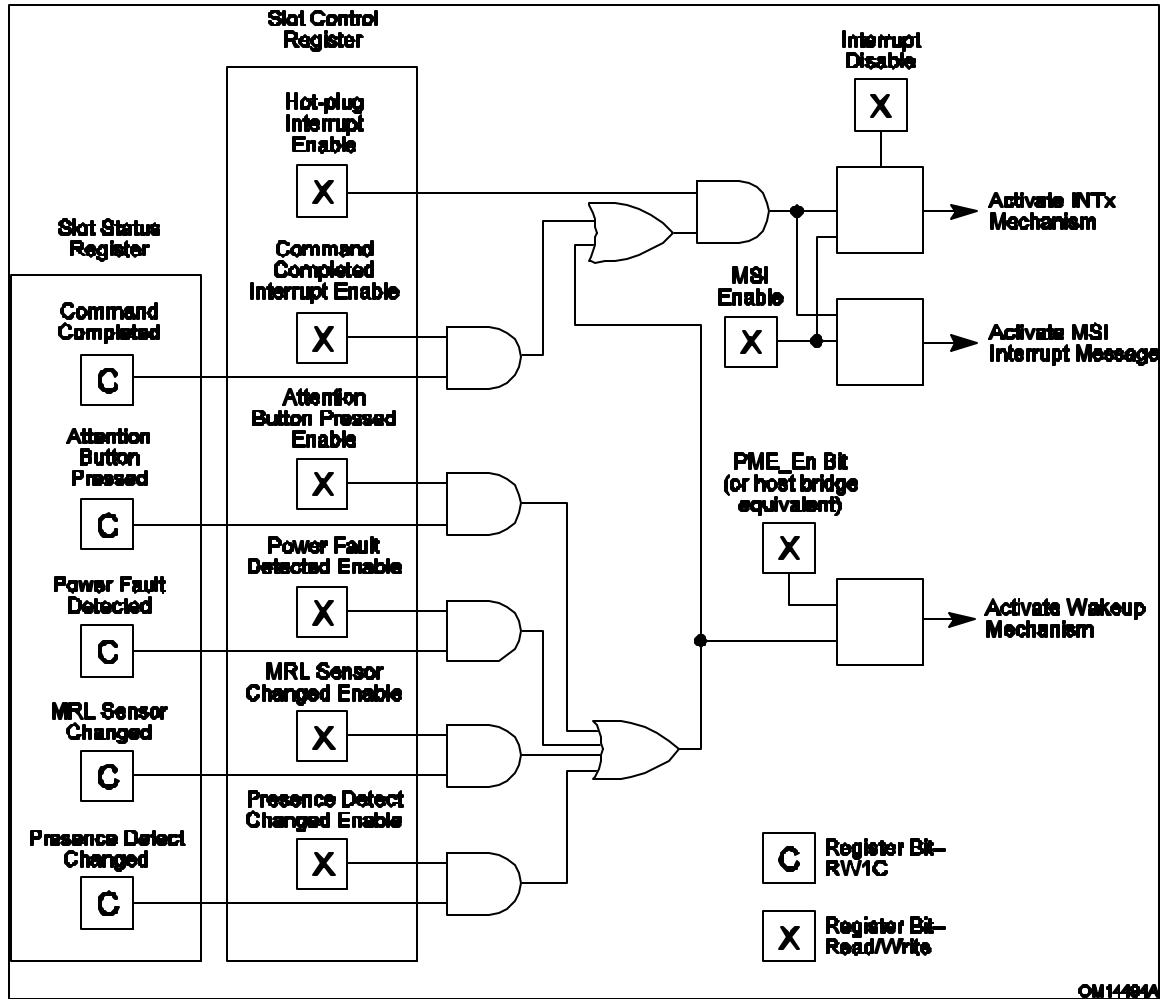
Revised Figure 7-36



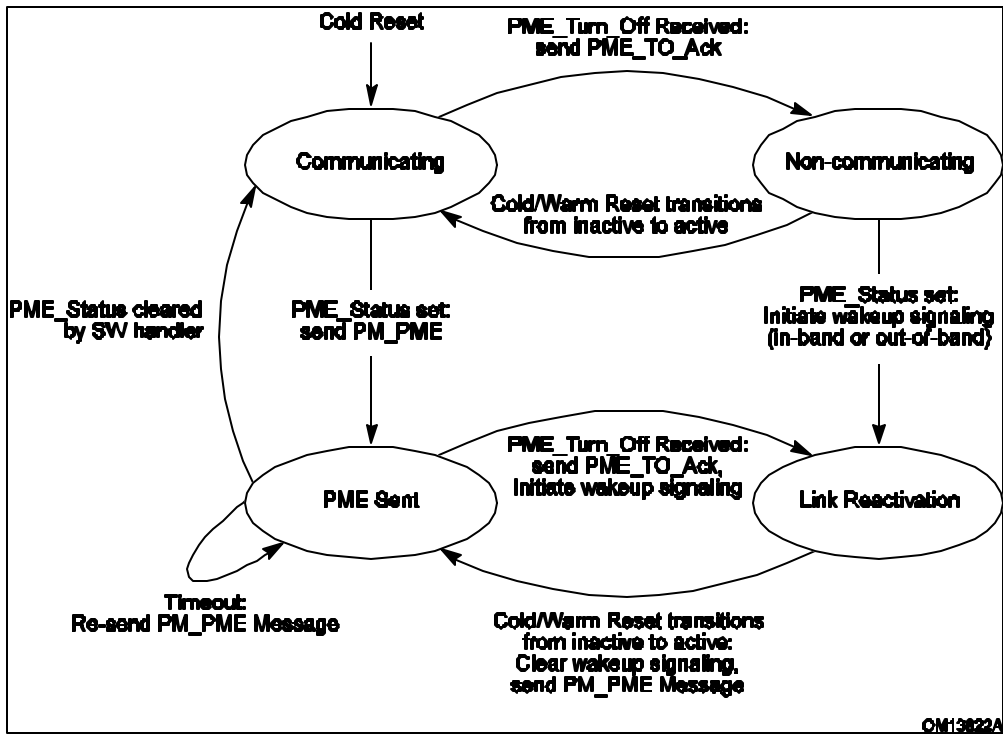
Revised Figure 7-1



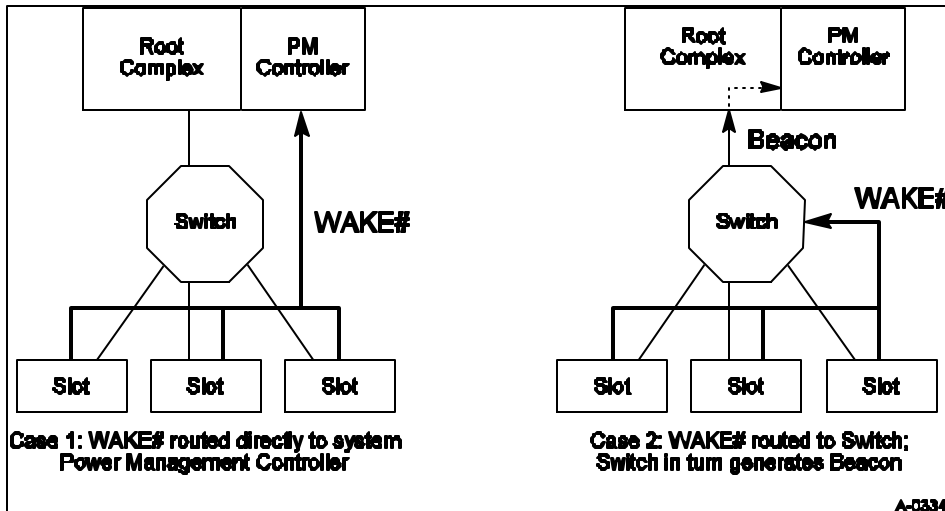
Revised Figure 5-7



Revised Figure 6-9



Revised Figure 5-4



New Figure