



PCI-SIG ENGINEERING CHANGE NOTICE

TITLE:	Link Bandwidth Notification
DATE:	20 April 2005, updated 28 June 2006
AFFECTED DOCUMENT:	PCI Express Base Specification Revision 1.1
SPONSOR:	David Harriman, Intel Corporation Joe Cowan, Hewlett-Packard Company

Part I

NOTE: This document includes highlighted text that *will not* be included in the final ECN to the 1.1 base spec, but *will* be included in the 2.0 base spec.

1 Summary of the Functional Changes

Adds a mechanism to signal, via an interrupt and associated status bit, that a link has changed width under hardware autonomous action, or Link retraining due to software direction has completed.

2 Benefits as a Result of the Changes

System software can be proactively notified of changes that affect link bandwidth.

3 Assessment of the Impact

Moderate hardware implementation/validation impact.

4 Analysis of the Hardware Implications

Implementation and validation of the associated interrupt & status mechanisms.

5 Analysis of the Software Implications

No impact to existing software. Mechanism for future software to be notified proactively in the case of a link width change.

Part II

Detailed Description of the change

In Section 4.2.6.3.3., modify as shown:

4.2.6.3.3. Configuration.Lanenum.Accept

4.2.6.3.3.1. Downstream Lanes

- ❑ If two consecutive TS1 ordered sets are received with non-PAD Link and non-PAD Lane numbers that match all the non-PAD Link and non-PAD Lane numbers (or reversed Lane numbers if Lane reversal is optionally supported) that are being transmitted in Downstream Lane TS1 ordered sets, the next state is Configuration.Complete.

- The 'Link Bandwidth Management Status' and 'Link Autonomous Bandwidth Status' bits of the Link Status register must be updated as follows on a link bandwidth change if the current transition to Configuration state was from the Recovery state:

- (i) If the link bandwidth was reduced and either the link autonomous bit (bit 6) of symbol 4 in the two consecutive TS1 ordered sets is reset to 0 or the bandwidth change was initiated by the upstream component due to reliability issues, the 'Link Bandwidth Management Status' bit is set to 1b

- (ii) else the 'Link Autonomous Bandwidth Status' bit is set to 1b.

In Section 4.2.6.3.6., modify as shown:

4.2.6.3.6. Configuration.Idle

...

- ❑ Next state is L0 if eight consecutive Symbol Times of Idle data received on all configured Lanes and 16 Idle data Symbols are sent after receiving one Idle data Symbol.

- If software has written a 1b to the 'Retrain Link' bit in the Link Control Register since the last transition to L0 from Recovery or Configuration, the upstream component must set the 'Link Bandwidth Management Status' bit of the Link Status register to 1b.

In Section 4.2.6.4.4., modify as shown:

4.2.6.4.4. Recovery.Idle

...

- ❑ Next state is L0 if eight consecutive Symbol Times of Idle data received on all configured Lanes and 16 Idle data Symbols are sent after receiving one Idle data Symbol.

- If software has written a 1b to the 'Retrain Link' bit in the Link Control Register since the last transition to L0 from Recovery or Configuration, the upstream component must set the 'Link Bandwidth Management Status' bit of the Link Status register to 1b.

In Section 7.8., modify as shown:

7.8.2. PCI Express Capabilities Register (Offset 02h)

...

Table 7-10: PCI Express Capabilities Register

Bit Location	Register Description	Attributes
13:9	Interrupt Message Number – This register must indicate which MSI/MSI-X vector is used for the interrupt message generated in association with <u>any of the status bits</u> in either the Slot Status register (see Section 6.7.3.4) or the Root Status register (see Section 6.1.6) of this capability structure.	RO
...	...	

...

7.8.6 Link Capabilities Register (Offset 0Ch)

...

Table 7-14: Link Capabilities Register

Bit Location	Register Description	Attributes
<u>21</u>	<u>Link Bandwidth Notification Capability</u> – A value of 1b indicates support for the Link Bandwidth Notification status and interrupt mechanisms. This capability is required for all Root Ports and Switch downstream ports supporting links wider than x1 and/or multiple link speeds . <u>This field is not applicable and is reserved for Endpoint devices, PCI Express to PCI/PCI-X bridges, and Upstream Ports of Switches.</u> <u>Devices that do not implement the Link Bandwidth Notification capability must hardwire this bit to 0b.</u>	<u>RO</u>

...

7.8.7. Link Control Register (Offset 10h)

...

Table 7-29: Link Control Register

Bit Location	Register Description	Attributes
10	<p>Link Bandwidth Management Interrupt Enable – When set to 1b this bit enables the generation of an interrupt to indicate that the Link Bandwidth Management Status bit has been set.</p> <p>Default value of this field is 0b.</p> <p>This bit is not applicable and is reserved for Endpoint devices, PCI Express to PCI/PCI-X bridges, and Upstream Ports of Switches.</p> <p>Devices that do not implement the Link Bandwidth Notification capability must hardwire this bit to 0b.</p>	RW

...

7.8.8 Link Status Register (Offset 12h)

...

Table 7-31: Link Status Register

Bit Location	Register Description	Attributes
14	<p>Link Bandwidth Management Status – This bit is set to 1b by hardware to indicate that either of the following has occurred without the port transitioning through DL_Down status:</p> <ul style="list-style-type: none"> A link retraining has completed following a write of 1b to the Retrain Link bit. <p>Note: This bit is set following any write of 1b to the Retrain Link bit, including when the link is in the process of retraining for some other reason.</p> <ul style="list-style-type: none"> Hardware has changed link speed or width to attempt to correct unreliable link operation, either through an LTSSM timeout or a higher level process. <p>This bit must be set if the Physical Layer reports a speed or width change was initiated by the downstream component that was not indicated as an Autonomous Change.</p> <p>The default value of this field is 0b.</p> <p>This bit is not applicable and is reserved for Endpoint devices, PCI Express to PCI/PCI-X bridges, and Upstream Ports of Switches.</p> <p>Devices that do not implement the Link Bandwidth Notification capability must hardwire this bit to 0b.</p>	RW1C