PCI-SIG ENGINEERING CHANGE NOTICE

TITLE:	SR-IOV Table Updates
DATE:	June 16, 2016
AFFECTED DOCUMENT:	Single Root I/O Virtualization and Sharing Revision 1.1, Readiness Notifications (RN) ECN
SPONSOR:	Gord Caruk, AMD

Part I

1. Summary of the Functional Changes

1. In SR-IOV specification, the PCI table of capabilities (Table 3-21) is updated to include all currently defined PCI Capability IDs.
    a. Allow presence of the Null Capability in PFs and VFs.
    b. Allow presence of the SATA Capability in PFs.
    c. Allow presence of the Enhanced Allocation Capability in PFs and disallow it in VFs.
2. In SR-IOV specification, the PCI-E table of capabilities (Table 3-22) is updated to include all currently defined PCI-E Extended Capability IDs.
    a. Allow presence of the Null Capability in PFs and VFs.
    b. Allow presence of Secondary PCI Express in PFs and disallow it in VFs.
    c. Allow presence of Protocol Multiplexing in PFs and disallow it in VFs.
    d. Allow presence of Process Address Space ID in PFs and disallow it in VFs.
    e. Allow presence of LN Requester in PFs and VFs.
    f. Disallow presence of Downstream Port Containment in PFs and VFs.
    g. Allow presence of L1 PM Substates in PFs and disallow it in VFs.
    h. Allow presence of Precision Time Measurement in PFs and disallow it in VFs.
    i. Allow presence of PCI Express over M-PHY in PFs and disallow it in VFs.
3. Since not all of the ECNs have been implemented into the latest draft of the SR-IOV specification, there is some minor spillover into the RN ECN.
4. As part of the 4.0 development effort, the *Single Root I/O Virtualization and Sharing Specification* is being incorporated into the base spec and other changes will come at that time.

2. Benefits as a Result of the Changes

Update SR-IOV specification to reflect current PCI Code and ID Assignment Specification, regarding PCI capabilities and PCI-E extended capabilities.

Clarify the requirements for VFs regarding the other Capabilities added by ECNs that should have updated the SR-IOV specification but did not.

3. Assessment of the Impact

In most cases, the changes formally state what already was stated in the Base spec or is implied by the basic definitions of what types of features a PF can support and what a VF can support.

4. Analysis of the Hardware Implications

It is not anticipated that any product changes are required by this ECN. All of the changes in this ECN have been derived from the various ECNs that added capabilities mentioned in this document. Products implementing those ECNs would have had to infer these requirements. Some difference in inference is possible.

5. Analysis of the Software Implications

It is not anticipated that any software changes are required by this ECN.

6. Analysis of the C&I Test Implications

1. PCI-SIG compliance test needs to conform to the updated table, with respect to whether or not a VF can implement the specific extended capability structures. Changes to the compliance test may be needed to check that the VF does or does not implement an extended capability structure in accordance with these clarifications.

Part II

Detailed Description of the change

In the Single Root I/O Virtualization and Sharing Specification, Section 3.6, make the following changes to Table 3-21.

Table 3-21: SR-IOV Usage of PCI Standard Capabilities

Capability ID	Description	PF Attributes	VF Attributes
00h	Null Capability	Base	Base
01h	PCI Power Management Interface	Base	Optional. See Chapter 6.
02h	AGP	n/a	n/a
03h	VPD	Base	Optional. See Section 3.6.1.
04h	Slot Identification	~~Base~~n/a	n/a
05h	MSI	Base	See ~~Chapter~~Section 5.1.1.
06h	CompactPCI* Hot Swap	n/a	n/a
07h	PCI-X™	n/a	n/a
08h	HyperTransport*	n/a	n/a
09h	Vendor Specific	Base	Base
0Ah	Debug Port	Base	Base
0Bh	CompactPCI Central Resource Control	n/a	n/a
0Ch	PCI Hot Plug	Base	n/a
0Dh	PCI Bridge Subsystem ID	n/a	n/a
0Eh	AGP 8x	n/a	n/a
0Fh	Secure Device	n/a	n/a
10h	PCI Express	Base	See Section 3.5.
11h	MSI-X	See ~~Chapter~~ Section 5.1.2 and Section 5.1.3.	See ~~Chapter~~ Section 5.1.2 and Section 5.1.3.
12h	Serial ATA Data/Index Configuration	~~n/a~~Base	n/a
13h	Advanced Features	n/a	n/a
14h	Enhanced Allocation	Base	Must not implement.

In the Single Root I/O Virtualization and Sharing Specification, Section 3.7, make the following changes to Table 3-22.

Table 3-22: SR-IOV Usage of PCI Express Extended Capabilities

Extended Capability ID	Description	PF Attributes	VF Attributes

0000h	Null Capability	Base	Base
0001h	Advanced Error Reporting (AER)	Base	See Section 4.2.
0002h	Virtual Channel (02h)	Base	Must not implement. ~~Not implemented.~~ See Section 3.7.1.
0003h	Device Serial Number	Base	~~Not implemented~~See Section 3.7.x1.
0004h	Power Budgeting	Base	Must not implement. ~~Not implemented.~~ See Section 3.7.x2.
0005h	Root Complex Link Declaration	n/a	n/a
0006h	Root Complex Internal Link Control	n/a	n/a
0007h	Root Complex Event Collector Endpoint Association	n/a	n/a
0008h	Multi Function Virtual Channel	Base	Must not implement. ~~Not implemented.~~ See Section 3.7.1.
0009h	Virtual Channel (09h)	Base	Must not implement. ~~Not implemented.~~ See Section 3.7.1.
000Ah	Root Complex Register Block Header ~~Capability~~ (RCRB)	n/a	n/a
000Bh	Vendor Specific	Base	Base
000Ch	Configuration Access Correlation ~~Access Capability~~	n/a	n/a
000Dh	Access Control Services (ACS)	~~Base.~~ See Section 3.7.2.	~~Base.~~ See Section 3.7.2.
000Eh	Alternative Routing ID Interpretation (ARI)	See Section 3.7.3.	See Section 3.7.3.
000Fh	Address Translation Services (ATS)	See Section 3.7.4.	See Section 3.7.4.
0010h	Single Root I/O Virtualization (SR-IOV)	See Section 3.3.	Must not implement. ~~Not implemented.~~ See Section 3.3.
0011h	Multi-Root I/O Virtualization (MR-IOV)	Must not implement. ~~Not implemented.~~ See Section 3.7.5.	Must not implement. ~~Not implemented.~~ See Section 3.7.5.
0012h	Multicast	See Section 3.7.6.	See Section 3.7.6.

0013h	~~Address Translation Services~~ Page Request Interface (PRI)	See Section 3.7.7.	See Section 3.7.7.
0014h	Reserved for AMD	Base	Base
~~0~~0015h	Resizable BAR	Base	Must not implement. ~~Not implemented.~~ See Section 3.7.x3.
0016h	Dynamic Power Allocation (DPA)	Base	Must not implement. ~~Not implemented.~~ See Section 3.7.8.
0017h	TLP Processing Hints (TPH)	See Section 3.7.9.	See Section 3.7.9.
0018h	Latency Tolerance Reporting (LTR)	Base	Must not implement. ~~Not implemented.~~ LTR is controlled using Function 0 which is never a VF.
0019h	Secondary PCI Express	Base	Must not implement. 8.0 GT/s is controlled using Function 0 which is never a VF.
001Ah	Protocol Multiplexing (PMUX)	Base	Must not implement. PMUX is controlled using Function 0 which is never a VF
001Bh	Process Address Space ID (PASID)	Base	See Section 3.7.y
001Ch	LN Requester (LNR)	Base	Base
001Dh	Downstream Port Containment (DPC)	n/a	n/a
001Eh	L1 PM Substates	Base	Must not implement. L1 PM Substates is controlled using Function 0 which is never a VF.
001Fh	Precision Time Measurement (PTM)	Base	Must not implement. PTM controls the port and must not be implemented in a VF.
0020h	PCI Express over M-PHY (M-PCIe)	Base	Must not implement. M-PCIe is controlled using Function 0 which is never a VF.

In the Single Root I/O Virtualization and Sharing Specification, after Section 3.7.1, add the following new sections. (note to technical writer: adjust all following section numbers, and ensure that all section references are updated.)

3.7.x1. Device Serial Number

The Device Serial Number Extended Capability may be present in PFs. If a PF contains the capability, its value applies to all associated VFs. VFs are permitted but not recommended to implement this capability. VFs that implement this capability must return the same Device Serial Number value as that reported by their associated PF.

3.7.x2. Power Budgeting

The Power Budgeting Extended Capability may be present in PFs, but VFs must not implement it. If a PF contains the capability, it must report values that cover all associated VFs.

In the Single Root I/O Virtualization and Sharing Specification, after Section 3.7.7, add the following new section. (note to technical writer: adjust all following section numbers, and ensure that all section references are updated.)

3.7.x3. Resizable BAR

The Resizable BAR Extended Capability may be present in PFs. Since VFs do not implement standard BARs the capability must not be present in a VF. The PF's Resizable BAR settings do not affect any settings in the SR-IOV capability.

In the Single Root I/O Virtualization and Sharing Specification, after Section 3.7.9, add the following new section.

3.7.y. PASID

An Endpoint device is permitted to support PASID. The PASID configuration of the single function (Function or PF) representing the device is also used by all VFs in the device. A PF is permitted to implement the PASID capability, but VFs must not implement it.

Even though the PASID configuration is shared between Functions, PFs and VFs, the device sends the requesting Function's ID (Function, PF or VF) in the Requester ID field of the TLP containing PASID.

In the Readiness Notifications (RN) ECN, Section 3.7, make the following changes to Table 3-22.

0021h	FRS Queuing	n/a	n/a
0022h	Readiness Time Reporting	Base	See Section 3.7.~~6~~z

In the Readiness Notifications (RN) ECN, Section 3.7.6, make the following changes.

3.7.~~6~~z. Readiness Time Reporting Extended Capability

The Readiness Time Reporting Extended Capability may be present in VFs, PFs, both or neither. If any VF associated with a given PF contains the capability, all VFs associated with that same PF must also support Readiness Time Reporting.

The **Reset Time** field contains the time required following setting of VF Enable (see Section 6.1).

The **DL_Up Time** field is RsvdP.

All VFs associated with ~~a~~the same PF shall report the same time values.