PCI-SIG ENGINEERING CHANGE NOTICE

TITLE:	Emergency Power Reduction mechanism with PWRBRK signal
DATE:	October 30, 2015
AFFECTED DOCUMENT:	PCI Express Card Electromechanical Specification, Revision 3.0 PCI Express Base Specification, Revision 3.1 Single Root I/O Virtualization and Sharing Specification Revision 1.1
SPONSOR:	NVIDIA

Part I

1. Summary of the Functional Changes

This ECR defines two sets of related changes to support an Emergency Power Reduction mechanism and to provide software visibility for this mechanism:

1. The Card Electromechanical Specification is updated to define an optional Emergency Power Reduction mechanism using RSVD pin B30.
2. The PCI Express Base Specification is updated to define an optional mechanism to indicate support for Emergency Power Reduction and to provide visibility as to the power reduction status of a Device.

2. Benefits as a Result of the Changes

This will standardize the pin location for the optional Emergency Power Reduction mechanism. For systems and cards that choose to leave B30 unconnected, there will be no impact to behavior.

This will provide a mechanism for software to detect support and to determine that the mechanism is operating correctly.

3. Assessment of the Impact

This helps to mitigate current conflicts between implementations that use this pin B30 for the optional Emergency Power Reduction mechanism and any other proprietary usage.

This provides an optional mechanism for software that is interested to detect support and to verify correct operation. No software changes are required, even for add-in cards that support the mechanism.

4. Analysis of the Hardware Implications

B30 will no longer be undefined RSVD pin, it will now be specified as PWRBRK pin.

Cards and Systems not implementing this specific Emergency Power Reduction mechanism must leave the signal disconnected and floating.

5. Analysis of the Software Implications

No software changes are required to use the Emergency Power Reduction mechanism.

System software can optionally use the new capability bits to determine support and to confirm correct operation of the mechanism.

6. Analysis of the Compatibility and Interoperability Test Implications

Some previously reserved fields are now defined. New test cases would need to be created to test functionality of this feature

Part II

Detailed Description of the change

In the PCI Express Card Electromechanical Specification, change Section 1.5 Electrical Overview, page 15 as follows:

- JTAG, optional
- SMBus, optional
- Wake (WAKE#), required only if the device/system supports wakeup and/or the OBFF mechanism
- +3.3Vaux, optional
- Power Brake (PWRBRK#), optional

REFCLK, JTAG, SMBus, PERST#, ~~and~~ WAKE#, and PWRBRK# are described in Chapter 2; +3.3Vaux is described in Chapter 4; and PRSNT1# and PRSNT2# are described in Chapter 3.

In the PCI Express Card Electromechanical Specification, change Section 2.0 Auxiliary Signals, page 21 as follows:

- PRSNT2# (required): add-in card presence detect pin. See Chapter 3 for a detailed description.
- PWRBRK# (optional): an open-drain, active low signal that is driven low by the system to signal an Emergency Power Reduction mechanism.

Note that the SMBus interface pins are collectively optional for both the add-in card and the system board. If the optional management features are implemented, SMBCLK and SMBDAT are both required. Similarly, the JTAG pins are collectively optional. If this test mode is implemented, all the JTAG pins are required. For additional system requirements related to these signals refer to the *PCI Local Bus Specification, Revision. 3.0*, Section 4.3.3.

In the PCI Express Card Electromechanical Specification, add Section 2.6 Power Brake Pin (Optional):

2.6 PWRBRK# Signal (Optional)

The PWRBRK# signal is an optional normative capability applicable to the CEM form factor. Only add-in cards that support Emergency Power Reduction connect to this pin. Likewise, only systems that support the Emergency Power Reduction mechanism connect to this pin. The assertion and de-assertion of PWRBRK# are asynchronous to any system clock. An add-in card that supports Emergency Power Reduction must provide a weak pull-up on PWRBRK# (minimum 95KΩ). A system that supports Emergency Power Reduction must provide a stronger pull-up on PWRBRK#. These pull-up resistor values must ensure meeting the rise time specification of T_{PWRBRK}.

The PWRBRK# signal is used to communicate requests to enter and exit the Emergency Power Reduction State. See PCI Express Base Specification, Section 6.x.

PWRBRK# is an open-drain, active low signal that is driven low by an external enclosure component. When asserted, add-in cards that support Emergency Power Reduction must quickly reduce their power consumption. When de-asserted, add-in cards that support Emergency Power Reduction are permitted to resume normal power consumption. The add-in card should debounce this input to avoid thrashing the system with rapid transitions in and out of the Emergency Power Reduction State.

This mechanism is an emergency fail-safe intended to be used to prevent system damage and is not intended to provide normal dynamic power management. The external enclosure should control how it asserts/deasserts this signal to avoid thrashing the system with rapid transitions in and out of the Emergency Power Reduction State. The amount of power consumed in the Emergency Power Reduction state is communicated through the Power Budgeting extended capability as defined in the PCI Express Base Specification (see Section 7.15). The time allowed to achieve this power reduction ($T_{PWRBRK\text{-}AIC\text{-}ENTER\text{-}LP\text{-}MODE}$) is defined in Table 2-4.

Add-in cards that support Emergency Power Reduction may contain any mix of Devices that advertise support for PWRBRK# or that don't advertise such support. Software is not required to configure or enable use of PWRBRK#. Software can optionally use the mechanisms defined in the PCI Express Base Specification to determine support for PWRBRK# and, if supported, to determine the associated power saving. Software can detect that a supporting Function has entered the Emergency Power Reduction State.

Electrical specifications for PWRBRK# at the CEM connector are defined in Table 2-3. Timing requirements are defined in Table 2-4 and Figure 2-15.

PWRBRK# may be bused to multiple PCI Express add-in card connectors, forming a single output connection at the external enclosure component, or individual connectors can have individual connections to the external enclosure component.

In the PCI Express Card Electromechanical Specification, change Table 2-3 in Section 2.6.1, page 39 as follows:

Table 2-3: Auxiliary Signal DC Specifications - PERST#, WAKE#, ~~and~~ SMBus, and PWRBRK#

Symbol	Parameter	Conditions	Min	Max	Unit	Notes
V_{IL1}	Input Low Voltage		-0.5	0.8	V	2, 6, 7
V_{IH1}	Input High Voltage		2.0	V_{cc3_3} + 0.5	V	2, 6, 7
V_{IL2}	Input Low Voltage		-0.5	0.8	V	4
V_{IH2}	Input High Voltage		2.1	V_{ccSus3_3} + 0.5	V	4
V_{OL1}	Output Low Voltage	4.0 mA		0.2	V	1, 3
V_{HMAX}	Max High Voltage			V_{cc3_3} + 0.5	V	3

Symbol	Parameter	Conditions	Min	Max	Unit	Notes
V_{OL2}	Output Low Voltage	4.0 mA		0.4	V	1, 4
I_{in}	Input Leakage Current	0 to 3.3 V	-10	+10	µA	2, 4, 7
I_{lkg}	Output Leakage Current	0 to 3.3 V	-50	+50	µA	3, 5
C_{in}	Input Pin Capacitance			7	pF	2, 7
C_{out}	Output (I/O) Pin Capacitance			30	pF	3, 4

Notes:

1. Open-drain output a pull-up is required on the system board. There is no V_{OH} specification for these signals. The number given is the maximum voltage that can be applied to this pin.
2. Applies to PERST#.
3. Applies to WAKE#.
4. Applies to SMBus signals SMBDATA and SMBCLK.
5. Leakage at the pin when the output is not active (high impedance).
6. Applies to WAKE# issued by Switch Downstream Ports and Root Complex for signaling of OBFF indications as received at the input of the Endpoint(s).
7. Applies to PWRBRK#.

In the PCI Express Card Electromechanical Specification, change Table 2-4 in Section 2.6.2, page 40 as follows:

Table 2-4: Power Sequencing and Reset Signal Timings

Symbol	Parameter	Min	Max	Units	Notes	Figure
T_{PVPERL}	Power stable to PERST# inactive	100		ms	1	Figure 2-10
$T_{PERST-CLK}$	REFCLK stable before PERST# inactive	100		µs	2	Figure 2-10
T_{PERST}	PERST# active time	100		µs		Figure 2-11
T_{FAIL}	Power level invalid to PERST# active		500	ns	3	Figure 2-13
T_{WKRF}	WAKE# rise – fall time		100	ns	4	Figure 2-14
$T_{WAKE-TX-MIN-PULSE}$	Minimum WAKE# pulse width; applies to both active-inactive-active and inactive-active-inactive cases	300		ns	5	
$T_{WAKE-FALL-FALL-CPU-ACTIVE}$	Time between two falling WAKE# edges when signaling CPU Active	700	1000	ns	5	

Symbol	Parameter	Min	Max	Units	Notes	Figure
T_{PWRBRK}	PWRBRK# rise – fall time		100	ns	6	
$T_{PWRBRK-FALL-RISE-ACTIVE}$	Time PWRBRK# is active	1		ms	6	
$T_{PWRBRK-RISE-FALL-INACTIVE}$	Time PWRBRK# is inactive	1		ms	6	
$T_{PWRBRK-AIC-ENTER-LP-MODE}$	Time for add-in card to enter low power mode		10	µs		

Notes:

1. Any supplied power is stable when it meets the requirements specified for that power supply.
2. A supplied reference clock is stable when it meets the requirements specified for the reference clock. The PERST# signal is asserted and de-asserted asynchronously with respect to the supplied reference clock.
3. The PERST# signal must be asserted within T_{FAIL} of any supplied power going out of specification.
4. Measured from WAKE# assertion/de-assertion to valid input level at the system PM controller. Since WAKE# is an open-drain signal, the rise time is dependent on the total capacitance on the platform and the system board pull-up resistor. It is the responsibility of the system designer to meet the rise time specification.
5. Refers to timing requirement for indicating an active window.
6. Refers to PWRBRK# timing diagram in Figure 2-15.

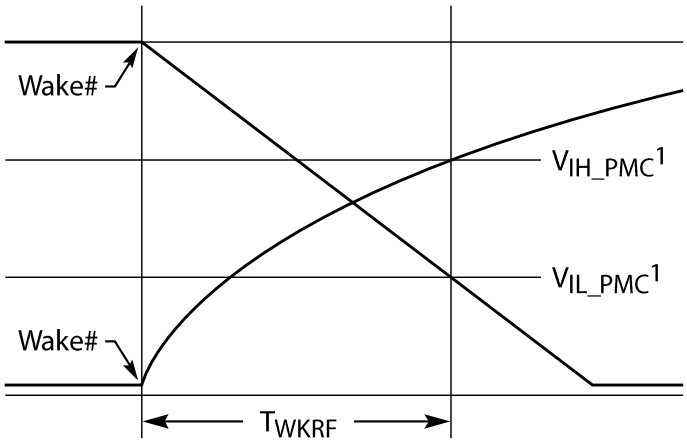

Note 1: Power Management Controller input switching levels are platform dependent and are not set by this specification.

A-0338

Figure 2-14: WAKE# Rise and Fall Time Measurement Points

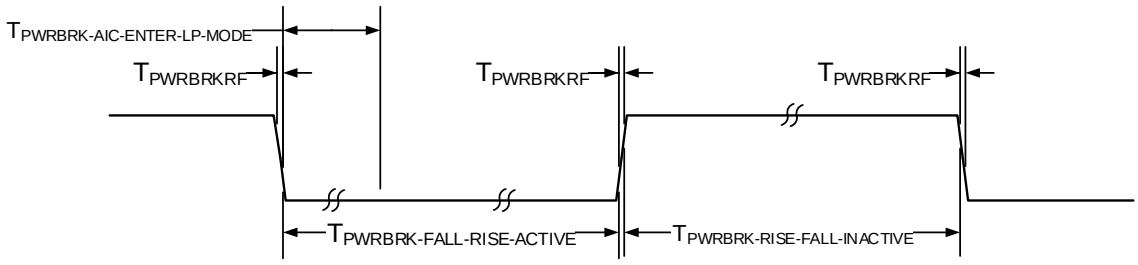

Figure 2-15: PWRBRK# Timing Requirement Diagram

In the PCI Express Card Electromechanical Specification, change Table 6-1 in Section 6.1 Connector Pinout, page 87 as follows:

19	PETp1	Transmitter differential pair, Lane 1	RSVD	
20	PETn1		GND	Ground
21	GND	Ground	PERp1	Receiver differential pair, Lane 1
22	GND	Ground	PERn1	
23	PETp2	Transmitter differential pair, Lane 2	GND	Ground
24	PETn2		GND	Ground
25	GND	Ground	PERp2	Receiver differential pair, Lane 2
26	GND	Ground	PERn2	
27	PETp3	Transmitter differential pair, Lane 3	GND	Ground
28	PETn3		GND	Ground
29	GND	Ground	PERp3	Receiver differential pair, Lane 3
30	PWRBRK	Emergency Power Reduction	PERn3	
31	PRSNT2#	Hot-Plug presence detect	GND	Ground
32	GND	Ground	RSVD	Reserved
End of the x4 connector				

In the PCI Express Base Specification, section 6.9, make the following changes:

PCI Express provides a mechanism for software controlled limiting of the maximum power per slot that a PCI Express adapter (associated with that slot) can consume. If supported, the Emergency Power Reduction State, over-rides the mechanisms listed here (see Section 6.x).

In the PCI Express Base Specification, section 6.16, make the following changes:

The DPA Capability provides a mechanism to allocate power dynamically for these types of devices. DPA is optional normative functionality applicable to Endpoint Functions that can benefit from the dynamic allocation of power and do not have an alternative mechanism. If supported, the Emergency Power Reduction State, over-rides the mechanisms listed here (see Section 6.x).

In the PCI Express Base Specification, add a new section to Chapter 6:

6.x Emergency Power Reduction State

Emergency Power Reduction State is an optional mechanism to request that Functions quickly reduce their power consumption. Emergency Power Reduction is a fail-safe mechanism intended to be used to prevent system damage and is not intended to provide normal dynamic power management.

If a Function implements Emergency Power Reduction State, it must also implement the Power Budgeting extended capability and must report Power Budgeting values for this state (see Section 7.15). Devices that are integrated on the system board are not required to implement the Power Budgeting extended capability, but if they do so, they must meet the preceding requirement.

Note to Reviewers:

This note is not part of the ECN.

For reference, Section 4.3 of the *PCI Express Card Electromechanical Specification, Revision 3.0* reads as follows:

4.3. Power Budgeting Capability

The Power Budget Capabilities register-set, as defined in the *PCI Express Base Specification, Revision 3.0*, shall be implemented for (1) cards capable of using more power than initially allowed at power-up (see Section 4.2), and (2) cards utilizing auxiliary power connections (150 W, 225 W, 300 W, etc.).

Sustained Thermal and Maximum Thermal values shall include all thermal power that is produced by the card.

Populated values shall include all power used by the card including power drawn from auxiliary power connections.

For multi-device cards, an instance of the Power Budget Capabilities register-set may report power at a device or card level. Thus for multi-device cards,

such as a switch with devices behind it, system software must aggregate all instances of the Power Budget Capabilities register-sets implemented at or beneath the base device of the card.

Functions enter and exit this state based on either autonomously or via external requests. External requests may be either following a signaling protocol defined in an applicable form factor specification, or by a vendor-specific method. Table 6-x defines how the Emergency Power Reduction Supported and Emergency Power Reduction Initialization Required fields determine the mechanisms that are allowed to trigger entry and exit from this state (see Section 7.8.15).

Table 6-x Emergency Power Reduction Supported values

Emergency Power Reduction Supported	Emergency Power Reduction Initialization Required	Entry / Exit permitted by		
		Form Factor Mechanism	Vendor Specific Mechanism(s)	Autonomous Mechanisms
00b	0	No	Yes	Yes
	1	No	No	No
01b	Any	No	Yes	Yes
10b	Any	Yes	Yes	Yes
11b	Reserved			

Functions may indicate that they require re-initialization on exit from this state:

- ❑ If the Emergency Power Reduction Initialization Required bit is Clear (see Section 7.8.15):
  - On entry to this state, the Function either operates normally (perhaps with reduced performance), or enters a device specific "power reduction dormant state". The Upstream Port of the Device remains operating. Outstanding requests initiated by or directed to the Function must complete normally.
  - On exit from this state, the Function operates normally (perhaps resuming normal performance). Functions that entered a "power reduction dormant state" exit that state. In either case, no software intervention is required.
- ❑ If the Emergency Power Reduction Initialization Required bit is Set (see Section 7.8.15):
  - On entry to this state, the Function ceases normal operation. The Upstream Port of the associated Device is permitted to enter DL_Down.
    - ♦ If the Upstream Port remains in DL_Up, outstanding requests directed to or initiated by the Function must complete normally.
    - ♦ If the Upstream Port enters DL_Down, outstanding request behavior is defined in Section 2.9.1. This transition may result in a Surprise Down error.
    - ♦ Sticky bits must be preserved in this state.
  - On exit from this state, software intervention is required to resume normal operation. The mechanism used to indicate to software when this is required is outside the scope of this specification (e.g., a device specific interrupt). If the Upstream Port entered

DL_Down, all Functions of the Device are reset and a full reconfiguration is required (see Section 2.9.2).

The following rules apply to the Emergency Power Reduction State:

- A Device supports Emergency Power Reduction State if at least one Function in the Upstream Port indicates support (i.e., Emergency Power Reduction Supported is non-zero).
- Emergency Power Reduction State is associated with a Device. All Functions in a Device that support it enter and exit this state at the same time.
- Functions where the Emergency Power Reduction Supported field is 00b are not affected by the Emergency Power Reduction State of the Device as long as the Upstream Port remains in DL_Up. The Emergency Power Reduction Detected bit is RsvdZ.
- Functions where the Emergency Power Reduction Supported field is 01b or 10b:
  - Set the Emergency Power Reduction Detected bit when the Device enters Emergency Power Reduction State.
  - Clear the Emergency Power Reduction Detected bit when requested if the Device has exited the Emergency Power Reduction State.
- For Switches, Downstream Switch Ports enter and exit Emergency Power Reduction State at the same time as the associated Upstream Switch Port. The corresponding fields in Configuration Space are reserved for Downstream Switch Ports.
- For SR-IOV Devices, VFs enter and exit Emergency Power Reduction State at the same time as their PF. The corresponding fields in Configuration Space are reserved for VFs.
- Encoding 10b shall not be used unless the associated form factor specification defines a mechanism for requesting Emergency Power Reduction.
- It is strongly recommended that the Emergency Power Reduction Supported field be initialized by hardware or firmware within the Function prior to initial device enumeration. This initialization is permitted to be deferred to device driver load when this is not practical (e.g., when there is no firmware ROM).

IMPLEMENTATION NOTE

Diagnostic Checking of Emergency Power Reduction Detected

The Emergency Power Reduction Detected bit permits system software to detect that Emergency Power Reduction State was entered, even momentarily. The Emergency Power Reduction Request bit can be used by software to request entry. Normally, software would use a system specific method to enter the Emergency Power Reduction State using external mechanisms.

IMPLEMENTATION NOTE

Emergency Power Reduction State: Example Add-in Card

Figure 6-x shows an example multi-Device add-in card supporting Emergency Power Reduction. Note that Device C does not support the Emergency Power Reduction State. Device C might be a Switch that fans out to Devices A and B.

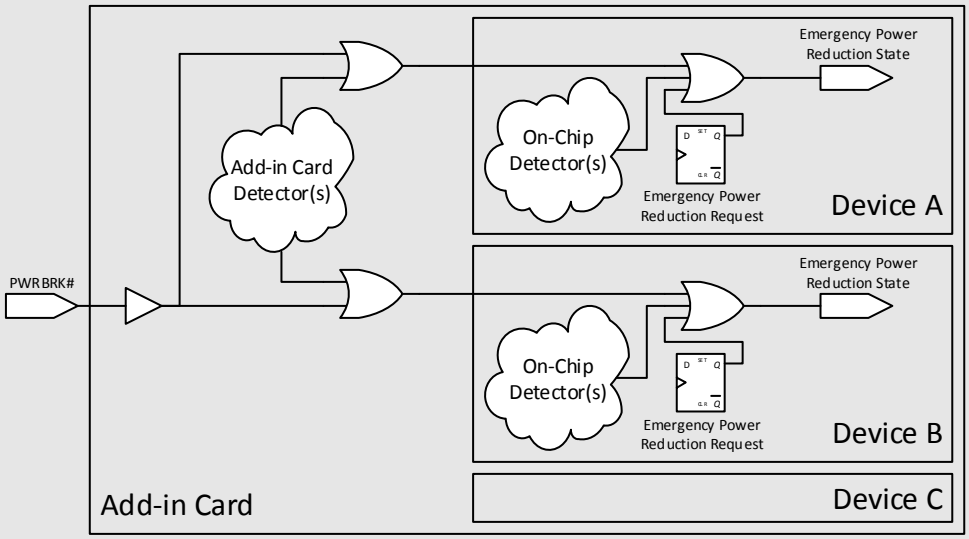

Figure 6-x Emergency Power Reduction State: Example add-in card

In the PCI Express Base Specification, Section 7.8.5, make the following changes:

7.8.5 Device Status Register (Offset 0Ah)

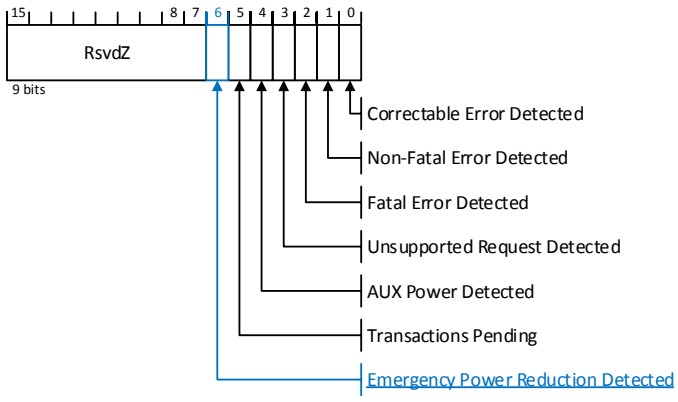

Figure 7-15: Device Status Register

Table 7-15: Device Status Register

Bit Location	Register Description	Attributes
6	**Emergency Power Reduction Detected** – This bit is Set when the Function is in the Emergency Power Reduction State. Whenever any condition is present that would cause the Emergency Power Reduction State to be entered, the Function remains in the Emergency Power Reduction State and writes to this bit have no effect. See Section 6.x for additional details. Multi-Function Devices associated with an Upstream Port must Set this bit in all Functions that support Emergency Power Reduction State. This bit is RsvdZ if the Emergency Power Reduction Supported field is 00b (see Section 7.8.15). This bit is RsvdZ in Functions that are not associated with an Upstream Port. Default value is 0b.	RW1C

In the PCI Express Base Specification, section 7.8.15, make the following changes:

7.8.15 Device Capabilities 2 Register (Offset 24h)

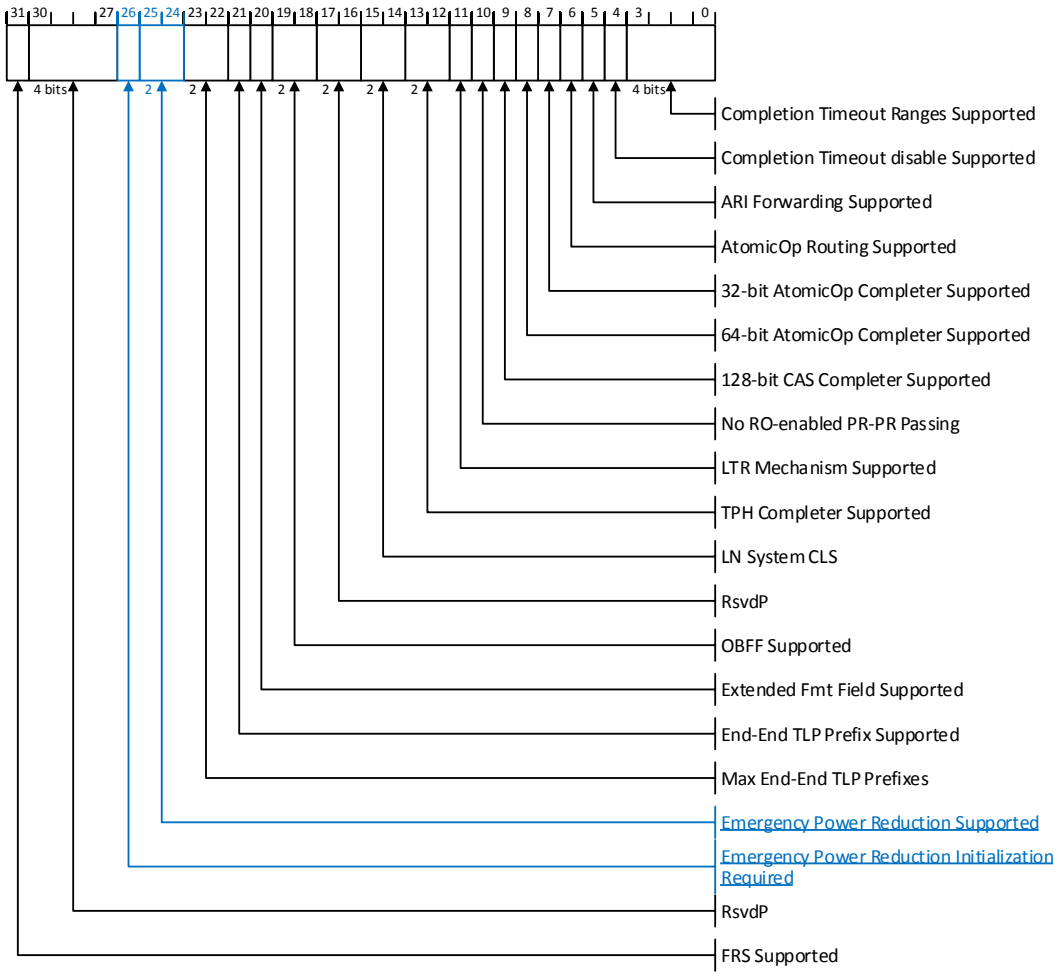

Figure 7-25: Device Capabilities 2 Register

Table 7-25: Device Capabilities 2 Register

Bit Location	Register Description	Attributes
25:24	**Emergency Power Reduction Supported** – Indicates support level of the optional Emergency Power Reduction State feature. A Function can enter Emergency Power Reduction State autonomously, or based on one of two mechanisms defined by the associated Form Factor Specification. Functions that are in the Emergency Power Reduction State consume less power. The Emergency Power Reduction mechanism permits a chassis to request add-in cards to rapidly enter Emergency Power Reduction State without involving system software. See Section 6.x for additional details. Values are: 00b Emergency Power Reduction State not supported 01b Emergency Power Reduction State is supported and is triggered by Device Specific mechanism(s) 10b Emergency Power Reduction State is supported and is triggered either by the mechanism defined in the corresponding Form Factor specification or by Device Specific mechanism(s) 11b Reserved This field is RsvdP in Functions that are not associated with an Upstream Port. For Multi-Function Devices associated with an Upstream Port, all Functions that report a non-zero value for this field, must report the same non-zero value for this field. Default value is 00b. After reset, once this field returns a non-zero value, it must continue to return the same non-zero value, until the next reset.	HwInit
26	**Emergency Power Reduction Initialization Required** – If Set, the Function requires complete or partial initialization upon exit from the Emergency Power Reduction State. If Clear, the Function requires no software intervention to return to normal operation upon exit from the Emergency Power Reduction State. See Section 6.x for additional details. For Multi-Function Devices associated with an Upstream Port, all Functions must report the same value for this bit. This bit is RsvdP in Functions that are not associated with an Upstream Port. Default value is 0b. After reset, when this field returns a non-zero value, it must continue to return the same non-zero value.	HwInit

In the PCI Express Base Specification, section 7.8.16, make the following changes:

7.8.16 Device Control 2 Register (Offset 28h)

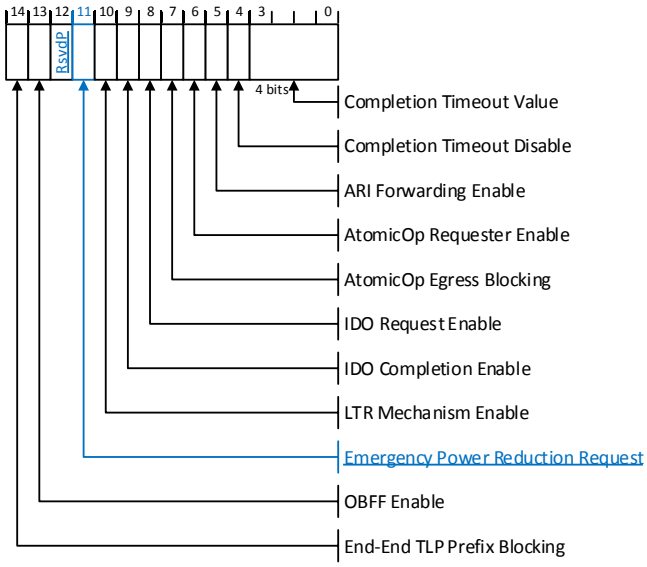

Figure 7-26: Device Control 2 Register

Table 7-26: Device Control 2 Register

Bit Location	Register Description	Attributes
11	**Emergency Power Reduction Request** – If Set, all Functions in the component that support Emergency Power Reduction State must enter the Emergency Power Reduction State. If Clear these Functions must exit the Emergency Power Reduction State if no other reasons exist to preclude exiting this state. See Section 6.x for additional details. This bit is implemented in the lowest numbered Function associated with an Upstream Port that has a non-zero value in the Emergency Power Reduction Supported field. This bit is RsvdP in all other Functions. Default is 0b.	RW/RsvdP (see description)

In the PCI Express Base Specification, section 7.15.3, make the following changes:

7.15.3. Data Register (Offset 08h)

…

Table 7-64: Power Budgeting Data Register

Bit Location	Register Description	Attributes
17:15	**Type** – Specifies the type of the operating condition being described. Defined encodings are: 000b PME Aux 001b Auxiliary 010b Idle 011b Sustained 100b Sustained – Emergency Power Reduction State (see Section 6.x) 101b Maximum – Emergency Power Reduction State (see Section 6.x) 111b Maximum All other encodings are Reserved.	RO

…

A device that implements the Power Budgeting Capability is required to provide data values for the D0 Maximum and D0 Sustained PM State and Type combinations for every power rail from which it consumes power; data for the D0 Maximum and D0 Sustained for Thermal must also be provided if these values are different from the sum of the values for an operating condition reported for D0 Maximum and D0 Sustained on the power rails.

Devices that support auxiliary power or PME from auxiliary power must provide data for the appropriate power type (Auxiliary or PME Aux).

If a device implements Emergency Power Reduction State, it must report Power Budgeting values for the following:

- Maximum Emergency Power Reduction State, PM State D0, all power rails used by the device
- Maximum Emergency Power Reduction State, PM State D0, Thermal (if different from the sum of the preceding values)
- Sustained Emergency Power Reduction State, PM State D0, all power rails used by the device
- Sustained Emergency Power Reduction State, PM State: D0, Thermal (if different from the sum of the preceding values)

In the Single Root I/O Virtualization and Sharing Specification, section 3.5.5, make the following changes:

3.5.5. Device Status Register (Offset 0Ah)

The Device Status register provides information about PCI Express device specific parameters. Figure 3-8 details allocation of register fields in the Device Status register; Table 3-16 provides the respective bit definitions.

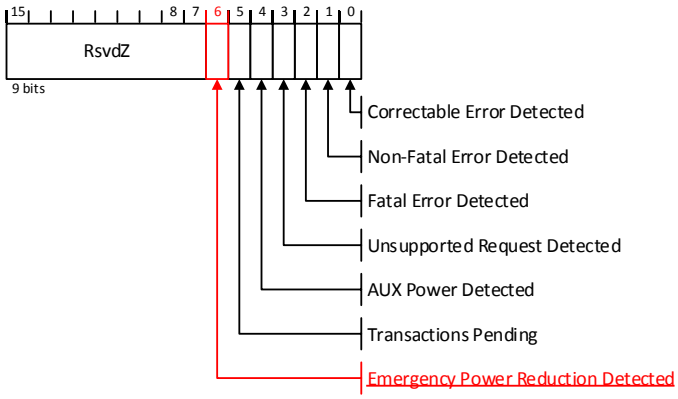

Figure 3-8: Device Status Register

PF and VF functionality is defined in the PCI Express Base Specification except where noted in Table 3-16.

Table 3-16: Device Status Register

Bit Location	PF and VF Register Differences From Base	PF Attributes	VF Attributes
4	AUX Power Detected	Base	0b
6	Emergency Power Reduction Detected	Base	0b

In the Single Root I/O Virtualization and Sharing Specification, section 3.5.9, make the following changes:

3.5.9. Device Capabilities 2 Register (Offset 24h)

PF and VF functionality is defined in the PCI Express Base Specification except as noted in Table 3-18.

Table 3-18: Device Capabilities 2 Register

Bit Location	PF and VF Register Differences From Base	PF Attributes	VF Attributes
3:0	Completion Timeout Ranges Supported VF value must be identical to PF value.	Base	Base
4	Completion Timeout Disable Supported VF value must be identical to PF value.	Base	Base
6	AtomicOp Routing Supported Not applicable to Endpoints.	RsvdP	RsvdP
7	32-bit AtomicOp Completer Supported VF value must be identical to PF value.	Base	Base
8	64-bit AtomicOp Completer Supported VF value must be identical to PF value.	Base	Base
9	128-bit CAS Completer Supported VF value must be identical to PF value.	Base	Base
25:24	Emergency Power Reduction Supported VF value must be identical to PF value.	Base	Base
26	Emergency Power Reduction Initialization Required VF value must be identical to PF value.	Base	Base

In the Single Root I/O Virtualization and Sharing Specification, section 3.5.10, make the following changes:

3.5.10. Device Control 2 Register (Offset 28h)

PF and VF functionality is defined in the PCI Express Base Specification except as noted in Table 3-19.

Table 3-19: Device Control 2 Register

Bit Location	PF and VF Register Differences From Base	PF Attributes	VF Attributes
3:0	Completion Timeout Ranges Value The PF value applies to all associated VFs.	Base	RsvdP
4	Completion Timeout Disable The PF value applies to all associated VFs.	Base	RsvdP
6	AtomicOp Requester Enable The PF value applies to all associated VFs.	Base	RsvdP
8	IDO Request Enable The PF value applies to all associated VFs.	Base	RsvdP
9	IDO Completion Enable The PF value applies to all associated VFs.	Base	RsvdP
11	Emergency Power Reduction Request This bit is only present in one Function associated with an Upstream Port. That Function can never be a VF.	Base	RsvdP

In the Single Root I/O Virtualization and Sharing Specification, add a new section 6.5 as follows:

6.5 VF Emergency Power Reduction State

If the Emergency Power Reduction Supported field in a VF is non-zero, that VF enters and exits the Emergency Power Reduction State at the same time as the associated PF. Software can use the Emergency Power Reduction Detected bit in the PF to emulate the corresponding bit in the VF.