ENGINEERING CHANGE NOTIFICATION

Date: 2/13/03

Name: Dave Walker / Prashant Sethi Email: dwalker@microsoft.com prashant.sethi@intel.com

Company: Microsoft / Intel Mailstop: FM5-60

Address: 1900 Prairie City Road

City: Folsom State/Province: CA

Country: USA Zip/Postal Code: 95630

Phone: (916) 356-4026 FAX:

Your Company SIG Voting Member Name:

Phone: Email:

PART I

Brief description of the functional changes proposed:
This ECN describes changes necessary for support of integrated PCI Express devices in PCI Express
Root Complexes. This ECN does not impact any existing functionality or devices.
Specification(s) this proposed change is against:
PCI Express Base Specification 1.0a
Benefits as a result of the proposed changes:
PCI Express devices can be integrated in root complexes for high integration systems. This capability is
already supported by PCI devices today.
An assessment of the impact to the existing revision and systems that currently conform to the PCI specification:
No impact to systems or peripherals that conform to PCI Express 1.0a specification.
An analysis of the hardware implications:
Allows PCI Express devices to be integrated in root complexes; hardware supporting such devices will have
to implement appropriate configuration space changes.
An analysis of the software implications:
Software supporting integrated endpoints will have to recognize new PCI Express Device Type for advanced
capabilities; software compliant with 1.0a specification limits capabilities to those provided

by conventional PCI devices

ENGINEERING CHANGE NOTIFICATION

PART II

THE FOLLOWING CRITERIA APPLIES TO DEFINING A RESERVED PIN ON THE PCI CONNECTOR OR TO REQUEST A NEW COMMAND:

? The function is deemed to be important for the Continued growth and long-term well being of PCI. (A shot-term fixes to a problem does not merit consideration, i.e., IDEIRZ14.)

? There is no other effective way for the function to be implemented except using a reserved pin or new command. (Is there a solution that can be implemented in configuration space?)

Describe how the new function (on an Addin Card) works with an existing system and other existing Addin Cards?
New function is not relevant to Addin Card
Describe how existing Addin Cards work when added into a system with the new function?
Not applicable as new function is not relevant to Addin Card; existing Addin cards continue to work as
before with no changes.
Are there any combinations not addressed by the previous two items? If yes, specify them and describe the interaction between new device and existing devices.
N/A

Description of changes:

Terms and Acronyms, p.~~24~~ 23 – replace definition of Root Complex with –

Root Complex An entity that includes a Host Bridge, zero or more Root Complex Integrated Endpoints, zero or more Root Complex Event Collectors and one or more Root Ports.

1.3.2, p.32 – replace last sentence of 1st paragraph –

Endpoints are classified as either legacy, PCI Express or Root Complex Integrated Endpoints.

1.3.2, p.32 - add a new section 1.3.2.3, entitled "Root Complex Integrated Endpoint "Rules

- A Root Complex Integrated Endpoint is implemented on internal logic of Root Complexes that contains the Root Ports.

- A Root Complex Integrated Endpoint must be a device with a Type 00h Configuration Space header.

- A Root Complex Integrated Endpoint must support Configuration Requests as a Completer

- A Root Complex Integrated Endpoint must not require I/O resources claimed through BAR(s).

- A Root Complex Integrated Endpoint must not generate I/O Requests.

- A Root Complex Integrated Endpoint must not support Locked Requests as a Completer or generate them as a Requestor. PCI Express-compliant software drivers and applications must be written to prevent the use of lock semantics when accessing a Root Complex Integrated Endpoint.

- A Root Complex Integrated Endpoint operating as the Requester of a Memory Transaction is required to be capable of generating addresses equal to or greater than the Host is capable of handling as a Completer.

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

- A Root Complex Integrated Endpoint is required to support MSI if an interrupt resource is requested, but is permitted to support either the 32-bit or 64-bit Message Address version of the MSI capability structure.

- A Root Complex Integrated Endpoint is permitted to support 32-bit addressing for Base Address registers that request memory resources.

- A Root Complex Integrated Endpoint must not implement Link Capabilities/Status/Control registers in the PCI Express Extended Capability.

- 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 one Root Complex Event Collector.

- A Root Complex Integrated Endpoint does not implement Active State Power Management.

- A Root Complex Integrated Endpoint may not be hot-plugged independent of the Root Complex as a whole.

- A Root Complex Integrated Endpoint must not appear in any of the hierarchy domains exposed by the Root Complex.

- A Root Complex Integrated Endpoint must not appear in Switches.

1.3.4 PCI Express - PCI Bridge, p.34, insert new section 1.3.4 entitled "Root Complex Event Collector" as follows, renumber existing 1.3.4 to 1.3.5 PCI Express – PCI Bridge

- 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 05h and Programming Interface 00h.

- A Root Complex Event Collector resides on the same logical bus as the Root Complex Integrated Endpoints it supports.

- 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.

2.5.3, p.~~106~~100, VC and TC Rules, change last bullet to –

Root Complex must support independent TC/VC mapping configuration for each RCRB, the associated Root Ports, and any Root Complex Integrated Endpoints.

5.3.3.1, p.235 – replace last paragraph:

PME indications are propagated to the Root Complex in the form of TLP messages. PM_PME messages include the logical location of the requesting agent within the Hierarchy (in the form of the Requester ID of the PME message header). Explicit identification within the PM_PME message is intended to facilitate quicker PME service routine response, and hence shorter resume time.

With:

PME indications that originate from PCI Express Endpoints or PCI Express Legacy Endpoints are propagated to the Root Complex in the form of TLP messages. PM_PME messages include the logical location of the requesting agent within the Hierarchy (in the form of the Requester ID of the PME message header). Explicit identification within the PM_PME message is intended to facilitate quicker PME service routine response, and hence shorter resume time.

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 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.

PME indications that originate from a Root Port itself are reported through the same Root Port.

5.4.1, p.~~242~~243, replace 1st sentence of section –

All PCI Express components other than Root Complex Integrated Endpoints are required to support the minimum requirements defined herein for Active State Link PM. Root Complex Integrated Endpoints do not have an associated link, and therefore do not support Active State Link PM.

6.1.6, p.261, add a new paragraph to the end of the section –

Root Complex Event Collectors provide support for the above described functionality for Root Complex Integrated Endpoints.

6.1.9, p.262, add a new paragraph to the end of the section –

Because Root Complex Integrated Endpoints are not contained in any of the hierarchy domains originated by Root Ports, these endpoints signal system software of a PME using the same mechanism used in present PCI systems. A Root Complex Event Collector, if implemented, enables the PCI Express Native PME model for Root Complex Integrated Endpoints.

6.2.3.2, p.266, Error Messages – add after 1st paragraph –

Error messages that originate from PCI Express or Legacy endpoints are sent to corresponding Root Ports. Errors that originate from a Root Port itself are reported through the same Root Port.

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 one Root Complex Event Collector.

6.2.4.1, p.268, add a new paragraph to the end of the section –

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 one Root Complex Event Collector.

7.1, p.311, add to 3rd paragraph –

Both a PCI Express Endpoint and a Legacy Endpoint are required to appear within one of the Hierarchy Domains originated by the Root Complex, meaning that they appear in configuration space in a tree that has a Root Port as its head. Root Complex Integrated Endpoints and Root Complex Event Collectors do not appear within one of the Hierarchy Domains originated by the Root Complex. These appear in configuration space as peers of the Root Ports.

7.2.2.1, p.315, Host Bridge Requirements, add text at end of paragraph –

A PCI Express Host Bridge is not required to signal errors through a Root Complex Event Collector. This support is optional for PCI Express Host Bridges.

7.2.3., p.~~314~~315, Root Complex Register Block, change first sentence to –

A Root Port or Root Complex Integrated Endpoint may be associated with an optional 4096 byte block of memory mapped registers referred to as the Root Complex Register Block (RCRB).

7.8, p. ~~331~~333, replace 2nd paragraph –

Figure 7-10 details allocation of register fields in the PCI Express Capability Structure. The PCI Express Capabilities, Device Capabilities and Device Status/Control registers are required for all PCI Express devices. The Link Capabilities and Link Status/Control are required for all Endpoints that are not Root Complex Integrated Endpoints. Endpoints are not required to implement registers other than those listed above and terminate the capability structure.

Root Complex Event Collectors implement the Root Status/Control registers in addition to PCI Express Capabilities, Device Capabilities and Device Status/Control registers.

Figure 7-10, p.~~332~~ 334 –

The bracket and label "All Devices" should extend only through the Device Control/Status registers. A new bracket and label "Devices With Links" should be created, extending through Link Control/Status.

A new bracket and label “Root Complex Event Collectors” should be created encompassing Root Control/Status registers.

7.8.2~~.~~, Table 7-10, Bits 7:4, p.~~334~~336, add new encoding –

Device/Port Type - 1001b - Root Complex Integrated Endpoint Device

Device/Port Type - 1010b - Root Complex Event Collector

Table 7-20, p.361, Root Control Register, add following to bit definitions for bits 0, 1 and 2:

Root Complex Event Collectors provide support for the above described functionality for Root Complex Integrated Endpoints.

7.10, p.364, Advanced Error Reporting Capability, replace 1st paragraph with –

The PCI Express Advanced Error Reporting capability is an optional extended capability that may be implemented by PCI Express devices supporting advanced error control and reporting. The Advanced Error Reporting capability structure definition has additional interpretation for Root Ports and Root Complex Event Collectors; software must interpret the PCI Express device/Port Type field (Section 7.8.1) in the PCI Express Capability Structure to determine the availability of additional registers for Root Ports and Root Complex Event Collectors.

Figure 7-26, p.365 –

Change “Root Ports Only” bracket at end to read “Root Ports and Root Complex Event Collectors”

7.10.9, p.373, Root Error Command Register, change text in parentheses in 1st paragraph:

(claimed by the Root Port)

To:

(claimed by the Root Port or Root Complex Event Collector)

Table 7-31, p.374, Root Error Command Register, add following to bit definitions for bits 0, 1 and 2:

Root Complex Event Collectors provide support for the above described functionality for Root Complex Integrated Endpoints.

7.10.10, p.374, Root Error Status Register, add at end of 1st paragraph:

Root Complex Event Collectors provide support for the above described functionality for Root Complex Integrated Endpoints (and for the Root Complex Event Collector itself).

Add new capability as follows:

<Ch 7 new material – New Section 7.15>

PCI Express Root Complex Event Collector Endpoint Association Capability

The PCI Express Root Complex Event Collector Endpoint Association Capability is implemented by Root Complex Event Collectors.

It declares the Root Complex Integrated Endpoints supported by the Root Complex Event Collector on the same logical bus on which the Root Complex Event Collector is located. A Root Complex Event Collector must implement the Root Complex Event Collector Endpoint Association Capability; no other PCI Express device may implement this capability.

The PCI Express Root Complex Event Collector Endpoint Association Capability, as shown in **Error! Reference source not found.**, consists of the PCI Express Enhanced Capability Header followed by a DWORD bitmap enumerating Root Complex Integrated Endpoints associated with the Root Complex Event Collector.

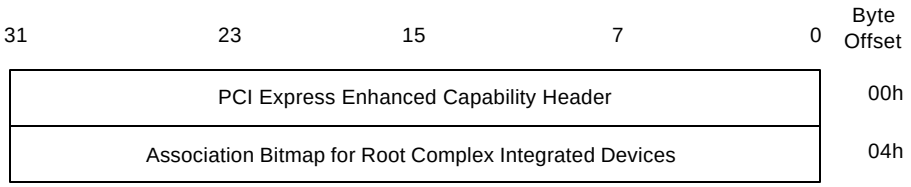

Figure 0-1 Root Complex Event Collector Endpoint Association Capability

<Section 7.15.1> Root Complex Event Collector Endpoint Association Enhanced Capability Header

The Extended Capability ID for the Root Complex Event Collector Endpoint Association Capability is 0007h.

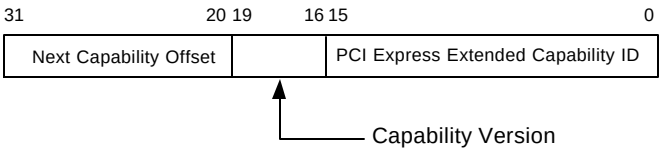

Figure 0-2 Root Complex Event Collector Endpoint Association Enhanced Capability Header

Table 0-1 Root Complex Event Collector Endpoint Association Enhanced Capability Header

Bit Location	Description	Register Attribute
15:0	**PCI Express Extended Capability ID** – This field is a PCI-SIG defined ID number that indicates the nature and format of the extended capability. Extended Capability ID for the Root Complex Event Collector Endpoint Association Capability is 0007h.	RO
19:16	**Capability Version** – This field is a PCI-SIG defined version number that indicates the version of the	RO

	capability structure present. Must be 1h for this version of the specification.	
31:20	**Next Capability Offset** – This field contains the offset to the next PCI Express capability structure or 000h if no other items exist in the linked list of capabilities. For Extended Capabilities implemented in device configuration space, this offset is relative to the beginning of PCI compatible configuration space and thus must always be either 000h (for terminating list of capabilities) or greater than 0FFh. The bottom two bits of this offset are reserved and must be implemented as 00b although software must mask them to allow for future uses of these bits.	RO

<Section 7.15.2> Association Bitmap for Root Complex Integrated Endpoint Devices

The Association Bitmap for Root Complex Integrated Endpoint Devices is a read-only field that sets the bits corresponding to the Device Numbers of Root Complex Integrated Endpoint Devices supported by the Root Complex Event Collector on the same logical bus as the Event Collector itself. The bit corresponding to the Device Number of the Root Complex Event Collector must always be set.