# PCI-SIG ENGINEERING CHANGE NOTICE

| | |
|---|---|
| **TITLE:** | Resizable BAR Capability |
| **DATE:** | Jan 22, 2008 – Updated and approved by PWG April 24, 2008 |
| **AFFECTED DOCUMENT:** | PCI Express Base Specification version 2.0 |
| **SPONSORS:** | Hewlett-Packard, Advanced Micro Devices |

## Part I

### 1. Summary of the Functional Changes

This optional ECN adds a capability for Functions with BARs to report various options for sizes of their memory mapped resources that will operate properly. Also added is an ability for software to program the size to configure the BAR to.

Together these allow the hardware to tell resource allocation software what are the sizes that may be used, and that software will attempt to allocate the resource to the largest of the reported sizes.

### 2. Benefits as a Result of the Changes

The Resizable BAR Capability allows system software to allocate all resources in systems where the total amount of resources requesting allocation plus the amount of installed system memory is larger than the supported address space.

Currently, resources are either: a) simply not allocated and left out of the system, or b) forced to report a smaller aperture in order to be allocated, but that aperture size in not optimal in all uses of the product.

This change allows the system to be configured with optimal resource settings.

### 3. Assessment of the Impact

This capability is applicable to all components that have Base Address registers and derive value from presenting two or more options for the sizes of the resources inferred by the BARs. It is thought that Functions using this capability will be limited to very large resources, for example, add in cards with very large local memories.

### 4. Analysis of the Hardware Implications

This is an optional capability that is useful for Functions 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 the resource capabilities, determine the optimal size to allocate to the resource, and size the BAR, prior to setting the base address.

## Part II

### Detailed Description of the change

*Modify Section 7.5.2.1. as shown*

… The minimum Memory Space address range requested by a BAR is 128 bytes. The attributes for some of the bits in the BAR are affected by the Resizable BAR Capability, if it is implemented.

*Add new Section 7.xx*

# 7.xx. Resizable BAR Capability

The Resizable BAR Capability is an optional capability that allows hardware to communicate resource sizes, and system software, after determining the optimal size, to communicate this optimal size back to the hardware. Hardware communicates the resource sizes that are acceptable for operation via the Resizable BAR Capability register. Software determines, through a proprietary mechanism, what the optimal size is for the resource, and programs that size via the BAR Size field of the Resizable BAR Control register. Hardware immediately reflects the size inference in the read-only bits of the appropriate Base Address register. Hardware must Clear any bits that change from R/W to read-only, so that subsequent reads return zero. Software must clear the Memory Space Enable bit in the Command register before writing the BAR Size field. After writing the BAR Size field, the contents of the corresponding BAR are undefined. To ensure that it contains a valid address after resizing the BAR, system software must reprogram the BAR, and Set the Memory Space Enable bit (unless the resource is not allocated).

The Resizable BAR Capability register is permitted to indicate the ability to operate at 4GB or greater only if the associated BAR is a 64b BAR.

This capability is applicable to Functions that have Base Address registers only. It is strongly recommended that a Function not advertise any supported BAR sizes in its Resizable BAR Capability register that are larger than the space it would effectively utilize if allocated.

---

## IMPLEMENTATION NOTE

### Using the Capability During Resource Allocation

System software that allocates resources can use this capability to resize the resources inferred by the Function's BAR's read-only bits. Previous versions of this software determined the resource size by writing FFFFh to the BAR, reading back the value, and determining the size by the number of bits that are Set. Following this, the base address is written to the BAR.

System software uses this capability in place of the above mentioned method of determining the resource size, and prior to assigning the base address to the BAR. Potential usable resource sizes are reported by the Function, and read, from the Resizable BAR Capability 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 Resizable BAR Command register for the appropriate BAR for the Function. Following this, the base address is written to the BAR.

For interoperability reasons, it is possible that hardware will set the default size of the BAR to a low size; a size lower than the largest reported in the 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.

With the Resizable BAR capability, the amount of address space consumed by a device can change. In a resource constrained environment, the allocation of more address space to a device may result in allocation of less of the address space to other memory-mapped hardware, like system RAM. System software responsible for allocating resources in this kind of environment is recommended to distribute the limited address space appropriately.

---

The 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 devices with this capability to be identified and controlled. A Capability and a Control register is implemented for each BAR that is resizable. Since a maximum of 6 BARs may be implemented by any Function, the Resizable BAR Capability structure can range from 12 bytes long (for a single BAR) to 52 bytes long (for all 6 BARs).

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

**Figure 7-x1 Resizable BAR Capability**

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

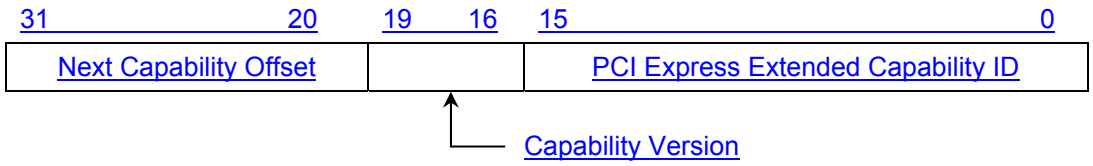


**Figure 7-x2 Resizable BAR Extended Capability Header**

**Table 7-x1 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.<br><br>PCI Express Extended Capability ID for the Resizable BAR Capability is 0015h. | RO |
| 19:16 | **Capability Version** – This field is a PCI-SIG defined version number that indicates the version of the capability structure present.<br><br>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. Resizable BAR Capability Register (Offset 04h)

| 31 - 24 | 23 | 22 | 21 | 20 | 19 | 18 | 17 | 16 | 15 | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 - 0 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| RsvdP | | | | | | | | | | | | | | | | | | | | | RsvdP |

Figure 7-x3 Resizable BAR Capability Register

Table 7-x2 Resizable BAR Capability Register

| Bit Location | Register Description | Attributes |
|---|---|---|
| 4 | When Set, indicates that the Function will operate with the BAR sized to 1MB | RO |
| 5 | When Set, indicates that the Function will operate with the BAR sized to 2MB | RO |
| 6 | When Set, indicates that the Function will operate with the BAR sized to 4MB | RO |
| 7 | When Set, indicates that the Function will operate with the BAR sized to 8MB | RO |
| 8 | When Set, indicates that the Function will operate with the BAR sized to 16 MB | RO |
| 9 | When Set, indicates that the Function will operate with the BAR sized to 32 MB | RO |
| 10 | When Set, indicates that the Function will operate with the BAR sized to 64 MB | RO |
| 11 | When Set, indicates that the Function will operate with the BAR sized to 128 MB | RO |
| 12 | When Set, indicates that the Function will operate with the BAR sized to 256 MB | RO |
| 13 | When Set, indicates that the Function will operate with the BAR sized to 512 MB | RO |
| 14 | When Set, indicates that the Function will operate with the BAR sized to 1 GB | RO |
| 15 | When Set, indicates that the Function will operate with the BAR sized to 2 GB | RO |
| 16 | When Set, indicates that the Function will operate with the BAR sized to 4 GB | RO |
| 17 | When Set, indicates that the Function will operate with the BAR sized to 8 GB | RO |
| 18 | When Set, indicates that the Function will operate with the BAR sized to 16 GB | RO |
| 19 | When Set, indicates that the Function will operate with the BAR sized to 32 GB | RO |
| 20 | When Set, indicates that the Function will operate with the BAR sized to 64 GB | RO |

| | | |
|---|---|---|
| 21 | When Set, indicates that the Function will operate with the BAR sized to 128 GB | RO |
| 22 | When Set, indicates that the Function will operate with the BAR sized to 256 GB | RO |
| 23 | When Set, indicates that the Function will operate with the BAR sized to 512 GB | RO |

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

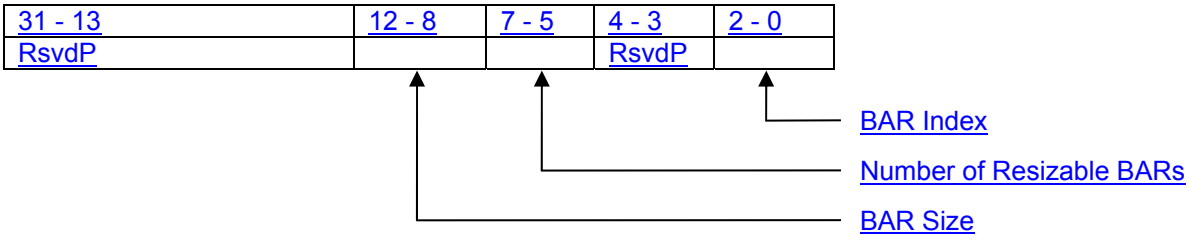


**Figure 7-x4 Resizable BAR Control Register**

**Table 7-x3 Resizable BAR Control Register**

| Bit Location | Register Description | Attributes |
|---|---|---|
| 2:0 | **BAR Index** – This encoded value points to the beginning of the BAR.<br>0 = BAR located at offset 10h<br>1 = BAR located at offset 14h<br>2 = BAR located at offset 18h<br>3 = BAR located at offset 1Ch<br>4 = BAR located at offset 20h<br>5 = BAR located at offset 24h<br>All other encodings are reserved.<br>For a 64-bit Base Address register, the BAR Index indicates the lower DWORD.<br>This value indicates which BAR supports a negotiable size. | RO |
| 7:5 | **Number of Resizable BARs** – Indicates the total number of resizable BARs in the capability structure for the Function. See Figure 7-x1.<br>The value of this field must be in the range of 01h to 06h. The field is valid in Resizable BAR Control register (0) (at offset 008h), and is RsvdP for all others. | RO/RsvdP |
| 12:8 | **BAR Size** – This is an encoded value.<br>0 = 1 MB<br>1 = 2 MB<br>2 = 4 MB<br>3 = 8 MB<br>…<br>19 = 512 GB<br>The default value of this field is equal to the default size of the address space that the BAR resource is requesting via the BAR's read-only bits.<br>When this register field is programmed, the value is | R/W |

| | | |
|---|---|---|
| | immediately reflected in the size of the resource, as encoded in the number of read-only bits in the BAR. | |