



PCI-SIG ENGINEERING CHANGE REQUEST

TITLE:	Hot Plug
DATE:	24 June 2004 [Updated 12 July 2004]
AFFECTED DOCUMENT:	PCI Express Base Specification Revision 1.0a
SPONSOR:	Mohan Kumar & David Harriman, Intel Corporation Dave Walker, Microsoft Carl Jackson, HP

Part I

1. Summary of the Functional Changes

- Removes the requirement for Hot plug messages described in Section 2.2.8.7
- Replace firmware interface description (OSHP) and replace with reference to the PCI Express Firmware specification
- Comprehend Electro Mechanical Interlock (EMI) mechanism
- Add a capability bit for command complete to support form factors (e.g. Express Card) that do not require command complete functionality
- Specify the mechanism to indicate to software when it is OK to initiate configuration transaction to a hot added device
- Specify the mechanism for how software determines that the power controller off command is complete

Note that this document comprehends preceding errata C21.

2. Benefits as a Result of the Changes

Improves hot plug mechanism by simplifying implementation requirements and adding new mechanisms & capabilities.

3. Assessment of the Impact

No negative impact.

4. Analysis of the Hardware Implications

No negative impact. New optional capabilities.

5. Analysis of the Software Implications

Hot Plug software operation is further clarified (beyond material included in earlier hot plug errata).

Part II

Detailed Description of the change

2.2.8.7 Hot-Plug Signaling Ignored Messages

The messages listed in Table 2-19 were previously used for a mechanism that is no longer supported. Transmitters are strongly encouraged not to transmit these messages, but if message transmission is implemented, it must conform to the requirements of 1.0a version of this specification.

Receivers are strongly encouraged to ignore receipt of these messages, but are allowed to process these messages in conformance with the requirements of 1.0a version of this specification.

Ignored messages listed in Table 2-19 are handled by the receiver as follows:

- The Physical and Data Link Layers must handle these messages identical to handling any other TLP
- The Transaction layer must account for flow control credit but take no other action in response to these messages

The hot-plug signaling messages (Table 2-19) are virtual signals between Switches/Root Ports used to support hot-plug event signaling and devices on adapters that support Removal Request functionality (doorbell mechanism) on the adapter. For more information see Section 6.7.

Table 2-19: Hot-Plug Signaling Ignored Messages

Name	Code[7:0]	Routing r[2:0]	Support				Req ID	Description/Comments
			R C	E p	S w	B r		
<u>Attention_Indicator_OnIgnored Message</u>	0100 0001	100	t	f	tf	f	BDF	This Message is issued by the Switch/Root Port when the Attention Indicator Control is set to 01b. The end device receiving the Message will terminate the Message and initiate appropriate action to cause the Attention Indicator located on the card to turn on. See note.
<u>Attention_Indicator_BlinkIgnored Message</u>	0100 0011	100	t	f	tf	f	BDF	This Message is issued by the Switch/Root Port when the Attention Indicator Control is set to 10b. The end device receiving the Message will terminate the Message and initiate appropriate action to cause the Attention Indicator located on the card to blink. See note.
<u>Attention_Indicator_OffIgnored Message</u>	0100 0000	100	t	f	tf	f	BDF	This Message is issued by the Switch/Root Port when the Attention Indicator Control is set to 11b. The end device receiving the Message will terminate the Message and initiate appropriate action to cause the Attention Indicator located on the card to turn off. See note.

Name	Code[7:0]	Routing r[2:0]	Support				Req ID	Description/Comments
			R C	E p	S w	B r		
<u>Power_Indicator_OnIgnored Message</u>	0100 0101	100	t	f	tf	f	BDF	This Message is issued by the Switch/Root Port when the Power Indicator Command is set to 01b. The end device receiving the Message will terminate the Message and initiate appropriate action to cause the Power Indicator located on the card to turn on. See note.
<u>Power_Indicator_BlinkIgnored Message</u>	0100 0111	100	t	f	tf	f	BDF	This Message is issued by the Switch/Root Port when the Power Indicator Control is set to 10b. The end device receiving the Message will terminate the Message and initiate appropriate action to cause the Power Indicator located on the card to blink. See note.
<u>Power_Indicator_OffIgnored Message</u>	0100 0100	100	t	f	tf	f	BDF	This Message is issued by the Switch/Root Port when the Power Indicator Command is set to 11b. The end device receiving the Message will terminate the Message and initiate appropriate action to cause the Power Indicator located on the card to turn off. See note.
<u>Attention_Button_Pressed Ignored Message</u>	0100 1000	100	All: f f As Required: t t t				BDF	This Message is issued by a device in a slot that implements an Attention Button on the card to signal the Switch/Root Port to generate the Attention Button Pressed Event. The Switch Switch/Root Port terminates the Message and sets the Attention Button Pressed register to 1b which may result in an interrupt being generated.

Note: If no indicators are present on the card, the Message is ignored.

~~Components must accept received hot-plug signaling messages as defined in the Support column of Table 2-19 without indicating an error. However, only components which implement or support the associated action must process the received Message and take the corresponding action. Components which do not implement or support the associated action discard the received Message.~~

...

2.3.1. Request Handling Rules

...

- ❑ If the Request is a Message, and the Message Code specifies a value that is undefined, or that corresponds to a Message not supported by the device (other than Vendor_Defined Type 1 which is not treated as an error – see Section 2.2.8.6), the Request is an Unsupported Request, and is reported according to Section 6.2
 - If the Message Code is a supported value, process the Message according to the corresponding Message processing rules; if the Message Code is an ignored Message, then ignore the Message without reporting any error (see Section 2.2.8.7).

...

6.6. PCI Express Reset – Rules

...

The second set of rules addresses requirements placed on the system:

- ❑ To allow components to perform internal initialization, system software must wait for at least 100 ms from the end of a reset of one or more device before it is permitted to issue Configuration Requests to those devices.
 - A system must guarantee that all components intended to be software visible at boot time are ready to receive Configuration Requests within 100 ms of the end of Fundamental Reset at the Root Complex – how this is done is beyond the scope of this specification

- ❑ The root complex and/or system software must allow 1.0s (+50%/-0%) after a reset of a device, before it may determine that a device which fails to return a Successful Completion status for a valid Configuration Request is a broken device

Note: This delay is analogous to the Trhfa parameter specified for PCI/PCI-X, and is intended to allow an adequate amount of time for devices which require self initialization.

- ❑ When attempting a Configuration access to devices on a PCI or PCI-X segment behind a PCI Express/PCI(-X) Bridge, the timing parameter Trhfa must be respected

For this second set of rules, if system software does not have direct visibility into the state of Fundamental Reset (eg, Hot-Plug, see Section 6.7), software must base these timing parameters on an event known to occur after the end of Fundamental Reset.

...

6.7.1 Elements of Hot-Plug

Table <6-5> lists the physical elements comprehended in PCI Express for support of hot-plug models. A PCI Express form-factor specification must define how these elements are used in that form-factor. For a given form-factor specification, it is possible that only some of the available hot-plug elements are required, or even that none of these elements are required. In all cases, the form-factor specification must define all assumptions and limitations placed on the system or the user by the choice of elements included. Silicon component implementations that are intended to be used only with selected form factors are permitted to support only those elements that are required by the associated form factor(s).

Table 6-5: ~~User Visible Elements of the Standard Usage Model~~ Elements of Hot-Plug

Element	Purpose
Indicators	Show the power and attention state of the slot
Manually-operated Retention Latch (MRL)	Holds adapter in place
MRL Sensor	Allows the Port and system software to detect the MRL being opened
Electromechanical Interlock	Prevents removal of adapter while from slot is powered
Attention Button	Allows user to request hot-plug operations
Software User Interface	Allows user to request hot-plug operations
Slot Numbering	Provides visual identification of slots
<u>Power Controller</u>	<u>Software-controlled Eelectronic component or components that control power to a slot or adapter and monitor that power for fault conditions</u>

6.7.1.2.1. Indicators

...

6.7.1.2.2. Manually-operated Retention Latch (MRL)

...

6.7.1.2.3. MRL Sensor

...

6.7.1.2.4. Electromechanical Interlock

An electromechanical interlock is a mechanism for physically locking the adapter or MRL in place until ~~the system software and Port~~ releases it. The state of the electromechanical interlock is set by software and must not change except in response to a subsequent software command. In particular, the state of the electromechanical

interlock must be maintained even when power to the hot-plug slot is removed. Requirements for the electromechanical interlock are called out in the associated form factor specification.

The current state of the electromechanical interlock must be reflected at all times in the Electromechanical Interlock Status field in the Slot Status register, which must be updated within 200 ms of any commanded change. Software must wait at least 1 second after issuing a command to toggle the state of the Electromechanical Interlock before another command to toggle the state can be issued. Systems may optionally expand control of interlocks to provide physical security of the adapter.

6.7.1.2.5. Attention Button

The Attention Button is a momentary-contact push button switch, located adjacent to each hot-plug slot or on the adapter that is pressed by the user to initiate a hot-plug operation at that slot. Regardless of the physical location of the button, the signal is processed and indicated to software by hot-plug hardware associated with the downstream port corresponding to the slot.

~~In form factors that electrically control the Attention Button by the chassis, the Attention Button must allow the user to initiate both hot add and hot remove operations regardless of the physical location of the button. In form factors that electrically control the Attention Button by the component on the adapter, the attention button can be used to initiate hot remove operations only.~~

If present, the Power Indicator provides visual feedback to the human operator (if the system software accepts the request initiated by the Attention Button) by blinking. Once the Power Indicator begins blinking, a 5-second abort interval exists during which a second depression of the Attention Button cancels the operation.

If an operation initiated by an Attention Button fails for any reason, it is recommended that system software present an error message explaining the failure via a software user interface or add the error message to a system log.

6.7.1.2.6. Software User Interface

...

6.7.1.2.7. Slot Numbering

...

6.7.1.2.8. Power Controller

The power controller is an element composed of one or more discrete components that acts under control of software to set the power state of either the hot-plug slot or the adapter as appropriate for the specific form factor. The power controller must also monitor the slot/adapter for power fault conditions (as defined in the associated form factor specification) that occur on the slot/adapter's main power rails and, if supported, auxiliary power rail.

If a power controller is not present, the power state of the hot-plug slot/adapter must be set automatically by the hot-plug controller in response to changes in the presence of an adapter in the slot.

The power controller monitors main and auxiliary power faults independently. If a power controller detects a main power fault on the hot-plug slot/adapter, it must automatically set its internal main power fault latch and remove main power from the hot-plug slot/adapter (without affecting auxiliary power). Similarly, if a power controller detects an auxiliary power fault on the hot-plug slot/adapter, it must automatically set its internal auxiliary power fault latch and remove auxiliary power from the hot-plug slot/adapter (without affecting main power). Power must remain off to the slot/adapter as long as the power fault condition remains latched, regardless of any writes by software to turn on power to the hot-plug slot/adapter. The main power fault latch is cleared when software turns off power to the hot-plug slot/adapter. The mechanism by which the auxiliary power fault latch is cleared is form factor specific but generally requires auxiliary power to be removed from the hot-plug slot/adapter. For example, one form factor may remove auxiliary power when the MRL for the slot is opened while another may require the adapter to be physically removed from the slot. See the associated form factor specifications for specific requirements.

Since the Power Controller Control field in the Slot Control register reflects the last value written and not the actual state of the power controller, this means there may be an inconsistency between the value of the Power Controller Control field and the state of the power to the slot/adapter in a power fault condition. To determine whether slot/adapter power is off due to a power fault, software must use the power fault software notification to detect power faults. **To determine that a requested power-up operation has otherwise failed, software must use** the hot-plug slot/adapter power-up time out mechanism described in section 6.7.7.3.

Software must not assume that writing to the Slot Control register to change the power state of a hot-plug slot/adapter causes an immediate power state transition. After turning power on, software must wait for a Data Link Layer State Changed event, as described in Section 6.7.7.3. After turning power off, software must wait for at least 1 second before taking any action that relies on power having been removed from the hot-plug slot/adapter. For example, software is not permitted to turn off the power indicator (if present) or attempt to turn on the power controller before completing the 1 second wait period.

6.7.3. Registers Grouped by Hot-Plug Element Association

...

6.7.3.1. Attention Button Registers

Attention Button Present (Slot Capabilities & Device Capabilities) – This bit indicates if an Attention Button is electrically controlled by the chassis (Slot Capabilities) or by the adapter (Device Capabilities).

Attention Button Pressed (Slot Status) – This bit is set when an Attention Button electrically controlled by the chassis is pressed. ~~This bit is also set upon receipt of an Attention_Button_Pressed Message from the adapter.~~

Attention Button Pressed Enable (Slot Control) – When set, this bit enables software notification on an Attention Button Pressed event (see Section 6.7.7.3).

6.7.3.2. Attention Indicator Registers

Attention Indicator Present (Slot Capabilities & Device Capabilities) – This bit indicates if an Attention Indicator is electrically controlled by the chassis (Slot Capabilities) or by the adapter (Device Capabilities).

Attention Indicator Control (Slot Control)– When written, sets an Attention Indicator electrically controlled by the chassis to the written state ~~and sends the appropriate Attention Indicator Message (determined by the decoding) to the device on the adapter.~~

6.7.3.3. Power Indicator Registers

Power Indicator Present (Slot Capabilities & Device Capabilities) – This bit indicates if a Power Indicator is electrically controlled by the chassis (Slot Capabilities) or by the adapter (Device Capabilities).

Power Indicator Control (Slot Control) – When written, sets a Power Indicator electrically controlled by the chassis to the written state ~~and sends the appropriate Power Indicator Message (determined by the decoding) to the device on the adapter.~~

6.7.3.4. Power Controller Registers

Power Controller Present (Slot Capabilities) – This bit indicates if a Power Controller is implemented.

Power Controller Control (Slot Control) – Turns the Power Controller on or off according to the value written.

Power Fault Detected (Slot Status) – This bit is set when a power fault is detected at the slot or the adapter.

Power Fault Detected Enable (Slot Control) – When set, this bit enables software notification on a power fault event (see Section 6.7.7.3).

6.7.3.5. Presence Detect Registers

Presence Detect State (Slot Status) – This bit indicates the presence of an adapter in the slot.

Presence Detect Changed (Slot Status) – This bit is set when a Presence Detect state change is detected.

Presence Detect Changed Enable (Slot Control) – When set, this bit enables software notification on a presence detect changed event (see Section 6.7.7.3)

6.7.3.6. MRL Sensor Registers

MRL Sensor Present (Slot Capabilities) – This bit indicates if an MRL Sensor is implemented.

MRL Sensor Changed (Slot Status) – This bit is set when the value of the MRL Sensor state changes.

MRL Sensor Changed Enable (Slot Control) – This bit when set enables software notification on a MRL Sensor changed event (see Section 6.7.7.3).

MRL Sensor State (Slot Status) – This register reports the status of the MRL Sensor if one is implemented.

6.7.3.7 Electromechanical Interlock Registers

Electromechanical Interlock Present (Slot Capabilities) – This bit indicates if an Electromechanical Interlock is implemented.

Electromechanical Interlock Status (Slot Status) – This bit reflects the current state of the Electromechanical Interlock.

Electromechanical Interlock Control (Slot Control) – This bit when set to 1b toggles the state of the Electromechanical Interlock.

6.7.3.87. Command Completed Registers

No Command Completed Support (Slot Capabilities) – This bit when set to 1b indicates that this slot does not generate software notification when an issued command is completed by the Hot-Plug Controller.

Command Completed (Slot Status) – This bit is set when the Hot-Plug Controller completes an issued command and is ready to accept the next command.

Command Completed Interrupt Enable (Slot Control) – This bit when set enables software notification (see Section 6.7.7.3) when a command is completed by the hot-plug control logic.

6.7.3.98. Port Capabilities and Slot Information Registers

Slot Implemented (PCI Express Capabilities) – When set, this bit indicates that the Link associated with this Downstream Port is connected to a slot.

Physical Slot Number (Slot Capabilities) – This hardware initialized field indicates the physical slot number attached to the Port.

Hot-Plug Capable (Slot Capabilities) – When set, this bit indicates that this slot is capable of supporting hot-plug.

Hot-Plug Surprise (Slot Capabilities) – When set, this bit indicates that adapter removal from the system without any prior notification is permitted for the associated form factor.

6.7.3.109. Hot-Plug Interrupt Control Register

Hot-Plug Interrupt Enable (Slot Control) – When set, this bit enables generation of the hot-plug interrupt on enabled hot-plug events.

<ed. note: The information in the following section has been comprehended above>

6.7.5. Messages

The Messages defined here allow for cards to implement indicators and buttons on the adapter without having to connect signals directly to the Port. Detailed explanation of each Message is located in Chapter 2.

6.7.5.1. Messages for the Attention Indicator

This series of Messages allows the Attention Indicator to be electrically controlled by the device on the adapter as opposed to the chassis. These Messages are sent by the Downstream Port to the device and instruct the device to set its Attention Indicator to the indicated state. The following Messages are used:

Attention_Indicator_On

Attention_Indicator_Blink

Attention_Indicator_Off

All Endpoint devices are required to accept without error the Attention Indicator Messages even if the device does not implement the indicator.

6.7.5.2. Messages for Power Indicator

This series of Messages allows the Power Indicator to be electrically controlled by the device on the adapter as opposed to the chassis. These Messages are sent by the Downstream Port to the device and instruct the device to set its Power Indicator to the indicated state. The following Messages are used:

Power_Indicator_On

Power_Indicator_Blink

Power_Indicator_Off

All Endpoint devices are required to accept without error the Power Indicator Messages even if the device does not implement the indicator.

6.7.5.3 Messages for Attention Button

Attention_Button_Pressed—This Message allows the attention button to be electrically controlled by the device on the adapter and informs the Port that the attention button has been pressed. Upon receipt of this Message the Port sets the Attention Button Pressed bit in the Slot Status register.

All downstream Ports of Switches and Root Ports are required to accept without error the Attention_Button_Pressed Message even if the slot is not hot plug capable.

6.7.6 Hot-Plug Slot Register Requirements

The following register fields are required for all slots that implement the hot plug capability:

Attention Button Pressed and Attention Button Pressed Enable

~~Attention Button Pressed Enabled (bit 0 of Slot Control register) and Attention Button Pressed (bit 0 of Slot Status register) are required to be implemented by all hot-plug capable slots to provide support for the Attention_Button_Pressed Message.~~

~~Attention Indicator Control~~

~~Attention Indicator Control field (bits 7:6 of the Slot Control register) must be implemented by all hot-plug capable slots to provide support for the Attention Indicator Messages.~~

~~Power Indicator Control~~

~~Power Indicator Control field (bits 9:8 of the Slot Control register) must be implemented by all hot-plug capable slots to provide support for the Power Indicator Messages.~~

~~Command Completed~~

~~All fields associated with Command Completion (Section 6.7.3.7) must be implemented by all hot-plug capable slots~~

~~Presence Detect~~

~~All fields associated with Presence Detect (Section 6.7.3.5) must be implemented by all hot-plug capable slots~~

6.7.7. PCI Express Hot-Plug Events

A Downstream Port with hot-plug capabilities supports the following hot-plug events:

Slot Events:

- Attention Button Pressed
- Power Fault Detected
- MRL Sensor Changed
- Presence Detect Changed

[Command Completed Events](#)

[Data Link Layer State Changed Events](#)

Each of these events has a status field, which indicates that an event has occurred but has not yet been processed by software, and an enable field, which indicates whether the event is enabled for software notification. Some events also have a capability field, which indicates whether the event type is supported on the port. The grouping of these fields by event type is listed in Section 6.7.3, and each individual field is described in Section 7.8.

6.7.7.1 Slot Events

...

6.7.7.2 Command Completed Events

Since changing the state of some hot-plug elements may not happen instantaneously, PCI Express supports hot-plug commands and command completed events. All hot-plug

capable ports are required to support hot-plug commands and if the capability is reported, command completed events.

Software issues a command to a hot-plug capable Downstream Port by issuing a write transaction that targets any portion of the port's Slot Control Register. A single write to the Slot Control register is considered to be a single command, even if the write affects more than one field in the Slot Control register. In response to this transaction, the port must carry out the requested actions and then set the associated status field for the command completed event. The port must process the command normally even if the status field is already set when the command is issued. If a single command results in more than one action being initiated, the order in which the actions are executed is unspecified. **All actions associated with a single cCommand** execution must not take longer than 1 second.

If command completed events are not supported as indicated by a value of 1b in the No Command Completed Support field of the Slot Capabilities register, a hot-plug capable port must process a write transaction that targets any portion of the port's Slot Control register without any dependency on previous Slot Control writes. Software is permitted to issue multiple Slot Control writes in sequence without any delay between the writes.

...

6.7.7.3 Data Link Layer State Changed Events

The Data Link Layer State Changed event provides an indication that the state of the Data Link Layer Link Active field in the Link Status register has changed. **Support for Data Link Layer State Changed events and software notification of these events are required for hot-plug capable downstream ports.** If this event is supported, the port sets the status field associated with the event when the value in the Data Link Layer Link Active field changes.

This event allows software to indirectly determine when power has been applied to a newly hot-plugged adapter. Software must wait for 100ms after the Data Link Layer Link Active field reads 1b before initiating a configuration access to the hot added device (see Section 6.6). Software must allow 1 second after the Data Link Layer Link Active field reads 1b before it is permitted to determine that a hot plugged device which fails to return a Successful Completion for a Valid Configuration Request is a broken device (see Section 6.6).

The Data Link Layer State Changed event must occur within 1 second of the event that initiates the hot-insertion. If a power controller is supported, the time out interval is measured from when software initiated a write to the Slot Control Register to turn on the power. If a power controller is not supported, the time out interval is measured from presence detect slot event. Software is allowed to time out on a hot add operation if the Data Link Layer State Changed event does not occur within 1 second. The action taken by software after such a timeout is implementation specific.

6.7.7.43 Software Notification of Hot-Plug Events

A hot-plug capable Downstream Port must support generation of an interrupt on a hot-plug event. As described in Sections 6.7.7.1 and 6.7.7.2, each type of hot-plug event has both an enable field which enables software notification of that type of event and a status field that indicates that an event has occurred but has not yet been processed by software.

If the enable field for an event indicates that software notification of the event is disabled, the hot-plug capable Downstream Port must not generate an interrupt on that event. Hot-plug interrupt generation is also globally enabled through the Hot-Plug Interrupt Enable field in the Slot Control register. If this field indicates that hot-plug interrupts are disabled, the hot-plug capable Downstream Port must not generate an ~~interrupt~~ interrupt on any hot-plug event.

The Downstream Port must generate an interrupt on a hot-plug event following the rules for interrupt generation described in the rest of this document provided these disable conditions are not present and the corresponding status field for the event is not set. If the status field is set, software has not yet handled a previous event; thus, a new interrupt must not be generated.

If the port is enabled for level-triggered interrupt generation using the INTx messages, the interrupt must remain asserted as long as global interrupt generation is enabled and at least one status field for an enabled hot-plug event remains set. If the port is enabled for edge-triggered interrupt generation using MSI or MSI-X, an MSI must be sent whenever global interrupt generation is enabled and the status field for an enabled hot-plug event transitions from not set to set. The port may optionally send an MSI when there are hot-plug events that occur while interrupt generation is disabled, and interrupt generation is subsequently enabled.

If wake generation is required by the associated form factor specification, a hot-plug capable Downstream Port must support generation of a wakeup event (using the PME mechanism) on hot-plug events that occur when the system is in a sleep state or the port is in device state D1, D2, or D3_{Hot}.

Software enables a hot-plug event to generate a wakeup event by enabling software notification of the event as described in Section 6.7.7.1. Note that in order for software to disable interrupt generation while keeping wakeup generation enabled, the Hot-Plug Interrupt Enable bit must be cleared. For form factors that support wake generation, a wakeup event must be generated if all three of the following conditions occur:

- The status register for an enabled event transitions from not set to set
- The port is in device state D1, D2, or D3_{Hot}, and
- The PME Enable bit in the port's Power Management Control/Status Register is set

Note that the Hot-Plug Controller generates the wakeup on behalf of the hot-plugged device, and it is not necessary for that device to have auxiliary (or main) power.

6.7.8. The Operating System Hot-Plug Method

Firmware support for Hot-Plug

Some systems that include hot-plug capable Root Ports and Switches that are released before ACPI-compliant operating systems with native hot-plug support are available, can use ACPI firmware for propagating hot-plug events. Firmware control of the hot-plug registers must be disabled if an operating system with native support is used. Platforms that provide ACPI firmware to propagate hot-plug events must also provide a [control method mechanism](#) to transfer control to the operating system. [The details of this mechanism are described in the PCI Firmware Specification.](#)

~~This method is called Operating System Hot-Plug (OSHP) and is provided for each Port that is hot-plug capable and being controlled by ACPI firmware.~~

~~Operating systems with native hot-plug support must execute the OSHP method, if present, for each hot-plug capable Port before accessing the hot-plug registers and when returning from a hibernated state. If a Port's OSHP method is executed multiple times, and the Switch to operating system control has already been achieved, the method must return successfully without doing anything. After the OSHP method is executed, the firmware must not access the Port's hot-plug registers. If any signals such as the System Interrupt or PME# have been redirected for servicing by the firmware, they must be restored appropriately for operating system control.~~

~~The following is an example of a namespace entry for an SHPC that is managed by firmware.~~

```
Device(PPB1){  
  
    ...  
  
    Method(OSHP, 0){  
        // Disable firmware access to SHPC and restore  
        // the normal System Interrupt and Wakeup mechanism.  
    }  
  
    ...  
}
```



IMPLEMENTATION NOTE

Controlling Hot-Plug by Using ACPI

When using ACPI to control the hot-plug events, the following should be considered:

Firmware should redirect the System Interrupt to a GPE so that ACPI can service the interrupts instead of the operating system. An appropriate *_Exx* GPE handler should be provided. When an operating system with native hot-plug support executes the OSHP

method, the firmware restores the normal System Interrupt so the interrupts can be serviced by the operating system.

...

7.8.3. Device Capabilities Register (Offset 04h)

The Device Capabilities register identifies PCI Express device specific capabilities. Figure 7-13 details allocation of register fields in the Device Capabilities register; Table 7-11 provides the respective bit definitions.

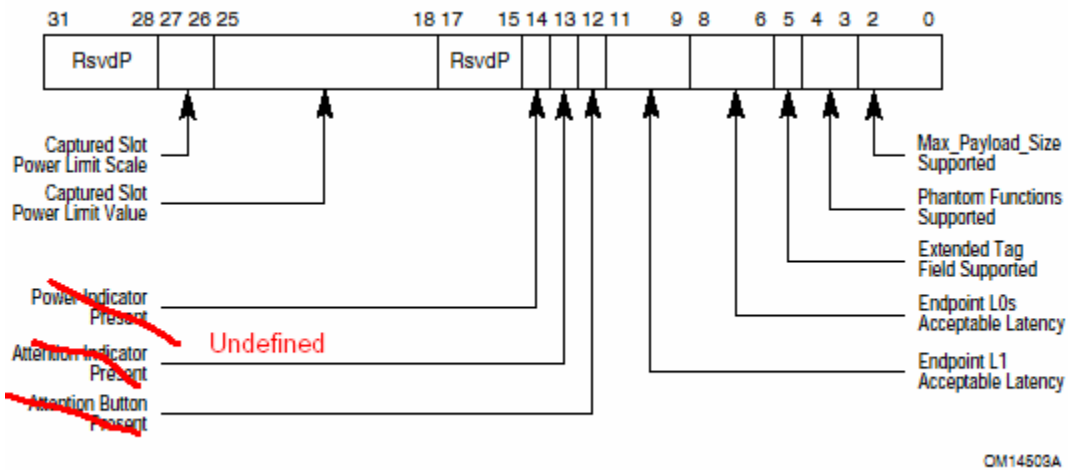


Figure 7-13: Device Capabilities Register

Table 7-11: Device Capabilities Register

Bit Location	Register Description	Attributes
...		
12	<p>The value read from this bit is undefined. In previous versions of this specification, this bit was used to indicate that an Attention Button is implemented on the adapter and electrically controlled by the component on the adapter. System software must ignore the value read from this bit. System software is permitted to write any value to this bit. Attention Button Present—When set to 1b, this bit indicates that an Attention Button is implemented on the adapter and is electrically controlled by the component on the adapter. Attention Button press events are reported using the Attention_Button_Pressed Message.</p> <p>This bit is valid for the following PCI Express device Types:</p> <ul style="list-style-type: none"> • PCI Express Endpoint device • Legacy PCI Express Endpoint device • Upstream Port of PCI Express Switch • PCI Express to PCI/PCI-X Bridge 	RO

Bit Location	Register Description	Attributes
13	<p>The value read from this bit is undefined. In previous versions of this specification, this bit was used to indicate that an Attention Indicator is implemented on the adapter and electrically controlled by the component on the adapter. System software must ignore the value read from this bit. System software is permitted to write any value to this bit. Attention Indicator Present—When set to 1b, this bit indicates that an Attention Indicator is implemented on the adapter and is electrically controlled by the component on the adapter using the Attention_Indicator_On, Attention_Indicator_Blink, and Attention_Indicator_Off Messages.</p> <p>This bit is valid for the following PCI Express device Types:</p> <ul style="list-style-type: none"> • PCI Express Endpoint device • Legacy PCI Express Endpoint device • Upstream Port of PCI Express Switch • PCI Express to PCI/PCI-X Bridge 	RO
14	<p>The value read from this bit is undefined. In previous versions of this specification, this bit was used to indicate that a Power Indicator is implemented on the adapter and electrically controlled by the component on the adapter. System software must ignore the value read from this bit. System software is permitted to write any value to this bit. Power Indicator Present—When set to 1b, this bit indicates that a Power Indicator is implemented on the adapter and is electrically controlled by the component on the adapter using the Power_Indicator_On, Power_Indicator_Blink, and Power_Indicator_Off Messages.</p> <p>This bit is valid for the following PCI Express device Types:</p> <ul style="list-style-type: none"> • PCI Express Endpoint device • Legacy PCI Express Endpoint device • Upstream Port of PCI Express Switch • PCI Express to PCI/PCI-X Bridge 	RO
...		

...

7.8.6. Link Capabilities Register (Offset 0Ch)

...

Table 7-14: Link Capabilities Register

Bit Location	Register Description	Attributes
...	...	

Bit Location	Register Description	Attributes
20	<p>Data Link Layer Link Active Reporting Capable – For a Downstream Port, this bit must be set to 1b if the component supports the optional capability of reporting the DL_Active state of the Data Link Control and Management State Machine. For a hot-plug capable downstream port (as indicated by Hot-Plug Capable field of Slot Capabilities Register), this bit must be set to 1b.</p> <p>For Upstream Ports and components that do not support this optional capability, this bit must be hardwired to 0b. <i>[Editorial Note: This change builds on Surprise down ECN]</i></p>	RO
...	...	

If L1 state is not supported for ASPM (as reported in the ASPM Support field), then the L1 Exit latency field is ignored.

7.8.7. Link Control Register (Offset 10h)

...

7.8.8. Link Status Register (Offset 12h)

...

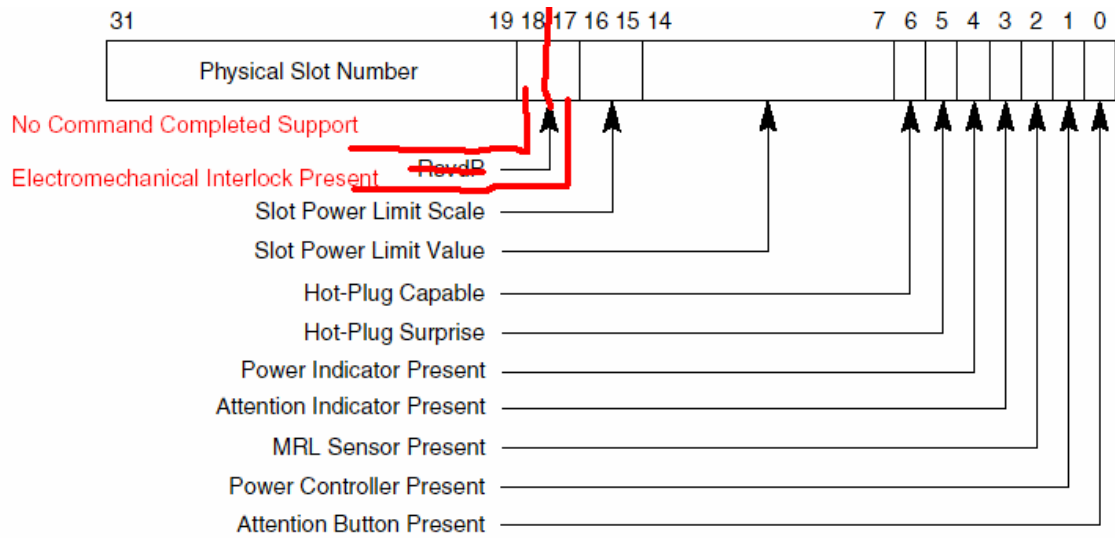
Table 7-16: Link Status Register

Bit Location	Register Description	Attributes
...	...	
13	<p>Data Link Layer Link Active (Optional) – Set to 1b when This bit indicates the status of the Data Link Control and Management State Machine. It returns a 1b to indicate is in the DL_Active state, 0b otherwise. If this optional status reporting capability is not implemented This bit must be implemented if the corresponding Data Link Layer Active Capability bit is implemented. Otherwise, this bit must be hardwired to 0b. <i>[Editorial Note: This change builds on Surprise down ECN]</i></p>	RO

...

7.8.9. Slot Capabilities Register (Offset 14h)

The Slot Capabilities register identifies PCI Express slot specific capabilities. Figure 7-19 details allocation of register fields in the Slot Capabilities register; Table 7-17 provides the respective bit definitions.



OM14509A

Figure 7-19: Slot Capabilities Register

Table 7-17: Slot Capabilities Register

Bit Location	Register Description	Attributes
...		
17	Electromechanical Interlock Present – When set to 1b, this bit indicates that an Electromechanical Interlock is implemented on the chassis for this slot.	Hwinit
18	No Command Completed Support – When set to 1b, this bit indicates that this slot does not generate software notification when an issued command is completed by the Hot-Plug Controller. This bit is only permitted to be set to 1b if the hot-plug capable port is able to accept writes to all fields of the Slot Control register without delay between successive writes.	Hwinit
...		

7.8.10. Slot Control Register (Offset 18h)

...

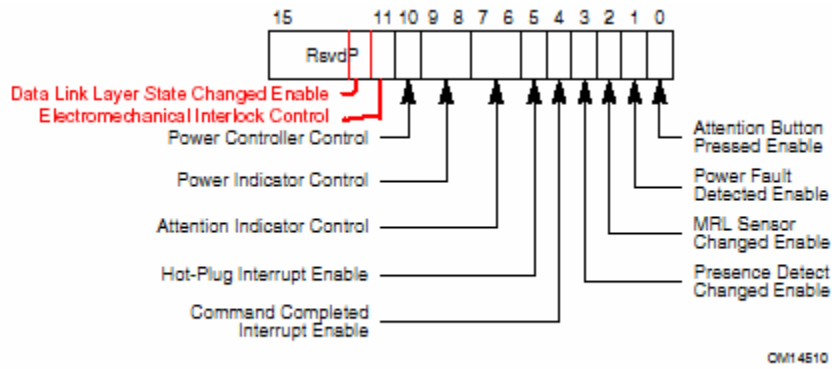


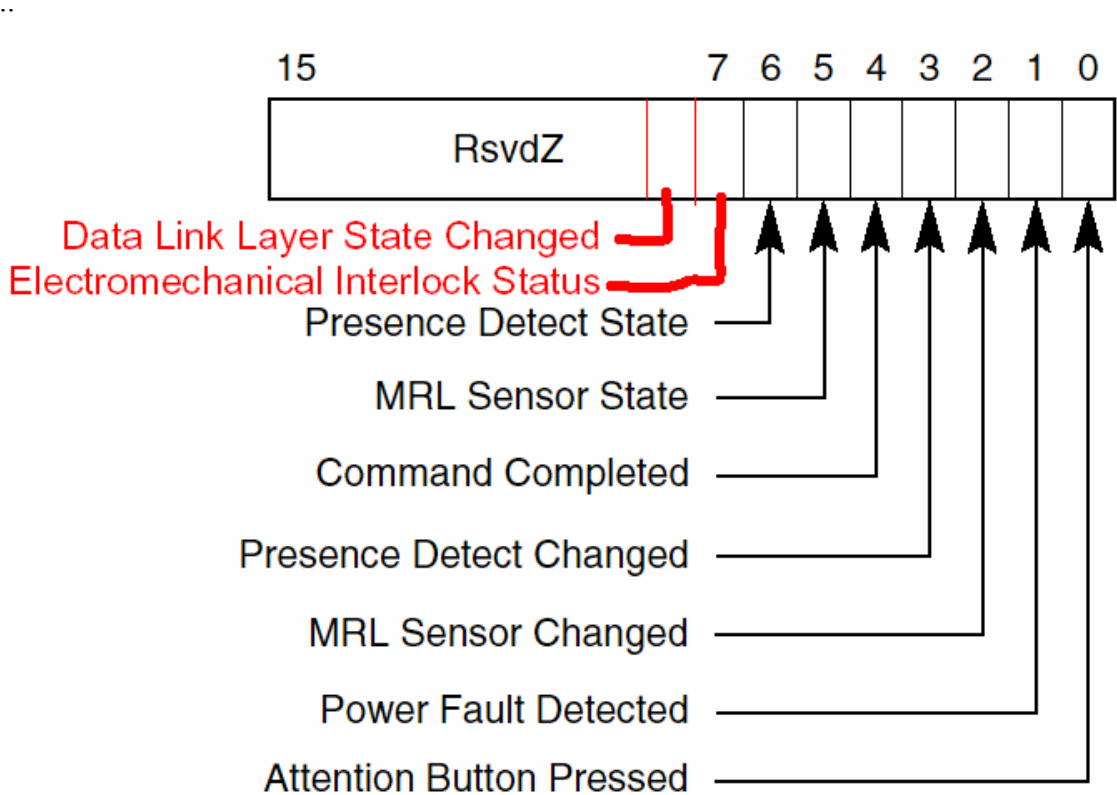
Figure 7-20: Slot Control Register

Table 7-18: Slot Control Register

Bit Location	Register Description	Attributes
0	Attention Button Pressed Enable – When set to 1b, this bit enables software notification on an attention button pressed event. See Section 6.7.7. Software may set this bit even if the Attention Button Supported bit in the Slot Capabilities register is 0b, so that Attention_Button_Pressed messages can cause software notification. Default value of this field is 0b.	RW
...		
4	Command Completed Interrupt Enable – If Command Completed notification is supported (as indicated by No Command Completed Support field of Slot Capabilities Register), When set to 1b, this bit enables software notification when a hot-plug command is completed by the Hot-Plug Controller. Default value of this field is 0b. If Command Completed notification is not supported this bit must be hardwired to 0b.	RW
...		

Bit Location	Register Description	Attributes								
7:6	<p>Attention Indicator Control – If an Attention Indicator is implemented, writes to this field set the Attention Indicator to the written state. Reads of this field must reflect the value from the latest write, even if the corresponding hot-plug command is not complete, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined.</p> <p>If the indicator is electrically controlled by chassis, the indicator is controlled directly by the downstream port through implementation specific mechanisms. If the indicator is electrically controlled by the adapter, the indicator is indirectly controlled by transmission of the appropriate Attention Indicator Message (determined by the value written). Regardless of the location of the indicator, the downstream port must transmit the appropriate Attention Indicator Message.</p> <p>Defined encodings are:</p> <table> <tr> <td>00b</td> <td>Reserved</td> </tr> <tr> <td>01b</td> <td>On</td> </tr> <tr> <td>10b</td> <td>Blink</td> </tr> <tr> <td>11b</td> <td>Off</td> </tr> </table> <p>Note: The default value of this field must be one of the non-Reserved values.</p>	00b	Reserved	01b	On	10b	Blink	11b	Off	RW
00b	Reserved									
01b	On									
10b	Blink									
11b	Off									
9:8	<p>Power Indicator Control – If a Power Indicator is implemented, writes to this field set the Power Indicator to the written state. Reads of this field must reflect the value from the latest write, even if the corresponding hot-plug command is not complete, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined.</p> <p>If the indicator is electrically controlled by chassis, the indicator is controlled directly by the downstream port through implementation specific mechanisms. If the indicator is electrically controlled by the adapter, the indicator is indirectly controlled by transmission of the appropriate Power Indicator Message (determined by the value written). Regardless of the location of the indicator, the downstream port must transmit the appropriate Power Indicator Message.</p> <p>Defined encodings are:</p> <table> <tr> <td>00b</td> <td>Reserved</td> </tr> <tr> <td>01b</td> <td>On</td> </tr> <tr> <td>10b</td> <td>Blink</td> </tr> <tr> <td>11b</td> <td>Off</td> </tr> </table> <p>Note: The default value of this field must be one of the non-Reserved values.</p>	00b	Reserved	01b	On	10b	Blink	11b	Off	RW
00b	Reserved									
01b	On									
10b	Blink									
11b	Off									
...										
<u>11</u>	<p>Electromechanical Interlock Control – If an <u>Electromechanical Interlock</u> is implemented, a write of 1b to this field causes the state of the interlock to toggle. A write of 0b to this field has no effect. A read to this register always returns a 0.</p>	<u>RW</u>								
<u>12</u>	<p>Data Link Layer State Changed Enable – If the <u>Data Link Layer Link Active</u> capability is implemented, when set to 1b, this field enables software notification when <u>Data Link Layer Link Active</u> field is changed</p>	<u>RW</u>								

7.8.11. Slot Status Register (Offset 1Ah)



OM14511

Figure 7-21: Slot Status Register

Table 7-19: Slot Status Register

Bit Location	Register Description	Attributes
0	Attention Button Pressed – If an Attention Button is implemented, this bit is set when the attention button is pressed (either as directly detected or by the receipt of an Attention_Button_Pressed Message). If an Attention Button is not supported and the form factor does not support the Attention_Button_Pressed Message , this bit must not be set.	RW1C
...		

Bit Location	Register Description	Attributes
4	<p>Command Completed – If Command Completed notification is supported (as indicated by No Command Completed Support field of Slot Capabilities Register), this bit is set when a hot-plug command has completed and the Hot-Plug Controller is ready to accept a subsequent command. The Command Completed status bit is set as an indication to host software that the Hot-Plug Controller has processed the previous command and is ready to receive the next command; it provides no guarantee that the action corresponding to the command is complete.</p> <p>If Command Completed notification is not supported, this bit must be hardwired to 0b.</p>	RW1C
...		
7	<p>Electromechanical Interlock Status – If an Electromechanical Interlock is implemented, this bit indicates the current status of the Electromechanical Interlock.</p> <p>Defined encodings are:</p> <p>0b Electromechanical Interlock Disengaged</p> <p>1b Electromechanical Interlock Engaged</p>	RO
8	<p>Data Link Layer State Changed - This bit is set when the value reported in Data Link Layer Link Active field of the Link Status register is changed</p> <p>In response to a Data Link Layer State Changed event, software must read Data Link Layer Link Active field of the Link Status register to determine if the link is active before initiating configuration cycles to the hot plugged device.</p>	RW1C
...		