



PCI-SIG ENGINEERING CHANGE REQUEST

TITLE:	PCIe Device Labeling under Operating Systems
DATE:	03/28/2010
AFFECTED DOCUMENT:	PCI Firmware Specification v3.0
SPONSOR:	Mukund Khatri & Vijay Nijhawan, Dell Inc

Part I

1. Summary of the Functional Changes

This ECR proposes a mechanism (new extension to _DSM) to make the device names/labels under Operating Systems deterministic. Currently, there is no well defined mechanism to consistently associate platform specific device names and instances of a device type under operating system. As a result, instance labels for specific device types under various operating systems (ex: ethx label for networking device instance under Linux OS) do not always map into the platform designated device labels. Additionally, the instance labels can change based on the system configuration. For example, under Linux operating systems, the "eth0" label does not necessarily map to the first embedded networking device as designed in a given platform. Depending on the hardware bus topology, current configuration including the number and type of networking adapters installed, the eth0 label assignment could change in a given platform.

2. Benefits as a Result of the Changes

With this proposed mechanism, a) it provides a means to make the device labels under operating system deterministic and b) one can match the device identification/ representation under OS with the names/labels listed on the platform chassis. This allows customers to directly associate OS identification of devices with the devices listed in the actual product. This interface could also be used by CLI tools to uniquely configure and manage specific end-point devices.

3. Assessment of the Impact

Support for this _DSM extension is optional. When supported, platform BIOS would provide support for this _DSM extension and ACPI-aware operating systems could take advantage of presence of this _DSM and use this device specific information for deterministic display in OS device manager or equivalent user interfaces/ representations.

4. Analysis of the Hardware Implications

None

5. Analysis of the Software Implications

See section 3.

Part II

Detailed Description of the change

Modify table 4-7 as follows:

Table 4-7: _DSM Definitions for PCI

UUID	Revision	Function	Description
E5C937D0-3553-4d7a-9117-EA4D19C3434D	42	1	PCI Express Slot Information
	2	2	PCI Express Slot Number
	2	3	Vendor-specific Token ID
	2	4	PCI Bus Capabilities
	42	5	Ignore PCI Boot Configuration
	2	57	Naming a PCI or PCI Express device under OS

Deleted: 1

Deleted: 1

Deleted: 4

Modified section 4.6.1 to reflect changes in Revision ID

4.6.1. _DSM for PCI Express Slot Information

Arguments:

Arg0: UUID: E5C937D0-3553-4d7a-9117-EA4D19C3434D

Arg1: Revision ID: 2

Arg2: Function Index: 1

Arg3: Empty Package

Formatted: Font color: Red

Formatted: Font color: Red

Above change also applies to section 4.6.2, 4.6.3 & 4.6.4.1.3,

Appended section 4.6.x in the PCI Firmware specification

4.6.x. _DSM for naming a PCI or PCI Express device under OS

This interface will return a package with the following information relative to the device or slot:

- First entry of the package is the instance number..
 - Instance number must be unique under _SB scope. This instance number does not have to be sequential in a given system configuration.
 - This entry is mandatory when this _DSM is implemented

Request Request Request Request Request Request Request Request

- In typical usage model, in a configuration with multiple devices of the same type, device with lower instance number would be assigned lower device label by the OS compared to another device with higher instance number.
- Second entry of the package is a string name
 - This string is optional even when this _DSM is implemented; When not implemented, this entry must return a null string.
 - When populated, this string object should match the label/name on the chassis. String should be specified in Unicode.

When implemented, this string must be unique for a given platform under _SB. This interface is defined for devices that are defined in ACPI-namespace for a given platform by the BIOS. For expansion slots, typically, platform BIOS will include entry for the PCI or PCI Express slot in its namespace and each of these expansion slots would be expected to include their respective _DSM entries. Assignment of specific device names to multi-function devices installed in expansion slots, and/or PCI or PCI Express devices that are hot-added to expansion slots in OS-environment would be handled in OS-specific manner, and is not specified via this specification.

Location:

This object will be placed under the ACPI object representing the embedded PCI or PCI Express device/function or under the ACPI PCI Slot description (as defined in Section 4.7) representing the add-in PCI or PCI Express Slot.

Arguments:

Arg0: UUID: E5C937D0-3553-4d7a-9117-EA4D19C3434D
Arg1: Revision ID: 2
Arg2: Function Index: 7
Arg3: None

Return:

A package of two elements.

```
Method (_DSM, 3) { // Assume GUID and revision match

If (Land( Lequal(Arg1, 2), Lequal(Arg2,7))) {
    Return ( Package(2){ // PCI Express Slot Parsing
        1, // Instance of the enumeration
        String1 //String name which matches the label on the
                chassis
    })
} // end of if (Land....
} // end of method
```

IMPLEMENTATION NOTE

When implementing support for this method, the following should be considered: System firmware should include this method for all applicable PCI or PCI Express devices, both embedded and in PCI or PCI Express slots. System firmware should assign appropriate instance numbers to each device so as to influence the device labeling under operating system accordingly for the

Request Request Request Request Request Request Request Request

given platform design. For an example system configuration that includes two networking devices, the following table shows a) the location for each device (Bus#, Device#, Function#), b) description of device, c) _DSM (function 7) for each device, and d) the resultant sample device label as assigned by an operating system that takes advantage of this method. In case of this sample configuration, operating system would sort the different networking devices in the system using the instance numbers provided via _DSM, and assign lower NIC-label (NIC1) to device which has lower instance number, as shown in the table below.

Table: 4-<x>: Example usage of device/slot enumeration ordering hints

<u>PCI/PCI Express device location</u>	<u>System device</u>	<u>DSM Fn 7 Package</u>	<u>OS Device Label</u>
Bus 2, Device 4, Function 0	Networking device in PCI Express Slot 1	(3, "PCIe Slot 1")	NIC 2
Bus 4, Device 8, Function 0	Embedded Networking device	(2, "Embedded NIC1")	NIC 1