
**Errata for the
PCI Express[®] Base Specification
Revision 3.1,
Single Root I/O Virtualization and
Sharing Revision 1.1,
Address Translation and Sharing
Revision 1.1, and M.2 Specification
Revision 1.0**

September 18, 2015



REVISION	REVISION HISTORY	DATE
1.0	First Release. Includes B208, A210, B211, B214 to B216, A217, B218 to B220, B224, B228 to B237, B242, B244, B246, B248 to B250, B252 to B264	2015-09-18

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 © 2011-2015 PCI-SIG

Contents

B208	Spread Spectrum Modulation Slew Rate	4
A210	M-PCIe Configuration.Update	6
B211	Root Complex Extended Capability Headers.....	7
B214	Inference of Electrical Idle	7
B215	ACS Source Validation and Bus Number 0	8
B216	Root Complex Integrated Endpoint Terminology	9
A217	L1 PM Substates: L1.2 Entry Conditions	9
B218	Root Complex and Routing Element Terms	9
B244	Device Terminology.....	10
B219	L1 PM Substates: Electrical Idle Terminology	10
B220	LTSSM Disabled to Detect Transition	10
B224	Malformed TLPs and Disabled VCs	11
B228	Link Speed Management	12
B229	Corrupt Received DLLP	13
B230	Loopback.....	13
B231	Flow Control Information Tracked by Receiver	13
B233	M.2 Specification: Reference to Link Capabilities register.....	14
B234	WAKE# and OBFF	14
B235	Readiness Time Reporting example times	14
B236	SR-IOV: Changing NumVFs or ARI Capable Hierarchy	15
B237	SR-IOV: Bus Numbering inside a Root Complex.....	15
B242	ATS: Terminology: PRG Response PASID Required	15
B246	Device Serial Number	16
B248	ATS: Page Request Interface and Dirty Bits	16
B249	SR-IOV and Access Control Services	18
B250	Enhanced Allocation.....	20
B252	Latency Tolerance Reporting (LTR).....	20
B253	LCRC and Sequence Number Rules (TLP Transmitter).....	21
B254	PCI Express Capability Structure.....	22
B255	Enhanced Allocation Part 2	22
B256	Enhanced Allocation Virtual Function BEI Ranges	23
B257	Lane Equalization Control Register Nomenclature	23
B258	Latency Tolerance Reporting	24
B259	Completion Timeout Mechanism.....	26
B260	Slot Power Limit	26
B263	AtomicOp Completion Rules	27
B264	PASID Completions.....	27

Note on Errata Numbering

Errata numbering is arbitrary and reflects numbers used during the development process.

B208 Spread Spectrum Modulation Slew Rate

In Section 4.3.7.3.3, page 427, make the following changes:

Table 4-32: Refclk Parameters for Common Refclk Rx Architecture at 5.0 GT/s

Symbol	Description	Limits		Units	Note
		Min	Max		
T _{REFCLK-HF-RMS}	> 1.5 MHz to Nyquist RMS jitter after applying Equation 4.3.3		3.1	ps RMS	1
T _{REFCLK-SSC-RES}	SSC residual		75	ps	1
T _{REFCLK-LF-RMS}	10 kHz - 1.5 MHz RMS jitter		3.0	ps RMS	2
T _{SSC-FREQ-DEVIATION}	SSC deviation		+0.0/-0.5	%	
T _{SSC-MAX-PERIOD-SLEW}	Maximum SSC df/dt rate of change of the clock frequency		0.75 1250	ps/μs ppm/μs	3

Notes:

1. T_{REFCLK-HF-RMS} is measured at the far end of the test circuit illustrated in Figure 4-89 after the filter function defined in Table 4-29 for Common Refclk Rx for >1.5 MHz jitter components has been applied.
2. T_{REFCLK-SSC-RES} and T_{REFCLK-LF-RMS} are measured after the filter function defined in Table 4-29 for Common Refclk Rx for >1.5 MHz jitter components has been applied.
3. Defined for a worst case SSC modulation profile such as Lexmark.

In Section 4.3.7.3.4, page 428, line 12, make the following changes:

In the data clocked Rx architecture, the amount of Refclk jitter propagated depends on the maximum PLL bandwidth (16 MHz, 3 dB of peaking). It is also the case that the Rx CDR must track the entirety of SSC, (20 ns) and the CDR must be capable of tracking SSC at a maximum slew rate corresponding to the largest ~~df/dt~~ [clock frequency change rate](#) SSC modulation profiles.

In Section 4.3.7.3.5, page 429, make the following changes:

Table 4-34: Refclk Parameters for Data Clocked Rx Architecture

Symbol	Description	Limits		Units	Note
		Min	Max		
T _{REFCLK-HF-RMS}	≥ 1.5 MHz to Nyquist RMS jitter after applying Equation 4.3.5		4.0	ps RMS	1
T _{REFCLK-SSC-RES}	Full SSC modulation corresponding to +0 – 0.5%		20	ns	1
T _{REFCLK-LF-RMS}	10 kHz - 1.5 MHz RMS jitter		7.5	ps RMS	2
T _{SSC-FREQ-DEVIATION}	SSC deviation		+0.0/-0.5	%	
T _{SSC-MAX-PERIOD-SLEW}	Max-SSC df/dt Maximum rate of change of the clock frequency		0.75 1250	ps/μs ppm/μs	3

Notes:

1. T_{REFCLK-HF-RMS} is measured at the far end of the test circuit illustrated in Figure 4-92 after the filter function defined in Table 4-29 for Common Refclk Rx for >1.5 MHz jitter components has been applied.

2. $T_{REFCLK-SSC-RES}$ and $T_{REFCLK-LF-RMS}$ are measured after the filter function defined in Table 4-29 for Common Refclk Rx for >1.5 MHz jitter components has been applied.
3. Defined for a worst case SSC modulation profile such as Lexmark.

In Section 4.3.7.5, page 433, add note reference for $T_{SSC-FREQ-DEVIATION}$ as follows:

Table 4-35: Parameters for Separate Refclk With Independent SSC Architecture (SRIS) at 5.0 GT/s

Symbol	Description	Limits		Units
		Min	Max	
F_{REFCLK}	Refclk frequency	99.97	100.03	MHz
$T_{REFCLK-RMS-SRIS}$	RMS Refclk jitter for Separate Refclk Independent SSC architecture ¹		2.0	ps RMS
F_{SSC}	SSC frequency range	30	33	kHz
$T_{SSC-FREQ-DEVIATION}$	SSC deviation ²		+0.0/-0.5	%

Notes:

1. The Refclk jitter is measured after applying PLL transfer function (16 MHz, 2 dB) and CDR jitter transfer function defined in Figure 4-95.
2. The maximum rate of change of the clock frequency is 1250 ppm/ μ s.

In Section 4.3.8.2, page 436, line 4, make the following changes:

Notes:

1. Before application of SSC.
2. ~~It is sufficient to define SSC deviation only without specifying anything about the shape of the modulation envelope. Envelopes with very large df/dt will fail the $T_{REFCLK-RMS-CO}$ parameter.~~ [The maximum rate of change of the clock frequency is 1250 ppm/ \$\mu\$ s.](#)
3. The Refclk jitter is measured after applying the jitter filtering function defined in Figure 4-96.

In Section 4.3.8.3, page 438, add note for $T_{SSC-FREQ-DEVIATION}$ as follows:

Table 4-38: Parameters for Data Clocked Rx Architecture at 8.0 GT/s

Symbol	Description	Limits		Units
		Min	Max	
F_{REFCLK}	Refclk frequency	99.97	100.03	MHz
$T_{REFCLK-RMS-SRIS}$	RMS Refclk jitter for data clocked Rx architecture ¹		1.0	ps RMS
F_{SSC}	SSC frequency range	30	33	kHz
$T_{SSC-FREQ-DEVIATION}$	SSC deviation ²		+0.0/-0.5	%

Notes:

1. The Refclk jitter is measured after applying jitter filtering function defined in Figure 4-97.
2. [The maximum rate of change of the clock frequency is 1250 ppm/ \$\mu\$ s.](#)

In Section 4.3.8.5, page 440, add note reference for $T_{SSC-FREQ-DEVIATION}$ as follows:

Table 4-39: Parameters for Separate Refclk With Independent SSC (SRIS) Architecture at 8.0 GT/s

Symbol	Description	Limits		Units
		Min	Max	
F _{REFCLK}	Refclk frequency	99.97	100.03	MHz
T _{REFCLK-RMS-SRIS}	RMS Refclk jitter for Separate Refclk Independent SSC architecture ¹		0.5	ps RMS
F _{SSC}	SSC frequency range	30	33	kHz
T _{SSC-FREQ-DEVIATION}	SSC deviation ²		+0.0/-0.5	%

Notes:

1. The Refclk jitter is measured after applying PLL transfer function (4 MHz, 2 dB) and CDR jitter transfer function defined in Figure 4-98.
2. The maximum rate of change of the clock frequency is 1250 ppm/μs.
3. The maximum rate of change of the clock phase is 3 ns/μs.

A210 M-PCIe Configuration.Update

In Section 8.4.9.3.3, page 963, line 15, make the following changes:

8.4.9.3.3. Configuration.Update

- On entry, all the TX-LANES and RX-LANES are in HIBERN8 state.
- The recovery_to_configuration variable must be set to 0b.
- The Configured TX-LANES must exit to STALL after residing in this state for THIBERN8.
- The next state is Configuration.Confirm if all Configured TX-LANES have resided in STALL for at least TACTIVATE, where TACTIVATE is measured after detecting DIF-N on all the RX-LANES that are part of the new LINK width.
- [Otherwise, the next state is Configuration.ExitToDetect after a 2 ms timeout.](#)

In Section 8.4.9.3.7, page 965, line 15, make the following changes:

8.4.9.3.7. Configuration.ExitToDetect

- The following steps must be followed in sequence:
 - ~~All the~~ Configured TX-LANES [that are in an HS-BURST](#) must send at least one EIOS, terminate the HS-BURST and enter STALL.
 - All Configured LANES initiate the process to enter HIBERN8 as described in Section 8.4.10-.
 - Local Reset must be asserted to all M-PHY MODULES associated with the LINK.
 - The next state is Detect

B211 Root Complex Extended Capability Headers

In Section 7.13.1, Table 7-56, page 789, line 10, make the following changes:

Table 7-56: Root Complex Link Declaration Extended Capability Header

Bit Location	Register Description	Attributes
15:0	<p>PCI Express Extended Capability ID – This field is a PCI-SIG defined ID number that indicates the nature and format of the Extended Capability.</p> <p>The Extended Capability ID for the Root Complex Link Declaration Extended Capability is 0005h.</p>	RO

In Section 7.13.2, page 790, line 1, make the following changes:

The Element Self Description register provides information about the Root Complex element containing the [Root Complex Link Declaration](#) Capability.

In Section 7.14.1, Table 7-60, page 796, line 6, make the following changes:

Table 7-60: Root Complex Internal Link Control Extended Capability Header

Bit Location	Register Description	Attributes
15:0	<p>PCI Express Extended Capability ID – This field is a PCI-SIG defined ID number that indicates the nature and format of the Extended Capability.</p> <p>The Extended Capability ID for the Root Complex Internal Link Control Extended Link Declaration Capability is 0006h.</p>	RO

B214 Inference of Electrical Idle

In Section 4.2.4.3, Page 257, line 4 make the following changes:

4.2.4.3. Inferring Electrical Idle

...

IMPLEMENTATION NOTE**Inference of Electrical Idle**

In the L0 state, ~~some number of one or more~~ Flow Control Update DLLPs are expected to be received in a 128 μ s window. Also in L0, ~~a one or more~~ SKP Ordered Sets ~~is are~~ expected to be received ~~in a 128 μ s window. on average every 1538 Symbols in 2.5 GT/s and 5.0 GT/s data rates and every 375 Blocks in 8.0 GT/s data rate.~~ As a simplification, it is permitted to use either one (or both) of these indicators to infer Electrical Idle. Hence, the absence of ~~either~~ a Flow Control Update

DLLP and/or ~~alternatively~~ a SKP Ordered Set in any 128 μ s window can be inferred as Electrical Idle.

B215 ACS Source Validation and Bus Number 0

In Section 6.12.1.1, page 588, line 12, make the following changes:

6.12.1.1. ACS Downstream Ports

...

- ACS Source Validation: must be implemented.

When enabled, the Downstream Port tests the Bus Number from the Requester ID of each Upstream Request received by the Port to determine if it is within the Bus Number “aperture” of the Port – the inclusive range specified by the Secondary Bus Number register and the Subordinate Bus Number register.

If the Bus Number from the Requester ID of the Request is not within this aperture, this is a reported error (ACS Violation) associated with the Receiving Port (see Section 6.12.4.)

Completions are never affected by ACS Source Validation.

IMPLEMENTATION NOTE:

Upstream Messages and ACS Source Validation

Functions are permitted to transmit Upstream Messages before they have been assigned a Bus Number. Such messages will have a Requester ID with a Bus Number of 00h. If the Downstream Port has ACS Source Validation enabled, these Messages (see Section 2.2.8.6 and Section 6.23.1) will likely be detected as an ACS Violation error.

- ACS Translation Blocking: must be implemented.

...

In Section 6.23.1, page 651, line 20, make the following changes:

6.23.1. Device Readiness Status (DRS)

...

Additional requirements for Root Ports and Switch Downstream Ports include:

- Implementation of the DRS Message Received bit, which indicates receipt of a DRS Message

IMPLEMENTATION NOTE:

DRS Messages and ACS Source Validation

Functions are permitted to transmit DRS Messages before they have been assigned a Bus Number. Such messages will have a Requester ID with a Bus Number of 00h. If the Downstream Port has ACS Source Validation enabled, these Messages (see Section 6.12.1.1) will likely be detected as an ACS Violation error.

B216 Root Complex Integrated Endpoint Terminology

Replace all 13 occurrences of '**RCIE**' with '[RCiEP](#)'.

Replace all 3 occurrences of '**RCIEs**' with '[RCiEPs](#)'.

In the Terms in Acronyms Section, page 42, line 15, add the following definition:

[RCiEP](#) Root Complex Integrated End Point.

A217 L1 PM Substates: L1.2 Entry Conditions

In Section 5.5.1, page 488, line 15, make the following changes:

5.5.1. Entry conditions for L1 PM Substates and L1.0 Requirements

...

- When in ASPM L1.0 and the ASPM L1.2 Enable bit is Set, the L1.2 substate must be entered when CLKREQ# is de-asserted and all of the following conditions are true:
 - ~~†~~The reported snooped LTR value [last sent or received by this Port](#) is greater than or equal to the value set by the LTR_L1.2_THRESHOLD Value and Scale fields, or there is no snoop service latency requirement.
 - ~~†~~The reported non-snooped LTR value [last sent or received by this Port](#) is greater than or equal to the value set by the LTR_L1.2_THRESHOLD Value and Scale fields, or there is no non-snoop service latency requirement.

B218 Root Complex and Routing Element Terms

In the Terms and Acronyms Section, page 42, line 36, make the following changes:

Terms and Acronyms

...

Root Complex, RC	A defined System Element that includes at least one Host Bridge, Root Port, or Root Complex Integrated Endpoint. zero or more Host Bridges, zero or more Root Complex Integrated Endpoints, zero or more Root Complex Event Collectors, and one or more Root Ports.
Root Complex Component	A logical aggregation of Root Ports, Root Complex Register Blocks, Root Complex Integrated Endpoints, and Root Complex Event Collectors.
Root Port	A PCI Express Port on a Root Complex that maps a portion of a Hierarchy through an associated virtual PCI-PCI Bridge.

RP PIO	Root Port Programmed I/O. See Section 6.2.10.3.
Routing Element	A term referring to a Root Complex, Switch, or Bridge in regard to its ability to route, multicast, or block TLPs.
RRAP	Remote Register Access Protocol (RRAP) is ...

B244 Device Terminology

In the Terms and Acronyms Section, page 38, line 15, make the following changes:

Device	A collection of one or more Functions within a single Hierarchy identified by common Bus Number and Device Number. An SR-IOV Device may have additional Functions accessed via additional Bus Numbers and/or Device Numbers configured through one or more SR-IOV Capability structures.
--------	--

B219 L1 PM Substates: Electrical Idle Terminology

In Section 5.5, page 485, line 1, make the following changes:

5.5. L1 PM Substates

...

L1.1 substate

...

- The Upstream and Downstream Ports are not required to be enabled to detect Electrical Idle [exit](#).

...

L1.2 substate

...

- The Upstream and Downstream Ports are not required to be enabled to detect ~~electrical idle~~ [\(E\) Electrical Idle exit](#).

B220 LTSSM Disabled to Detect Transition

In Section 4.2.6.9, page 326, line 17, make the following changes:

4.2.6.9. Disabled

...

- ❑ If an EIOS was transmitted (one if the current Link speed is 2.5 GT/s or 8.0 GT/s and two consecutive ones if the current Link speed is 5.0 GT/s) and an EIOS was received on any Lane (even while transmitting TS1 with the Disable Link bit asserted), then:
 - LinkUp = 0b (False)
 - At this point, the Lanes are considered Disabled.
 - For Upstream Ports: All Receivers must meet the Z_{RX-DC} specification for 2.5 GT/s within 1 ms (see Table 4-24).
 - For Upstream Ports: The next state is Detect when Electrical Idle ~~E_{exit}~~ is detected on at least one Lane ~~at the Receiver~~.

B224 Malformed TLPs and Disabled VCs

In Section 2.3, page 123, line 22, make the following changes:

2.3. Handling of Received TLPs

...

- Received Malformed TLPs that are ambiguous with respect to which buffer to release or are mapped to an uninitialized or disabled Virtual Channel must be discarded without updating Receiver Flow Control information.

In Section 5.3.2.1, page 454, line 10, make the following changes:

5.3.2.1. Entry into the L1 State

...

3. The Downstream component must then wait until it accumulates at least the minimum number of credits required to send the largest possible packet for any FC type for all enabled VCs (if it does not already have such credits). All Transaction Layer TLP scheduling is then suspended.

In Section 5.4.1.2.1, page 472, line 13, make the following changes:

5.4.1.2.1. Entry into the L1 State

...

- ❑ The Downstream component must not initiate ASPM L1 entry until it accumulates at least the minimum number of credits required to send the largest possible packet for any FC type for all enabled VCs.

B228 Link Speed Management

In Section 6.11, page 585, line 7, make the following changes:

6.11. Link Speed Management

...

The Target Link Speed field in the Link Control 2 register [in the Downstream Port](#) sets the upper bound for the Link speed. Except as described below, the Upstream component must attempt to maintain the Link at the Target Link Speed, or at the highest speed supported by both components on the Link (as reported by the values in the training sets – see Section 4.2.4.1), whichever is lower.

~~If Any Upstream Port or Downstream Port with~~ the Hardware Autonomous Speed Disable bit in the Link Control 2 register ~~is clear, the component~~ is permitted to autonomously [change-adjust](#) the Link speed using implementation specific criteria.

In Section 6.11, page 586, line 20, make the following changes:

Software is permitted to cause a Link to transition to the Polling.Compliance LTSSM state [at a particular speed](#) by writing [the Link Control 2 register in both components with the same value in](#) ~~to~~ the Target Link Speed field and ~~s~~Setting the Enter Compliance bit ~~in the Link Control 2 register in both components~~, and then initiating a Hot Reset on the Link ~~(through the Downstream Port~~Upstream component). ~~Software is required to write the same value into the Target Link Speed field in both the Upstream and Downstream components.~~

In Section 7.8.19, Table 7-28, page 742, Bits 3:0, make the following changes:

Table 7-28: Link Control 2 Register

Bit Location	Register Description	Attributes
3:0	<p>Target Link Speed – For Downstream Ports, this field sets an upper limit on Link operational speed by restricting the values advertised by the Upstream component in its training sequences.</p> <p>...</p> <p>For both Upstream and Downstream Ports, this field is used to set the target compliance mode speed when software is using the Enter Compliance bit to force a Link into compliance mode.</p> <p>For Upstream Ports, if the Enter Compliance bit is clear, this field is permitted to have no effect.</p> <p>For a Multi-Function device associated with an Upstream Port, the field 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 field is of type RsvdP.</p> <p>Components that support only the 2.5 GT/s speed are permitted to hardwire this field to 0000b.</p>	RWS/RsvdP (see description)

B229 Corrupt Received DLLP

In Section 3.5.2.2, page 200, line 14, make the following changes:

3.5.2.2. Handling of Received DLLPs

...

- A corrupt received DLLP is discarded. [This is a Bad DLLP error and](#), ~~and~~ is a reported error associated with the Port (see Section 6.2).

B230 Loopback

In Section 4.2.2.5, page 238, line 10, make the following changes:

4.2.2.5. Loopback with 128b/130b Code

When using 128b/130b encoding, Loopback Masters must transmit Blocks with the defined 01b and 10b Sync Headers. However, they are not required to transmit an SDS Ordered Set when transitioning from Ordered Set Blocks to Data Blocks, nor are they required to transmit an EDS Token when transitioning from Data Blocks to Ordered Set Blocks. Masters must transmit SKP Ordered Sets periodically as defined in Section 4.2.7, and they must be capable of processing received (looped-back) SKP Ordered Sets of varying length. Masters are permitted to transmit Electrical Idle Exit Ordered Sets (EIEOS) as defined in Section 4.2.2.2.1. Masters are permitted to transmit any payload in Data Blocks and Ordered Set Blocks that they expect to be looped-back. ~~However, Ordered Set Blocks that match the definition of SKP OS, EIEOS, or EIOS should be avoided since they have defined purposes while in Loopback.~~ [If the Loopback Master transmits an Ordered Set Block whose first symbol matches the first symbol of SKP OS, EIEOS, or EIOS, that Ordered Set Block must be a complete and valid SKP OS, EIEOS, or EIOS.](#)

When using 128b/130b encoding, Loopback Slaves must ...

In Section 4.2.6.10.2, page 329, line 24, make the following changes:

4.2.6.10.2. Loopback.Active

- The loopback master must send valid encoded data. The loopback master [must not transmit](#) ~~should avoid sending~~ EIOS as data until it wants to exit Loopback. [When operating with 128b/130b encoding, loopback masters must follow the requirements of Section 4.2.2.5.](#)

B231 Flow Control Information Tracked by Receiver

In Section 2.6.1.2, page 161, line 3, make the following changes:

2.6.1.2. FC Information Tracked by Receiver

...

- ❑ When the Link is in the L0 or L0s Link state, Update FCPs for each enabled type of non-infinite FC credit must be scheduled for transmission at least once every 30 μ s (-0%/+50%), except when the Extended Sync bit of the ~~Control~~ Link [Control](#) register is set, in which case the limit is 120 μ s (-0%/+50%).

B233 M.2 Specification: Reference to Link Capabilities register

In the M.2 Specification, Revision 1.0, Section 3.1.3.3, page 100, line 5, make the following changes:

3.1.3.3. Clock Request Support Reporting and Enabling

Support for the CLKREQ# dynamic clock protocol should be reported using bit 18 in the PCI Express [Link Capabilities](#) register (offset ~~0C4h~~ [00Ch](#)). To enable dynamic clock management, bit 8 of the Link Control register (offset 010h) is provided.

B234 WAKE# and OBF#

In the Section 6.19, page 625, line 5, make the following changes:

An OBF# Message received at a Port that does not implement OBF# or when OBF# is not enabled must be handled as an Unsupported Request (UR). [This is a reported error associated with the receiving Port \(see Section 6.2\). If a Port has OBF# enabled using WAKE# signaling, and that Port receives an OBF# Message, the behavior is undefined.](#)

B235 Readiness Time Reporting example times

In Section 7.35, page 912, line 12, make the following changes:

...

Registers and fields in the Readiness Time Reporting Extended Capability are shown in Figure 7-179. Time values are encoded in floating point as shown in Figure 7-180. The actual time value is $Value \times Mult[Scale]$. For example, the value A1Eh represents about 1 second (actually 1.006 sec) and the value ~~40Ah~~ [80Ah](#) represents about 10 ms (actually 10.240 ms).

In Section 7.35.3, page 915, Table 7-158, make the following changes:

Table 7-158: Readiness Time Reporting 2 Register

Bit Location	Register Description	Attributes
23:12	D3hot to D0 Time – If Immediate_Readiness_on_Return_to_D0 is Clear, D3hot to D0 Time is the time that the Function requires after it is directed from D3hot to D0 before it is Configuration-Ready and has returned to either D0uninitialized or D0active state (see the PCI Bus Power Management Interface Specification).	HwInit/RsvdP

	<p>This field is RsvdP if the Immediate_Readiness_on_Return_to_D0 bit is Set.</p> <p>This field is undefined when the Valid bit is Clear.</p> <p>This field must be less than or equal to the encoded value 40Ah80Ah.</p>	
--	--	--

B236 **SR-IOV: Changing NumVFs or ARI Capable Hierarchy**

In the SR-IOV Specification, Section 2.1.2, page 27, line 3, make the following changes:

IMPLEMENTATION NOTE:

NumVFs and ARI Capable Hierarchy

After configuring NumVFs and ARI Capable Hierarchy [where applicable](#), software may read First VF Offset and VF Stride to determine how many Bus Numbers would be consumed by the PF's VFs. The additional Bus Numbers, if any, are not actually used until VF Enable is Set

B237 **SR-IOV: Bus Numbering inside a Root Complex**

In the SR-IOV Specification, Section 2.1.2, page 27, line 16, make the following changes:

As in the PCI Express Base Specification, SR-IOV capable Devices [that are associated with an Upstream Port](#) capture the Bus Number from any Type 0 Configuration [Write](#) Request. SR-IOV capable Devices do not capture the Bus Number from any Type 1 Configuration Write Requests. [SR-IOV capable Root Complex Integrated Endpoints use an implementation specific mechanism to assign their Bus Numbers.](#)

B242 **ATS: Terminology: PRG Response PASID Required**

In the PASID ECN to the ATS Specification, Part I, line 15, make the following changes:

In addition, three new bits are defined in the ATS Translation Completion Data Entry (Global Mapping, Execute Permitted, Privileged Mode Access), one new bit is defined in the ATS Invalidation Request Message (Global Invalidate), one new bit is defined in the ATS Extended Capability (Global Invalidate Supported), and one new bit is defined in the Page Request Extended Capability (PRG Response PASID [Required](#)~~Enable~~).

In the PASID ECN to the ATS Specification, in the added Section 4.2.2, page 13, line 29, make the following changes:

*Note: **Green** text was inserted by the ECN and is unchanged by this errata.*

4.2.2. PASID TLP Prefix Usage

If a Page Request has a PASID TLP Prefix, the corresponding PRG Response Message may optionally contain one as well.

If the PRG Response PASID ~~Required~~Enable bit is Clear, PRG Response Messages do not have a PASID TLP Prefix.

If the PRG Response PASID ~~Required~~Enable bit is Set, PRG Response Messages have a PASID TLP Prefix if the Page Request also had one. The Function is permitted to use the PASID value from the prefix in conjunction with the PRG Index to match requests and responses.

In PASID TLP Prefixes attached to PRG Response Messages, the Execute Requested and Privileged Mode Requested bits are Reserved and the PASID value is copied from the PASID value of the Page Request.

B246 Device Serial Number

In Section 7.12, page 785, line 8, make the following changes:

It is permitted but not recommended for Root Complex Integrated Endpoints to implement this Capability.

All multi-Function devices that implement this Capability must implement it for Function 0; other Functions that implement this Capability must return the same Device Serial Number value as that reported by Function 0.

A PCI Express ~~multi-device~~ component containing multiple Devices such as a PCI Express Switch that implements this Capability must return the same Device Serial Number for each device.

B248 ~~ATS:~~ Page Request Interface and Dirty Bits

In the ATS Specification, Section 1.2, page 19, line 22, make the following changes:

1.2. Page Request Interface Extension

...

The ability to page, begs the question of page table status flag management. Typical TAs associate flags (e.g. dirty and access indications) with each untranslated address. Without any additional hints about how to manage pages mapped to an ~~I/O device~~ Function, ~~an RC~~ such TAs would need to conservatively assume that ~~if an I/O device can write to a page, it has written to the page when they grant a Function permission to read or write a page, that Function will use the permission.~~ ~~I/O~~ Such writable pages would need to be marked as dirty before their translated addresses are made available to a ~~device~~ Function.

This conservative dirty-on-write-permission-grant behavior is generally not a significant issue for [Functions that do not support paging-non-pageable devices](#), where I/O pages are pinned and the cost of saving a clean page to memory will seldom be paid. However, ~~devices~~ [Functions](#) that support the Page Request Interface could pay a significant penalty if all writable pages are treated as dirty, since such ~~devices~~ [Functions](#) operate without pinning their accessible memory footprints and may issue speculative page requests for performance. The cost of saving clean pages (instead of just discarding them) in such systems can diminish the value of otherwise attractive paging techniques. [This can cause significant performance issues and risk functional issues in circumstances where the backing store is unable to be written, such as a CD-ROM.](#)

The No Write (NW) flag in Translation Requests indicates [that a Function](#) ~~whether a page should be marked as dirty if it is writable, or whether a device~~ is willing to restrict its usage to only reading the page, independent of the access rights that would otherwise have been granted.

If a device chooses to request only read access [by issuing a Translation Request with the NW flag Set](#) and ~~later then~~ determines that it ~~does wish~~ [needs](#) to write to the page, then the device ~~would need to~~ [must](#) issue a new Translation Request.

[Upon receiving a Translation Request with the NW flag Clear, TAs are permitted to mark the associated pages dirty. It is strongly recommended that Functions not issue such Requests unless they have been given explicit write permission. An example of write permission is where the host issues a command to a Function to load data from a storage device and write that data into memory.](#)

...

In the ATS Specification, Section 2.2.5, page 24, line 23, make the following changes:

2.2.5. No Write (NW) Flag

The No Write flag, when Set, indicates that the Function is requesting read-only access for this translation.

The TA may ignore the No Write Flag, however, if the TA responds with a translation marked as read-only then the Function must not issue MemWrite transactions using that translation. In this case, the Function may issue another translation request with the No Write flag Clear, which may result in a new translation completion with or without the W (Write) bit Set.

[Upon receiving a Translation Request with the NW flag Clear, TAs are permitted to mark the associated pages dirty. It is strongly recommended that Functions not issue such Requests unless they have been given explicit write permission.](#)

In the ATS Specification, Section 4.1, Table 4-1, page 41, make the following changes:

4.1. Page Request Message

...

Table 4-1: Page Request Message Data Fields

Field	Meaning
W	Write Access Requested – This field, when Set, indicates that the requesting Function seeks write access and/or zero-length read access to the associated page. When

Clear, this field indicates that the requesting Function will not write to the associated page.

Upon receiving a Page Request Message with the W field Set, the host is permitted to mark the associated page dirty. Thus, Functions must not issue such Requests unless the Function has been given explicit write permission.

B249 SR-IOV and Access Control Services

In Section 6.12, page 586, line 30, make the following changes:

6.12. Access Control Services (ACS)

ACS defines a set of control points within a PCI Express topology to determine whether a TLP ~~should~~ **is to** be routed normally, blocked, or redirected. ACS is applicable to RCs, Switches, and multi-Function devices.⁸⁹ For ACS requirements, single-Function devices that are SR-IOV capable must be handled as if they were multi-Function devices, since they essentially behave as multi-Function devices after their Virtual Functions (VFs) are enabled.

In Section 6.12.1.2, page 590, line 17, make the following changes:

6.12.1.2. ACS Functions in SR-IOV Capable and Multi-Function Devices

This section applies to multi-Function device ACS Functions, with the exception of Downstream Port Functions, which are covered in the preceding section. For ACS requirements, single-Function devices that are SR-IOV capable must be handled as if they were multi-Function devices.

- ACS Source Validation: must not be implemented.
- ACS Translation Blocking: must not be implemented.
- ACS P2P Request Redirect: must be implemented by Functions that support peer-to-peer traffic with other Functions. This includes SR-IOV Virtual Functions (VFs).

ACS P2P Request Redirect is subject to interaction with the ACS P2P Egress Control and ACS Direct Translated P2P mechanisms (if implemented). Refer to Section 6.12.3 for more information.

When ACS P2P Request Redirect is enabled in a multi-Function or SR-IOV capable device, peer-to-peer Requests (between Functions of the device) must be redirected Upstream towards the RC.

Completions are never affected by ACS P2P Request Redirect.

- ACS P2P Completion Redirect: must be implemented by Functions that implement ACS P2P Request Redirect.

The intent of ACS P2P Completion Redirect is to avoid ordering rule violations between Completions and Requests when Requests are redirected. Refer to Section 6.12.5 for more information.

ACS P2P Completion Redirect does not interact with ACS controls that govern Requests.

When ACS P2P Completion Redirect is enabled in a multi-Function [or SR-IOV capable](#) device, peer-to-peer Read Completions that do not have the Relaxed Ordering bit set must be redirected Upstream towards the RC. Otherwise, peer-to-peer Completions must be routed normally.

Requests are never affected by ACS P2P Completion Redirect.

- ❑ ACS Upstream Forwarding: must not be implemented.
- ❑ ACS P2P Egress Control: implementation is optional; is based on Function Numbers or Function Group Numbers; controls peer-to-peer Requests between the different Functions within the multi-Function [or SR-IOV capable](#) device.

ACS P2P Egress Control is subject to interaction with the ACS P2P Request Redirect and ACS Direct Translated P2P mechanisms (if implemented). Refer to Section 6.12.3 for more information.

Each Function within a multi-Function [or SR-IOV capable](#) device that supports ACS P2P Egress Control can be selectively enabled to block peer-to-peer communication with other Functions or Function Groups⁹³ within the device. This is configured on a per Function basis.

With ACS P2P Egress Control in multi-Function [or SR-IOV capable](#) devices, controls in the “sending” Function determine if the Request is blocked, and if so, the “sending” Function handles the ACS Violation error per Section 6.12.4.

When ACS Function Groups are enabled in an ARI Device, ACS P2P Egress Controls are enforced on a per Function Group basis instead of a per Function basis. See Section 6.13.

Completions are never affected by ACS P2P Egress Control.

- ❑ ACS Direct Translated P2P: must be implemented if the multi-Function [or SR-IOV capable](#) device Function supports Address Translation Services (ATS) and also peer-to-peer traffic with other Functions.

When ACS Direct Translated P2P is enabled in a multi-Function [or SR-IOV capable](#) device Function, peer-to-peer Memory Requests whose Address Type (AT) field indicates a Translated address must be routed normally (“directly”) to the peer Function, regardless of ACS P2P Request Redirect and ACS P2P Egress Control settings. All other peer-to-peer Requests must still be subject to ACS P2P Request Redirect and ACS P2P Egress Control settings.

Completions are never affected by ACS Direct Translated P2P.

In Section 6.12.1.3, page 591, line 31, make the following changes:

6.12.1.3. Functions in Single-Function Devices

This section applies to single-Function device Functions, with the exception of Downstream Port Functions [and SR-IOV capable Functions](#), which are covered in a preceding sections. [For ACS requirements, single-function devices that are SR-IOV capable must be handled as if they were multi-Function devices.](#)

No ACS capabilities are applicable, and the Function must not implement an ACS Extended Capability structure.

B250 Enhanced Allocation

In the Enhanced Allocation (EA) ECN, Section 6.2.5.1, page 4, line 8, make the following changes:

*Note: **Green** text was inserted by the ECN and is unchanged by this errata.*

6.2.5.1. Address Maps

After determining this information, power-up software can map the I/O controllers into reasonable locations and proceed with system boot. In order to do this mapping in a device independent manner, the base registers for this mapping are placed in the predefined header portion of Configuration Space. **It is strongly recommended that power-up firmware/software also support the optional Enhanced Allocation mechanism (see Section 6.9). Resources allocated by a specific Function using Enhanced Allocation must be handled by system firmware/software as being unavailable for use by any other entity, regardless of the state of that specific Function, including cases where that Function is disabled, and without regard to the state of that Function's Memory Space Enable, IO Space Enable bits in the Command register, or any function-specific mechanisms.**

...

Devices that map control functions into I/O Space must not consume more than 256 bytes per I/O Base Address register **or per each I/O Space entry in the Enhanced Allocation capability.**

...

B252 Latency Tolerance Reporting (LTR)

In Section 6.18, page 619, line 15, make the following changes:

6.18. Latency Tolerance Reporting (LTR) Mechanism

...

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 one of the following applies.**

- If the bit was cleared due to a Configuration Write to the Device Control 2 register, the device must send a new LTR Message with all the Requirement bits clear.**
- If the bit was cleared due to an FLR, it is strongly recommended that the device send a new LTR Message with all the Requirement bits clear.**

When a Downstream Port goes to DL_Down status, any previous latencies recorded for that Port must be treated as invalid.

An LTR Message from a device ...

In Section 6.18, page 621, line 3, make the following changes:

- ❑ ~~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 transmitted Upstream if the conglomerated latencies are changed as a [result of DL_Down invalidating the previous latencies recorded for that Port](#).

B253 LCRC and Sequence Number Rules (TLP Transmitter)

In Section 3.5.2.1, page 196, line 26, make the following changes:

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 3-7, Table 3-8, and Table ~~3-6~~.

B254 PCI Express Capability Structure

In Section 3.5.2.1, page 655, Figure 7-3, change text as following:

PCI Express
 Capability Structure
~~Capability needed by BIOS
 or by driver software on non
 PCI Express aware operating
 systems~~

so the figure becomes the following:

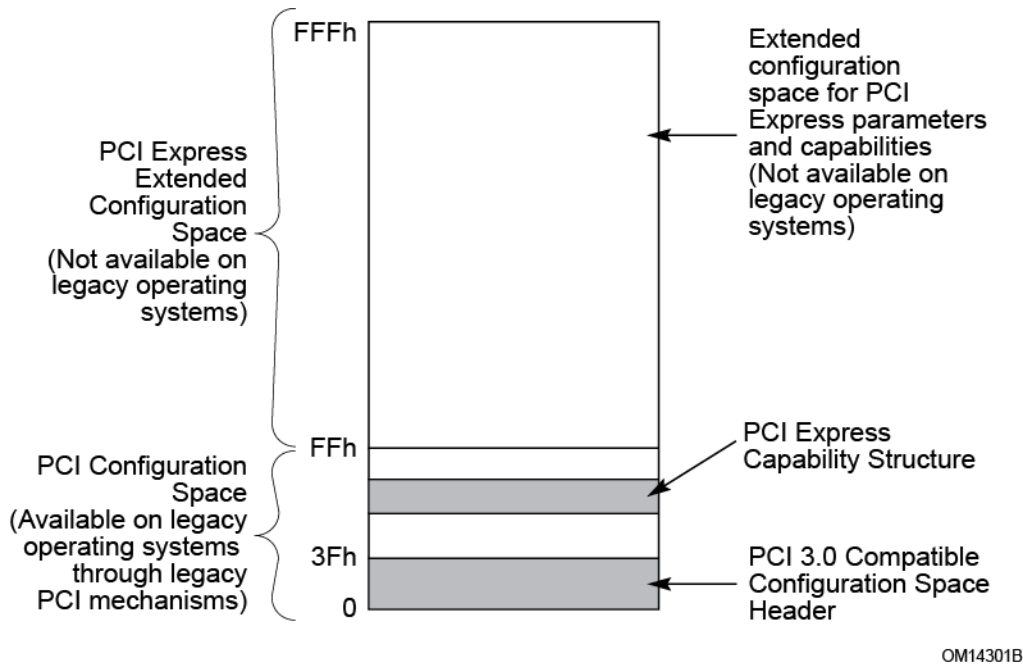


Figure 7-3: PCI Express Configuration Space Layout

B255 Enhanced Allocation Part 2

In the Enhanced Allocation (EA) ECN, page 13, Table 6-1, make the following changes:

Value (h)	Resource Definition
...	...
08-FC E	<p>Reserved for future use; System firmware/software must not write to this entry, and must not attempt to interpret this entry or to use this resource.</p> <p>When software reads a Primary Properties value that is within this range, is it strongly recommended that software treat this resource according to the value in the Secondary Properties</p>

B256 Enhanced Allocation Virtual Function BEI Ranges

In the Enhanced Allocation (EA) ECN, page 11, line 14, make the following changes:

- For Memory ~~or I/O~~ BARs where the Primary or Secondary ~~Property~~ Properties is 03h or 04h it is permitted to assign the same BEI in the range of 0-5 once for a range where Base + MaxOffset is below 4GB, and again for a range where Base + MaxOffset is greater than 4GB; It is not otherwise permitted to assign the same BEI in the range 0-5 for more than one entry.
- For Virtual Function BARs where the Primary or Secondary ~~Property~~ Properties is 03h or 04h it is permitted to assign the same BEI in the range of ~~0-5~~ 9-14 once for a range where Base + MaxOffset is below 4GB, and again for a range where Base + MaxOffset is greater than 4GB; It is not otherwise permitted to assign the same BEI in the range ~~0-5~~ 9-14 for more than one VF entry.

B257 Lane Equalization Control Register Nomenclature

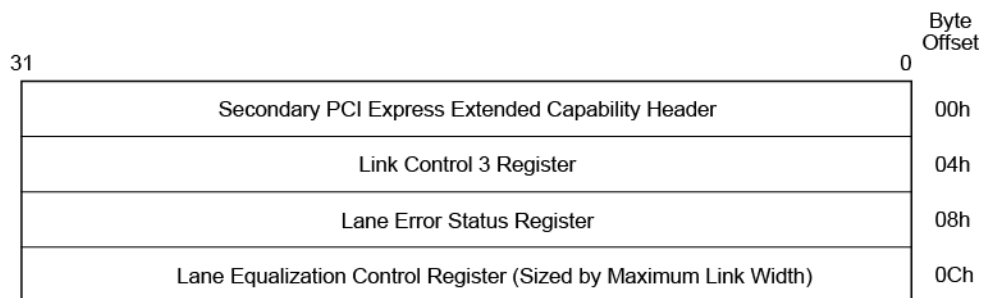
In Section 6.6.2, page 562, line 31, make the following changes:

- ~~Link Lane~~ Equalization Control register in the Secondary PCI Express Extended Capability structure

In Section 7.27, page 857, Figure 7-132, change the name of the field at offset 0Ch as follows:

Lane Equalization Control Register (Sized by Maximum Link Width)

so that the figure looks like:



A-0798A

Figure 7-132: Secondary PCI Express Extended Capability Structure

In Section 7.27.4, page 860, line 4, make the following changes:

The Lane Equalization Control register consists of of control fields required for per Lane equalization and the number of entries in this register are sized by Maximum Link Width (see Section 7.8.6). Each entry contains the values for the Lane with the corresponding default Lane

number which is invariant to Link width and Lane reversal negotiation that occurs during Link training.

In Section 7.27.4, page 861, line 1, make the following changes:

Figure 7-137: Lane ~~((Maximum Link Width—1):0)~~ Equalization Control Register [Entry](#)

Table 7-117: Lane ~~((Maximum Link Width—1):0)~~ Equalization Control Register [Entry](#)

In Section 4.2.3, page 241, line 1, make the following changes:

rate. These preset values are derived from the Upstream Port Transmitter Preset and Upstream Port Receiver Preset Hint fields of each Lane's Equalization Control register [entry](#). After the data rate change to 8.0 GT/s, the Upstream Port transmits TS1 Ordered Sets with the preset values it received. The preset values must be within the operable range defined in Section 4.3.3.5.2 if reduced swing will be used by the Transmitter.

Phase 1: Both components make the Link operational enough at 8.0 GT/s data rate to be able to exchange TS1 Ordered Sets to complete remaining phases for the fine-tuning their Transmitter/Receiver pairs. It is expected that the Link will operate at a BER of less than 10^{-4} before the component is ready to move on to the next Phase.

The Downstream Port initiates Phase 1 by transmitting TS1 Ordered Sets with EC=01b (indicating Phase 1) to the Upstream Port using the preset values in the Downstream Port Transmitter Preset and, optionally, the Downstream Port Receiver Preset Hint fields of each Lane's Equalization Control register [entry](#). The Upstream Port, after ...

In Section 4.2.6.4.4, page 310, line 1, make the following changes:

The Downstream Port must transmit EQ TS2 Ordered Sets (TS2 Ordered Sets with Symbol 6 bit 7 set to 1b) with the Transmitter Preset and Receiver Preset Hint fields set to the values specified by the Upstream Port Transmitter Preset and the Upstream Port Receiver Preset Hint fields of each Lane's Equalization Control register [entry](#) if all of the following conditions are satisfied: ...

B258 Latency Tolerance Reporting

In section 7.25.2, page 851, Table 7-108, make the following changes:

Bit Location	Register Description	Attributes
9:0	<p>Max Snoop LatencyValue – Along with the Max Snoop LatencyScale field, this register specifies the maximum snoop latency that a device is permitted to request. Software should set this to the platform's maximum supported latency or less. It is strongly recommended that any updates to this field are reflected in LTR Message(s) sent by the device within 1 ms.</p> <p>The default value for this field is 00 0000 0000b.</p>	RW
12:10	<p>Max Snoop LatencyScale – This register provides a scale for the value contained within the Maximum Snoop LatencyValue field. Encoding is the same as the LatencyScale fields in the LTR Message.</p>	RW

	<p>See Section 6.18. It is strongly recommended that any updates to this field are reflected in LTR Message(s) sent by the device within 1 ms.</p> <p>The default value for this field is 000b.</p> <p>Hardware operation is undefined if software writes a Not Permitted value to this field.</p>	
--	--	--

In section 7.25.3, page 852, Table 7-109, make the following changes:

Bit Location	Register Description	Attributes
9:0	<p>Max No-Snoop LatencyValue – Along with the Max No-Snoop LatencyScale field, this register specifies the maximum no-snoop latency that a device is permitted to request. Software should set this to the platform’s maximum supported latency or less. It is strongly recommended that any updates to this field are reflected in LTR Message(s) sent by the device within 1 ms.</p> <p>The default value for this field is 00 0000 0000b.</p>	RW
12:10	<p>Max No-Snoop LatencyScale – This register provides a scale for the value contained within the Max No-Snoop LatencyValue field. Encoding is the same as the LatencyScale fields in the LTR Message. See Section 6.18. It is strongly recommended that any updates to this field are reflected in LTR Message(s) sent by the device within 1 ms.</p> <p>The default value for this field is 000b.</p> <p>Hardware operation is undefined if software writes a Not Permitted value to this field.</p>	RW

In Section 6.18, page 623, line 13, change the implementation note as follows:

IMPLEMENTATION NOTE:

Optimal Use of LTR

...

Note that the RC may delay processing of device Request TLPs, provided it satisfies the device's service requirements. If, for example, an Endpoint connected to Root Port 1 reports a latency tolerance of 100 μ s, and an Endpoint on Root Port 2 report a value of 30 μ s, the RC might implement a policy of stalling an initial Request following an idle period from Root Port 1 for 70 μ s before servicing the Request with a 30 μ s latency, thus providing a perceived service latency to the first Endpoint of 100 μ s. This RC behavior provides the RC the ability to batch together Requests for more efficient servicing.

[It is possible that, after it is determined that the RC can service snoop and no-snoop Requests from all Endpoints within the maximum snoop and maximum no-snoop time intervals, this information may be communicated to Endpoints by updating the Max Snoop LatencyValue, Max Snoop](#)

[LatencyScale and Max No-Snoop LatencyValue, Max No-Snoop LatencyScale fields. The intention of this communication would be to prevent Endpoints from sending needless LTR updates.](#)

[When an Endpoint’s LTR value for snoop Requests changes to become larger \(looser\) than the value indicated in the Max Snoop LatencyValue/Scale fields, it is recommended that the Endpoint send an LTR message with the snoop LTR value indicated in the Max Snoop LatencyValue/Scale fields. Likewise, when an Endpoint’s LTR value for no-snoop Requests changes to become larger \(looser\) than the value indicated in the Max No-Snoop LatencyValue/Scale fields, it is recommended that the Endpoint send an LTR message with the no-snoop LTR value indicated in the Max No-Snoop LatencyValue/Scale fields.](#)

It is recommended that Endpoints buffer Requests as much as possible, and then use the full Link bandwidth in bursts that are as long as the Endpoint can practically support, as this will generally lead to the best overall platform power efficiency.

B259 Completion Timeout Mechanism

In Section 2.8, page 170, line 26, make the following changes:

The Completion Timeout mechanism may be disabled by configuration software. The Completion Timeout limit is set in the Completion Timeout Value field of the Device Control 2 register. ~~Refer to Section 2.2.8.10.~~ A Completion Timeout is a reported error associated with the Requester Function (see Section 6.2).

B260 Slot Power Limit

In Section 7.15.3, page 804, Table 7-65, make the following changes:

Bit Location	Register Description	Attributes
7:0	<p>Base Power – Specifies in watts the base power value in the given operating condition. This value must be multiplied by the data scale to produce the actual power consumption value except when the Data Scale field equals 00b (1.0x) and Base Power exceeds EFh, the following alternative encodings are used:</p> <p>F0h = greater than 239 W and less than or equal to 250 W Slot Power Limit</p> <p>F1h = greater than 250 W and less than or equal to 275 W Slot Power Limit</p> <p>F2h = greater than 275 W and less than or equal to 300 W Slot Power Limit</p> <p>F3h to FFh = Reserved for Slot Power Limit values above greater than 300 W</p>	RO

B263 AtomicOp Completion Rules

In Section 6.15.2, page 610, line 20, make the following changes:

6.15.2 AtomicOp Transaction Protocol Summary

...

- ❑ Unless there's a higher precedence error, an [AtomicOp-aware](#) 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 Section 2.7.2.2. The value of the target location must remain unchanged.
- ❑ If the Completer of an AtomicOp Request encounters an uncorrectable error accessing the target location or carrying out the Atomic operation, the Completer must handle it as a Completer Abort (CA). The subsequent state of the target location is implementation specific.
- ❑ [AtomicOp-aware](#) Completers are required to handle any properly formed AtomicOp Requests with types or operand sizes they don't support as an Unsupported Request (UR). If the Length field in an AtomicOp Request contains an unarchitected value, the Request must be handled [by an AtomicOp-aware Completer](#) as a Malformed TLP. See Section 2.2.7.
- ❑ If any Function in a multi-Function device supports AtomicOp Completer or AtomicOp routing capability, all Functions with Memory Space BARs in that device must decode properly formed AtomicOp Requests and handle any they don't support as an Unsupported Request (UR). Note that in such devices, Functions lacking AtomicOp Completer capability are forbidden to handle properly formed AtomicOp Requests as Malformed TLPs.
- ❑ If an RC has any Root Ports that support AtomicOp routing capability, all Root Complex Integrated Endpoints in the RC reachable by forwarded AtomicOp Requests must decode properly formed AtomicOp Requests and handle any they don't support as an Unsupported Request (UR).
- ❑ With an AtomicOp Request having a supported type and operand size, the [AtomicOp-aware](#) Completer is required either to carry out the Request or handle it as Completer Abort (CA) for any location in its target Memory Space. Completers are permitted to support AtomicOp Requests on a subset of their target Memory Space as needed by their programming model (see Section 2.3.1). Memory Space structures defined or inherited by PCI Express (e.g., the MSI-X Table structure) are not required to be supported as AtomicOp targets unless explicitly stated in the description of the structure.

B264 PASID Completions

In Section 6.20.2.1, page 630, line 15, make the following changes:

For Endpoints, the following rules apply:

- ❑ The Endpoint is not permitted to send TLPs with a PASID value greater than or equal to $2^{\text{Max PASID Width}}$.
- ❑ The Endpoint is optionally permitted to signal an error when it receives a [TLPRequest](#) with a PASID value greater than or equal to $2^{\text{Max PASID Width}}$. ~~For Requests, this error~~ This is an [an](#)

Unsupported Request ~~error associated with the Receiving Port (see Section 6.2) and for Completions, this error is Unexpected Completion.~~

For Root Complexes, the following rules apply:

- ❑ A Root Complex is not permitted to send a TLP with a PASID value greater than it supports.
- ❑ A Root Complex is optionally permitted to signal an error when it receives a ~~TLPRequest~~ with a PASID value greater than it supports. ~~For Requests, this error~~ This is ~~an~~ Unsupported Request ~~error associated with the Receiving Port (see Section 6.2) and for Completions, this error is Unexpected Completion.~~