



PCI-SIG ENGINEERING CHANGE NOTICE

TITLE:	End-End TLP Prefix Changes for RCs
DATE:	May 26, 2010
AFFECTED DOCUMENTS:	TLP Prefix ECN, PCI Express Base Specification Rev 2.1, PCI Express Base Specification Rev 3.0
SPONSORS:	Intel Corporation, Hewlett-Packard Company, Advanced Micro Devices

Part I

1. Summary of the Functional Changes

This change allows for all Root Ports with the End-End TLP Prefix Supported bit Set to have different values for the Max End-End TLP Prefixes field in the Device Capabilities 2 register. It also changes and clarifies error handling for a Root Port receiving a TLP with more End-End TLP Prefixes than it supports.

2. Benefits as a Result of the Changes

This change enables End-End TLP Prefix to be enabled on multi-component root ports that may differ in values of Max End-End TLP prefix size capability. This change is applicable to Root Ports only.

3. Assessment of the Impact

This change is applicable to Root Ports only. This change will impact software that relies on all root ports that support End-End TLP Prefix to have the same values for the Max End-End TLP Prefix field in the Device Capabilities 2 Register.

4. Analysis of the Hardware Implications

Root Ports set the Max End-End TLP prefix in Device Capabilities 2 register according to component capabilities.

5. Analysis of the Software Implications

Software must consider this change while enabling TLP Prefix. For Peer-Peer access via Root Ports software must ensure that path from source to destination supports the required End-End TLP Prefix sizes before enabling TLP Prefix usage.

6. Analysis of the C&I Test Implications

C & I Test must comprehend this change in protocol testing. This change is applicable to Root Ports only.

Part II

Detailed Description of the change

Change Section 2.2.10.2 as follows

2.2.10.2. End-End TLP Prefix Processing

...

- ❑ Functions indicate how many End-End TLP Prefixes they support by the Max_End_End TLP Prefixes field in the Device Capabilities 2 register (see Section 7.8.15).

- For Root Ports, the Max End-End TLP Prefixes field is permitted to return a value indicating support for fewer End-End TLP Prefixes than what the Root Port hardware actually implements; however, the error handling semantics must still be based on the value contained in the field. TLPs received that contain more End-End TLP Prefixes than are supported by the Root Port must be handled as follows. It is recommended that Requests be handled as Unsupported Requests, but otherwise they must be handled as Malformed TLPs. It is recommended that Completions be handled as Unexpected Completions, but otherwise they must be handled as Malformed TLPs. For TLPs received by the Ingress Port, this is a reported error associated with the Ingress Port. For TLPs received internally to be transmitted out the Egress Port, this is a reported error associated with the Egress Port. See Section 6.2.

- For all other Function types, TLPs received that contain more End-End TLP Prefixes than are supported by a Function must be handled as Malformed TLPs. This is a reported error associated with the Receiving Port (see Section 6.2).

AER logging (if supported) occurs as specified in Section 6.2.4.4. ~~This is a reported error associated with the Receiving Port (see Section 6.2).~~

...

- ❑ ~~Root Complexes must support forwarding of TLPs with up to Max End-End TLP Prefixes if the End-End TLP Prefix Supported bit is Set in both the Ingress and Egress Ports and forwarding of TLPs is supported between those Ports.⁴⁸~~

- ❑ Different Root Ports with the End-End TLP Prefix Supported bit Set are permitted to report different values for Max End-End TLP Prefixes.

⁴⁸ ~~Root Port indication of End-End TLP Prefix Supported does not imply any particular level of peer-to-peer support by the RC, or that peer-to-peer traffic is supported at all (see Section 2.2.10.2.2).~~

Change Section 2.2.10.2.2 as follows

2.2.10.2.2. Root Ports with End-End TLP Prefix Supported

Support for peer-to-peer routing of TLPs containing End-End TLP Prefixes between Root Ports is optional and implementation dependent. If an RC supports End-End TLP Prefix routing capability between two or more Root Ports, it must indicate that capability in each associated Root Port via the End-End TLP Prefix Supported bit in the Device Capabilities 2 register.

An RC is not required to support End-End TLP Prefix routing between all pairs of Root Ports that have the End-End TLP Prefix Supported bit Set. A Request with End-End TLP Prefixes that would require routing between unsupported pairs of Root Ports must be handled as an Unsupported Request (UR). A Completion with End-End TLP Prefixes that would require routing between unsupported pairs of Root Ports must be handled as an Unexpected Completion (UC). In both cases, this error is reported by the “sending” Port.

The End-End TLP Prefix Supported bit must be Set for any Root Port that supports forwarding of TLPs with End-End TLP Prefixes initiated by host software or Root Complex Internal Endpoints. The End-End TLP Prefix Supported bit must be Set for any Root Ports that support forwarding of TLPs with End-End TLP Prefixes received on their Ingress Port to Root Complex Integrated Endpoints.

~~All Root Ports with the End-End TLP Prefix Supported bit Set must have the same value for the Max End-End TLP Prefixes field in the Device Capabilities 2 register.~~

Different Root Ports with the End-End TLP Prefix Supported bit Set are permitted to report different values for Max End-End TLP Prefixes.

An RC that splits a TLP into smaller TLPs when performing peer-to-peer routing between Root Ports must replicate the original TLP’s End-End TLP Prefixes in each of the smaller TLPs (see Section 1.3.1).

Change Section 6.2.4.4 as follows

6.2.4.4. TLP Prefix Logging

...

- ❑ A Function that receives a TLP containing more End-End TLP Prefixes than are indicated by the Function’s Max End-End TLP Prefixes field must handle the TLP as ~~a Malformed TLP~~ an error (see Section 2.2.10.2 for specifics) and store the first overflow End-End TLP Prefix in the 1st DW of the Header Log register with the remainder of the Header Log register being undefined.

Change Section 7.8.15 as follows

7.8.15. Device Capabilities 2 Register (Offset 24h)

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 See Section 2.2.10.2 for important details). Values are:</p> <table><tr><td>01b</td><td>1 End-End TLP Prefix</td></tr><tr><td>10b</td><td>2 End-End TLP Prefixes</td></tr><tr><td>11b</td><td>3 End-End TLP Prefixes</td></tr><tr><td>00b</td><td>4 End-End TLP Prefixes</td></tr></table> <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-Different Root Ports that have the End-End TLP Prefix Supported bit Set must contain the same value are permitted to report different values 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>	01b	1 End-End TLP Prefix	10b	2 End-End TLP Prefixes	11b	3 End-End TLP Prefixes	00b	4 End-End TLP Prefixes
01b	1 End-End TLP Prefix								
10b	2 End-End TLP Prefixes								
11b	3 End-End TLP Prefixes								
00b	4 End-End TLP Prefixes								

Change Section 7.10.12 as follows

7.10.12. TLP Prefix Log Register (Offset 38h)

...

The First TLP Prefix Log Register contains the first End-End TLP Prefix from the TLP (see Section 6.2.4.4). The Second TLP Prefix Log register contains the second End-End TLP Prefix and so forth. If the TLP contains fewer than four End-End TLP Prefixes, the remaining TLP Prefix Log Registers contain zero. A TLP that contains more End-End TLP Prefixes than are indicated by the Function's Max End-End TLP Prefixs field must be handled as ~~a Malformed TLP (see Section 2.2.10.2)~~ ~~an error (see Section 2.2.10.2 for specifics)~~. To allow software to detect this condition, the supported number of End-End TLP Prefixes are logged in this register, the first overflow End-End TLP Prefix is logged in the first DW of the Header Log register and the remaining DWs of the Header Log Register are undefined (see Section 6.2.4.4).