PCI-SIG ENGINEERING CHANGE NOTICE

TITLE:	VF Resizable BARs
DATE:	July 6, 2016
AFFECTED DOCUMENT:	Single Root I/O Virtualization and Sharing Specification version 1.1, PCI Code and ID Assignment Specification version 1.7
SPONSORS:	Advanced Micro Devices

Part I

1. Summary of the Functional Changes

Similar to, and based on, the Resizable BAR and Expanded Resizable BAR ECNs, this optional ECN adds a capability for PFs to be able to resize their VF BARs.

This ECN is written with the expectation that the Expanded Resizable BAR ECN will have been released prior to this ECN's release. This ECN supports all of the BAR sizes defined by both the Resizable BAR and Expanded Resizable BAR ECNs.

2. Benefits as a Result of the Changes

The VF Resizable BAR Capability allows VF BARs to be sized and allocated with the same flexibility and algorithms that regular BARs have, and allowing software to adjust the size of all BARs in a similar way. The benefits of this ECN are the same as for the Resizable BAR and Expanded Resizable BAR ECNs. Namely that all resources in the system can be allocated to their maximum level of efficiency within the available resources.

3. Assessment of the Impact

This capability is applicable only to PFs (Functions with an SR-IOV Capability) that have VF BARs. It is thought that PFs using this capability will limit its use to very large resources, for example, very large local memories.

4. Analysis of the Hardware Implications

This is an optional normative capability that is useful for PFs that have resources requiring large amounts of address space. Components that this does not apply to or do not wish to add this capability are not required to. The hardware implications for a component that does implement this capability are limited to some new Configuration Space registers and a modification to their BARs.

5. Analysis of the Software Implications

There is no impact to current software. The registers associated with this capability default to a benign state so that current resource allocation algorithms continue to operate as they currently do.

Software that allocates resources into the address space of the system is recommended to add a utility that collects resource capabilities and sizes all BARs; regular BARs as well as VF BARs.

6. Analysis of the Compatibility and Interoperability Test Implications

New test cases would need to be created to test the functionality of this feature.

Part II

Detailed Description of the change to the Single Root I/O Virtualization and Sharing Specification

Modify Section 2.1.1.1, page 25, line 18, as shown

…

The behavior of VF BARs is the same as the ***PCI Local Bus Specification*** normal PCI Memory Space BARs, except that a VF BAR describes the aperture for each VF, whereas a PCI BAR describes the aperture for a single Function~~:~~. The attributes for some of the bits in the VF BARs are affected by the VF Resizable BAR capability (see Section 3.7.x4) if it is implemented.

Modify Section 2.2.2, page 30, line 1, as shown

VFs must support Function Level Reset (FLR).

Note: Software may use FLR to reset a VF. FLR to a VF affects a VF's state but does not affect its existence in PCI Configuration Space or PCI Bus address space. The VFs BARn values (see Section 3.3.14) and VF MSE (see Section 3.3.3.4) in the PF's SR-IOV extended capability, and the VF Resizable BAR capability values (see Section 3.7.x4) are unaffected by FLRs issued to the VF.

Modify Section 3.3.14, page 48, line 4, as shown

These fields must define the VF's Base Address Registers (BARs). These fields behave as normal PCI BARs, as described in the ***PCI Local Bus Specification***. They can be sized by writing all 1s and reading back the contents of the BARs as described in the ***PCI Local Bus Specification***, complying with the low order bits that define the BAR type fields.
These fields may have their attributes affected by the VF Resizable BAR capability (see Section 3.7.x4) if it is implemented.
The amount of address space decoded by each BAR shall be an integral multiple of System Page Size.
…

The algorithm to determine the amount of address space mapped by a VF BAR*n* differs from the ~~base specification~~ standard BAR algorithm as follows:
1. Resize the BAR via the VF Resizable BAR capability (see Section 3.7.x4) if it is implemented.

2. After reading the low order bits to determine if the BAR …

In Section 3.7, page 69, add a row at the end of Table 3-22 as follows:

0024h	VF Resizable BAR	See Section 3.7.x4	Not Implemented. See Section 3.7.x4

Add new Section 3.7.x4, page 74, line 11. (Note to Editor: preferably this follows the Section entitled 'Resizable BAR'.)

3.7.x4. VF Resizable BAR Capability

The VF Resizable BAR capability is permitted to be implemented only in PFs that implement at least one VF BAR, and affects the size and base of a PF's VF BARs. Since VFs do not implement the BARs themselves the capability must not be present in a VF. A PF may implement both a VF Resizable BAR capability and a Resizable BAR capability, as each capability operates independently.

The VF Resizable BAR capability is an optional capability that permits PFs to be able to have their VF's BARs resized. The VF Resizable BAR capability permits hardware to communicate the resource sizes that are acceptable for operation via the VF Resizable BAR Capability and Control registers and system software to communicate the optimal size back to the hardware via the VF BAR Size field of the VF Resizable BAR Control register.

Hardware immediately reflects the size inference in the read-only bits of the appropriate VF BAR. The size inferred is the greater of the values decoded from the System Page Size and VF BAR Size fields. Hardware must Clear any bits that change from read-write to read-only, so that subsequent reads return zero. Software must clear the VF MSE bit in the SR-IOV Control register before writing the VF BAR Size field. After writing the VF BAR Size field, the contents of the corresponding VF BAR are undefined. To ensure that it contains a valid address after resizing the VF BAR, system software must reprogram the VF BAR, and Set the VF MSE bit (unless the resource is not allocated).

The VF Resizable BAR Capability and Control registers are permitted to indicate the ability to operate at 4 GB or greater only if the associated VF BAR is a 64-bit BAR.

It is strongly recommended that a Function not advertise any supported VF BAR sizes in its VF Resizable BAR Capability and Control registers that are larger than the space it would effectively utilize if allocated.

IMPLEMENTATION NOTE

Using the Capability During Resource Allocation

System software uses this capability in a similar way to the Resizable BAR capability. System software must first configure the System Page Size register (see Section 2.1.1.1). Potential usable memory aperture sizes are reported by the PF, and read, from the VF Resizable BAR Capability and Control registers. It is intended that the software allocate the largest of the reported sizes that it can, since allocating less address space than the largest reported size can result in lower performance. Software then writes the size to the VF Resizable BAR Control register for the appropriate VF BAR for the Function. Following this, the base address is written to the VF BAR.

For interoperability reasons, it is possible that hardware will set the default size of the VF BAR to a low size; a size lower than the largest reported in the VF Resizable BAR Capability register. Software that does not use this capability to size resources will likely result in sub-optimal resource allocation, where the resources are smaller than desirable, or not allocatable because there is no room for them.

It is recommended that system software responsible for allocating resources in a resource constrained environment distribute the limited address space to all memory-mapped hardware, including system RAM, appropriately.

The VF Resizable BAR Capability structure defines a PCI Express Extended Capability which is located in PCI Express Extended Configuration Space, that is, above the first 256 bytes, and is shown below in Figure 7-x1. This structure allows PFs with this capability to be identified and controlled. A Capability register and a Control register are implemented for each VF BAR that is resizable. Since a maximum of 6 VF BARs may be implemented by any PF, the VF Resizable BAR Capability structure can range from 12 bytes long (for a single VF BAR) to 52 bytes long (for all 6 VF BARs).

31 0	Byte Offset
VF Resizable BAR Extended Capability Header	000h
VF Resizable BAR Capability Register(0)	004h
VF Resizable BAR Control Register(0)	008h
…	
VF Resizable BAR Capability Register(n)	(n*8+4)
VF Resizable BAR Control Register(n)	(n*8+8)

Figure 7-x1 VF Resizable BAR Capability

7.xx.1. VF Resizable BAR Extended Capability Header (Offset 00h)

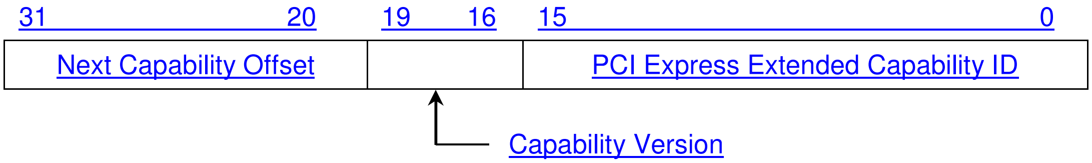

Figure 7-x2 VF Resizable BAR Extended Capability Header

Table 7-x1 VF Resizable BAR Extended Capability Header

Bit Location	**Register Description**	**Attributes**
15:0	**PCI Express Extended Capability ID** – This field is a PCI-SIG defined ID number that indicates the nature and format of the extended capability. PCI Express Extended Capability ID for the VF Resizable BAR Capability is 0024h.	RO
19:16	**Capability Version** – This field is a PCI-SIG defined version number that indicates the version of the capability structure present. Must be 1h for this version of the specification.	RO
31:20	**Next Capability Offset** – This field contains the offset to the next PCI Express Extended Capability structure or 000h if no other items exist in the linked list of capabilities.	RO

7.xx.2. VF Resizable BAR Capability Register (Offset 04h)

The VF Resizable BAR Capability register field descriptions are the same as the definitions in the Resizable BAR Capability register in *PCI Express Base Specification*, Table 7-96. Where those descriptions say 'BAR', this register's description is for 'VF BAR'. Where those descriptions say 'Function', this register's description is for 'PF'. Otherwise the field descriptions, the number of bits, their positions, and their attributes are the same. Consequently *PCI Express Base Specification*, Figure 7-110 similarly allocates the register fields in this register.

7. xx.3. VF Resizable BAR Control Register (Offset 08h)

The VF Resizable BAR Control register bits 31:16 follow the same definitions as the Resizable BAR Control register in *PCI Express Base Specification*, Table 7-97.

31 - 16	15 - 14	13 - 8	7 - 5	4 - 3	2 - 0
Follows allocation in *PCI Express Base Specification*, Figure 7-111	RsvdP			RsvdP	

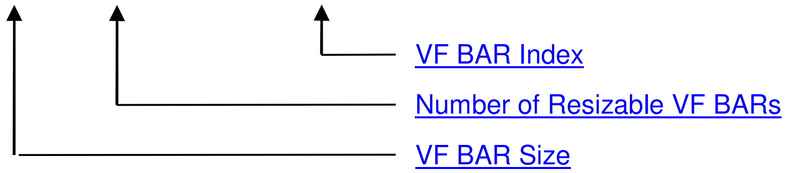

Figure 7-x4 VF Resizable BAR Control Register

Table 7-x3 VF Resizable BAR Control Register

Bit Location	**Register Description**	**Attributes**
2:0	**VF BAR Index** – This encoded value points to the beginning of this particular VF BAR located in the SR-IOV Capability. 0 = VF BAR located at offset 24h 1 = VF BAR located at offset 28h 2 = VF BAR located at offset 2Ch 3 = VF BAR located at offset 30h 4 = VF BAR located at offset 34h 5 = VF BAR located at offset 38h All other encodings are reserved. For a 64-bit Base Address register, the VF BAR Index indicates the lower DWORD. This value indicates which VF BAR supports a negotiable size.	RO
7:5	**Number of VF Resizable BARs** – Indicates the total number of resizable VF BARs in the capability structure for the Function. See Figure 7-x1. The value of this field must be in the range of 01h to 06h. The field is valid in VF Resizable BAR Control register (0) (at offset 008h), and is RsvdP for all others.	RO/RsvdP
13:8	**VF BAR Size** – This is an encoded value. 0 = 1 MB (2^{20}) 1 = 2 MB (2^{21}) 2 = 4 MB (2^{22}) 3 = 8 MB (2^{23}) … 43 = 8 EB (2^{63}) The default value of this field is equal to the default size of the address space that the VF BAR resource is requesting via the VF BAR's read-only bits. Behavior is undefined if a value is written in this field and the corresponding supported size bit is not Set in the VF Resizable BAR Capability or VF Resizable BAR Control registers. When this register field is programmed, the value is immediately reflected in the size of the resource, as encoded in the number of read-only bits in the VF BAR.	RW
31:16	The descriptions for these bits are the same as the definitions in the Resizable BAR Control register in *PCI Express Base Specification*, Table 7-97. Where those descriptions say 'BAR', this register's description is for 'VF BAR'. Where those descriptions say 'Function', this register's description is for 'PF'.	From *PCI Express Base Specification*, Table 7-97

Detailed Description of the change to the Codes and ID Specification

In Section 3, page 25, add a row at the end of Table 3-1 as follows:

0024h	VF Resizable BAR