

Errata for the
PCI Express[®] Base Specification
Revision 2.1

November 18, 2010



Errata for the PCI Express Base Specification, Revision 2.1

Revision	Revision History	Date
	Initial release.	06/08/2010
	Added errata A14-A21, B40-B91, and C1-C2.	11/18/2010

PCI-SIG® disclaims all warranties and liability for the use of this document and the information contained herein and assumes no responsibility for any errors that may appear in this document, nor does PCI-SIG make a commitment to update the information contained herein.

Contact the PCI-SIG office to obtain the latest revision of this specification.

Questions regarding the PCI Express Base Specification or membership in PCI-SIG may be forwarded to:

Membership Services

www.pcisig.com

E-mail: administration@pcisig.com

Phone: 503-619-0569

Fax: 503-644-6708

Technical Support

techsupp@pcisig.com

DISCLAIMER

This PCI Express Base Specification Errata document is provided “as is” with no warranties whatsoever, including any warranty of merchantability, noninfringement, fitness for any particular purpose, or any warranty otherwise arising out of any proposal, specification, or sample. PCI-SIG disclaims all liability for infringement of proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

PCI, PCI Express, PCIe, and PCI-SIG are trademarks or registered trademarks of PCI-SIG.
All other product names are trademarks, registered trademarks, or servicemarks of their respective owners.

Copyright © 2010 PCI-SIG

Contents

A1	Rules for Use of Data Poisoning	6
A2	Memory, I/O, and Configuration Request Rules	7
A3	Vendor_Defined Messages	7
A4	Error Message Forwarding and PCI Mapping for Bridge – Rules.....	8
A5	Configuration.Linkwidth.Start	8
A6	Capability Version	9
A7	Uncorrectable Error Severity Register	9
A8	Rules for Transmitters and Receivers.....	10
A9	Transaction Descriptor – Transaction ID Field.....	10
A10	Latency Tolerance Reporting (LTR) Mechanism	11
A11	Enable Clock Power Management.....	12
A12	Request Handling Rules	13
A13	Prefetchable BARs	13
A14	Memory, I/O, and Configuration Request Rules	15
A15	Conventional Reset.....	15
A16	Multicast	16
A17	ASPM Configuration.....	17
A18	Power Management Messages.....	18
A19	Multicast Capability	18
A20	Advanced Error Reporting Capability.....	19
A21	ACS Capability Register.....	19
B1	Latency Tolerance Reporting(LTR) Mechanism	21
B2	Mapping Between PCI 3.0 and PCI Express Registers	21
B3	IDO Completion Enable	22
B4	Physical Layer Services	22
B5	Detect.Active	22
B6	Terms and Acronyms	23
B7	Address Based Routing Rules and Memory, I/O, and Configuration Request Rules...	29
B8	First/Last DW Byte Enables Rules	30
B9	Memory, I/O, and Configuration Request Rules	30
B10	End-End TLP Prefix Processing.....	30
B11	ST Modes of Operation	30
B12	Latency Tolerance Mechanism	31
B13	ID Based Routing Rules	31
B14	ST Modes	31
B15	Latency Tolerance Mechanism	32
B16	Device Capabilities 2 Register	33
B17	ACS Control Register	33
B18	Latency Tolerance Reporting Capability	34
B19	TPH Requester Capability Structure	34
B20	TPH Requester Capability Register—ST Table Location	34

Errata for the PCI Express Base Specification, Revision 2.1

B21	TPH Requester Capability Register—ST Table Size	35
B22	TPH Requester Control Register—ST Model Select	35
B23	TPH Requester Control Register—TPH Requester Enable.....	36
B24	TPH ST Table.....	36
B25	TPH ST Table—ST Upper and ST Lower	36
B26	Common Packet Header Fields	37
B27	Compliance Pattern.....	37
B28	Modified Compliance Pattern	38
B29	Polling.Compliance.....	39
B30	OBFF ECN Update [affects PCI Express CEM Specification, Revision 2.0]	40
B31	MC_Base_Address Register	40
B32	D3 State	41
B33	Trusted Configuration Space	41
B34	Latency Tolerance Reporting Mechanism	41
B35	Latency Tolerance Reporting Mechanism	42
B36	Additional References to Tables	42
B37	Capitalizing “Reserved”	43
B38	L0.....	43
B39	The Set of Lanes in Polling.Active	44
B40	Completion Rules	45
B41	Completion Handling Rules.....	45
B42	Flow Control Initialization Protocol.....	45
B43	Reduce Swing Corrections.....	46
B44	Transmitter De-emphasis.....	48
B45	Transmitter and Receiver Termination.....	48
B46	Procedure for Channel Measurement and Margin Extraction.....	48
B47	Data Register	49
B48	Flow Control Rules	49
B49	Receiver Loopback.....	49
B50	Optional Errors	50
B51	Link Status 2 Register	52
B52	Multicast Capability Register.....	52
B53	Multicast TLP Processing.....	53
B54	Data Poisoning	53
B55	LCRC and Sequence Number Rules	53
B56	Device Capabilites 2 Register	54
B57	Link Control 2 Register.....	55
B58	I/O Transactions	56
B59	First/Last DW Byte Enables Rules	56
B60	Request Handling Rules	56
B61	Flow Control Rules	57
B62	Data Link Control and Management State Machine Rules	57
B63	LCRC and Sequence Number Rules (TLP Transmitter).....	58
B64	Transmitter De-emphasis.....	58

Errata for the PCI Express Base Specification, Revision 2.1

B65	Receiver Specification.....	59
B66	Link Training and Status State Machine (LTSSM) Descriptions.....	59
B67	Link Training and Status State Rules.....	59
B68	Recovery.Speed and Loopback.Entry.....	60
B69	Device Capabilities Register.....	60
B70	Device Control Register.....	61
B71	Link Status Register.....	61
B72	Slot Capabilities Register.....	62
B73	Slot Control Register.....	63
B74	Link Address for Link Type 1.....	64
B75	Data Select Register/Data Register.....	65
B76	Various Register Corrections.....	66
B77	Egress Control Vector.....	67
B78	Link State Power Management.....	67
B79	Entry into the L0s State.....	68
B80	L1 ASPM State.....	68
B81	ASPM Configuration.....	68
B82	Root Complex Event Collector.....	68
B83	Power Controller.....	70
B84	Link Speed Management.....	71
B85	Configuration Register Types.....	71
B86	Cache Line Size Register.....	72
B87	Base Address Registers.....	72
B88	Device Capabilities Register.....	72
B89	Link Control Register.....	73
B90	Root Complex Event Collector.....	73
B91	Slot Power Limit Value.....	74
C1	Scrambling Spectrum Clarification.....	75
C2	Capitalization Corrections.....	75

A1 Rules for Use of Data Poisoning

In Section 2.7.2.2, page 147, lines 18-38, make the following edits:

- ~~❑ Receipt of a Poisoned TLP is a reported error associated with the Receiving Function (see Section 6.2)~~
- ~~❑ A Poisoned Configuration Write Request must be discarded by the Completer, and a Completion with a Completion Status of UR is returned (see Section 2.2.9).~~
 - ~~● A Switch must route a Poisoned Configuration Write Request in the same way it would route a non-Poisoned Request, unless the Request addresses the Configuration Space of the Switch itself, in which case the Switch is the Completer for the Request and must follow the above rule.~~
- ~~❑ A Poisoned I/O or Memory Write Request, or a Message with data (except for vendor-defined Messages), that addresses a control register or control structure in the Completer must be handled as an Unsupported Request (UR) by the Completer (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 the same Request if it were not poisoned, unless the Request addresses a control register or control structure of the Switch itself, in which case the Switch is the Completer for the Request and must follow the above rule.~~
- ~~❑ Unless there's a higher precedence error, a Completer must handle a Poisoned AtomicOp Request as a Poisoned TLP Received error, and must also return a Completion with a Completion Status of Unsupported Request (UR). See Sections 6.2.3.2.3, 6.2.3.2.4, and 6.2.5. The value of the target location must remain unchanged.~~
 - ~~● A Switch must route a Poisoned AtomicOp Request the same way it would route the same Request if it were not poisoned, unless the Request targets a Memory Space location in the Switch itself, in which case the Switch is the Completer for the Request and must follow the above rule.~~
- ~~❑ The following Requests with Poisoned data must not modify the value of the target location:~~
 - ~~● Configuration Write Request~~
 - ~~● Any of the following that target a control register or control structure in the Completer: I/O Write Request, Memory Write Request, or non-vendor-defined Message with data~~
 - ~~● AtomicOp Request~~

~~Unless there is a higher precedence error, a Completer must handle these Requests as a Poisoned TLP Received error^[footnote] and the Completer must also return a Completion with a Completion Status of Unsupported Request (UR) if the Request is Non-Posted. See Section 6.2.3.2.3, Section 6.2.3.2.4, and Section 6.2.5.~~

~~A Switch must route the Request the same way it would route the same Request if it were not Poisoned, unless the Request targets a location in the Switch itself, in which case the Switch is the Completer for the Request and must follow the above rules.~~

^[footnote] ~~Due to ambiguous language in earlier versions of this specification, a component is permitted to handle this error as an Unsupported Request, but this is strongly discouraged~~

A2 Memory, I/O, and Configuration Request Rules

In Section 2.2.7, page 76, line 23, add the following text:

- ❑ Requests must not specify an Address/Length combination which causes a Memory Space access to cross a 4-KB 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)
 - For AtomicOp Requests, the mandatory Completer check for natural alignment of the Address (see above) already guarantees that the access will not cross a 4-KB boundary so a separate 4-KB boundary check is not necessary.
 - Note that if a 4-KB boundary check is performed for AtomicOp CAS Requests, this check must comprehend that the TLP Length value is based on the size of two operands, whereas the access to Memory Space is based on the size of one operand.

A3 Vendor_Defined Messages

In Section 2.2.8.6, page 92, lines 10 and 12, make the following edits:

Receivers Completers silently discard Vendor_Defined Type 1 Messages which they are not designed to receive – this is not an error condition.

Receivers Completers handle the receipt of an unsupported Vendor_Defined Type 0 Message as an Unsupported Request, and the error is reported according to Section 6.2.

In Section 2.9.1, page 150, line 2, make the following edits:

- ❑ The Port must terminate any PME Turn_Off handshake Requests targeting the Port in such a way that the Port is considered to have acknowledged the PME_Turn_Off request (see the Implementation Note in Section 5.3.3.2.1).
- ❑ The Port must handle Vendor Defined Message Requests as described in Section 2.2.8.6 (e.g., silently discard Vendor Defined Type 1 Messages Requests that it is not designed to receive) since the DL_Down prevents the Request from reaching its targeted Function.
- ❑ for all other Posted Requests, discarding the Requests
 - This is a reported error associated with the Function for the (virtual) Bridge associated with the Port (see Section 6.2), and must be reported as an Unsupported Request. For Root Ports, the reporting of this error is optional.
 - For a Posted Request already being processed by the Transaction Layer, the Port is permitted not to report ~~it~~ the error.

A4 Error Message Forwarding and PCI Mapping for Bridge – Rules

In Section 6.2.8.1, page 385, Line 8, make the following edits:

...

- Poisoned TLPs are forwarded according to the same rules as non-Poisoned TLPs

...

- When forwarding a Poisoned Completion Downstream:
 - ◆ Set the Detected Parity Error bit in the Status register
 - ◆ Set the Master Data Parity Error bit in the ~~Secondary~~ Status register if the Parity Error Response ~~Enable~~ bit in the ~~Bridge Control Command~~ register is set

A5 Configuration.Linkwidth.Start

In Section 4.3.6.3.1.1, page 222, Line 11, make the following edit:

- In the optional case where a crosslink is supported, the next state is Disable after all Lanes that ~~detected a Receiver during Detect~~ are transmitting TS1 Ordered Sets receive two consecutive TS1 Ordered Sets with the Disable Link bit asserted.

In Section 4.3.6.3.1.1, page 223, Line 1, make the following edit:

- Next state is Loopback after all Lanes that ~~detected a Receiver during Detect~~ are transmitting TS1 Ordered Sets, that are also receiving Ordered Sets, receive the Loopback bit asserted in two consecutive TS1 Ordered Sets.

In Section 4.3.6.3.1.2, page 225, Lines 5, 8, and 10, make the following edits:

- Next state is Disable after any Lanes that ~~detected a Receiver during Detect~~ are transmitting TS1 Ordered Sets, receive two consecutive TS1 Ordered Sets with the Disable Link bit asserted.
 - 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~~ are transmitting TS1 Ordered Sets, that are also receiving TS1 Ordered Sets, receive the Disable Link bit asserted in two consecutive TS1 Ordered Sets.
- Next state is Loopback after all Lanes that ~~detected a Receiver during Detect~~ are transmitting TS1 Ordered Sets, that are also receiving TS1 Ordered Sets, receive the Loopback bit asserted in two consecutive TS1 Ordered Sets.

A6 Capability Version

In Section 7.10.1, page 556, Table 7-29, Bits 19:16, make the following edit:

Bit Location	Register Description	Attributes
19:16	<p>Capability Version – This field is a PCI-SIG defined version number that indicates the version of the Capability structure present.</p> <p>Must be 2h for this version of the specification. This field must be 2h if the End-End TLP Prefix Supported bit (see Section 7.8.15) is Set and must be 1h or 2h otherwise.</p>	RO

In Section 7.10.12, page 571, line 6 add the following text:

The TLP Prefix Log Registers beyond the number supported by the Function are hardwired to zero. For example, if a Functions, Max End-End TLP Prefixes field contains 10b (indicating 2 DW of buffering) then the third and fourth TLP Prefix Log Registers are hardwired to zero. ~~If the End-End TLP Prefix Supported bit (see Section 7.8.15) is Clear, the TLP Prefix Log register is not required to be implemented.~~

A7 Uncorrectable Error Severity Register

In Section 7.10.4, page 560, line 5, make the following edit:

The Uncorrectable Error Severity register controls whether an individual error is reported as a Nonfatal or Fatal error. An error is reported as fatal when the corresponding error bit in the severity register is Set. If the bit is Clear, the corresponding error is considered non-fatal. Refer to Section 6.2 for further details. Register fields for bits not implemented by the Function are hardwired to ~~the specified default value an implementation specific value.~~

A8 Rules for Transmitters and Receivers

In Section 4.2.7.1, page 259, line 9, make the following edit:

- ❑ Any and all time spent in ~~any lower power Link state (L0s, L1, L2) or in any other a~~ state when the Transmitter is electrically idle ~~(such as Detect and Recovery Speed)~~ does not count in ~~(and may optionally reset)~~ the 1180 to 1538 Symbol Time interval used to schedule the transmission of SKP ordered sets.
- ❑ ~~During all lower power Link states any counter(s) or other mechanisms used to schedule SKP Ordered Sets must be reset. It is recommended that any counter(s) or other mechanisms used to schedule SKP Ordered Sets be reset any time when the Transmitter is electrically idle.~~

In Section 4.2.7.2, page 259, line 21, add the following text:

- ❑ Receivers shall be tolerant to receive and process SKP Ordered Sets at an average interval between 1180 to 1538 Symbol Times.
 - Note: Since Transmitters in electrical idle are not required to reset their mechanism for time-based scheduling of SKP Ordered Sets, Receivers shall be tolerant to receive the first time-scheduled SKP Ordered Set following electrical idle in less than the average time interval between SKP Ordered Sets.

A9 Transaction Descriptor – Transaction ID Field

In Section 2.2.6.2, page 71, lines 3 and 10, make the following edits:

- ❑ Tag[7:0] is ~~an~~ 8-bit field generated by each Requester, and it must be unique for all outstanding Requests that require a Completion for that Requester
 - Receivers/Completers must handle correctly 8-bit Tag values regardless of the setting of their Extended Tag Field Enable bit. Refer to the PCI Express to PCI/PCI-X Bridge Specification for details on the bridge handling of extended tags.
 - If the Extended Tag Field Enable bit (see Section 7.8.4) is clear, the maximum number of outstanding Requests per Function shall be limited to 32, and only the lower 5 bits of the Tag field are used with the remaining upper 3 bits required to be 000b
 - If the Extended Tag Field Enable bit (see Section 7.8.4) is set, the maximum is increased to 256, and the entire Tag field is used
 - The initial value of the Extended Tag Field Enable bit (see Section 7.8.4) is implementation specific
 - Receiver/Completer behavior is undefined if multiple uncompleted Requests are issued non-unique Tag values

A10 Latency Tolerance Reporting (LTR) Mechanism

In Section 6.18, page 461, lines 11, 21, and 25, make the following edits:

When directed to a non-D0 state by a Write to the PMCSR register, if a device's ~~had most recently previously transmitted LTR message (since the last DL_Down to DL_Up transition)~~ reported one or both latency fields with ~~the any~~ Requirement bit set, ~~then~~ it must send a new LTR Message with both ~~of the~~ Requirement bits clear prior to transitioning to the non-D0 state.

When the LTR Mechanism Enable bit is cleared, if a device's most recently sent LTR Message (since the last DL_DOWN to DL_Up transition) reported latency tolerance values with any Requirement bit set, then it must send a new LTR Message with all the Requirement bits clear.

An LTR Message from a device reflects the tolerable latency from the perspective of the device, for which the platform must consider the service latency itself, plus the delay added by the use of Clock Power Management (CLKREQ#), if applicable. The service latency itself is defined as follows:

- ❑ If the device issues a Read Request, latency is measured as delay from transmission of the END symbol in the Request TLP to the receipt of the STP symbol in the first Completion TLP ~~for that Request~~.
- ❑ If the device issues one or more Write Requests such that it cannot issue another Write Request due to Flow Control backpressure, the latency is measured from the transmission of the END symbol of the TLP that exhausts the FC credits to the receipt of the SDP symbol of the DLLP returning more credits ~~for that Request type~~.

A11 Enable Clock Power Management

In Section 7.8.7, page 521, Table 7-16, Bit 8, make the following edit:

Bit Location	Register Description	Attributes
8	<p>Enable Clock Power Management – Applicable only for Upstream Ports and with form factors that support a “Clock Request” (CLKREQ#) mechanism, this bit operates as follows:</p> <p>0b Clock power management is disabled and device must hold CLKREQ# signal low.</p> <p>1b When this bit is Set, the device is permitted to use CLKREQ# signal to power manage Link clock according to protocol defined in appropriate form factor specification.</p> <p>For a non-ARI multi-Function device, power-management-configuration software must only Set this bit if all Functions of the multi-Function device indicate a 1b in the Clock Power Management bit of the Link Capabilities register. The component is permitted to use the CLKREQ# signal to power manage Link clock only if this bit is Set for all Functions.</p> <p>For ARI Devices, Clock Power Management is enabled solely by the setting in Function 0. The settings in the other Functions always return whatever value software programmed for each, but otherwise are ignored by the component.</p> <p>Downstream Ports and components that do not support Clock Power Management (as indicated by a 0b value in the Clock Power Management bit of the Link Capabilities register) must hardwire this bit to 0b.</p> <p>Default value of this bit is 0b, <u>unless specified otherwise by the form factor specification.</u></p>	RW

A12 Request Handling Rules

In Section 2.3.1, page 107, lines 15-19, reformat the text in as shown in red:

- For Configuration Requests only, following reset it is possible for a device to terminate the request but indicate that it is temporarily unable to process the Request, but will be able to process the Request in the future – in this case, the Configuration Request Retry Status (CRS) Completion Status is used (see Section 6.6). Valid reset conditions after which a device is permitted to return CRS are:
 - ◆ Cold, Warm, and Hot Link Resets
 - ◆ FLRs
 - ◆ A reset initiated in response to a $D3_{hot}$ to $D0_{uninitialized}$ device state transition.

A device Function is explicitly not permitted to return CRS following a software-initiated reset (other than an FLR) of the device, e.g., by the device's software driver writing to a device-specific reset bit. Additionally, a device Function is not permitted to return CRS after having previously returned a Successful Completion without an intervening valid reset (i.e., FLR or Conventional Reset) condition.

A13 Prefetchable BARs

In Section 1.3.2.2, page 41, line 20, make the following edit:

- 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 Function does not tolerate write merging. [See Section 7.5.2.1 for additional guidance on having the Prefetchable bit Set.](#)

In Section 2.2.5, page 67, line 16, make the following edit:



IMPLEMENTATION NOTE

Read Request with TPH to Non-Prefetchable Space

Memory Read Requests with the TH bit Set and [that](#) target Non-Prefetchable Memory Space should only be issued when it can be guaranteed that completion of such reads will not create undesirable side effects. [See Section 7.5.2.1 for consideration of certain BARs that may have the Prefetchable bit Set even though they map some locations with read side-effects.](#)

In Section 7.5.2.1, page 488, after line 3, insert the following implementation note:



IMPLEMENTATION NOTE

Additional Guidance on the Prefetchable Bit in Memory Space BARs

PCI Express adapters with Memory Space BARs that request a large amount of non-prefetchable Memory Space (e.g., over 64 MB) may cause shortages of that Space on certain scalable platforms, since many platforms support a total of only 1 GB or less of non-prefetchable Memory Space. This may limit the number of such adapters that can be supported on those platforms. For this reason, it is especially encouraged for BARs requesting large amounts of Memory Space to have their Prefetchable bit Set, since prefetchable Memory Space is more bountiful on most scalable platforms.

While a Memory Space BAR is required have its Prefetchable bit Set if none of the locations within its range have read side effects and all of the locations tolerate write merging, there are system configurations where having the Prefetchable bit Set will still allow correct operation even if those conditions are not met. For those cases, it may make sense for the adapter to have the Prefetchable bit Set in certain candidate BARs so that the system can map those BARs into prefetchable Memory Space in order to avoid non-prefetchable Memory Space shortages.

On PCI Express systems that meet the criteria enumerated below, setting the Prefetchable bit in a candidate BAR will still permit correct operation even if the BAR's range includes some locations that have read side-effects or cannot tolerate write merging. This is primarily due to the fact that PCI Express Memory Reads always contain an explicit length, and PCI Express Switches never prefetch or do byte merging. Generally only 64-bit BARs are good candidates, since only Legacy Endpoints are permitted to set the Prefetchable bit in 32-bit BARs, and most scalable platforms map all 32-bit Memory BARs into non-prefetchable Memory Space regardless of the Prefetchable bit value.

Here are criteria that are sufficient to guarantee correctness for a given candidate BAR:

- The entire path from the host to the adapter is over PCI Express.
- No conventional PCI or PCI-X devices do peer-to-peer reads to the range mapped by the BAR.
- The PCI Express Host Bridge does no byte merging. (This is believed to be true on most platforms.)
- Any locations with read side-effects are never the target of Memory Reads with the TH bit Set. See Section 2.2.5.

The above criteria are a simplified set that are sufficient to guarantee correctness. Other less restrictive but more complex criteria may also guarantee correctness, but are outside the scope of this specification.

A14 Memory, I/O, and Configuration Request Rules

In Section 2.2.7, page 78, line 3, make the following edit:

...

Receivers may optionally check for violations of these rules (but must not check reserved bits). If a Receiver implementing these checks determines that a TLP violates these rules, the TLP is a Malformed TLP.

In Section 2.2.7, page 78, line 20, make the following edit:

...

Receivers may optionally check for violations of these rules (but must not check reserved bits). If a Receiver implementing these checks determines that a TLP violates these rules, the TLP is a Malformed TLP.

A15 Conventional Reset

In Section 6.6.1, page 405, line 11, make the following edit:

- There are three distinct types of Conventional Reset: cold, warm, and hot:
 - A Fundamental Reset 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 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. This document does not specify a means for generating a warm reset.
 - There is an in-band mechanism for propagating Conventional Reset across a Link. This is called a hot reset and is described in Section 4.2.4.7.
 - There is an in-band mechanism for software to force a Link into Electrical Idle, “disabling” the Link. The Disabled LTSSM state is described in Section 4.2.5.9 and the Link Disable control bit is described in Section 7.8.7. Disabling a Link causes Downstream components to undergo a hot reset
 - Note also that the Data Link Layer reporting DL_Down status is in some ways identical to a hot reset – see Section 2.9.

In Section 6.6.1, page 406, line 10, make the following edit:

The second set of rules addresses 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 Conventional Reset of one or more devices before it is permitted to issue Configuration Requests to 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 Conventional Reset at the Root Complex – how this is done is beyond the scope of this specification
 - Note: Software should use 100 ms wait periods only if software enables CRS Software Visibility. Otherwise, Completion timeouts, platform timeouts, or lengthy processor instruction stalls may result. See the Configuration Request Retry Status Implementation Note in Section 2.3.1.
- ❑ The Root Complex and/or system software must allow at least 1.0 s after a Conventional Reset 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. This period is independent of how quickly Link training completes.
 - Note: This delay is analogous to the T_{rhta} parameter specified for PCI/PCI-X, and is intended to allow an adequate amount of time for devices which require self initialization.

A16 Multicast

In Section 6.14.5, page 446, following line 16, edit as shown below:

6.14.5. MC_Overlay Mechanism

The MC Overlay mechanism is provided to allow a single BAR in an Endpoint that doesn't contain a Multicast Capability structure to be used for both Multicast and unicast TLP reception. Software can configure the MC_Overlay mechanism to affect this by setting the MC_Overlay_BAR in a Downstream Port so that the Multicast address range, or a portion of it, is remapped (overlaid) onto the Memory Space range accepted by the Endpoint's BAR. At the Upstream Port of a Switch, the mechanism can be used to overlay a portion of the Multicast address range onto a Memory Space range associated with host memory.

A Downstream Port's MC_Overlay mechanism applies to TLPs exiting that Port. An Upstream Port's MC_Overlay mechanism applies to TLPs exiting the Switch heading Upstream. A Port's MC_Overlay mechanism does not apply to TLPs received by the Port, to TLPs targeting memory space within the Port, or to TLPs routed Peer-to-Peer between Functions in a Multi-Function Upstream Port.

In Section 6.14.3, page 445, following line 28, edit as shown below:

6.14.3. Multicast Capability Structure Field Updates

Some fields of the Multicast Capability structure may be changed at any time. Others cannot be changed with predictable results unless the MC_Enable bit is Clear in every Function of the component. The latter group includes MC_Base_Address and MC_Index_Position.

Fields which software may change at any time include MC_Enable, MC_Num_Group, MC_Receive, MC_Block_All, and MC_Block_Untranslated. Updates to these fields must themselves be ordered. Consider, for example, TLPs A and B arriving in that order at the same Ingress Port and in the same TC. If A uses value X for one of these fields, then B must use the same value or a newer value.

For Multi-Function Upstream Switch Ports Multicast TLPs received by one Switch or transmitted by one Endpoint Function are presented to the other parallel Endpoint Functions and the Downstream Switch Ports of the other parallel Switches (Functions are considered to be parallel if they are in the same Device. A single Multicast TLP is forwarded Upstream when any of the Upstream Switch Functions has the appropriate MC Receive bit Set.

A17 ASPM Configuration

In Section 5.4.1.3, page 352, line 4, edit as shown below:

...

ASPM Control = 00b

Port's Transmitter must not ~~bring a Link into center~~ L0s-state.

Ports connected to the Downstream end of the Link must not issue a PM_Active_State_Request_L1 DLLP on its Upstream Link.

Ports connected to the Upstream end of the Link receiving a L1 request must respond with negative acknowledgement.

...

A18 Power Management Messages

In Section 5.6, page 355, line 4, edit as shown below:

Power Management Messages follow the general rules for all Messages. Power Management Message fields follow the following rules:

- Length field is reserved.
- Attribute field must be set to the default values (all 0's).
- Address field is reserved.
- Requester ID – see Table 2-21 in Section 2.2.8.2.

Traffic Class field must use the default class (TC0).

A19 Multicast Capability

In Section 7.21, page 633, line 1, edit as shown below:

7.21. Multicast Capability

Multicast is an optional normative functionality that is controlled by the Multicast Capability structure. The Multicast Capability is applicable to Root Ports, RCRBs, Switch Ports, Endpoint Functions, and Root Complex Integrated Endpoints. It is not applicable to PCI Express to PCI/PCI-X Bridges.

In the cases of a Switch or Root Complex or a component that contains multiple Functions, multiple copies of this Capability structure are required – one for each Endpoint Function, Switch Port, or Root Port that supports Multicast. To provide implementation efficiencies, certain fields within each of the Multicast Capability structures within a component must be programmed the same and results are indeterminate if this is not the case. The fields and registers that must be configured with the same values include MC_Enable, MC_Num_Group, MC_Base_Address and MC_Index_Position. These same fields in an Endpoint's Multicast Capability structure must match those configured into a Multicast Capability structure of the Switch or Root Complex above the Endpoint or in which the Root Complex Integrated Endpoint is integrated.

A20 Advanced Error Reporting Capability

In Section 7.10, page 554, line 12, edit as shown below:

...

Note that if an error reporting bit field is marked as optional in the error registers, the bits must be implemented or not implemented as a group across the Status, Mask and Severity registers. In other words, a Function is required to implement the same error bit fields in corresponding Status, Mask and Severity registers. Bits corresponding to bit fields that are not implemented must be hardwired to 0, unless otherwise specified.

Except for Root Ports and Root Complex Event Collectors, if the End-End TLP Prefix Supported bit is Set, the Root Error Command and Error Source Identification registers must be RsvdP and the Root Error Status register must be RsvdZ.

A21 ACS Capability Register

In Section 7.16.2, page 609, make the following edits to Table 7-66:

Bit Location	Register Description	Attributes
6	<p>ACS Direct Translated P2P (T) – Required for Root Ports that support Address Translation Services (ATS) and also support peer-to-peer traffic with other Root Ports; required for Switch Downstream Ports; <u>required for multi-Function device Functions that support Address Translation Services (ATS) and also support peer-to-peer traffic with other Functions;</u> must be hardwired to 0b otherwise. If 1b, indicates that the component implements ACS Direct Translated P2P.</p>	RO

B1 Latency Tolerance Reporting(LTR) Mechanism

In Section 6.18, page 460, line 20, make the following edits:

When enabling the LTR mechanism in a hierarchy, devices closest to the Root Port must be enabled first, ~~then moving downwards towards the leaf Endpoints.~~

If an LTR Message is received at a ~~Root~~ Downstream Port that does not support LTR or if LTR is not enabled, the Message must be treated as an Unsupported Request.

B2 Mapping Between PCI 3.0 and PCI Express Registers

In Section 7.5.1.1, page 481, line 2, add the following text:

Table 7-3 establishes the mapping between PCI 3.0 and PCI Express for the PCI 3.0 Configuration Space Command register. For PCI Express to PCI/PCI-X Bridges, refer to the *PCI Express to PCI/PCI-X Bridge Specification* for additional requirements for some bits in this register.

In Section 7.5.1.2, page 483, line 2, add the following text:

Table 7-4 establishes the mapping between PCI 3.0 and PCI Express for PCI 3.0 Configuration Space Status register. For PCI Express to PCI/PCI-X Bridges, refer to the *PCI Express to PCI/PCI-X Bridge Specification* for additional requirements for some bits and fields in this register.

In Section 7.5.3.3, page 489, line 13, add the following text:

This register does not apply to PCI Express. It must be read-only and hardwired to 00h. For PCI Express to PCI/PCI-X Bridges, refer to the *PCI Express to PCI/PCI-X Bridge Specification* for additional requirements for this register.

In Section 7.5.3.4, page 489, line 15, add the following text:

Table 7-5 establishes the mapping between PCI 3.0 and PCI Express for PCI 3.0 Configuration Space Secondary Status register. For PCI Express to PCI/PCI-X Bridges, refer to the *PCI Express to PCI/PCI-X Bridge Specification* for additional requirements for some bits and fields in this register.

In Section 7.5.3.6, page 491, line 2, add the following text:

Table 7-6 establishes the mapping between PCI 3.0 and PCI Express for the PCI 3.0 Configuration Space Bridge Control register. For PCI Express to PCI/PCI-X Bridges, refer to the *PCI Express to PCI/PCI-X Bridge Specification* for additional requirements for some bits in this register.

B3 IDO Completion Enable

In Section 7.8.16, page 546, Table 7-25, Bit 9, edit as follows:

Bit Location	Register Description	Attributes
9	<p>IDO Completion Enable – If this bit is Set, the Function is permitted to set the ID-Based Ordering (IDO) bit (Attribute[2]) of Completions it returns (see Section 2.2.6.3 and Section 2.4). Endpoints, including RC Integrated Endpoints, and Root Ports are permitted to implement this capability.</p> <p>A Function is permitted to hardwire this bit to 0b if it never sets the IDO attribute in Requests<u>Completions</u>.</p> <p>Default value of this bit is 0b.</p>	RW

B4 Physical Layer Services

In Section 1.5.4.3, page 49, line 11, edit as follows:

Interface initialization, maintenance control, and status tracking:

- Reset/Hot-Plug control/status
- Interconnect power management
- Width and Lane mapping negotiation
- ~~Polarity reversal~~Lane polarity inversion

B5 Detect.Active

In Section 4.2.6.1.2, page 214, line 20, make the following edit:

- The Transmitter performs a Receiver Detection sequence on all unconfigured Lanes that can form one or more Links (see Section ~~4.3.1.84.3.5.7~~ for more information).

B6 Terms and Acronyms

On page 30 make the following edit:

invariant A field of a TLP header **or TLP Prefix** that contains a value that cannot legally be modified as the TLP flows through the PCI Express fabric.

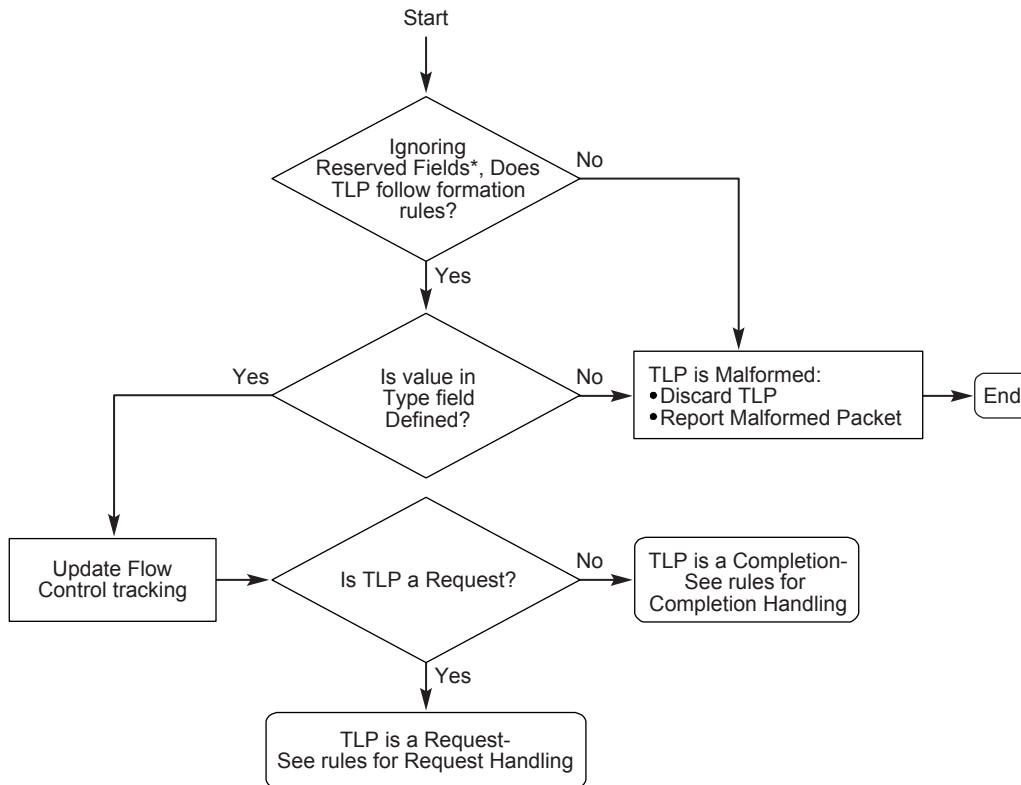
On page 31 make the following edits:

Packet A fundamental unit of information transfer consisting of **an optional TLP Prefix, followed by** a header ~~that~~**and**, in some cases, ~~is~~ **followed by** a data payload.

On page 33 make the following edit:

variant A field of a TLP header **or TLP Prefix** that contains a value that is subject to possible modification according to the rules of this specification as the TLP flows through the PCI Express fabric.

In Section 2.3, page 102, replace Figure 2-30: Flowchart for Handling of Received TLPs with the following:



*TLP fields which are marked Reserved are not checked at the Receiver

In Section 2.6.1.2, page 140, line 21, make the following edit:

TLP Overhead	Represents the additional TLP components which consume Link bandwidth (<u>TLP Prefix</u> , header, LCRC, framing Symbols) and is treated here as a constant value of 28 Symbols
--------------	--

In Section 3.5.2.1, page 172, line 2, make the following edit:

TLP Overhead	represents the additional TLP components which consume Link bandwidth (<u>TLP Prefix</u> , header, LCRC, framing Symbols) and is treated here as a constant value of 28 Symbols
--------------	--

In Section 3.5.3.1, page 182, line 32, make the following edit:

TLP Overhead	represents the additional TLP components which consume Link bandwidth (<u>TLP Prefix</u> , header, LCRC, framing Symbols) and is treated here as a constant value of 28 Symbols.
--------------	---

In Section 2.7.1, page 142, lines 21 and 26, make the following edits:

A 32-bit ECRC is calculated for the ~~entire~~ TLP (End-End TLP Prefixes, header, and data payload) using the following algorithm and appended to the end of the TLP (see Figure 2-3):

- The ECRC value is calculated using the following algorithm (see Figure 2-37).
- The polynomial used has coefficients expressed as 04C1 1DB7h
- The seed value (initial value for ECRC storage registers) is FFFF FFFFh
- All invariant fields of the TLP's End-End TLP Prefixes (if present), header, and the entire data payload (if present) are included in the ECRC calculation, all bits in variant fields must be set to 1b for ECRC calculations.


In Section 6.2.4.2, page 374, line 36, make the following edit:

Since an implementation only has the ability to record a finite number of headers, it is important that software services the First Error Pointer ~~and Header Log, Header Log, and TLP Prefix Log~~ registers in a timely manner, to limit the risk of missing this information for subsequent errors.

In Section 6.2.4.3, page 375, line 25, make the following edits:

If the “corresponding” uncorrectable error bit in the Uncorrectable Error Mask register is clear and the error is one that requires header logging, then the prefix and header ~~is~~are recorded, subject to the availability of resources.

In Section 6.2.5 , page 377, replace Figure 6-2: Flowchart Showing Sequence of Device Error Signaling and Logging Operations with the following:



OM14546D

In Section 6.2.7, page 379, make the following edits to Table 6-2:

Error Name	Error Type (Default Severity)	Detecting Agent Action	References
Corrected Internal Error	Correctable (masked by default)	<i>Component:</i> Send ERR_COR to Root Complex.	Section 6.2.9
Uncorrectable Internal Error	Uncorrectable (Fatal and masked by default)	<i>Component:</i> Send ERR_FATAL to Root Complex. Optionally, log the prefix /header of the first TLP associated with the error.	Section 6.2.9
Header Log Overflow	Correctable (masked by default)	<i>Component:</i> Send ERR_COR to Root Complex.	Section 6.2.4.2

In Section 6.2.7, page 381, make the following edits to Table 6-5:

Error Name	Error Type (Default Severity)	Detecting Agent Action	References
Poisoned TLP Received	Uncorrectable (Non-Fatal)	<i>Receiver:</i> Send ERR_NONFATAL to Root Complex or ERR_COR for the Advisory Non-Fatal Error cases described in Sections 6.2.3.2.4.2 and 6.2.3.2.4.3. Log the prefix /header of the Poisoned TLP.	Section 2.7.2.2
ECRC Check Failed		<i>Receiver (if ECRC checking is supported):</i> Send ERR_NONFATAL to Root Complex or ERR_COR for the Advisory Non-Fatal Error case described in Section 6.2.3.2.4.2. Log the prefix /header of the TLP that encountered the ECRC error.	Section 2.7.1

Errata for the PCI Express Base Specification, Revision 2.1

Error Name	Error Type (Default Severity)	Detecting Agent Action	References
Unsupported Request (UR)		<p><i>Request Receiver:</i> Send ERR_NONFATAL to Root Complex or ERR_COR for the Advisory Non-Fatal Error case described in Section 6.2.3.2.4.1. Log the prefix/header of the TLP that caused the error.</p>	Section 2.2.8.6, Section 2.3.1, Section 2.3.2, Section 2.7.2.2., Section 2.9.1, Section 5.3.1, Section 6.2.3.1, Section 6.2.6, Section 6.2.8.1, Section 6.5.7, Section 7.3.1, Section 7.3.3, Section 7.5.1.1, Section 7.5.1.2
Completion Timeout		<p><i>Requester:</i> Send ERR_NONFATAL to Root Complex or ERR_COR for the Advisory Non-Fatal Error case described in Section 6.2.3.2.4.4.</p>	Section 2.8
Completer Abort		<p><i>Completer:</i> Send ERR_NONFATAL to Root Complex or ERR_COR for the Advisory Non-Fatal Error case described in Section 6.2.3.2.4.1. Log the prefix/header of the Request that encountered the error.</p>	Section 2.3.1
Unexpected Completion		<p><i>Receiver:</i> Send ERR_COR to Root Complex. This is an Advisory Non-Fatal Error case described in Section 6.2.3.2.4.5. Log the prefix/header of the Completion that encountered the error.</p>	Section 2.3.2
ACS Violation		<p><i>Receiver (if checking):</i> Send ERR_NONFATAL to Root Complex. Log the prefix/header of the Request TLP that encountered the error.</p>	
MC Blocked TLP		<p><i>Receiver (if checking):</i> Send ERR_NONFATAL to Root Complex. Log the prefix/header of the Request TLP that encountered the error.</p>	Section 6.14.4

Errata for the PCI Express Base Specification, Revision 2.1


Error Name	Error Type (Default Severity)	Detecting Agent Action	References
AtomicOp Egress Blocked		<i>Egress Port:</i> Send ERR_COR to Root Complex. This is an Advisory Non-Fatal Error case described in Section 6.2.3.2.4.1. Log the <u>prefix</u> /header of the AtomicOp Request that encountered the error.	Section.6.15.2
Receiver Overflow	Uncorrectable (Fatal)	<i>Receiver (if checking):</i> Send ERR_FATAL to Root Complex.	Section 2.6.1.2
Flow Control Protocol Error		<i>Receiver (if checking):</i> Send ERR_FATAL to Root Complex.	Section 2.6.1
Malformed TLP		<i>Receiver:</i> Send ERR_FATAL to Root Complex. Log the <u>prefix</u> /header of the TLP that encountered the error.	Section 2.2.2, Section 2.2.3, Section 2.2.5, Section 2.2.7, Section 2.2.8.1, Section 2.2.8.2, Section 2.2.8.3, Section 2.2.8.4, Section 2.2.8.5, Section 2.2.9, Section 2.2.10, Section 2.2.10.1, Section 2.2.10.2, Section 2.3, Section 2.3.1, Section 2.3.1.1, Section 2.3.2, Section 2.5, Section 2.5.3, Section 2.6.1, Section 2.6.1.2, Section 6.2.4.4, Section 6.3.2
TLP Prefix Blocked		<i>Egress Port:</i> Send ERR_NONFATAL to Root Complex or ERR_COR for the Advisory Non-Fatal Error case described in Section 6.2.3.2.4.1. Log the <u>prefix</u> /header of the TLP that encountered the error.	Section 2.2.10.2

In Section 6.12.4, page 435, line 17, make the following edits:

The AER prefix/header logging and the Prefix Log/Header Log register registers may be used to determine the prefix/header of the offending Request.


B7 Address Based Routing Rules and Memory, I/O, and Configuration Request Rules

In Section 2.2.4.1, page 64, replace Figure 2-7: 64-bit Address Routing with the following:



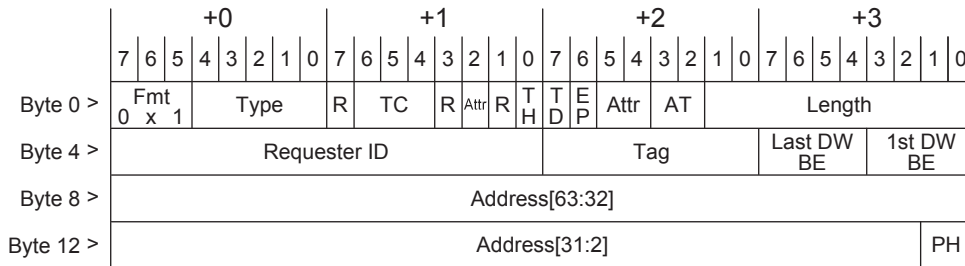
OM14544C

In Section 2.2.4.1, page 64, replace Figure 2-8: 32-bit Address Routing with the following:




OM14543C

In Section 2.2.7, page 77, replace Figure 2-15: Request Header Format for 64-bit Addressing of Memory with the following:



OM13764C

In Section 2.2.7, page 77, replace Figure 2-16: Request Header Format for 32-bit Addressing of Memory with the following:



OM13763C

B8 First/Last DW Byte Enables Rules

In Section 2.2.5, page 67, line 4, make the following edits:

For Memory Read Requests that have the TH bit Set, the Byte Enable fields are repurposed to carry the ST[7:0] field ([Refer to Section 2.2.7.1 for details](#)), and values for the Byte Enables are implied as defined below. ~~Such Requests~~ The TH bit must only be ~~issued-set in Memory Read Requests~~ when it is acceptable to complete ~~the-those~~ Requests as if all bytes for ~~the~~ requested ~~payload-data~~ were enabled.

- For Memory ~~Reads~~ Read Requests that have the TH bit Set, the following values are implied for the Byte Enables.

B9 Memory, I/O, and Configuration Request Rules

In Section 2.2.7, page 75, line 21, make the following edits:

- Last DW BE[3:0] and 1st DW BE[3:0]. For Memory Read Requests and AtomicOp Requests with the TH bit Set, the byte location for the Last DW BE[3:0] and 1st DW BE [3:0] fields in the header are repurposed to carry ST[7:0] field. ~~and the values for the DW BE fields are implied to be reserved. Otherwise, the DW BE fields are reserved. For Memory Read Requests with the TH bit Clear, see Section 2.2.5 for First/Last DW Byte Enable Rules. For AtomicOp Requests with TH bit Set, the values for the DW BE fields are implied to be reserved. For AtomicOp Requests with TH bit clear, the DW BE fields are reserved.~~

B10 End-End TLP Prefix Processing

In Section 2.2.10.2, page 100, line 15, make the following edit:

- For routing elements, the End-End TLP Prefix ~~Egress~~-Blocking bit in each Egress Port determines whether TLPs containing End-End TLP Prefixes can be transmitted via that Egress Port (see Section 7.8.16).

B11 ST Modes of Operation

In Section 6.17.3, page 459, line 3, make the following edits:

A Function that is capable of generating TPH Requests is required to support the No ST Mode of operation, ~~and support~~. Support for other ~~ST Modes~~~~modes~~ of ~~operations~~operation is optional. Only one ~~ST Modemode~~ of operation can be selected at a time by programming ST Mode Select.

B12 Latency Tolerance Mechanism

In Section 6.18, page 464, line 9, make the following edits:

Time D for a Read transaction is the time between the transmission of the END symbol in the Request TLP to the receipt of the STP symbol in the Completion TLP for that Request. Time D for a Write transaction is the time between the transmission of the END symbol of the TLP that exhausts the FC credit to the receipt of the SDP symbol of the DLLP returning more credits for that Request type.

B13 ID Based Routing Rules

In Section 2.2.4.2, page 65, line 13, make the following edit:

This specification defines Vendor_Defined Messages that are ID Routed (Section 2.2.8.6). Other specifications are permitted to define additional ID Routed Messages.

B14 ST Modes

In Section 6.17.3, page 458, make the following edits to Table 6-12:

ST Mode Select [2:0]	ST Mode Name	Description
000	No ST Mode	The Function must use a value of all zeroes for all Steering Tags.
001	Interrupt Vector Mode	Each Steering Tag is selected by an MSI/MSI-X interrupt vector number. The Function is required to use the Steering Tag value from an ST Table entry that can be indexed by a valid MSI/MSI-X interrupt vector number.
010	Device Specific Mode	It is recommended for the Function to use a Steering Tag value from an ST Table entry, but it is not required.
<u>All other encodings</u>	<u>All other- encodingsReserved</u>	Reserved for future use.

B15 Latency Tolerance Mechanism

In Section 6.18, page 462, lines 13-39, make the following edits:

Switches must collect the Messages from Downstream Ports that have the LTR mechanism enabled and transmit a “conglomerated” Message Upstream according to the following rules:

- ❑ If a Switch supports the LTR feature, it must support the feature on its Upstream Port and all Downstream Ports.
- ❑ A Switch Upstream Port is permitted to transmit LTR Messages only when its LTR Mechanism Enable bit is Set or shortly after software clears its LTR Mechanism Enable bit as described earlier in this section.
- ~~❑ The acceptable latency values for the Message sent Upstream by the Switch must reflect the lowest values received from any Downstream Port.~~
 - ~~• When any Downstream Port reports a field with the Requirement bit set to 1 and a LatencyValue of all 0’s (regardless of the LatencyScale value), the Message sent Upstream must report a LatencyScale of 000b and a LatencyValue of all 0’s.~~
- ❑ The acceptable latency values for the Message sent Upstream by the Switch must be calculated as follows:
 - If none of the Downstream Ports receive an LTR Message containing a requirement for a certain type of traffic (snoop/no-snoop), ~~the~~ then any LTR Message sent by the Switch must not set the Requirement bit corresponding to that type of traffic.
 - Define LTRdnport[N] as the value reported in the LTR Message received at Downstream Port N, with these adjustments if applicable:
 - ◆ LTRdnport[N] is effectively infinite if the Requirement bit is clear or if a Not Permitted LatencyScale value is used
 - ◆ LTRdnport[N] must be 0 if the Requirement bit is 1 and the LatencyValue field is all 0’s regardless of the LatencyScale value
 - Define LTRdnportMin as the minimum value of LTRdnport[N] across all Downstream Ports
 - Define Lswitch as all latency induced by the Switch
 - ◆ If Lswitch dynamically changes based on the Switch’s operational mode, the Switch must not allow Lswitch to exceed 20% of LTRdnportMin, unless Lswitch for the Switch’s lowest latency mode is greater, in which case the lowest latency state must be used
 - Calculate the value to transmit upstream, LTRconglomerated, as LTRdnportMin – Lswitch, unless this value is less than 0 in which case LTRconglomerated is 0
 - If LTRconglomerated is 0, both the LatencyValue and LatencyScale fields must be all 0’s in the conglomerated LTR message
 - ~~• Any additional latency induced by the Switch must be accounted for in the conglomerated Message. A Switch must ensure that its Link and internal power management and other internal operation shall not cause its conglomerated latency to be reduced by more than 20% of the lowest received latency.~~

- ~~☐ If any Downstream Port reports a field(s) for which the Requirement bit is clear, or uses a Not-Permitted LatencyScale value, that Port must not be considered when determining the corresponding field(s) reported in the Message sent Upstream.~~
 - ~~• Valid, permitted values for other fields must still be considered.~~
- ☐ When a Switch Downstream Port goes to DL_Down status, the previous latencies recorded for that Port must be treated as invalid. A new LTR message must be ~~and the latencies to be transmitted Upstream updated and a new conglomerated Message~~ transmitted Upstream if the conglomerated latencies are changed as a result.

B16 Device Capabilities 2 Register

In Section 7.8.15, page 542, make the following edits to Table 7-24:

Bit Location	Register Description	Attributes								
13:12	<p>TPH Completer Supported – <u>Value indicates Completer support for TPH or Extended TPH.</u> Applicable only to Root Ports and Endpoints. must be 00b for other Function types. For all other Functions, this field is Reserved.</p> <p>Defined Encodings are:</p> <table style="margin-left: 20px;"> <tr> <td>00b</td> <td>TPH and Extended TPH Completer not supported.</td> </tr> <tr> <td>01b</td> <td>TPH Completer supported; Extended TPH Completer not supported.</td> </tr> <tr> <td>10b</td> <td>Reserved.</td> </tr> <tr> <td>11b</td> <td>Both TPH and Extended TPH Completer supported.</td> </tr> </table> <p>See Section 6.17 for details.</p>	00b	TPH and Extended TPH Completer not supported.	01b	TPH Completer supported; Extended TPH Completer not supported.	10b	Reserved.	11 b	Both TPH and Extended TPH Completer supported.	RO
00b	TPH and Extended TPH Completer not supported.									
01b	TPH Completer supported; Extended TPH Completer not supported.									
10b	Reserved.									
11 b	Both TPH and Extended TPH Completer supported.									

B17 ACS Control Register

In Section 7.16.3, page 610, make the following edits to Table 7-67:

Bit Location	Register Description	Attributes
6	<p>ACS Direct Translated P2P Enable (T) – When Set, overrides the ACS P2P Request Redirect and ACS P2P Egress Control mechanisms with peer-to-peer Memory Requests whose Address Translation (AT) field indicates a Translated address (see Section 6.12.3).</p> <p>This bit is ignored if ACS Translation Blocking <u>Enable</u> (B) is 1b.</p> <p>Default value of this bit is 0b. Must be hardwired to 0b if the ACS Direct Translated P2P functionality is not implemented.</p>	RW

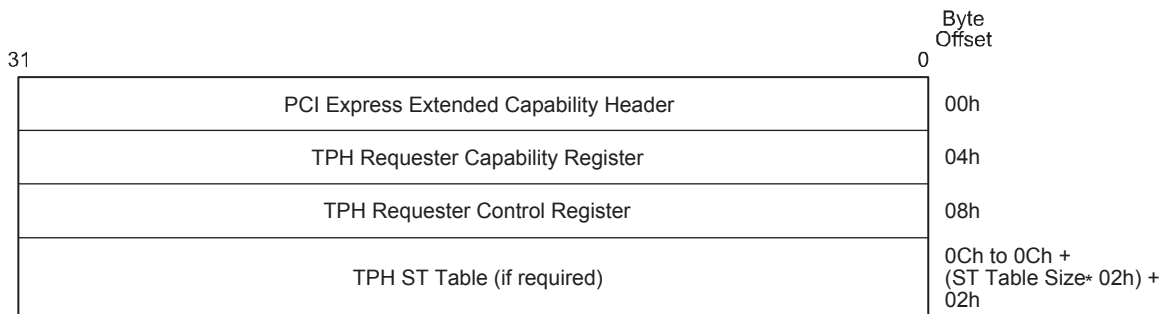
B18 Latency Tolerance Reporting Capability

In Section 7.25, page 651, line 3, make the following edits:

The PCI Express Latency Tolerance Reporting (LTR) Capability is an optional Extended Capability that allows software to provide platform latency information to components with Upstream Ports (Endpoints and Switches), and is required [for Switch Upstream Ports and Endpoints](#) if the [component Function](#) supports the LTR mechanism. It is not applicable to Root Ports, Bridges, or ~~Downstream Ports in a~~ Switch [Downstream Ports](#).

B19 TPH Requester Capability Structure

In Section 7.26, page 653, replace Figure 7-126: TPH Extended Capability Structure with the following:



A-0779A

B20 TPH Requester Capability Register—ST Table Location

In Section 7.26.2, page 655, make the following edits to Table 7-109:

Bit Location	Register Description	Attributes
10:9	<p>ST Table Location – Value indicates if and where the ST Table is located.</p> <p>Defined Encodings are:</p> <p>00b – ST Table is not present.</p> <p>01b – ST Table is located in the TPH Requester Capability structure.</p> <p>10b – ST Table is located in the MSI-X Table structure.</p> <p>11b – Reserved</p> <p>A Function that only supports the No ST Mode of operation must have a value of 00b in this field.</p> <p>A Function may report a value of 10b only if it implements an MSI-X Capability structure.</p>	RO

B21 TPH Requester Capability Register—ST Table Size

In Section 7.26.2, page 655, make the following edits to Table 7-109:

Bit Location	Register Description	Attributes
26:16	<p>ST Table Size – <u>Value indicates the maximum number of ST Table entries the Function may use.</u> Software reads this field to determine the ST Table Size N, which is encoded as N-1. For example, a returned value of <u>“0000000011”000 0000 0011b</u> indicates a table size of <u>4-four entries.</u></p> <p>There is an upper limit of 64 entries when the ST Table is located in the TPH Requester Capability structure.</p> <p>There is an upper limit of 2-K<u>2048</u> entries when the ST Table is located in the MSI-X Table.</p> <p>This field is only applicable for Functions that implement an ST Table as indicated by the ST Table Location field. Otherwise, the value in this field is undefined.</p>	RO

B22 TPH Requester Control Register—ST Model Select

In Section 7.26.3, page 656, make the following edits to Table 7-110:

Bit Location	Register Description	Attributes
2:0	<p>ST Mode Select – selects the ST mode<u>Mode</u> of operation.</p> <p>Defined encodings are:</p> <p>000b – No ST Mode</p> <p>001b – Interrupt Vector Mode</p> <p>010b – Device Specific Mode</p> <p>Others<u>all other encodings</u> – reserved for future use</p> <p>Functions that support only the No ST Mode of operation must hardwire this field to 000b.</p> <p><u>Function operation is undefined if software enables a mode of operation that does not correspond to a mode supported by the Function.</u></p> <p>The default value of this field is 000b.</p> <p>See Section 6-17-26.17.3 for details on ST modes of operation.</p>	RW

B23 TPH Requester Control Register—TPH Requester Enable

In Section 7.26.3, page 656, make the following edits to Table 7-110:

Bit Location	Register Description	Attributes
9:8	<p>TPH Requester Enable – <u>Controls the ability to issue Request TLPs using either TPH or Extended TPH.</u></p> <p>Defineddefined encodings are:</p> <p>00b – Function operating as a Requester is not permitted to issue Requests with TPH or Extended TPH.</p> <p>01b – Function operating as a Requester is permitted to issue Requests with TPH and is not permitted to issue Requests with Extended TPH.</p> <p>10b – Reserved.</p> <p>11b – Function operating as a Requester is permitted to issue Requests with TPH and Extended TPH.</p> <p>The default value of this field is 00b.</p>	RW

B24 TPH ST Table

In Section 7.26.4, page 657, lines 1-5, make the following edits:

The ST Table must be implemented in the TPH Requester Capability ~~Structure~~structure if the value of the ST Table Location field is 01b. For all other values, the ST Entry ~~Registers~~registers must not be implemented. Each implemented ST Entry is 16 bits. The number of ST Entry ~~Registers~~registers implemented must be equal to the number of ST Table entries supported by the Function, which is the value of the ST Table Size field plus one.

B25 TPH ST Table—ST Upper and ST Lower

In Section 7.26.4, page 657, make the following edits to Table 7-111:

Bit Location	Register Description	Attributes
7:0	<p>ST Lower – If the Function implements a TPH Requester Capability structure, and the ST Table Location indicates a value of 01b, then this<u>This</u> field contains the lower 8 bits of a Steering Tag.</p> <p>Default value of this field is <u>00h</u>.</p>	RW
15:8	<p>ST Upper – If the Function implements a TPH Requester Capability structure, and the ST Table Location indicates a value of 01b, and the Function's Extended TPH Requester Supported bit is Set, then this field contains the upper 8 bits of a Steering Tag. Otherwise, this field is RsvdP.</p> <p>Default value of this field is <u>00h</u>.</p>	RW

B26 Common Packet Header Fields

In Section 2.2.1, page 57, make the following edits to Table 2-2:

Table 2-2: Fmt[24:0] Field Values

Fmt[24:0]	Corresponding TLP Format
000b	3 DW header, no data
001b	4 DW header, no data
010b	3 DW header, with data
011b	4 DW header, with data
100b	TLP Prefix
	All encodings not shown above are reserved (see Section 2.3).

In Section 2.2.1, page 58, make the following edit to the title of Table 2.3:

Table 2-3: Fmt[24:0] and Type[4:0] Field Encodings

B27 Compliance Pattern

In Section 4.2.8, page 260, lines 3-9, make the following edits:

During Polling, the Polling.Compliance substate must be entered from Polling.Active based on the conditions described in Section 4.2.6.2.1. The compliance pattern consists of the sequence of 8b/10b Symbols K28.5, D21.5, K28.5, and D10.2 repeating. The ~~compliance~~Compliance sequence is as follows:

Symbol	K28.5	D21.5	K28.5	D10.2
Current Disparity	Negative	Positive	Positive	Negative
Pattern	0011111010	1010101010	1100000101	0101010101

The first Compliance sequence Symbol must have negative disparity. It is permitted to create a disparity error to align the running disparity to the negative disparity of the first Compliance sequence Symbol.

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 Compliance Pattern sequence. Note: A x1 device, or a xN device operating a Link in x1 mode, is permitted to include the Delay sSymbols with the Compliance Pattern.

This delay sequence on every eighth Lane is then:

Symbol:	D	D	K28.5	D21.5	K28.5	D10.2	D	D
----------------	---	---	-------	-------	-------	-------	---	---

Where D is a K28.5 Symbol. The first D Symbol has negative disparity to align the delay disparity with the disparity of the Compliance sequence.

In Section 4.2.1.3, page 190, lines 3-10, make the following edits:

A Transmitter is permitted to pick any disparity, unless otherwise required, 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 Symbol used to obtain Symbol lock. Disparity may also be re-initialized 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.

B28 Modified Compliance Pattern

In Section 4.2.9, page 261, line 5, make the following edits:

The Modified Compliance Pattern consists of the same basic Compliance Pattern sequence (see Section 4.2.8) with one change. Two identical error status Symbols followed by two K28.5 are appended to the basic Compliance sequence of 8b/10b Symbols (K28.5, D21.5, K28.5, and D10.2) to form the Modified Compliance Sequence of (K28.5, D21.5, K28.5, D10.2, error status Symbol, error status Symbol, K28.5, K28.5). The first Modified Compliance Sequence Symbol must have negative disparity. It is permitted to create a disparity error to align the running disparity to the negative disparity of the first Modified Compliance Sequence Symbol. For any given device that has multiple Lanes, every eighth Lane is moved by a total of eight Symbols. Four Symbols of K28.5 occurs at the beginning and another four Symbols of K28.7 occurs at the end of the eight Symbol Modified Compliance Pattern sequence. The first D Symbol has negative disparity to align the delay disparity with the disparity of the Modified Compliance Sequence. After the 16 Symbols are sent, the delay Symbols are advanced to the next Lane, until the delay Symbols have been sent on all eight lanes. Then the delay Symbols cycle back to Lane 0, and the process is repeated. It is permitted to advance the delay sequence across all eight lanes, regardless of the number of lanes detected or supported. Note: A x1 device, or a xN device operating a Link in x1 mode, is permitted to include the Delay symbols with the Modified Compliance Pattern.

B29 Polling.Compliance

In Section 4.2.6.2.2, page 220, lines 30-38, reformat the text as shown in red:

- Else (i.e., entry to Polling.Compliance is not due to the Compliance Receive bit assertion with Loopback bit deassertion in the received consecutive TS Ordered Sets and not due to the Enter Compliance bit and the Enter Modified Compliance bit in the Link Control 2 register set to 1b)
 - (a) Transmitter sends out the compliance pattern on all Lanes that detected a Receiver during Detect at the data rate and de-emphasis level determined above.
 - (b) Next state is Polling.Active if any of the following three conditions are true:
 1. Electrical Idle exit has been detected at the Receiver of any Lane that detected a Receiver during Detect, and the Enter Compliance bit in the Link Control 2 register is 0b.
 2. The Enter Compliance bit in the Link Control 2 register has changed to 0b (from 1b) since entering Polling.Compliance.
 3. The component is a Downstream component, the Enter Compliance bit in the Link Control 2 register is set to 1b and an EIOS has been detected on any Lane. The Enter Compliance bit is reset to 0b when this condition is true.

If the Transmitter is transmitting at a data rate other than 2.5 GT/s, or the Enter Compliance bit in the Link Control 2 register was set to 1b during entry to Polling.Compliance, the Transmitter sends eight consecutive EIOSs and enters Electrical Idle prior to transitioning to Polling.Active. During the period of Electrical Idle, the data rate is changed to 2.5 GT/s and stabilized. The period of Electrical Idle is greater than 1 ms but must not exceed 2 ms.

Note: Sending multiple EIOSs provides enough robustness such that the other component detects at least one EIOS and exits Polling.Compliance substate when the configuration register mechanism was used for entry.

B30 OBFF ECN Update [affects PCI Express CEM Specification, Revision 2.0]

On page 13 of the ECN, make the following edits (in red) to the Power Sequencing and Reset Signal Timings table:

Table 2-4: Power Sequencing and Reset Signal Timings

Symbol	Parameter	Min	Max	Units	Notes	Figure
T _{PVPERL}	Power stable to PERST# inactive	100		ms	1	Figure 2-10
T _{PERST-CLK}	REFCLK stable before PERST# inactive	100		μs	2	Figure 2-10
T _{PERST}	PERST# active time	100		μs		Figure 2-11
T _{FAIL}	Power level invalid to PERST# active		500	ns	3	Figure 2-13
T _{WKRF}	WAKE# rise – fall time		100	ns	4	Figure 2-14
<u>T_{WAKE-TX-MIN}</u> / <u>T_{WAKE-TX-MAX}</u> / <u>T_{WAKE-TX-PULSE}</u>	<u>Minimum/Maximum WAKE# pulse width; applies to both active-inactive-active and inactive-active-inactive cases</u>	<u>300</u>	<u>500</u>	<u>ns</u>	<u>5</u>	
<u>T_{WAKE-FALL-FALLCPU-ACTIVE}</u>	<u>Time between two falling WAKE# edges when signaling CPU Active</u>	<u>700</u>	<u>1000</u>	<u>ns</u>	<u>5</u>	

Notes:

1. Any supplied power is stable when it meets the requirements specified for that power supply.
2. A supplied reference clock is stable when it meets the requirements specified for the reference clock. The PERST# signal is asserted and de-asserted asynchronously with respect to the supplied reference clock.
3. The PERST# signal must be asserted within T_{FAIL} of any supplied power going out of specification.
4. Measured from WAKE# assertion/de-assertion to valid input level at the system PM controller. Since WAKE# is an open-drain signal, the rise time is dependent on the total capacitance on the platform and the system board pull-up resistor. It is the responsibility of the system designer to meet the rise time specification.
5. [Refers to timing requirement for indicating an active window.](#)

B31 MC_Base Address Register

In Section 7.21.4, page 636, make the following edits:

7.21.4 ~~Multicast Base~~MC_Base Address Register (Offset 08h)

Table 7-88: ~~Multicast Base~~MC_Base Address Register

B32 D3 State

In Section 5.3.1.4, page 321, line 5, make the following edit:

D3 support is required, (both the D3_{cold} and the D3_{hot} states). Functions supporting PME generation from D3 must support it for both D3_{cold} and the D3_{hot} states.

Functional context is required to be maintained by Functions in the D3_{hot} state if the No_Soft_Reset field in the PMCSR is Set. In this case, software is not required to re-initialize the Function after a transition from D3_{hot} to D0 (the Function will be in the D0_{initializedactive} state). If the No_Soft_Reset bit is Clear, functional context is not required to be maintained by the Function in the D3_{hot} state. As a result, in this case software is required to fully re-initialize the Function after a transition to D0 as the Function will be in the D0_{uninitialized} state.

The Function will be reset if the Link state has transitioned to the L2/L3 Ready state regardless of the value of the No_Soft_Reset bit.

B33 Trusted Configuration Space

In Section 2.2.1, page 58, footnote 3, make the following edit:

³ Deprecated TLP Types: previously used for [Trusted Configuration Space \(TCS\)](#), which is no longer supported by this specification. If a Receiver does not implement TCS, the Receiver must treat such Requests as Malformed Packets

B34 Latency Tolerance Reporting Mechanism

In Section 6.18, page 463, line 12, make the following edits:

When the latency requirement is updated during a series of Requests, it is required that the updated latency figure be comprehended by the RC prior to servicing no later than the larger of either (a) ~~waiting as long as the previously indicated latency or (b) following the servicing of~~ a subsequent Request. In all cases the updated latency value must take effect within a time period equal to or less than the previously reported latency requirement. It is permitted for the RC to comprehend the updated latency figure earlier than this limit.

B35 Latency Tolerance Reporting Mechanism

In Section 6.18, page 464, lines 6-24, make the following edits:

Time C is the period during which the transition from L1 ASPM to L0 takes place. ~~This value will not exceed the maximum L1 exit latency reported by the Endpoint.~~

Time D for a Read transaction is the time between the transmission of the END symbol in the Request TLP to the receipt of the STP symbol in the Completion TLP. Time D for a Write transaction is the time between the transmission of the END symbol of the TLP that exhausts the FC credits to the receipt of the SDP symbol in the DLLP returning more credits. This value will not exceed the latency in effect.

Time E is the period where the data path from the Endpoint to system memory is open, and data transactions are not subject to the leadoff latency.

The LTR latency semantic reflects the tolerable latency seen by the device as measured by one or both of the following:

Case 1: the device may or may not support Clock PM, but has not deasserted its CLKREQ# signal – The latency observed by the device is represented in Figure 6-16 as the sum of times C and D.

Case 2: the device supports Clock PM and has deasserted CLKREQ#- The latency observed by the device is represented as the sum of times A, C, and D.

To effectively use the LTR mechanism in conjunction with Clock PM, the device will know or be able to measure times B and C, so that it knows when to assert CLKREQ#. The actual values of Time A, Time C, and Time D may vary dynamically, and it is the responsibility of the platform to ensure the sum will not exceed the latency.

B36 Additional References to Tables

In Section 2.6.1.2, page 140, line 27, make the following edit:

UpdateFactor	Represents the number of maximum size TLPs sent during the time between UpdateFC receptions, and is used to balance Link bandwidth efficiency and receive buffer sizes – the value varies according to Max_Payload_Size and Link width, and is included in Table 2-38 <u>and Table 2-39</u> .
--------------	---

In Section 3.5.2.1, page 171, line 20, make the following edit:

- If the Transmit Retry Buffer contains TLPs for which no Ack or Nak DLLP has been received, and (as indicated by REPLAY_TIMER) no Ack or Nak DLLP has been received for a period exceeding the time indicated in Table 3-4 for 2.5 GT/s mode operation and Table 3-5 for 5.0 GT/s mode operation, the Transmitter initiates a replay.

In Section 3.5.2.1, page 172, line 8, make the following edit:

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 included <u>defined</u> in Table 3-6 <u>and Table 3-7</u> .
-----------	---

In Section 3.5.3.1, page 183, line 5, make the following edit:

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 Table 3-6 <u>and Table 3-7</u> .
-----------	--

B37 Capitalizing “Reserved”

Excluding page 101 footnote 19, page 118 line 15, page 415 line 28, and page 665 line 24:

Change all occurrences of “reserved” to “Reserved”.

B38 L0

In Section 4.2.6.5, page 246, lines 10 and 26, make the following edits:

- 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 the EIOS(s) is met. A transition to L1 can be initiated by PCI-PM (see Section 5.4.3.2.1), or by ASPM (see Section 5.4.1.2.1).
- 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 the EIOS(s) is met (see Section 5.4.3.2.3 for more details).

B39 The Set of Lanes in Polling.Active

In Section 4.2.6.2.1, page 217, line 1, make the following edits:

...

- ii At least a predetermined ~~number~~set of Lanes that detected a Receiver during Detect have detected an exit from Electrical Idle at least once since entering Polling.Active.
 - Note: This ~~may~~ 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. The exact ~~number~~set of predetermined Lanes -is implementation specific. Note that up to the 1.1 specification this predetermined ~~number~~set was ~~equal to~~ the total ~~number~~set of Lanes that detected a Receiver.
 - Note: Any Lane that receives eight consecutive TS1 or TS2 Ordered Sets should have detected an exit from Electrical Idle at least once since entering Polling.Active.
- Else Polling.Compliance if either (a) or (b) is true:
 - (a) ~~less than~~not all Lanes from the predetermined ~~number~~set of Lanes from (ii) above have detected an exit from Electrical Idle since entering Polling.Active.
 - (b) any Lane that detected a Receiver during Detect received eight consecutive TS1 Ordered Sets (or their complement) with the Lane and Link numbers set to PAD (K23.7), the Compliance Receive bit (bit 4 of Symbol 5) is 1b, and the Loopback bit (bit 2 of Symbol 5) is 0b.

...

In Section 4.2.6.2.2, page 220, line 21, make the following edits:

...

- (b) Next state is Polling.Active if any of the following three conditions are true:
 1. Electrical Idle exit has been detected at the Receiver of any Lane that detected a Receiver during Detect, and the Enter Compliance bit in the Link Control 2 register is 0b.

Note: the “any Lane” condition supports the usage model described in the “Compliance Load Board Usage to Generate Compliance Patterns” Implementation Note earlier in this section.
 2. The Enter Compliance bit in the Link Control 2 register has changed to 0b (from 1b) since entering Polling.Compliance.
 3. The component is a Downstream component, the Enter Compliance bit in the Link Control 2 register is set to 1b and an EIOS has been detected on any Lane. The Enter Compliance bit is reset to 0b when this condition is true.

...

B40 Completion Rules

In Section 2.2.9, page 96, line 23, add the following text:

- ❑ Completion headers must supply the same values for the Attribute as were supplied in the header of the corresponding Request, except as explicitly allowed when IDO is used (see Section 2.2.6.4).
- ❑ [the TH bit is reserved for Completions.](#)
- ❑ AT[1:0] must be 00b. Receivers are not required or encouraged to check this.

B41 Completion Handling Rules

In Section 2.3.2, page 118, line 26, make the following edits:

...

Alternatively, Root Complex behavior may be managed through the CRS Software Visibility Enable bit in ~~an~~[the RCRB Control register](#)~~Header Capability~~ as described in Section ~~7.8.127.20.4~~, permitting the behavior of one or more Root Ports or Root Complex Integrated Endpoints to be controlled by a single Enable bit.

...

B42 Flow Control Initialization Protocol

In Section 3.3, page 156, line 8, make the following edit:

...

In addition, when additional Virtual Channels (VCs) are enabled, the Flow Control initialization process must be completed for each newly enabled VC before it can be used (see Section ~~2.4.22.6.1~~).

...

B43 Reduce Swing Corrections


In Section 4.3.3.2, page 266, make the following edits:

4.3.3.2 ~~Low~~Reduced and Full Swing Transmitter Output Levels

Both the 2.5 GT/s and 5.0 GT/s PCI Express specifications define two voltage swing levels: full swing and ~~low~~reduced swing. Full swing signaling implements de-emphasis, while ~~low~~reduced swing does not. Typically, ~~low~~reduced swing is specified for power sensitive applications where a shorter channel is acceptable. The requirement as to whether a Transmitter need support full swing, ~~low~~reduced swing, or both modes, is dependent on its usage model. The method by which the output mode is selected is not explicitly defined in this specification, and may be implementation dependent. Note: All PCI Express device Transmitters must support full swing signaling, while support for ~~half~~reduced swing signaling is optional.

While two different Transmitter output signaling levels are defined, only a single Receiver specification is defined; this implies that margins (as specified at the Receiver) are identical regardless of the Transmitter's output swing capabilities. It also implies that the channel's characteristics need to be matched to the Transmitter output swing. Typically, ~~low~~reduced swing output is utilized for short channels, such as would occur in mobile platforms.

In Section 4.3.3.3, page 267, replace Figure 4-22 with the following:



A-0574A

Figure 4 22: Transmitter Margining Voltage Levels and Codes

In Section 4.3.3.3, page 267, make the following edits on lines 13 and 14:

Two modes of Tx margining are defined. All 5.0 GT/s capable Transmitters must implement full-swing margining, and those 5.0 GT/s capable Transmitters supporting ~~half-reduced~~ swing signaling must additionally implement ~~halfreduced~~ swing margining.

In Section 4.3.3.5, page 271, make the following edit on line 18:

9. ~~LowReduced~~ swing output, defined by $V_{TX-DIFF-PP-LOW}$ must be implemented as shown in Figure 4 27 with no de-emphasis.

In Section 4.3.3.8, page 273, make the following edit on line 16:

...

Otherwise, the data is driven to a partial swing. De-emphasis is not implemented for low-power signaling; in this case, the Transmitter drives a ~~halfreduced~~ swing output at all times for all transitions.

In Section 4.3.3.8, page 274, make the following edit the title of Figure 4-27:

Figure 4 27: ~~LowReduced~~ Swing Tx Parameters

In Section 4.3.6.2.6, page 304, change the section heading as follows:

4.3.6.2.6. Specifying the Channel for the ~~Low VoltageReduced~~ Swing Option

In Section 4.3.6.2.6, page 304, make the following edits in lines 2 and 3:

Certain PCI Express applications that are power sensitive, such as mobile, may be implemented using ~~lowreduced~~ swing Transmitter option. This involves the use of a ~~halfreduced~~ swing transmit signal with no de-emphasis. The procedure for specifying a channel for the ~~lowreduced~~ swing Transmitter is identical to that used for the full swing Transmitter, with the exception that the worst case behavioral Tx characteristics must reflect the ~~lowreduced~~ swing and lack of de-emphasis.

B44 Transmitter De-emphasis

In Section 4.3.3.9, page 276, make the following edits on line 2:

...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 minimum and maximum values as specified by [the appropriate \$V_{TX-DIFF-PP}\$ or \$V_{TX-DIFF-PP-LOW}\$ parameters](#) in Table 4-9.

B45 Transmitter and Receiver Termination

In Section 4.3.5.4, page 293, make the following edit on line 14:

- The Transmitter is required to meet $R_{LTX-DIFF}$ and R_{LTX-CM} (see Table 4-9 and Figure [4-234-34](#)) any time functional differential signals are being transmitted.

B46 Procedure for Channel Measurement and Margin Extraction

In Section 4.3.6.2.1, page 293, make the following edits on line 9:

A channel may be characterized to specification by means of the following procedure.

1. s-parameters for the channel under test are measured. These include insertion loss, return loss, and (if applicable) crosstalk.
2. The s-parameters are imported into a simulation environment along with a reference load and a behavioral Transmitter model.
3. Worst case Transmitter corners and data patterns are driven through the channel model into a reference load. Then the resulting eye margins are post processed to account for Refclk and Transmitter jitter effects that were not included in the simulation.
4. The resulting eye margins are compared to T_{RX-EYE} ~~and V_{RX-EYE}~~ in Table 4-12 ~~and V_{RX-EYE} in~~ [Table 4-10 and Table 4-11](#) to determine if the channel is within specification.

B47 Data Register

In Section 7.15.3, page 606, make the following edit to Table 7-63:

Table 7 63: Power Budgeting Data Register

Bit Location	Register Description	Attributes
20:18	<p>Power Rail – Specifies the power rail of the operating condition being described.</p> <p>Defined encodings are:</p> <p>000b Power (12V)</p> <p>001b Power (3.3V)</p> <p>010b Power (<u>1.5V or</u> 1.8V)</p> <p>111b Thermal</p> <p>All other encodings are reserved.</p>	RO

B48 Flow Control Rules

In Section 2.6.1, page 134, make the following edits on line 22:

- If only the Data or header advertisement (but not both) for a given type (~~NP~~, NP, or CPL) has been made with infinite credits during initialization, 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.

B49 Receiver Loopback

In Section 4.3.4.9, page 292, make the following edits on line 3:

4.3.4.9. Receiver Loopback

Receivers ~~capable of generating the Modified Compliance Pattern; whether operating at 2.5 GT/s or 5.0 GT/s,~~ must implement ~~a loopback Receiver Error Count~~ (described in Section ~~4.2.9.4.2.6.2.2~~) when placed in the appropriate compliance mode. ~~In this mode Receivers must verify that the incoming data is free of errors, and if not, an error count is incremented and returned via the corresponding Receiver Port. This mechanism permits an easy monitoring of the cumulative error count by observing the returned data stream and decoding the appropriate fields.~~

~~Some details of the error check and loopback circuits are implementation dependent and are not covered in this specification. However, the basic requirements are listed below:~~

- ~~The error check circuit must increment the error count by one and only one for each error detected.~~

- ~~☐ In order to guarantee no degradation in signal integrity, the returned error count must obey the run length and ones/zeros disparity required of 8b/10b coding.~~
- ~~☐ Error check/return must be implemented on a per-Lane basis.~~
- ~~☐ Error codes are defined in Section 4.2.9.~~

B50 Optional Errors

In Section 2.2.5, page 69, make the following edit to line 12:

- ☐ Receivers may optionally check for violations of the Byte Enables rules specified in this section. If a Receiver implementing such checks determines that a TLP violates one or more Byte Enables rules, the TLP is a Malformed TLP. These checks are independently optional (see Section 6.2.3.4).
 - If Byte Enables rules are checked, a violation is a reported error associated with the Receiving Port (see Section 6.2)

In Section 2.2.7, page 78, make the following edit to line 3:

- ☐ I/O Requests have the following restrictions:

...

- Length[9:0] must be 00 0000 0001b
- Last DW BE[3:0] must be 0000b

Receivers may optionally check for violations of these rules. These checks are independently optional (see Section 6.2.3.4). If a Receiver implementing these checks determines that a TLP violates these rules, the TLP is a Malformed TLP.

- ◆ If checked, this is a reported error associated with the Receiving Port (see Section 6.2).

In Section 2.2.7, page 78, make the following edit to line 20:

- ☐ Configuration Requests have the following restrictions:

...

- Last DW BE[3:0] must be 0000b

Receivers may optionally check for violations of these rules. These checks are independently optional (see Section 6.2.3.4). If a Receiver implementing these checks determines that a TLP violates these rules, the TLP is a Malformed TLP.

- ◆ If checked, this is a reported error associated with the Receiving Port (see Section 6.2).

In Section 2.3, page 103, make the following edit to line 32:

- It is an error to receive any Msg/MsgD Request other than a PME_TO_Ack that specifies 101b routing. It is an error to receive a PME_TO_Ack at the Upstream Port of a Switch. Switches may optionally check for violations of these rules. These checks are independently optional (see Section 6.2.3.4).~~TLPs in violation are Malformed TLPs.~~ If checked, violations are Malformed TLPs, and are this is a reported errors associated with the Receiving Port (see Section 6.2)

In Section 4.2.2, page 192, make the following edit to line 19:

- ❑ Receivers may optionally check for violations of the rules of this section. These checks are independently optional (see Section 6.2.3.4). ~~If checked, Any such~~ violations are is a Receiver Errors, and ~~is are~~ reported errors associated with the Port (see Section 6.2).

In Section 6.2.3, page 371, add the following new section on line 26:

6.2.3.4 Optional Error Checking

This specification contains a number of optional error checks. Unless otherwise specified, behavior is undefined if an optional error check is not performed and the error occurs.

When an optional error check involves multiple rules, unless otherwise specified, each rule is independently optional. An implementation may check against all of the rules, none of them or any combination.

Unless otherwise specified, implementation specific criteria are used in determining whether an optional error check is performed.

In Section 6.2.4.2, page 374, make the following edits to line 37:

Since an implementation only has the ability to record a finite number of headers, it is important that software services the First Error Pointer and Header Log registers in a timely manner, to limit the risk of missing this information for subsequent errors. ~~If an error requiring header logging is detected but the header cannot be recorded, then a Header Log Overflow error is reported by the Function. This~~ A Header Log Overflow occurs when an error that requires header logging is detected and either the number of recorded headers supported by an implementation has been reached, or the Multiple Header Recording Enable bit is not Set and the First Error Pointer register is valid.

Implementations may optionally check for this condition and report a Header Log Overflow error. This is a reported error associated with the detecting Function.

The setting of Multiple Heading Recording Capable and the checking for Header Log Overflow are independently optional.

In Section 6.2.9, page 386, make the following edits to line 13:

Reporting of Corrected Internal Errors and Uncorrectable Internal Errors is independently optional. If either is Internal Errors are reported, then AER must be implemented.

B51 Link Status 2 Register

In Section 7.8.20, page 551, make the following edit to Table 7-27:

Table 7-27: Link Status 2 Register

Bit Location	Register Description	Attributes
0	<p>Current De-emphasis Level – When the Link is operating at 5 GT/s speed, this bit reflects the level of de-emphasis.</p> <p>Encodings:</p> <p>1b -3.5 dB</p> <p>0b -6 dB</p> <p>The value in this bit is undefined when the Link is operating at 2.5 GT/s speed.</p> <p>Components that support only the 2.5 GT/s speed are permitted to hardwire this field to 0b.</p> <p><u>For components that support speeds greater than 2.5 GT/s, Multi-Function devices associated with an Upstream Port must report the same value in this field for all Functions of the Port.</u></p>	RO

B52 Multicast Capability Register

In Section 7.21.2, page 635, make the following edits to Table 7-86:

Table 7-86: Multicast Capability Register

Bit Location	Register Description	Attributes
15	<p>MC_ECRC_Regeneration_Supported – If Set, indicates that ECRC regeneration is supported.</p> <p>This bit must not be Set unless the Function supports Advanced Error Reporting, and the ECRC Check Capable bit in the Advanced Error Capabilities and Control register is also Set. However, if ECRC regeneration is supported, its operation is not contingent upon the setting of the ECRC Check Enable bit in the Advanced Error Capabilities and Control register. <u>This bit is applicable to Switch and Root Ports and is RsvdP in all other Functions.</u></p>	RO/ <u>RsvdP</u>

B53 Multicast TLP Processing

In Section 6.14.1, page 442, make the following edits to line 16:

In this step, each Switch Ingress Port and other components use values of MC_Enable, MC_Base_Address, MC_Index_Position, and MC_Num_Group from any one of their Functions. Software is required to configure all Functions [of a Switch and all Functions of a Multi-Function Upstream Port](#) to have the same values in each of these fields and results are indeterminate if this is not the case.

B54 Data Poisoning

In Section 2.7.2.2, page 147, make the following edits to line 8:

- Data poisoning applies only to the data within a Write Request (Posted or Non-Posted), [a Message with Data](#), an AtomicOp Request, a Read Completion, or an AtomicOp Completion.
 - Poisoning of a TLP is indicated by a 1b value in the EP field
 - Transmitters are permitted to set the EP field to 1b only for TLPs that include a data payload. The behavior of the receiver is not specified if the EP bit is set for any TLP that does not include a data payload.

B55 LCRC and Sequence Number Rules

In Section 3.5.3.1, page 179, make the following edit to line 10:

This is a [Bad TLP error and is a](#) reported error associated with the Port (see Section 6.2).

In Section 3.5.3.1, page 179, make the following edit to line 20:

This is a [Bad TLP error and is a](#) reported error associated with the Port (see Section 6.2).

In Section 3.5.3.1, page 179, make the following edits to lines 31 and 32:

...

- 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 immediately
 - set the NAK_SCHEDULED flag
 - ~~report TLP missing~~

This is a Bad TLP error and is a reported error associated with the Port (see Section 6.2).

Regardless of the state of the NAK_SCHEDULED flag, it is permitted for this to be a reported error associated with the Port (see Section 6.2), and this permitted behavior is illustrated in Figure 3-17. However, in order to prevent error pollution it is recommended that the Port only report such an error when the NAK_SCHEDULED flag is clear.

- If the TLP Sequence Number is equal to the expected value stored in NEXT_RCV_SEQ:

...

B56 Device Capabilities 2 Register

In Section 7.8.15, pages 542 and 543, make the following edits to Table 7-24:

Bit Location	Register Description	Attributes
21	<p>End-End TLP Prefix Supported – Indicates whether End-End TLP Prefix support is offered by a Function. Values are:</p> <p>0b No Support</p> <p>1b Support is provided to receive TLPs containing End-End TLP Prefixes.</p> <p>This bit is Hwlnit for Root Ports and is RO for all other Functions.</p> <p>All Ports of a Switch must have the same value for this bit.</p>	<p>RO/Hwlnit (see description)</p>

Bit Location	Register Description	Attributes
23:22	<p>Max End-End TLP Prefixes – Indicates the maximum number of End-End TLP Prefixes supported by this Function. TLPs received by this Function that contain more End-End TLP Prefixes than are supported must be handled as Malformed TLPs (see Section 2.2.10.2). Values are:</p> <p>01b 1 End-End TLP Prefix 10b 2 End-End TLP Prefixes 11b 3 End-End TLP Prefixes 00b 4 End-End TLP Prefixes</p> <p>If End-End TLP Prefix Supported is Clear, this field is RsvdP.</p> <p>This field is HwInit for Root Ports and is RO for all other Functions. All Root Ports that have the End-End TLP Prefix Supported bit Set must contain the same value for this field.</p> <p>For Switches where End-End TLP Prefix Supported is Set, this field must be 00b indicating support for up to four End-End TLP Prefixes (see Section 2.2.10.2).</p>	<u>HwInt</u>

B57 Link Control 2 Register

In Section 7.8.19, page 548, make the following edits to Table 7-26:

Bit Location	Register Description	Attributes
4	<p>Enter Compliance – Software is permitted to force a Link to enter Compliance mode (at the speed indicated in the Target Link Speed field <u>and the de-emphasis level indicated by the Compliance De-emphasis bit</u>) by setting this bit to 1b in both components on a Link and then initiating a hot reset on the Link.</p> <p>Default value of this bit following Fundamental Reset is 0b.</p> <p>For a Multi-Function device associated with an Upstream Port, the bit in Function 0 is of type RWS, and only Function 0 controls the component's Link behavior. In all other Functions of that device, this bit is of type RsvdP.</p> <p>Components that support only the 2.5 GT/s speed are permitted to hardwire this bit to 0b.</p> <p><u>This bit is intended for debug, compliance testing purposes only. System firmware and software is allowed to modify this bit only during debug or compliance testing. In all other cases, the system must ensure that this bit is set to the default value.</u></p>	RWS/RsvdP (see description)

B58 I/O Transactions

In Section 2.1.1.2, page 53, make the following edits to line 10:

PCI Express supports I/O Space for compatibility with legacy devices which require their use. Future revisions of this specification ~~are expected to may~~ deprecate the use of I/O Space. I/O Transactions include the following types:

- Read Request/Completion
- Write Request/Completion

I/O Transactions use a single address format:

- Short Address Format: 32-bit address

B59 First/Last DW Byte Enables Rules

In Section 2.2.5, page 67, make the following edits to line 8:

Byte Enables are included with Memory, I/O, and Configuration Requests. This section defines the corresponding rules. Byte Enables, when present in the Request header, are located in byte 7 of the header (see Figure 2-11). For Memory Read Requests that have the TH bit Set, the Byte Enable fields are repurposed to carry the ST[7:0] field, and values for the Byte Enables are implied as defined below. Such Requests must only be issued when it is acceptable to complete the Requests as if all bytes for requested payload were enabled.

- For Memory Reads that have the TH bit Set, the following values are implied for the Byte Enables. [See Section 2.2.7 for additional requirements.](#)
 - If the Length field for this Request indicates a length of 1 DW, then the value for the 1st DW Byte Enables is implied to be 1111b and the value for the Last DW Byte Enables is implied to be 0000b.
 - If the Length field for this Request indicates a length of greater than 1 DW, then the value for the 1st DW Byte Enables and the Last DW Byte Enables is implied to be 1111b.

B60 Request Handling Rules

In Section 2.3.1, page 106, make the following edits to line 6:

- If the Request is a Message, and the Message Code specifies a value that is undefined, or that corresponds to a Message not supported by the device Function, (other than Vendor_Defined Type 1 which is not treated as an error – see Section 2.2.8.6), the Request is an Unsupported Request, and is reported according to Section 6.2
 - If the Message Code is a supported value, process the Message according to the corresponding Message processing rules; if the Message Code is an ~~Ignored Message~~, [and the receiver is ignoring it](#), ignore the Message without reporting any error (see Section 2.2.8.7)

B61 Flow Control Rules

In Section 2.6.1, page 131, make the following edits to line 13:

...

- For headers, ~~the unit of Flow Control credit is one maximum-size header plus TLP digest:~~
 - The unit of Flow Control credit for Receivers that don't support TLP Prefixes is the sum of one maximum-size Header and TLP Digest.
 - The unit of Flow Control credit for Receivers that support End-End TLP Prefixes is sum of one maximum-size Header, TLP Digest and the maximum number of End-End TLP Prefixes permitted in a TLP.
 - The management of Flow Control for Receivers that support Local TLP Prefixes is dependent on the Local TLP Prefix type.

B62 Data Link Control and Management State Machine Rules

In Section 3.2.1, page 155, make the following edits to line 23:

...

- ◆ Upstream components are optionally permitted to treat this transition from DL_Active to DL_Inactive as a Surprise Down error, except in the following cases where this error detection is blocked:
 - If the Secondary Bus Reset in Bridge Control register has been set to 1b by software, then the subsequent transition to DL_Inactive must not be considered an error.
 - If the Link Disable bit has been set to 1b by software, then the subsequent transition to DL_Inactive must not be considered an error.
 - If a PME_Turn_Off Message has been sent through this Port, then the subsequent transition to DL_Inactive must not be considered an error.

Note that the DL_Inactive transition for this condition will not occur until a power off, a reset, or a request to restore the Link is sent to the Physical layer.

Note also that in the case where the PME_Turn_Off/PME_TO_Ack handshake fails to complete successfully, a Surprise Down error may be detected.

- If the Port is associated with a hot-pluggable slot (Hot-Plug Capable bit in the Slot Capabilities register set to 1b), and the Hot-Plug Surprise bit in the Slot Capabilities register is set to 1b, then any transition to DL_Inactive must not be considered an error.

...

B63 LCRC and Sequence Number Rules (TLP Transmitter)

In Section 3.5.2.1, page 166, make the following edits to line 18:

...

□ The following timer is used:

- REPLAY_TIMER - Counts time that determines when a replay is required, according to the following rules:
 - ◆ 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
 - ◆ ~~Resets and R~~estarts for each Ack DLLP received while there are ~~more~~ 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

B64 Transmitter De-emphasis

In Section 4.3.3.9, page 275, make the following edits to line 12:

For full swing signaling only, 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 ± 0.5 dB at 2.5 GT/s and either -3.5 ± 0.5 dB or -6 ± 0.5 dB at 5.0 GT/s) below the first bit. At 5.0 GT/s de-emphasis is selectable ~~as a~~ via ~~various methods including~~ configuration register bits. The two de-emphasis values defined for 5.0 GT/s operation permit optimum equalization for both short, reflection dominated channels, and long, loss dominated ones. 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 minimum and maximum values as specified by $V_{TX-DIFF-PP}$ in Table 4-9.

The de-emphasis level is defined via configuration register bits in Chapter 7 and Section 4.2.

The only exception pertains to transmitting the Beacon (see Section 4.3.5.8).

B65 Receiver Specification

In Section 4.3.4, page 281, make the following edits to line 7:

For 2.5 GT/s, the parameters defined in Table 4-12 are defined at the Receiver pins. As such, they do not directly measure a Receiver's performance. There is instead an implied correlation between the margins observed at the Receiver's pins and the BER, but no prescribed methodology for measuring it.

At 5.0 GT/s it becomes necessary to define a performance-based methodology for tolerancing a Receiver. Tolerancing is a two step procedure in which a test apparatus is calibrated to yield the worst case signal margins defined in Table 4-12 or Table 4-11 into a test load. The margins applied to a Receiver are dependent on the clocking architecture with which the Receiver will operate. Therefore two different sets of margining parameters are defined for the common Refclk clock and data clocked Rx architectures. The Receiver under test then replaces the test load, and its BER is observed. The following sections describe the Receiver tolerancing procedure in detail.

B66 Link Training and Status State Machine (LTSSM) Descriptions

In Section 4.2.5, page 208, make the following edits to line 26:

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/~~ Fundamental Reset. All counter values must be set to the specified values after ~~power-on/~~ Fundamental Reset.

B67 Link Training and Status State Rules

In Section 4.2.6, page 212, make the following edits to line 3:

Various Link status bits are monitored through software with the exception of LinkUp which is monitored by the Data Link Layer. Table 4-7 describes how the Link status bits must be handled throughout the LTSSM (for more information, see Section 3.4.3.2 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).

B68 Recovery.Speed and Loopback.Entry

In Section 4.2.6.4.2, page 238, make the following edits to line 36:

- The Transmitter enters Electrical Idle and stays there until the Receiver Lanes have entered Electrical Idle, and then additionally remains there for at least 800 ns on a successful speed negotiation (i.e., successful_speed_negotiation = 1b) or at least 6 μs on an unsuccessful speed negotiation (i.e., successful_speed_negotiation = 0b), but stays there no longer than an additional 1 ms. The frequency of operation is permitted to be changed to the new data rate only after the Receiver Lanes have entered Electrical Idle. If the negotiated data rate is 5.0 GT/s and if operating in full swing mode, -6 dB de-emphasis level must be selected for operation if the select_deemphasis variable is 0b and -3.5 dB de-emphasis level must be selected for operation if the select_deemphasis variable is 1b. Note that if the link is already operating at the highest data rate supported by both components, Recovery.Speed is executed but the data rate is not changed.

...

In Section 4.2.6.10.1, page 254, make the following edits to line 21:

- The loopback slave must send one EIOS if the current Link speed is 2.5 GT/s or two consecutive EIOSs if the current Link speed is greater than 2.5 GT/s and then immediately go to Electrical Idle for 2 ms. During this time the speed must be changed to the highest common transmitted and received Link speed advertised on two consecutive TS1 Ordered Sets. The select_deemphasis variable must be set equal to the Selectable De-emphasis bit (bit 6 of Symbol 4) in the two consecutive TS1 or TS2 Ordered Sets prior to entry to this substate. If the speed is 5 GT/s, and if operating in full swing mode, the de-emphasis level must be chosen to be -3.5 dB if select_deemphasis is 1b else the de-emphasis level must be -6 dB.

B69 Device Capabilities Register

In Section 7.8.3, page 503, make the following edits to Table 7-12:

Bit Location	Register Description	Attributes
25:18	<p>Captured Slot Power Limit Value (Upstream Ports only) – In combination with the Captured Slot Power Limit Scale value, specifies the upper limit on power <u>supplied by slot available to the adapter</u>.</p> <p>Power limit (in Watts) <u>is</u> calculated by multiplying the value in this field by the value in the Captured Slot Power Limit Scale field <u>except when the Captured Slot Power Limit Scale field equals 00b (1.0x) and the Captured Slot Power Limit Value exceeds EFh, then alternative encodings are used (see Section 7.8.9)</u>.</p> <p>This value is set by the Set_Slot_Power_Limit Message or hardwired to 00h (see Section 6.9). The default value is 00h.</p>	RO

B70 Device Control Register

In Section 7.8.4, page 505, make the following edits to Table 7-13:

Bit Location	Register Description	Attributes
3	<p>Unsupported Request Reporting Enable – This bit, in conjunction with other bits, controls the signaling of Unsupported Requests Errors by sending error Messages (see Section 6.2.5 and Section 6.2.6 for details). For a multi-Function device, this bit controls error reporting for each Function from point-of-view of the respective Function.</p> <p>A Root Complex Integrated Endpoint that is not associated with a Root Complex Event Collector is permitted to hardwire this bit to 0b.</p> <p>Default value of this bit is 0b.</p>	RW

B71 Link Status Register

In Section 7.8.8, page 526, make the following edits to Table 7-17:

Bit Location	Register Description	Attributes
13	<p>Data Link Layer Link Active – This bit indicates the status of the Data Link Control and Management State Machine. It returns a 1b to indicate the DL_Active state, 0b otherwise.</p> <p>This bit must be implemented if the corresponding-Data Link Layer Link Active Reporting Capablecapability bit is 1implemented. Otherwise, this bit must be hardwired to 0b.</p>	RO

B72 Slot Capabilities Register

In Section 7.8.9, page 529, make the following edits to Table 7-18:

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 <u>the slot (see Section 6.9) or by other means to the adapter.</u></p> <p>Power limit (in Watts) <u>is</u> calculated by multiplying the value in this field by the value in the Slot Power Limit Scale field except when the Slot Power Limit Scale field equals 00b (1.0x) and Slot Power Limit Value exceeds EFh, the following alternative encodings are used:</p> <p style="padding-left: 40px;">F0h = 250 W Slot Power Limit</p> <p style="padding-left: 40px;">F1h = 275 W Slot Power Limit</p> <p style="padding-left: 40px;">F2h = 300 W Slot Power Limit</p> <p style="padding-left: 40px;">F3h to FFh = reserved</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

B73 Slot Control Register

In Section 7.8.10, page 531, make the following edits to Table 7-19:

Bit Location	Register Description	Attributes
7:6	<p>Attention Indicator Control – If an Attention Indicator is implemented, writes to this field set the Attention Indicator to the Written<u>written</u> state.</p> <p>Reads of this field must reflect the value from the latest write, even if the corresponding hot-plug command is not complete, unless software issues a write without waiting, <u>if required to</u>, for the previous command to complete in which case the read value is undefined.</p> <p>Defined encodings are:</p> <p>00b Reserved</p> <p>01b On</p> <p>10b Blink</p> <p>11b Off</p> <p>Note: The default value of this field must be one of the non-Reserved values. If the Attention Indicator Present bit in the Slot Capabilities register is 0b, this bit is permitted to be read-only with a value of 00b.</p>	RW
9:8	<p>Power Indicator Control – If a Power Indicator is implemented, writes to this field set the Power Indicator to the Written<u>written</u> state. Reads of this field must reflect the value from the latest write, even if the corresponding hot-plug command is not complete, unless software issues a write without waiting, <u>if required to</u>, for the previous command to complete in which case the read value is undefined.</p> <p>Defined encodings are:</p> <p>00b Reserved</p> <p>01b On</p> <p>10b Blink</p> <p>11b Off</p> <p>Note: The default value of this field must be one of the non-Reserved values. If the Power Indicator Present bit in the Slot Capabilities register is 0b, this bit is permitted to be read-only with a value of 00b.</p>	RW

Bit Location	Register Description	Attributes
10	<p>Power Controller Control – If a Power Controller is implemented, this bit when written sets the power state of the slot per the defined encodings. Reads of this bit must reflect the value from the latest write, even if the corresponding hot-plug command is not complete, unless software issues a write, <u>if required to</u>, without waiting for the previous command to complete in which case the read value is undefined.</p> <p>Depending on the form factor, the power is turned on/off either to the slot or within the adapter.–Note that in some cases the power controller may autonomously remove slot power or not respond to a power-up request based on a detected fault condition, independent of the Power Controller Control setting.</p> <p>The defined encodings are:</p> <p>0b Power On</p> <p>1b Power Off</p> <p>If the Power Controller Implemented<u>Present</u> bit in the Slot Capabilities register is Clear, then writes to this bit have no effect and the read value of this bit is undefined.</p>	RW
12	<p>Data Link Layer State Changed Enable – If the Data Link Layer Link Active Reporting Ccapability is implemented, when Set 1b, this bit enables software notification when Data Link Layer Link Active Reporting bit is changed.</p> <p>If the Data Link Layer Link Active Reporting <u>Capable</u> bit<u>Capability</u> is not implemented<u>0b</u>, this bit is permitted to be read-only with a value of 0b.</p> <p>Default value of this bit is 0b.</p>	RW

B74 Link Address for Link Type 1

In Section 7.13.3.2.2, page 594, make the following edits to line 7:

For a Link pointing to the Configuration Space of a Root Complex element (Link Type bit = 1), bits in the first DWORD specify the Bus, Device, and Function Number of the target element. As shown in Figure ~~7-627-64~~, bits 2:0 (N) encode the number of bits n associated with the Bus Number, with N = 000b specifying $n = 8$ and all other encodings specifying $n = \langle \text{value of N} \rangle$. Bits 11:3 are reserved and hardwired to 0. Bits 14:12 specify the Function Number, and bits 19:15 specify the Device Number. Bits $(19 + n):20$ specify the Bus Number, with $1 \leq n \leq 8$.

B75 Data Select Register/Data Register

In Section 7.15.2, page 603, make the following edits to line 6:

This read-write register indexes the Power Budgeting Data reported through the Data register and selects the DWORD of Power Budgeting Data that ~~should~~is to appear in the Data register. Index values for this register start at zero to select the first DWORD of Power Budgeting Data; subsequent DWORDs of Power Budgeting Data are selected by increasing index values.

In Section 7.15.3, page 604, make the following edits to line 7:

This read-only register returns the DWORD of Power Budgeting Data selected by the Data Select register. Each DWORD of the Power Budgeting Data describes the power usage of the device in a particular operating condition. Power Budgeting Data for different operating conditions is not required to be returned in any particular order, as long as incrementing the Data Select register causes information for a different operating condition to be returned. If the Data Select register contains a value greater than or equal to the number of operating conditions for which the device provides power information, this register ~~should~~must return all zeros. Figure 7-72 details allocation of register fields in the Power Budgeting Data register; Table 7-63 provides the respective bit definitions.

In Section 7.15.3, page 606, make the following edits to Table 7-63:

Bit Location	Register Description	Attributes
20:18	<p>Power Rail – Specifies the thermal load or power rail of the operating condition being described.</p> <p>Defined encodings are:</p> <p>000b Power (12V)</p> <p>001b Power (3.3V)</p> <p>010b Power (1.8V)</p> <p>111b Thermal</p> <p>All other encodings are reserved.</p>	RO

In Section 7.15.3, page 606, make the following edits to line 2:

A device that implements the Power Budgeting Capability is required to provide data values for the D0 Maximum and D0 Sustained PM State ~~/and~~ Type combinations for every power rail from which it consumes power; data for the D0 Maximum ~~Thermal~~ and D0 Sustained for Thermal ~~combinations~~ must also be provided if these values are different from the sum of the values for an operating condition reported for D0 Maximum and D0 Sustained on the power rails.

Devices that support auxiliary power or PME from auxiliary power must provide data for the appropriate power type (Auxiliary or PME Aux).

B76 Various Register Corrections

In Section 7.16.4, page 612, make the following edits to Table 7-68:

Bit Location	Register Description	Attributes
N-1:0	Egress Control Vector – An N-bit bit-array configured by software, <u>where N is given by the value in the Egress Control Vector Size field</u> . When a given bit is set, peer-to-peer Requests targeting the associated Port, Function, or Function Group are blocked or redirected (if enabled) (see Section 6.12.3). Default value of <u>this field each bit</u> is 0b.	RW

In Section 7.21.5, page 637, make the following edits to Table 7-89:

Bit Location	Register Description	Attributes
MC_Max_Group:0	MC_Receive – For each bit that's Set, this Function gets a copy of any Multicast TLPs for the associated Multicast Group. Bits above MC_Num_Group are ignored by hardware. Default <u>is 0 value of each bit is 0b</u> .	RW
All other bits	Reserved	RsvdP

In Section 7.21.6, page 637, make the following edits to Table 7-90:

Bit Location	Register Description	Attributes
MC_Max_Group:0	MC_Block_All – For each bit that is Set, this Function is blocked from sending TLPs to the associated Multicast Group. Bits above MC_Num_Group are ignored by hardware. Default <u>is 0 value of each bit is 0b</u> .	RW
All other bits	Reserved	RsvdP

In Section 7.21.7, page 638, make the following edits to Table 7-91:

Bit Location	Register Description	Attributes
MC_Max_Group:0	MC_Block_Untranslated – For each bit that is Set, this Function is blocked from sending TLPs containing Untranslated Addresses to the associated MCG. Bits above MC_Num_Group are ignored by hardware. Default <u>is 0 value of each bit is 0b</u> .	RW
All other bits	Reserved	RsvdP

B77 Egress Control Vector

In Section 7.16.4, page 612, make the following edits to line 6:

The following examples illustrate how the vector might be configured:

- ❑ For an 8-Port Switch, each Port will have a separate vector indicating which Downstream Egress Ports it may forward Requests to.
 - Port 1 being not allowed to communicate with any other Downstream Ports would be configured as: 1111 1100b with ~~0b indicating in~~-bit 0 ~~corresponding~~~~corresponds~~ to the Upstream Port (~~hardwired to 0b~~) and ~~a 0b in~~-bit 1 ~~represents~~~~corresponding to~~ the Ingress Port (~~hardwired to 0b~~)~~as well~~.
 - Port 2 being allowed to communicate with Ports 3, 5, and 7 would be configured as: 0101 0010b.
- ❑ For a 4-Function device, each Function will have a separate vector that indicates which Function it may forward Requests to.
 - Function 0 being not allowed to communicate with any other Functions would be configured as: 1110b with ~~0b in~~-bit 0 corresponding to Function 0 (~~hardwired to 0b~~).
 - Function 1 being allowed to communicate with Functions 2 and 3 would be configured as: 0001b with ~~a 0b in~~-bit 1 corresponding to Function 1 (~~hardwired to 0b~~)~~as well~~.

B78 Link State Power Management

In Section 5.2, page 315, make the following edits to line 9:

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

L0s support is required for ASPM. It is not applicable to PCI-PM compatible power management.

All main power supplies, component reference clocks, and components’ internal PLLs must be active at all times during L0s. TLP and DLLP transmission is disabled for a Port whose Link is in L0sTx L0s.

The Physical Layer provides mechanisms for quick transitions from this state to the L0 state. When common (distributed) reference clocks are used on both sides of a Link, the transition time from L0s to L0 is typically less than 100 Symbol Times.

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.

B79 Entry into the L0s State

In Section 5.4.1.1.1, page 339 make the following edits to line 12:

Definition of Idle

The definition of an “idle” Upstream Port varies with device Function category. An Upstream Port of a multi-Function device is considered idle only when all of its Functions are idle.

A ~~non-Switch Endpoint Function or a Root~~ Port is determined to be idle if the following conditions are met:

...

B80 L1 ASPM State

In Section 5.4.1.2, page 340 make the following edits to line 31:

A component may optionally support the ASPM L1 state; a state that provides greater power savings at the expense of longer exit latency. L1 exit latency is visible to software, and reported via the ~~Link Capabilities register defined in Section 7.8.6~~L1 Exit Latency field.

When supported, L1 entry is disabled by default in the ASPM Control configuration field. Software must enable ASPM L1 on the Downstream component only if it is supported by both components on a Link. Software must sequence the enabling and disabling of ASPM L1 such that the Upstream component is enabled before the Downstream component and disabled after the Downstream component.

B81 ASPM Configuration

In Section 5.4.1.3, page 350 make the following edits to line 2:

...

Each component must support the Common Clock Configuration bit in their Capability structure’s Link ~~StatusControl~~ register. Software writes to this register bit to indicate to the device whether it is sharing the same clock source as the device on the other end of the Link.

B82 Root Complex Event Collector

In Section 1.3.2.3, page 42 make the following edits to line 20:

- A Root Complex Integrated Endpoint must signal PME and error conditions through the same mechanisms used on PCI systems. If a Root Complex Event Collector is implemented, a Root Complex Integrated Endpoint may optionally signal PME and error conditions through a Root Complex Event Collector. In this case, a Root Complex Integrated Endpoint must be associated with ~~exactly no more than~~ one Root Complex Event Collector.

In Section 5.3.1.1, page 329 make the following edits to line 22:

...

If a Root Complex Event Collector is implemented, PME indications that originate from a Root Complex Integrated Endpoint may optionally be reported in a Root Complex Event Collector residing on the same Logical Bus as the Root Complex Integrated Endpoint. The Root Complex Event Collector must explicitly declare supported Root Complex Integrated Endpoints as part of its capabilities; each Root Complex Integrated Endpoint must be associated with exactly no more than one Root Complex Event Collector. Root Complex Event Collectors explicitly identify the logical location of the requesting agent to facilitate quicker PME service routine response.

In Section 6.2.3.2, page 365 make the following edits to line 11:

If a Root Complex Event Collector is implemented, errors that originate from a Root Complex Integrated Endpoint may optionally be sent to the corresponding Root Complex Event Collector. Errors that originate from a Root Complex Integrated Endpoint are reported in a Root Complex Event Collector residing on the same Logical Bus as the Root Complex Integrated Endpoint. The Root Complex Event Collector must explicitly declare supported Root Complex Integrated Endpoints as part of its capabilities; each Root Complex Integrated Endpoint must be associated with exactly no more than one Root Complex Event Collector.

In Section 6.2.4.1.1, page 373 make the following edits to line 11:

If a Root Complex Event Collector is implemented, errors from a Root Complex Integrated Endpoint may optionally be reported in a Root Complex Event Collector residing on the same Logical Bus as the Root Complex Integrated Endpoint. The Root Complex Event Collector must explicitly declare supported Root Complex Integrated Endpoints as part of its capabilities. Each Root Complex Integrated Endpoint must be associated with exactly no more than one Root Complex Event Collector.

B83 Power Controller

In Section 6.7.1.8, page 416 make the following edits:

The power controller is an element composed of one or more discrete components that acts under control of software to set the power state of ~~either the hot-plug slot or the adapter~~ as appropriate for the specific form factor. The power controller must also monitor the slot/~~adapter~~ for power fault conditions (as defined in the associated form factor specification) that occur on the slot/~~adapter~~'s main power rails and, if supported, auxiliary power rail.

If a power controller is not present, the power state of the hot-plug slot/~~adapter~~ must be set automatically by the hot-plug controller in response to changes in the presence of an adapter in the slot.

The power controller monitors main and auxiliary power faults independently. If a power controller detects a main power fault on the hot-plug slot/~~adapter~~, it must automatically set its internal main power fault latch and remove main power from the hot-plug slot/~~adapter~~ (without affecting auxiliary power). Similarly, if a power controller detects an auxiliary power fault on the hot-plug slot/~~adapter~~, it must automatically set its internal auxiliary power fault latch and remove auxiliary power from the hot-plug slot/~~adapter~~ (without affecting main power). Power must remain off to the slot/~~adapter~~ as long as the power fault condition remains latched, regardless of any writes by software to turn on power to the hot-plug slot/~~adapter~~. The main power fault latch is cleared when software turns off power to the hot-plug slot/~~adapter~~. The mechanism by which the auxiliary power fault latch is cleared is form factor specific but generally requires auxiliary power to be removed from the hot-plug slot/~~adapter~~. For example, one form factor may remove auxiliary power when the MRL for the slot is opened while another may require the adapter to be physically removed from the slot. Refer to the associated form factor specifications for specific requirements.

Since the Power Controller Control bit in the Slot Control register reflects the last value written and not the actual state of the power controller, this means there may be an inconsistency between the value of the Power Controller Control bit and the state of the power to the slot/~~adapter~~ in a power fault condition. To determine whether slot/~~adapter~~ power is off due to a power fault, software must use the power fault software notification to detect power faults. To determine that a requested power-up operation has otherwise failed, software must use the hot-plug slot/~~adapter~~ power-up time out mechanism described in Section 6.7.3.3.

Software must not assume that writing to the Slot Control register to change the power state of a hot-plug slot/~~adapter~~ causes an immediate power state transition. After turning power on, software must wait for a Data Link Layer State Changed event, as described in Section 6.7.3.3. After turning power off, software must wait for at least 1 second before taking any action that relies on power having been removed from the hot-plug slot/~~adapter~~. For example, software is not permitted to turn off the power indicator (if present) or attempt to turn on the power controller before completing the 1 second wait period.

B84 Link Speed Management

In Section 6.11, page 428 make the following edits to line 3:

This section describes how Link speed management is coordinated between the LTSSM (Section 4.2.6) and the software Link observation and control mechanisms (Sections 7.8.6, 7.8.7, ~~and 7.8.8, 7.8.18, 7.8.19, and 7.8.20~~).

B85 Configuration Register Types

In Section 6.11, page 478 make the following edits to Table 7-2:

Register Attribute	Description
ROS	<p>Sticky - Read-only – Register bits are read-only and cannot be altered by software. Bits are neither initialized nor modified by hot reset or FLR.</p> <p>Where noted, devices that consume AUX power must preserve sticky register values when AUX power consumption (via either AUX power<u>Aux Power PM Enable</u> or PME Enable) is enabled. In these cases, registers are neither initialized nor modified by hot, warm, or cold reset (see Section 6.6).</p>
RWS	<p>Sticky - Read-Write – Register bits are read-write and are Set or Cleared by software to the desired state. Bits are neither initialized nor modified by hot reset or FLR.</p> <p>Where noted, devices that consume AUX power must preserve sticky register values when AUX power consumption (via either AUX power<u>Aux Power PM Enable</u> or PME Enable) is enabled. In these cases, registers are neither initialized nor modified by hot, warm, or cold reset (see Section 6.6).</p>
RW1CS	<p>Sticky - Write-1-to-clear status – Registers indicate status when read. A Set bit indicates a status event which is Cleared by writing a 1b. Writing a 0b to RW1CS bits has no effect. Bits are neither initialized nor modified by hot reset or FLR.</p> <p>Where noted, devices that consume AUX power must preserve sticky register values when AUX power consumption (via either AUX power<u>Aux Power PM Enable</u> or PME Enable) is enabled. In these cases, registers are neither initialized nor modified by hot, warm, or cold reset (see Section 6.6).</p>

B86 Cache Line Size Register

In Section 7.5.1.3, page 485 make the following edits to line 5:

The Cache Line Size register is set by the system firmware or the operating system to system cache line size. However, note that legacy PCI 3.0 software may not always be able to program this field correctly especially in the case of Hot-Plug devices. This field is implemented by PCI Express devices as a read-write field for legacy compatibility purposes but has no effect on any PCI Express device behavior. For PCI Express to PCI/PCI-X Bridges, refer to the *PCI Express to PCI/PCI-X Bridge Specification* for additional requirements for this register. The default value of this register is 00h.

B87 Base Address Registers

In Section 7.5.3.1, page 489 make the following edits to line 8:

A PCI Express Function 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 Function does not tolerate write merging. It is strongly encouraged that memory-mapped resources be designed as prefetchable whenever possible. PCI Express device Functions 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. The attributes for some of the bits in the BAR are affected by the Resizable BAR Capability, if it is implemented (see Section 7.22).

B88 Device Capabilities Register

In Section 7.8.3, page 501 make the following edits to Table 7-12:

Bit Location	Register Description	Attributes
5	<p>Extended Tag Field Supported – This bit indicates the maximum supported size of the Tag field as a Requester.</p> <p>Defined encodings are:</p> <p>0b 5-bit Tag field supported</p> <p>1b 8-bit Tag field supported</p> <p>Note that 8-bit Tag field generation must be enabled by the Extended Tag Field Enable bit in the Device Control register <u>of the Requester Function</u> before 8-bit Tags can be generated by the Requester.</p>	RO

B89 Link Control Register

In Section 7.8.7, page 520 make the following edits to Table 7-16:

Bit Location	Register Description	Attributes
7	<p>Extended Synch – When Set, this bit forces the transmission of additional Ordered Sets when exiting the L0s state (see Section 4.2.4.5) and when in the Recovery state (see Section 4.2.6.4.1). This mode provides external devices (e.g., logic analyzers) monitoring the Link time to achieve bit and Symbol lock before the Link enters the L0 state and resumes communication.</p> <p>For multi-Function devices if any Function has this bit Set, then the component must transmit the additional Ordered Sets when exiting L0s <u>or when in Recovery</u>.</p> <p>Default value for this bit is 0b.</p>	RW

B90 Root Complex Event Collector

In Section 1.3.4, page 44 make the following edits to line 9:

- A Root Complex Event Collector provides support for terminating error and PME messages from Root Complex Integrated Endpoints.
- A Root Complex Event Collector must follow all rules for a Root Complex Integrated Endpoint.
- A Root Complex Event Collector is not required to decode any memory or IO resources.
- A Root Complex Event Collector has the Base Class 08h, Sub-Class 06h and Programming Interface 00h.
- A Root Complex Event Collector resides on the same Logical Bus as the Root Complex Integrated Endpoints it supports.
- Multiple Root Complex Event Collectors are permitted to reside on a single Logical Bus.
- A Root Complex Event Collector explicitly declares supported Root Complex Integrated Endpoints through the Root Complex Event Collector Endpoint Association Capability.
- Root Complex Event Collectors are optional.

B91 Slot Power Limit Value


In Section 7.8.9, page 529 make the following edits to Table 7-18:

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 except when the Slot Power Limit Scale field equals 00b (1.0x) and Slot Power Limit Value exceeds EFh, the following alternative encodings are used:</p> <p>F0h = 250 W Slot Power Limit</p> <p>F1h = 275 W Slot Power Limit</p> <p>F2h = 300 W Slot Power Limit</p> <p>F3h to FFh = reserved <u>for Slot Power Limit values above 300 W</u></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

C1 Scrambling Spectrum Clarification

In Section C1, page 685, make the following edits on line 1:

At 2.5 GT/s, scrambling produces the power spectrum (in the 10-bit domain) shown in Figure C-1.



OM14292A

Figure C-1: Scrambling Spectrum at 2.5 GT/s for Data Value of 0

C2 Capitalization Corrections

In Section 4.3.1.2, page 263, make the following edits starting on line 10:

4.3.1.2. Component Interfaces

PCI Express components from different manufacturers must interoperate reliably together. At the electrical level, this is achieved by specifying a set of parameters and the interfaces at which those parameters must be met. For 5.0 GT/s PCI Express components, this interface is defined to be at the pins of the Receiver and Transmitter devices and the corresponding locations at either end of the channel. At 2.5 GT/s the requirement to reference all measurements to the Tx or Rx pins is less explicitly stated in the specification, but still holds true. The pin location was also chosen because it permits a component to be characterized independently of others. For example, a Transmitter can be measured without needing any other PCI Express components.

