



PCI-SIG ENGINEERING CHANGE NOTICE

TITLE:	PCIe hot plug
DATE:	10/7/13
AFFECTED DOCUMENT:	PCI Firmware Specification Ver 3.1
SPONSOR:	Mohan J. Kumar, Intel

Part I

1. Summary of the Functional Changes

Clarifications to support PCIe surprise hot plug to allow certain errors to be suppressed by platform if software stack ensured error containment. Also removes the implementation note in section 4.5.2.4 which is not representative of OS behavior.

2. Benefits as a Result of the Changes

Enables uniform and better support for PCIe surprise hot plug in ecosystem. MMIO accesses targeting a surprise removed PCIe device cause Completion Timeouts that result in uncorrected error signaling that potentially brings the system down. If software ensures error containment by preventing interpreting all 1s response due to completion timeout as valid data, then platform can suppress the uncorrected error signaling. This ECR provides a mechanism for OS and platform to communicate their capabilities with respect to handling of surprise remove related errors.

3. Assessment of the Impact

Changes to platform and OS behavior.

4. Analysis of the Hardware Implications

None

5. Analysis of the Software Implications

Change to BIOS and OS

Part II

Detailed Description of the change

From

4.5.2.4. Dependencies Between _OSC Control Bits

Because handling of hot-plug events, power management events, and advanced error reporting all require the modification of PCI Express Capability registers, the operating system is required to claim control over the PCI Express Capability (bit 4 of the Control field) in conjunction with claiming control over PCI Express Native Hot Plug, PCI Express Native Power Management Events, or PCI Express Advanced Error Reporting (bits 0, 2, and 3 of the Control field). If the operating system attempts to claim control of any of these features without also claiming control over the PCI Express Capability, the firmware is required to refuse control of the feature being illegally claimed and mask the corresponding bit.

IMPLEMENTATION NOTE

Some operating systems may require the platform to grant control over all of PCI Express Native Hot Plug, PCI Express Native Power Management Events, PCI Express Advanced Error Reporting and the PCI Express Capability (bits 0, 2, 3, and 4) in an “ALL or NOTHING” manner. If control of any one of these capabilities is denied by the platform, such operating systems may choose not to claim control of any of these features. Platforms that support these operating systems need to take *this into account*.

To

4.5.2.4. Dependencies Between _OSC Control Bits

Because handling of hot-plug events, power management events, and advanced error reporting all require the modification of PCI Express Capability registers, the operating system is required to claim control over the PCI Express Capability (bit 4 of the Control field) in conjunction with claiming control over PCI Express Native Hot Plug, PCI Express Native Power Management Events, or PCI Express Advanced Error Reporting (bits 0, 2, and 3 of the Control field). If the operating system attempts to claim control of any of these features without also claiming control over the PCI Express Capability, the firmware is required to refuse control of the feature being illegally claimed and mask the corresponding bit.

Operating systems must comprehend that platforms may not grant control of Native PCI Express Advanced Error Reporting feature.

Operating System determines that slots in the platform support surprise hot plug by consulting the slot capabilities register of corresponding root port or bridge.

IMPLEMENTATION NOTE

~~Some operating systems may require the platform to grant control over all of PCI Express Native Hot Plug, PCI Express Native Power Management Events, PCI Express Advanced Error Reporting and the PCI Express Capability (bits 0, 2, 3, and 4) in an “ALL or NOTHING” manner. If control of any one of these capabilities is denied by the platform, such operating systems may choose not to claim control of any of these features. Platforms that support these operating systems need to take *this into account*.~~

Interpretation of the _OSC Control Field, Passed in via Arg3

Control Field Bit Offset	Interpretation
6	OS sets this bit to 1 to indicate that all certified IO drivers will never interpret as

	valid data an all 1's response to a read from a Function's Memory Mapped IO space (MMIO BAR memory), IO space, or Configuration space. An all 1's read response is returned by hardware to indicate certain error conditions including surprise removal of a PCIe device, and reliable detection of those error cases requires that an all 1's response must never interpreted as valid by an IO driver OS sets this bit to 1 to indicate that all certified IO drivers will handle all 1s response caused by accesses to surprise removed PCIe device.
7-31	Reserved

Interpretation of the _OSC Control

Field, Returned Value

Control Field Bit Offset	Interpretation
6	Firmware sets this bit to 1 to indicate that it will suppress error notification caused by surprise hot remove (Completion Time Out errors).
7-31	Reserved