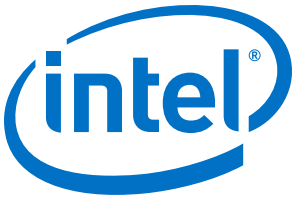

Intel® Xeon® Processor E3-1200v4 Product Family

Datasheet – Volume 2 of 2

June 2015

Order No.: 332375-001

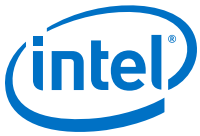

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or visit http://www.intel.com/design/literature.htm.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at http://www.intel.com/ or from the OEM or retailer.

No computer system can be absolutely secure.

Intel® Hyper-Threading Technology (Intel® HT Technology) is available on select Intel® Core™ processors. It requires an Intel® HT Technology enabled system. Consult your PC manufacturer. Performance will vary depending on the specific hardware and software used. Not available on Intel® Core™ i5-750. For more information including details on which processors support Intel® HT Technology, visit http://www.intel.com/info/hyperthreading.

Intel® High Definition Audio (Intel® HD Audio) requires an Intel® HD Audio enabled system. Consult your PC manufacturer for more information. Sound quality will depend on equipment and actual implementation. For more information about Intel® HD Audio, refer to http://www.intel.com/design/chipsets/hdaudio.htm.

Intel® 64 architecture requires a system with a 64-bit enabled processor, chipset, BIOS and software. Performance will vary depending on the specific hardware and software you use. Consult your PC manufacturer for more information. For more information, visit http://www.intel.com/content/www/us/en/architecture-and-technology/microarchitecture/intel-64-architecture-general.html.

Intel® Virtualization Technology (Intel® VT) requires a computer system with an enabled Intel® processor, BIOS, and virtual machine monitor (VMM). Functionality, performance or other benefits will vary depending on hardware and software configurations. Software applications may not be compatible with all operating systems. Consult your PC manufacturer. For more information, visit http://www.intel.com/go/virtualization.

The original equipment manufacturer must provide TPM functionality, which requires a TPM-supported BIOS. TPM functionality must be initialized and may not be available in all countries.

For Enhanced Intel SpeedStep® Technology, see the Processor Spec Finder at http://ark.intel.com/ or contact your Intel representative for more information.

Intel® AES-NI requires a computer system with an AES-NI enabled processor, as well as non-Intel software to execute the instructions in the correct sequence. AES-NI is available on select Intel® processors. For availability, consult your reseller or system manufacturer. **For more information, see** http://software.intel.com/en-us/articles/intel-advanced-encryption-standard-instructions-aes-ni/.

Intel® Active Management Technology (Intel® AMT) should be used by a knowledgeable IT administrator and requires enabled systems, software, activation, and connection to a corporate network. Intel AMT functionality on mobile systems may be limited in some situations. Your results will depend on your specific implementation. Learn more by visiting Intel® Active Management Technology.

No computer system can provide absolute security under all conditions. Intel® Trusted Execution Technology (Intel® TXT) requires a computer with Intel® Virtualization Technology, an Intel TXT-enabled processor, chipset, BIOS, Authenticated Code Modules and an Intel TXT-compatible measured launched environment (MLE). Intel TXT also requires the system to contain a TPM v1.s. For more information, visit http://www.intel.com/technology/security.

Requires a system with Intel® Turbo Boost Technology. Intel Turbo Boost Technology and Intel Turbo Boost Technology 2.0 are only available on select Intel® processors. Consult your PC manufacturer. Performance varies depending on hardware, software, and system configuration. For more information, visit https://www-ssl.intel.com/content/www/us/en/architecture-and-technology/turbo-boost/turbo-boost-technology.html.

Intel® Advanced Vector Extensions (Intel® AVX) are designed to achieve higher throughput to certain integer and floating point operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at less than the rated frequency and b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software, and system configuration and you should consult your system manufacturer for more information. Intel® Advanced Vector Extensions refers to Intel® AVX, Intel® AVX2 or Intel® AVX-512. For more information on Intel® Turbo Boost Technology 2.0, visit https://www-ssl.intel.com/content/www/us/en/architecture-and-technology/turbo-boost/turbo-boost-technology.html

Intel, Intel Core, Intel SpeedStep, and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

Contents

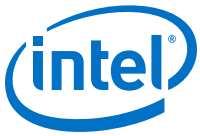

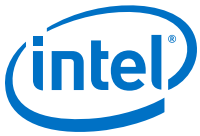

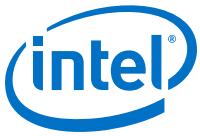

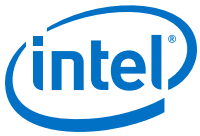

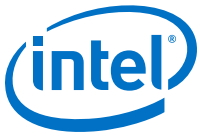

Figures

Tables

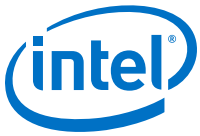

Revision History

Revision	Description	Date
001	• Initial Release	June 2015

1.0 Introduction

This is Volume 2 of the Datasheet for the following product family. Volume 2 provides register information for the processor.

- Intel® Xeon® processor E3-1200 v4 product family

Note: Throughout this document, the product family listed above may be referred to simply as "processor".

For Volume 1, refer to the following document.

- Server H-Processor Line. Refer to document number 332374 for the *Intel® Xeon® processor E3-1200 v4 Product Family (LGA) Entry Server and Workstation Processor Datasheet – Volume 1 of 2*

The processor contains one or more PCI devices within a single physical component. The configuration registers for these devices are mapped as devices residing on the PCI Bus assigned for the processor socket. This document describes these configuration space registers or device-specific control and status registers only.

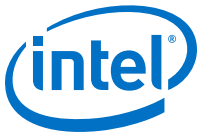

2.0 Processor Configuration Register Definitions and Address Ranges

This chapter describes the processor configuration register I/O and memory address ranges. The chapter provides register terminology. PCI Devices and Functions are described.

2.1 Register Terminology

Register Attributes and Terminology table lists the register-related terminology and access attributes that are used in this document. Register Attribute Modifiers table provides the attribute modifiers.

Table 1. Register Attributes and Terminology

Item	Description
RO	**Read Only:** These bits can only be read by software, writes have no effect. The value of the bits is determined by the hardware only.
RW	**Read / Write:** These bits can be read and written by software.
RW1C	**Read / Write 1 to Clear:** These bits can be read and cleared by software. Writing a '1' to a bit will clear it, while writing a '0' to a bit has no effect. Hardware sets these bits.
RW0C	**Read / Write 0 to Clear:** These bits can be read and cleared by software. Writing a '0' to a bit will clear it, while writing a '1' to a bit has no effect. Hardware sets these bits.
RW1S	**Read / Write 1 to Set:** These bits can be read and set by software. Writing a '1' to a bit will set it, while writing a '0' to a bit has no effect. Hardware clears these bits.
RsvdP	**Reserved and Preserved:** These bits are reserved for future RW implementations and their value must not be modified by software. When writing to these bits, software must preserve the value read. When SW updates a register that has RsvdP fields, it must read the register value first so that the appropriate merge between the RsvdP and updated fields will occur.
RsvdZ	**Reserved and Zero:** These bits are reserved for future RW1C implementations. Software must use 0 for writes.
WO	**Write Only:** These bits can only be written by software, reads return zero. *Note:* Use of this attribute type is deprecated and can only be used to describe bits without persistent state.
RC	**Read Clear:** These bits can only be read by software, but a read causes the bits to be cleared. Hardware sets these bits. *Note:* Use of this attribute type is only allowed on legacy functions, as side-effects on reads are not desirable
RSW1C	**Read Set / Write 1 to Clear:** These bits can be read and cleared by software. Reading a bit will set the bit to '1'. Writing a '1' to a bit will clear it, while writing a '0' to a bit has no effect.
RCW	**Read Clear / Write:** These bits can be read and written by software, but a read causes the bits to be cleared. *Note:* Use of this attribute type is only allowed on legacy functions, as side-effects on reads are not desirable.

Table 2. Register Attribute Modifiers

Attribute Modifier	Applicable Attribute	Description
S	RO (w/ -V)	**Sticky** : These bits are only re-initialized to their default value by a "Power Good Reset". *Note:* Does not apply to RO (constant) bits.
	RW	
	RW1C	
	RW1S	
-K	RW	**Key**: These bits control the ability to write other bits (identified with a 'Lock' modifier)
-L	RW	**Lock**: Hardware can make these bits "Read Only" using a separate configuration bit or other logic. *Note:* Mutually exclusive with 'Once' modifier.
	WO	
-O	RW	**Once**: After reset, these bits can only be written by software once, after which they become "Read Only". *Note:* Mutually exclusive with 'Lock' modifier and does not make sense with 'Variant' modifier.
	WO	
-FW	RO	**Firmware Write**: The value of these bits can be updated by firmware (PCU, TAP, and so on).
-V	RO	**Variant**: The value of these bits can be updated by hardware. *Note:* RW1C and RC bits are variant by definition and therefore do not need to be modified.

2.2 PCI Devices and Functions

The processor contains the following PCI devices within a single component. The configuration registers for the devices are mapped as devices residing on PCI Bus 0.

- Device 0: Host Bridge / DRAM Controller / LLC Controller0 – Logically this device appears as a PCI device residing on PCI bus 0. Device 0 contains the standard PCI header registers, PCI Express base address register, DRAM control (including thermal/throttling control), configuration for the DMI, and other processor specific registers.
- Device 1: Host-PCI Express* Bridge - Logically this device appears as a "virtual" PCI-to-PCI bridge residing on PCI bus 0, and is compliant with the *PCI-to-PCI Bridge Architecture Specification, Revision 1.2*. Device 1 is a multi-function device (MFD) consisting of three functions (0, 1, and 2). Device 1 contains the standard PCI-to-PCI bridge registers and the standard PCI Express/PCI configuration registers.
- Device 2: Integrated Graphics Device – Logically, this device appears as a PCI device residing on PCI bus 0. Physically, Device 2 contains the configuration registers for 3D, 2D, and display functions. In addition, Device 2 is located in two separate physical locations – GT and Display Engine.
- Device 3: High Definition Audio controller. This device contains registers used as control and status for integrated audio controller. Previous implementation of this controller was in the PCH.

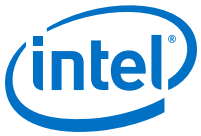

Table 3. PCI Devices and Functions for H-Processor Line

Description	DID (UP Server)	Bus	Device	Functions
HOST Bridge	0x1618	0	0	0
PCI Express* Controller (x16 PCIe)	0x1601	0	1	0
PCI Express* Controller (x8 PCIe)	0x1605	0	1	1
PCI Express* Controller (x4 PCIe)	0x1609	0	1	2
Integrated Graphics Device	N/A - GT1	0	2	0
	N/A - GT2			
	0x162A - GT3			
Audio Controller	0x160C	0	3	0

Notes: 1. Not all devices are enabled in all configurations.
2. DID values in this table differ from what is listed at the register ID description tables. The values listed in this table are the correct DID default values.

From a configuration standpoint, the DMI is logically PCI bus 0. As a result, all devices internal to the processor and the PCH appear to be on PCI bus 0.

For H-Processor Line 2-chip LGA/BGA, the PCI Express controllers (PEG10, PEG11, and PEG12) appear to system software to be real PCI buses behind PCI-to-PCI bridges that are devices resident on PCI bus 0. This is shown in Figure 1 on page 20.

Table 4. PCI Device Enumeration

Segment	Bus ID [7:0]	Device ID [4:0]	Function ID [2:0]	Endpoint
H-Processor Line	0x00	00000b (0)	000 (0)	Host Bridge
H-Processor Line	0x00	00010b (2)	000 (0)	Integrated Graphics Device
H-Processor Line	0x00	00011b (3)	000 (0)	Audio Controller

Figure 1. Conceptual Platform PCI Configuration Diagram for H-Processor Line

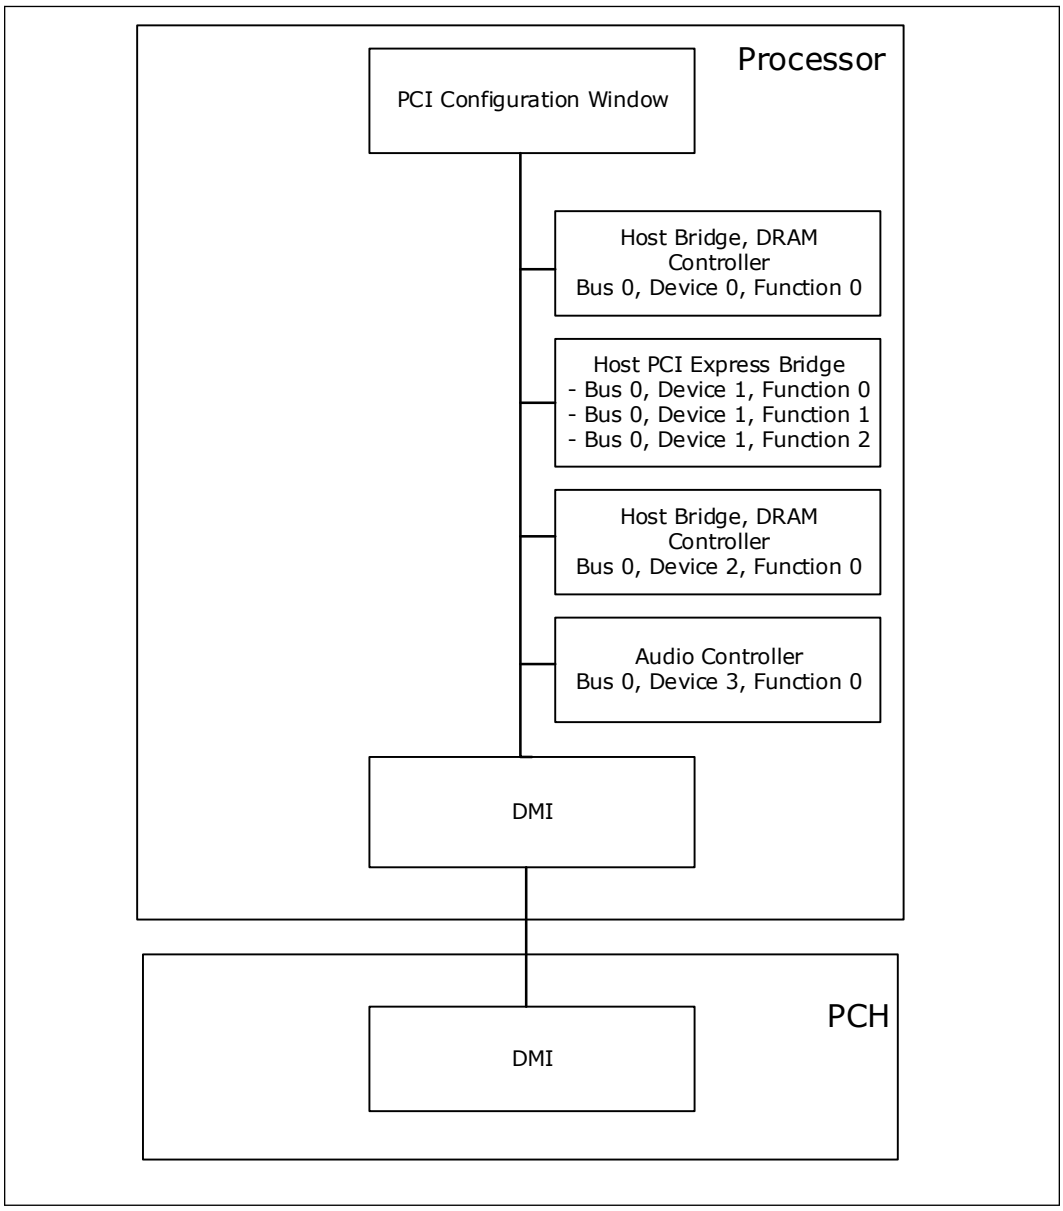

2.3 System Address Map

The processor supports 512 GB (39 bits) of addressable memory space and 64 KB+3 of addressable I/O space.

This section focuses on how the memory space is partitioned and how the separate memory regions are used. I/O address space has simpler mapping and is explained towards the end of this chapter.

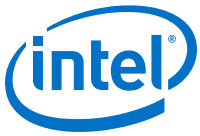

The processor supports a maximum of 32 GB of DRAM. No DRAM memory will be accessible above 32 GB. DRAM capacity is limited by the number of address pins available. There is no hardware lock to prevent more memory from being inserted than is addressable.

When running in internal graphics mode, processor initiated TileX/TileY/linear reads/ writes to GMADR range are supported. Write accesses to GMADR linear regions are supported from DMI. GMADR write accesses to TileX and TileY regions (defined using fence registers) are not supported from the DMI port. GMADR read accesses are not supported from DMI.

In the following sections, it is assumed that all of the compatibility memory ranges reside on the DMI Interface. The exception to this rule is VGA ranges, which may be mapped to PCI Express*, DMI, or to the internal graphics device (IGD). In the absence of more specific references, cycle descriptions referencing PCI should be interpreted as the DMI Interface/PCI, while cycle descriptions referencing PCI Express or IGD are related to the PCI Express bus or the internal graphics device respectively. The processor does not remap APIC or any other memory spaces above TOLUD (Top of Low Usable DRAM). The TOLUD register is set to the appropriate value by BIOS. The remapbase/remaplimit registers remap logical accesses bound for addresses above 4 GB onto physical addresses that fall within DRAM.

The Address Map includes a number of programmable ranges:

- Device 0:
  - PXPEPBAR – PxP egress port registers. (4 KB window)
  - MCHBAR – Memory mapped range for internal MCH registers. (32 KB window)
  - DMIBAR –This window is used to access registers associated with the processor/PCH Serial Interconnect (DMI) register memory range. (4 KB window)
  - GGC.GMS – Graphics Mode Select. Used to select the amount of main memory that is pre-allocated to support the internal graphics device in VGA (non-linear) and Native (linear) modes. (0 – 512 MB options).
  - GGC.GGMS – GTT Graphics Memory Size. Used to select the amount of main memory that is pre-allocated to support the Internal Graphics Translation Table. (0 – 2 MB options).
- For each of the following device functions
- Device 2, Function 0: (Integrated Graphics Device (IGD))
  - IOBAR – I/O access window for internal graphics. Through this window address/data register pair, using I/O semantics, the IGD and internal graphics instruction port registers can be accessed. This allows accessing the same registers as GTTMMADR. The IOBAR can be used to issue writes to the GTTMMADR or the GTT Table.
  - GMADR – Internal graphics translation window (128 MB, 256 MB, 512 MB window).
  - GTTMMADR – This register requests a 4 MB allocation for combined Graphics Translation Table Modification Range and Memory Mapped Range. GTTADR will be at GTTMMADR + 2 MB while the MMIO base address will be the same as GTTMMADR
- Device 3, Function 0: (Audio Controller)

The rules for the above programmable ranges are:

1. For security reasons, the processor will now positively decode (FFE0_0000h to FFFF_FFFFh) to DMI. This ensures the boot vector and BIOS execute off the PCH.
2. ALL of these ranges MUST be unique and NON-OVERLAPPING. It is the BIOS or system designer's responsibility to limit memory population so that adequate PCI, PCI Express, High BIOS, PCI Express Memory Mapped space, and APIC memory space can be allocated.
3. In the case of overlapping ranges with memory, the memory decode will be given priority. This is an Intel® Trusted Execution Technology (Intel® TXT) requirement. It is necessary to get Intel TXT protection checks, avoiding potential attacks.
4. There are NO Hardware Interlocks to prevent problems in the case of overlapping ranges.
5. Accesses to overlapped ranges may produce indeterminate results.
6. The only peer-to-peer cycles allowed below the Top of Low Usable memory (register TOLUD) are DMI Interface to PCI Express VGA range writes. Peer-to-peer cycles to the Internal Graphics VGA range are not supported.

Figure 2. System Address Range Example

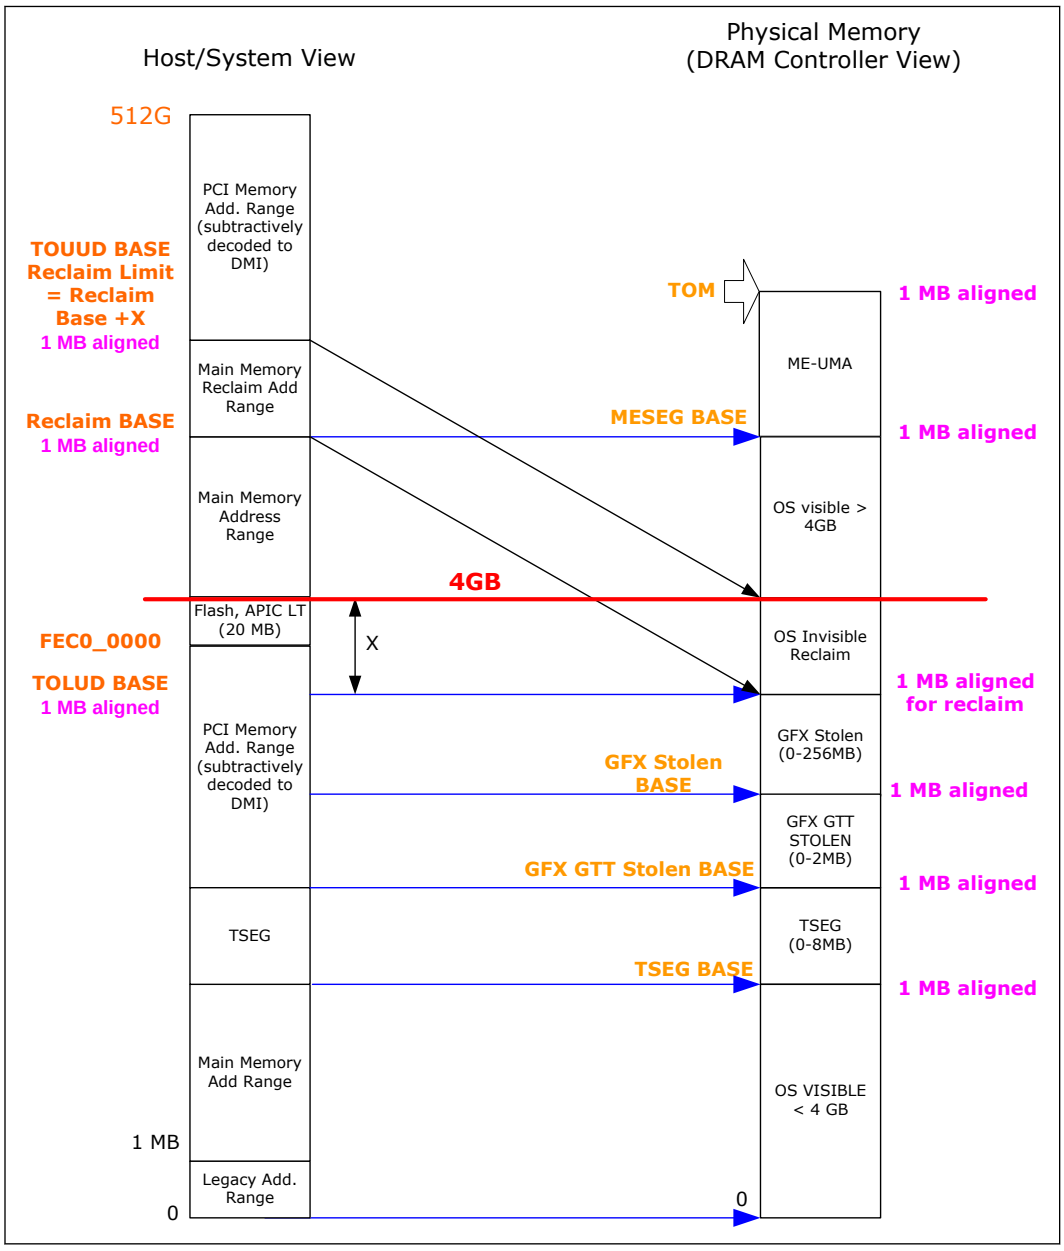

2.4 Legacy Address Range

The memory address range from 0 to 1 MB is known as Legacy Address. This area is divided into the following address regions:

- 0 – 640 KB - DOS Area
- 640 – 768 KB - Legacy Video Buffer Area
- 768 – 896 KB in 16 KB sections (total of 8 sections) – Expansion Area
- 896 – 960 KB in 16 KB sections (total of 4 sections) – Extended System BIOS Area

- 960 KB – 1 MB Memory, System BIOS Area

The area between 768 KB – 1 MB is also collectively referred to as PAM (Programmable Address Memory). All accesses to the DOS and PAM ranges from any device are sent to DRAM. However, access to the legacy video buffer area is treated differently.

Assumption: GT never sends requests in the Legacy Address Range; thus, there is no blocking of GT requests to this range in the System Agent.

Figure 3. DOS Legacy Address Range

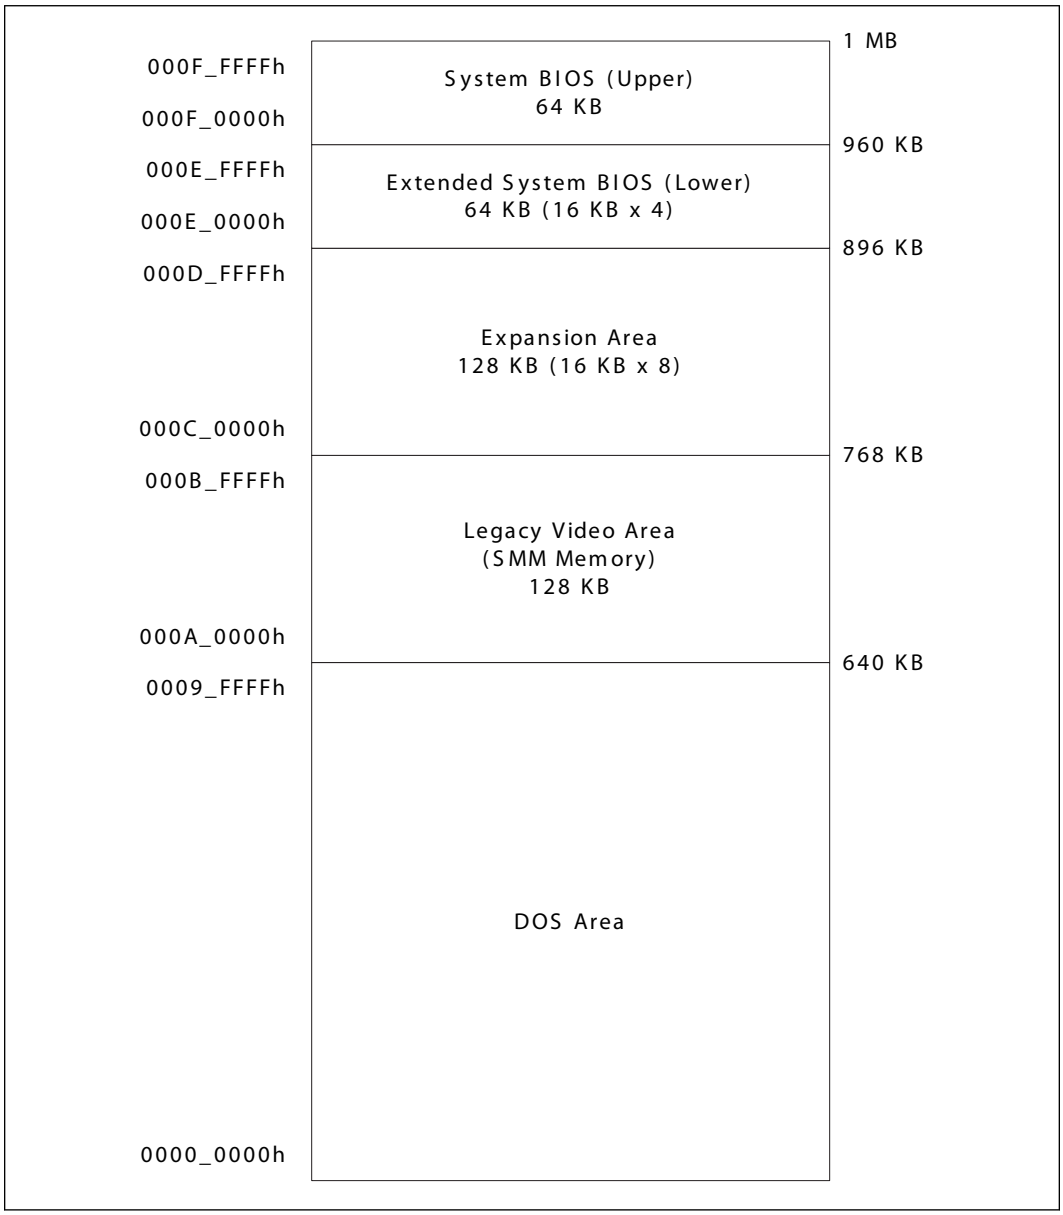

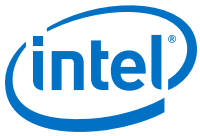

DOS Range (0h – 9_FFFFh)

The DOS area is 640 KB (0000_0000h – 0009_FFFFh) in size and is always mapped to the main memory.

Legacy Video Area / Compatible SMRAM Area (A_0000h – B_FFFFh)

The same address region is used for both Legacy Video Area and Compatible SMRAM.

- Legacy Video Area: The legacy 128 KB VGA memory range, frame buffer, at 000A_0000h – 000B_FFFFh, can be mapped to IGD (Device 2), to PCI Express (Device 1), and/or to the DMI Interface.
- Monochrome Adapter (MDA) Range: Legacy support requires the ability to have a second graphics controller (monochrome) in the system. The monochrome adapter may be mapped to IGD, PCI Express or DMI. Like the Legacy Video Area, decode priority is given first to IGD, then to PCI Express, and finally to DMI.
- Compatible SMRAM Address Range:

Legacy Video Area

The legacy 128 KB VGA memory range, frame buffer at 000A_0000h – 000B_FFFFh, can be mapped to IGD (Device 2), to PCI Express (Device 1), and/or to the DMI Interface.

Monochrome Adapter (MDA) Range

Legacy support requires the ability to have a second graphics controller (monochrome) in the system. The monochrome adapter may be mapped to IGD, PCI Express or DMI. Like the Legacy Video Area, decode priority is given first to IGD, then to PCI Express, and finally to DMI.

Compatible SMRAM Address Range

When compatible SMM space is enabled, SMM-mode CBO accesses to this range route to physical system DRAM at 00_000A_0000h – 00_000B_FFFFh.

Non-SMM mode CBO accesses to this range are considered to be to the Video Buffer Area as described above. PCI Express and DMI originated cycles to SMM space are not supported and are considered to be to the Video Buffer Area.

The processor always positively decodes internally mapped devices, namely the IGD and PCI Express. Subsequent decoding of regions mapped to PCI Express or the DMI Interface depends on the Legacy VGA configuration bits (VGA Enable and MDAP). This region is also the default for SMM space.

Programmable Attribute Map (PAM) (C_0000h – F_FFFFh)

PAM is a legacy BIOS ROM area in MMIO. It is overlaid with DRAM and used as a faster ROM storage area. It has a fixed base address (000C_0000h) and fix size of 256 KB. The 13 sections from 768 KB to 1 MB comprise what is also known as the PAM Memory Area. Each section has Read enable and Write enable attributes.

Figure 4. PAM Region Space

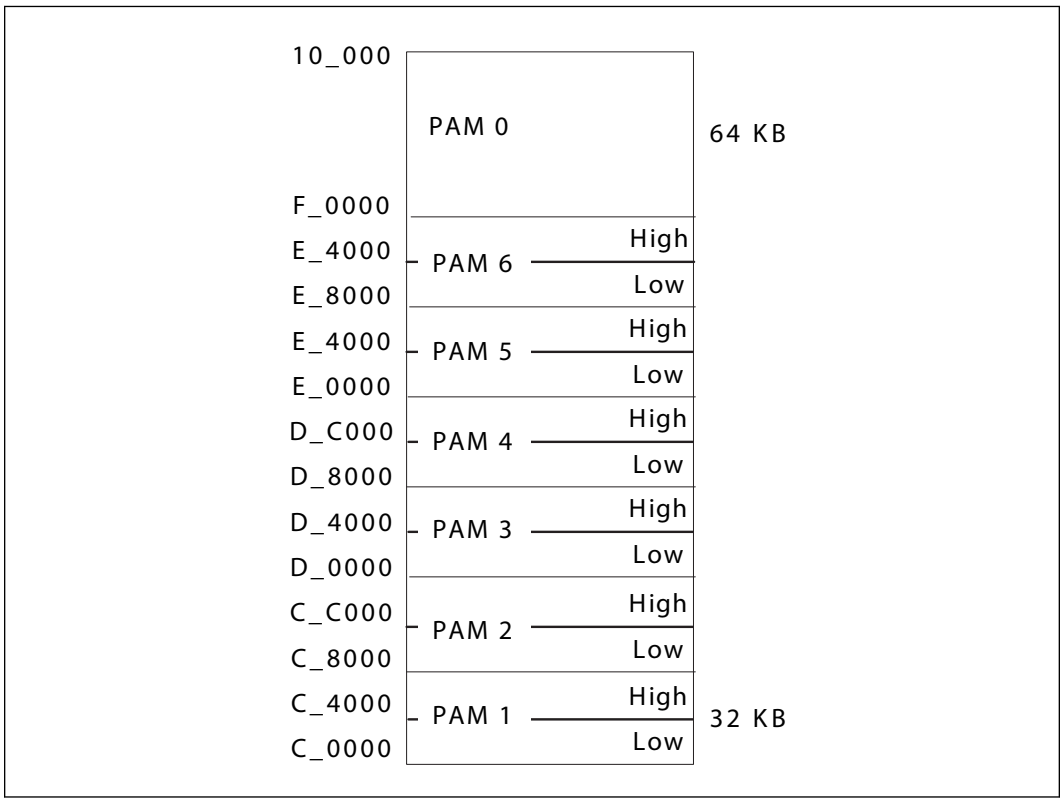

The PAM registers are mapped in Device 0 configuration space.

- ISA Expansion Area (C_0000h – D_FFFFh)
- Extended System BIOS Area (E_0000h – E_FFFFh)
- System BIOS Area (F_0000h – F_FFFFh)

The processor decodes the Core request, then routes to the appropriate destination (DRAM or DMI).

Snooped accesses from PCI Express or DMI to this region are snooped on processor Caches.

Non-snooped accesses from PCI Express or DMI to this region are always sent to DRAM.

Graphics translated requests to this region are not allowed. If such a mapping error occurs, the request will be routed to C_0000h. Writes will have the byte enables de-asserted.

2.5 Main Memory Address Range (1 MB – TOLUD)

This address range extends from 1 MB to the top of Low Usable physical memory that is permitted to be accessible by the processor (as programmed in the TOLUD register). The processor will route all addresses within this range to the DRAM unless it falls into the optional TSEG, optional ISA Hole, or optional IGD stolen VGA memory.

This address range is divided into two sub-ranges:

- 1 MB to TSEGMB
- TSEGMB to TOULUD

TSEGMB indicates the TSEG Memory Base address.

Figure 5. Main Memory Address Range

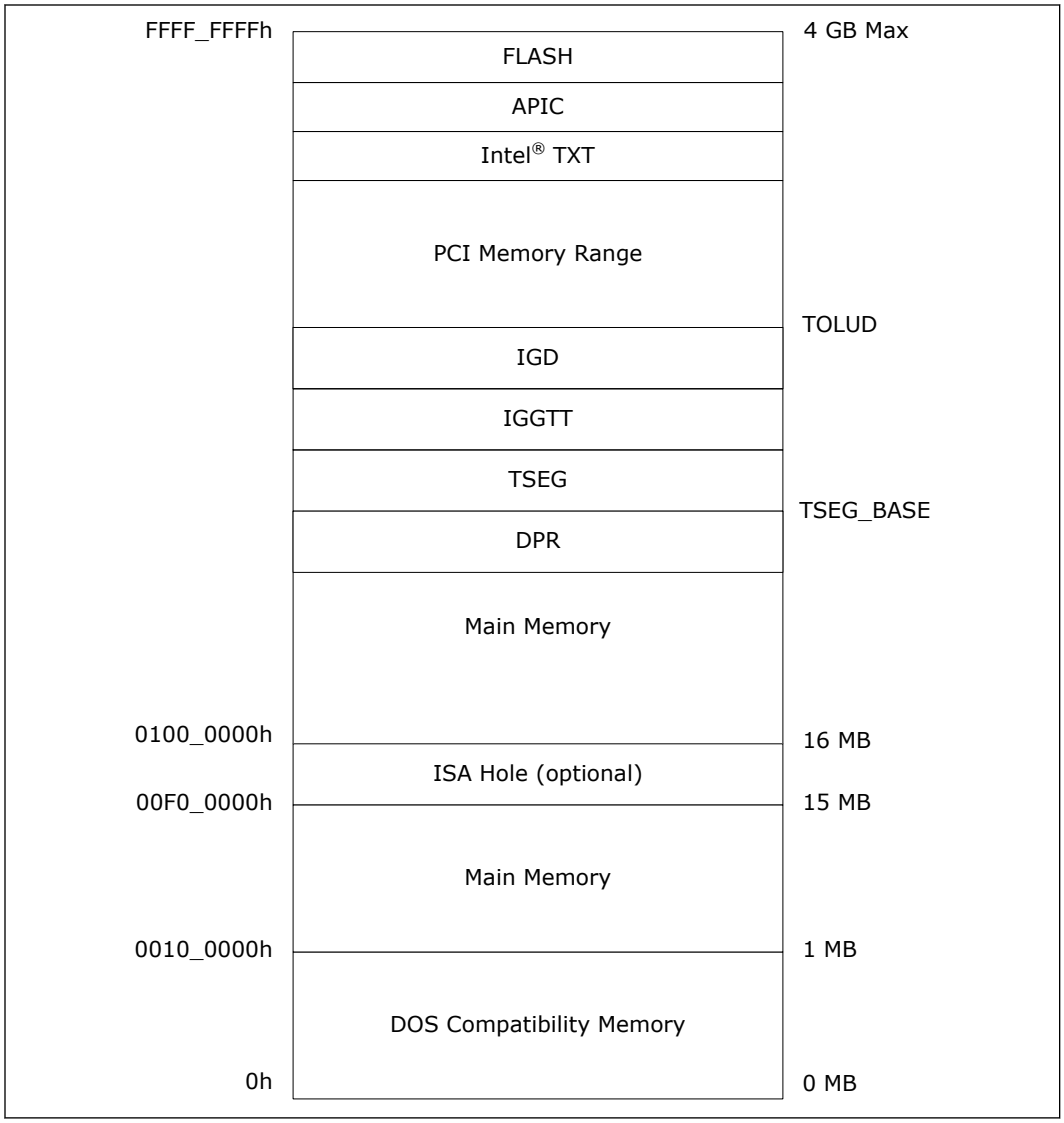

ISA Hole (15 MB –16 MB)

The ISA Hole (starting at address F0_0000h) is enabled in the Legacy Access Control Register in Device 0 configuration space. If no hole is created, the processor will route the request to DRAM. If a hole is created, the processor will route the request to DMI, since the request does not target DRAM. These downstream requests will be sent to DMI (subtractive decoding).

Graphics translated requests to the range will always route to DRAM.

1 MB to TSEGMB

Processor access to this range will be directed to memory, unless the ISA Hole is enabled.

TSEG

For processor initiated transactions, the processor relies on correct programming of SMM Range Registers (SMRR) to enforce TSEG protection.

TSEG is below IGD stolen memory, which is at the Top of Low Usable physical memory (TOLUD). BIOS will calculate and program the TSEG BASE in Device 0 (TSEGMB), used to protect this region from DMA access. Calculation is:

TSEGMB = TOLUD – DSM SIZE – GSM SIZE – TSEG SIZE

SMM-mode processor accesses to enabled TSEG access the physical DRAM at the same address.

When the extended SMRAM space is enabled, processor accesses to the TSEG range without SMM attribute or without WB attribute are handled by the processor as invalid accesses.

Non-processor originated accesses are not allowed to SMM space. PCI-Express, DMI, and Internal Graphics originated cycles to enabled SMM space are handled as invalid cycle type with reads and writes to location C_0000h and byte enables turned off for writes.

Protected Memory Range (PMR) - (programmable)

For robust and secure launch of the MVMM, the MVMM code and private data need to be loaded to a memory region protected from bus master accesses. Support for protected memory region is required for DMA-remapping hardware implementations on platforms supporting Intel TXT, and is optional for non-Intel TXT platforms. Since the protected memory region needs to be enabled before the MVMM is launched, hardware must support enabling of the protected memory region independently from enabling the DMA-remapping hardware.

As part of the secure launch process, the SINIT-AC module verifies the protected memory regions are properly configured and enabled. Once launched, the MVMM can setup the initial DMA-remapping structures in protected memory (to ensure they are protected while being setup) before enabling the DMA-remapping hardware units.

To optimally support platform configurations supporting varying amounts of main memory, the protected memory region is defined as two non-overlapping regions:

- **Protected Low-memory Region**: This is defined as the protected memory region below 4 GB to hold the MVMM code/private data, and the initial DMA-remapping structures that control DMA to host physical addresses below 4 GB. DMA-remapping hardware implementations on platforms supporting Intel TXT are required to support protected low-memory region 5.
- **Protected High-memory Region**: This is defined as a variable sized protected memory region above 4 GB, enough to hold the initial DMA-remapping structures for managing DMA accesses to addresses above 4 GB. DMA-remapping hardware

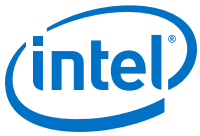

implementations on platforms supporting Intel TXT are required to support protected high-memory region 6, if the platform supports main memory above 4 GB.

Once the protected low/high memory region registers are configured, bus master protection to these regions is enabled through the Protected Memory Enable register. For platforms with multiple DMA-remapping hardware units, each of the DMA-remapping hardware units must be configured with the same protected memory regions and enabled.

DRAM Protected Range (DPR)

This protection range only applies to DMA accesses and GMADR translations. It serves a purpose of providing a memory range that is only accessible to processor streams. The range just below TSEGMB is protected from DMA accesses.

The DPR range works independent of any other range, including the PMRC checks in Intel VT-d. It occurs post any Intel VT-d translation. Therefore, incoming cycles are checked against this range after the Intel VT-d translation and faulted if they hit this protected range, even if they passed the Intel VT-d translation.

The system will set up:

- 0 to (TSEG_BASE – DPR size – 1) for DMA traffic
- TSEG_BASE to (TSEG_BASE – DPR size) as no DMA.

After some time, software could request more space for not allowing DMA. It will get some more pages and make sure there are no DMA cycles to the new region. DPR size is changed to the new value. When it does this, there should not be any DMA cycles going to DRAM to the new region.

If there were cycles from a rogue device to the new region, then those cycles could use the previous decode until the new decode can ensure PV. No flushing of cycles is required. On a clock-by-clock basis proper decode with the previous or new decode needs to be ensured.

All upstream cycles from 0 to (TSEG_BASE – 1 – DPR size), and not in the legacy holes (VGA), are decoded to DRAM.

Because Bus Master cycles can occur when the DPR size is changed, the DPR size needs to be treated dynamically.

Pre-allocated Memory

Voids of physical addresses that are not accessible as general system memory and reside within the system memory address range (< TOLUD) are created for SMM-mode, legacy VGA graphics compatibility, and GFX GTT stolen memory. **It is the responsibility of BIOS to properly initialize these regions**.

2.6 Main Memory Address Space (4 GB to TOUUD)

The maximum main memory size supported is 32 GB total DRAM memory.

A hole between TOLUD and 4 GB occurs when main memory size approaches 4 GB or larger. As a result, TOM and TOUUD registers and REMAPBASE/REMAPLIMIT registers become relevant.

The remap configuration registers exist to remap lost main memory space. The greater than 32-bit remap handling will be handled similar to other MCHs.

Upstream read and write accesses above 39-bit addressing will be treated as invalid cycles by PEG and DMI.

Top of Memory (TOM)

The "Top of Memory" (TOM) register reflects the total amount of populated physical memory. This is NOT necessarily the highest main memory address (holes may exist in main memory address map due to addresses allocated for memory mapped IO above TOM).

On Front Side Bus (FSB) chipsets, the TOM was used to allocate the Intel Management Engine (Intel ME) stolen memory. The Intel ME stolen size register reflects the total amount of physical memory stolen by the Intel ME. The Intel ME stolen memory is located at the top of physical memory. The Intel ME stolen memory base is calculated by subtracting the amount of memory stolen by the Intel ME from TOM.

Top of Upper Usable DRAM (TOUUD)

The Top of Upper Usable DRAM (TOUUD) register reflects the total amount of addressable DRAM. If remap is disabled, TOUUD will reflect TOM minus Intel ME stolen size. If remap is enabled, then it will reflect the remap limit. When there is more than 4 GB of DRAM and reclaim is enabled, the reclaim base will be the same as TOM minus Intel ME stolen memory size to the nearest 1 MB alignment.

Top of Low Usable DRAM (TOLUD)

TOLUD register is restricted to 4 GB memory (A[31:20]), but the processor can support up to 32 GB, limited by DRAM pins. For physical memory greater than 4 GB, the TOUUD register helps identify the address range between the 4 GB boundary and the top of physical memory. This identifies memory that can be directly accessed (including remap address calculation) that is useful for memory access indication and early path indication. TOLUD can be 1 MB aligned.

TSEG_BASE

The "TSEG_BASE" register reflects the total amount of low addressable DRAM, below TOLUD. BIOS will calculate memory size and program this register; thus, the system agent has knowledge of where (TOLUD) – (Gfx stolen) – (Gfx GTT stolen) – (TSEG) is located. I/O blocks use this minus DPR for upstream DRAM decode.

Memory Re-claim Background

The following are examples of Memory Mapped IO devices that are typically located below 4 GB:

- High BIOS
- TSEG
- GFX stolen
- GTT stolen
- XAPIC
- Local APIC
- MSI Interrupts

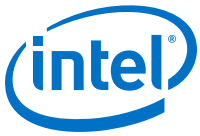

- Mbase/Mlimit
- Pmbase/PMlimit
- Memory Mapped IO space that supports only 32B addressing

The processor provides the capability to re-claim the physical memory overlapped by the Memory Mapped IO logical address space. The MCH re-maps physical memory from the Top of Low Memory (TOLUD) boundary up to the 4 GB boundary to an equivalent sized logical address range located just below the Intel ME stolen memory.

Indirect Accesses to MCHBAR Registers

Similar to prior chipsets, MCHBAR registers can be indirectly accessed using:

- Direct MCHBAR access decode:
  - Cycle to memory from processor
  - Hits MCHBAR base, AND
  - MCHBAR is enabled, AND
  - Within MMIO space (above and below 4 GB)
- GTTMMADR (10000h – 13FFFh) range -> MCHBAR decode:
  - Cycle to memory from processor, AND
  - Device 2 (IGD) is enabled, AND
  - Memory accesses for device 2 is enabled, AND
  - Targets GFX MMIO Function 0, AND
  - MCHBAR is enabled or cycle is a read. If MCHBAR is disabled, only read access is allowed.
- MCHTMBAR -> MCHBAR (Thermal Monitor)
  - Cycle to memory from processor, AND
  - Targets MCHTMBAR base
- IOBAR -> GTTMMADR -> MCHBAR.
  - Follows IOBAR rules. See GTTMMADR information above as well.

Memory Remapping

An incoming address (referred to as a logical address) is checked to see if it falls in the memory re-map window. The bottom of the re-map window is defined by the value in the REMAPBASE register. The top of the re-map window is defined by the value in the REMAPLIMIT register. An address that falls within this window is re-mapped to the physical memory starting at the address defined by the TOLUD register. The TOLUD register must be 1 MB aligned.

Hardware Remap Algorithm

The following pseudo-code defines the algorithm used to calculate the DRAM address to be used for a logical address above the top of physical memory made available using re-claiming.

```
IF (ADDRESS_IN[38:20] >= REMAP_BASE[35:20]) AND
(ADDRESS_IN[38:20] <= REMAP_LIMIT[35:20]) THEN
     ADDRESS_OUT[38:20] = (ADDRESS_IN[38:20] - REMAP_BASE[35:20]) +
0000000b & TOLUD[31:20]
     ADDRESS_OUT[19:0] = ADDRESS_IN[19:0]
```

2.7 PCI Memory Address Range (TOLUD – 4 GB)

Top of Low Usable DRAM (TOLUD) – TOLUD is restricted to 4 GB memory (A[31:20]), but the System Agent may support up to a much higher capacity, which is limited by DRAM pins.

This address range from the top of low usable DRAM (TOLUD) to 4 GB is normally mapped to the DMI Interface.

Device 0 exceptions are:

1. Addresses decoded to the egress port registers (PXPEPBAR)

2. Addresses decoded to the memory mapped range for internal MCH registers (MCHBAR)

3. Addresses decoded to the registers associated with the MCH/PCH Serial Interconnect (DMI) register memory range. (DMIBAR)

For each PCI Express* port, there are two exceptions to this rule:

4. Addresses decoded to the PCI Express Memory Window defined by the MBASE, MLIMIT registers are mapped to PCI Express.

5. Addresses decoded to the PCI Express prefetchable Memory Window defined by the PMBASE, PMLIMIT registers are mapped to PCI Express.

In integrated graphics configurations, there are exceptions to this rule:

6. Addresses decode to the internal graphics translation window (GMADR)

7. Addresses decode to the internal graphics translation table or IGD registers. (GTTMMADR)

In an Intel VT enable configuration, there are exceptions to this rule:

8. Addresses decoded to the memory mapped window to Graphics Intel VT remap engine registers (GFXVTBAR)

9. Addresses decoded to the memory mapped window to DMI VC1 Intel VT remap engine registers (DMIVC1BAR)

10. Addresses decoded to the memory mapped window to PEG/DMI VC0 Intel VT remap engine registers (VTDPVC0BAR)

11. TCm accesses (to Intel ME stolen memory) from PCH do not go through Intel VT remap engines.

Some of the MMIO Bars may be mapped to this range or to the range above TOUUD.

There are sub-ranges within the PCI memory address range defined as APIC Configuration Space, MSI Interrupt Space, and High BIOS address range. The exceptions listed above for internal graphics and the PCI Express ports ***MUST NOT*** *overlap with these ranges.*

Figure 6. PCI Memory Address Range

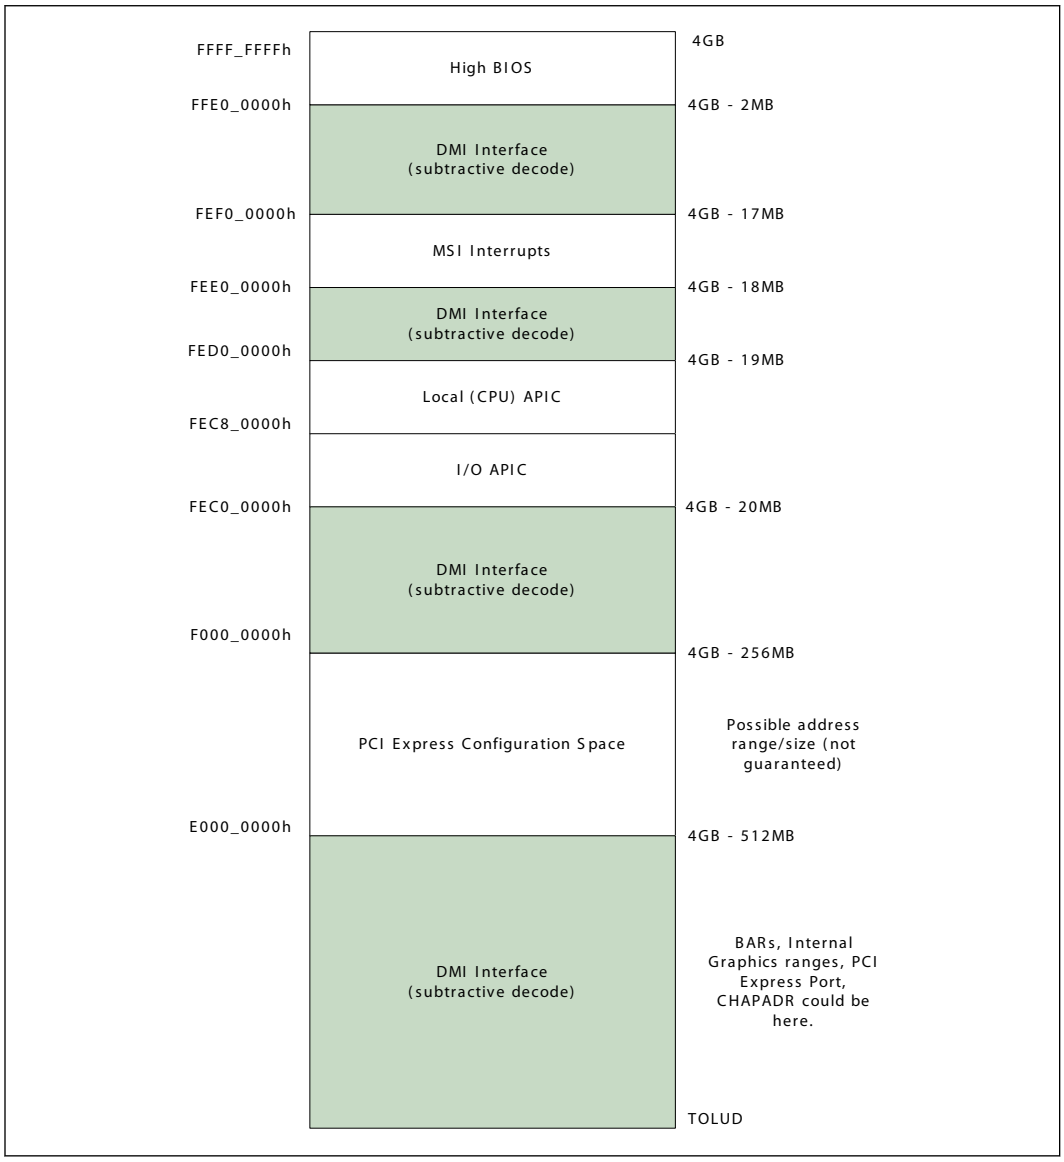

APIC Configuration Space (FEC0_0000h – FECF_FFFFh)

This range is reserved for APIC configuration space. The I/O APIC(s) usually reside in the PCH portion of the chipset, but may also exist as stand-alone components like PXH.

The IOAPIC spaces are used to communicate with IOAPIC interrupt controllers that may be populated in the system. Since it is difficult to relocate an interrupt controller using plug-and-play software, fixed address decode regions have been allocated for them. Processor accesses to the default IOAPIC region (FEC0_0000h to FEC7_FFFFh) are always forwarded to DMI.

The processor optionally supports additional I/O APICs behind the PCI Express* "Graphics" port. When enabled using the APIC_BASE and APIC_LIMIT registers (mapped PCI Express* Configuration space offset 240h and 244h), the PCI Express* port(s) will positively decode a subset of the APIC configuration space.

Memory requests to this range would then be forwarded to the PCI Express* port. This mode is intended for the entry Workstation/Server SKU of the PCH, and would be disabled in typical Desktop systems. When disabled, any access within the entire APIC Configuration space (FEC0_0000h to FECF_FFFFh) is forwarded to DMI.

HSEG (FEDA_0000h – FEDB_FFFFh)

This decode range is not supported on this processor platform.

MSI Interrupt Memory Space (FEE0_0000h – FEEF_FFFFh)

Any PCI Express* or DMI device may issue a Memory Write to 0FEEx_xxxxh. This Memory Write cycle does not go to DRAM. The system agent will forward this Memory Write along with the data to the processor as an Interrupt Message Transaction.

High BIOS Area

For security reasons, the processor will positively decode this range to DMI. This positive decode ensures any overlapping ranges will be ignored. This ensures that the boot vector and BIOS execute off the PCH.

The top 2 MB (FFE0_0000h – FFFF_FFFFh) of the PCI Memory Address Range is reserved for System BIOS (High BIOS), extended BIOS for PCI devices, and the A20 alias of the system BIOS.

The processor begins execution from the High BIOS after reset. This region is positively decoded to DMI. The actual address space required for the BIOS is less than 2 MB. However, the minimum processor MTRR range for this region is 2 MB; thus, the full 2 MB must be considered.

2.8 PCI Express* Configuration Address Space

PCIEXBAR is located in Device 0 configuration space as in Front Side Bus (FSB) platforms. The processor detects memory accesses targeting PCIEXBAR. BIOS must assign this address range such that it will not conflict with any other address ranges.

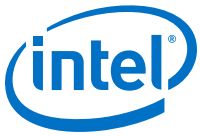

2.9 PCI Express* Graphics Attach (PEG)

The processor can be programmed to direct memory accesses to a PCI Express interface. Addresses are within either of two ranges specified using registers in each PEG(s) configuration space.

- The first range is controlled using the Memory Base Register (MBASE) and Memory Limit Register (MLIMIT) registers.
- The second range is controlled using the Pre-fetchable Memory Base (PMBASE) and Pre-fetchable Memory Limit (PMLIMIT) registers.

Conceptually, address decoding for each range follows the same basic concept. The top 12 bits of the respective Memory Base and Memory Limit registers correspond to address bits A[31:20] of a memory address . For the purpose of address decoding, the processor assumes that address bits A[19:0] of the memory base are zero and that address bits A[19:0] of the memory limit address are F_FFFFh. This forces each memory address range to be aligned to 1 MB boundary and to have a size granularity of 1 MB.

The processor positively decodes memory accesses to PCI Express memory address space as defined by the following equations

- Memory_Base_Address ≤ Address ≤ Memory_Limit_Address
- Prefetchable_Memory_Base_Address ≤ Address ≤ Prefetchable_Memory_Limit_Address

The window size is programmed by the plug-and-play configuration software. The window size depends on the size of memory claimed by the PCI Express device. Normally these ranges will reside above the Top-of-Low Usable-DRAM and below High BIOS and APIC address ranges. They MUST reside above the top of low memory (TOLUD) if they reside below 4 GB and MUST reside above top of upper memory (TOUUD) if they reside above 4 GB or they will steal physical DRAM memory space.

It is essential to support a separate Pre-fetchable range in order to apply USWC attribute (from the processor point of view) to that range. The USWC attribute is used by the processor for write combining.

The processor memory range registers described above are used to allocate memory address space for any PCI Express devices on PCI Express that require such a window.

The PCICMD register can override the routing of memory accesses to PCI Express. In other words, the memory access enable bit must be set to enable the memory base/ limit and pre-fetchable base/limit windows.

The upper PMUBASE/PMULIMIT registers are implemented for PCI Express Specification compliance. The processor locates MMIO space above 4 GB using these registers.

2.10 Graphics Memory Address Ranges

The integrated memory controller can be programmed to direct memory accesses to the IGD when addresses are within any of the ranges specified using registers in MCH Device 2 configuration space.

- The Graphics Memory Aperture Base Register (GMADR) is used to access graphics memory allocated using the graphics translation table.

- The Graphics Translation Table Base Register (GTTADR) is used to access the translation table and graphics control registers. This is part of the GTTMMADR register.

These ranges can reside above the Top-of-Low-DRAM and below High BIOS and APIC address ranges. They MUST reside above the top of memory (TOLUD) and below 4 GB so they do not take any physical DRAM memory space.

Alternatively, these ranges can reside above 4 GB, similar to other BARs that are larger than 32 bits in size.

GMADR is a Prefetchable range in order to apply USWC attribute (from the processor point of view) to that range. The USWC attribute is used by the processor for write combining.

IOBAR Mapped Access to Device 2 MMIO Space

Device 2, integrated graphics device, contains an IOBAR register. If Device 2 is enabled, IGD registers or the GTT table can be accessed using this IOBAR. The IOBAR is composed of an index register and a data register.

MMIO_Index: MMIO_INDEX is a 32-bit register. A 32-bit (all bytes enabled) I/O write to this port loads the offset of the MMIO register or offset into the GTT that needs to be accessed. An I/O Read returns the current value of this register. I/O read/write accesses less than 32 bits in size (all bytes enabled) will not target this register.

MMIO_Data: MMIO_DATA is a 32-bit register. A 32-bit (all bytes enabled) I/O write to this port is re-directed to the MMIO register pointed to by the MMIO-index register. An I/O read to this port is re-directed to the MMIO register pointed to by the MMIO-index register. I/O read/write accesses less than 32 bits in size (all bytes enabled) will not target this register.

The result of accesses through IOBAR can be:

- Accesses directed to the GTT table. (that is, route to DRAM)
- Accesses to internal graphics registers with the device.
- Accesses to internal graphics display registers now located within the PCH. (that is, route to DMI).

Note: GTT table space writes (GTTADR) are supported through this mapping mechanism.

This mechanism to access internal graphics MMIO registers MUST NOT be used to access VGA I/O registers that are mapped through the MMIO space. VGA registers must be accessed directly through the dedicated VGA I/O ports.

Trusted Graphics Ranges

Trusted graphics ranges are NOT supported.

2.11 System Management Mode (SMM)

Unlike Front Side Bus (FSB) platforms, the Core handles all SMM mode transaction routing. The platform does not support HSEG, and the processor will does not allow I/O devices access to CSEG/TSEG/HSEG ranges.

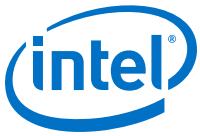

DMI Interface and PCI Express* masters are Not allowed to access the SMM space.

Table 5. SMM Regions

SMM Space Enabled	Transaction Address Space	DRAM Space (DRAM)
Compatible (C)	000A_0000h to 000B_FFFFh	000A_0000h to 000B_FFFFh
TSEG (T)	(TOLUD – STOLEN – TSEG) to TOLUD – STOLEN	(TOLUD – STOLEN – TSEG) to TOLUD – STOLEN

2.12 SMM and VGA Access Through GTT TLB

Accesses through GTT TLB address translation SMM DRAM space are not allowed. Writes will be routed to memory address 000C_0000h with byte enables de-asserted and reads will be routed to Memory address 000C_0000h. If a GTT TLB translated address hits SMM DRAM space, an error is recorded in the PGTBL_ER register.

PCI Express* and DMI Interface originated accesses are **never** allowed to access SMM space directly or through the GTT TLB address translation. If a GTT TLB translated address hits enabled SMM DRAM space, an error is recorded in the PGTBL_ER register.

PCI Express and DMI Interface write accesses through the GMADR range will not be snooped. Only PCI Express and DMI assesses to GMADR linear range (defined using fence registers) are supported. PCI Express and DMI Interface tileY and tileX writes to GMADR are not supported. If, when translated, the resulting physical address is to enable SMM DRAM space, the request will be remapped to address 000C_0000h with de-asserted byte enables.

PCI Express and DMI Interface read accesses to the GMADR range are not supported; therefore, there are no address translation concerns. PCI Express and DMI Interface reads to GMADR will be remapped to address 000C_0000h. The read will complete with UR (unsupported request) completion status.

GTT fetches are always decoded (at fetch time) to ensure fetch is not in SMM (actually, anything above base of TSEG or 640 KB - 1 MB). Thus, the fetches will be invalid and go to address 000C_0000h. This is not specific to PCI Express or DMI; it also applies to processor or internal graphics engines.

2.13 Intel® Management Engine (Intel® ME) Stolen Memory Accesses

There are two ways to validly access Intel ME stolen memory:

- PCH accesses mapped to VCm will be decoded to ensure only Intel ME stolen memory is targeted. These VCm accesses will route non-snooped directly to DRAM. This is the means by which the Intel ME (located within the PCH) is able to access the Intel ME stolen range.
- The display engine is allowed to access Intel ME stolen memory as part of Intel® KVM technology flows. Specifically, display-initiated HHP reads (for displaying a Intel KVM technology frame) and display initiated LP non-snoop writes (for display writing an Intel KVM technology captured frame) to Intel ME stolen memory are allowed.

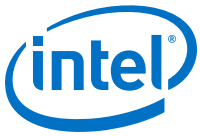

2.14 I/O Address Space

The system agent generates either DMI Interface or PCI Express* bus cycles for all processor I/O accesses that it does not claim. The Configuration Address Register (CONFIG_ADDRESS) and the Configuration Data Register (CONFIG_DATA) are used to generate PCI configuration space access.

The processor allows 64K+3 bytes to be addressed within the I/O space. The upper 3 locations can be accessed only during I/O address wrap-around when address bit 16 is asserted. Address bit 16 is asserted on the processor bus whenever an I/O access is made to 4 bytes from address 0FFFDh, 0FFFEh, or 0FFFFh. Address bit 16 is also asserted when an I/O access is made to 2 bytes from address 0FFFFh.

A set of I/O accesses are consumed by the internal graphics device if it is enabled. The mechanisms for internal graphics I/O decode and the associated control is explained in following sub-sections.

The I/O accesses are forwarded normally to the DMI Interface bus unless they fall within the PCI Express I/O address range as defined by the mechanisms explained below. I/O writes are NOT posted. Memory writes to PCH or PCI Express are posted. The PCI Express devices have a register that can disable the routing of I/O cycles to the PCI Express device.

The processor responds to I/O cycles initiated on PCI Express or DMI with an UR status. Upstream I/O cycles and configuration cycles should never occur. If one does occur, the transaction will complete with an UR completion status.

Similar to Front Side Bus (FSB) processors, I/O reads that lie within 8-byte boundaries but cross 4-byte boundaries are issued from the processor as one transaction. The reads will be split into two separate transactions. I/O writes that lie within 8-byte boundaries but cross 4-byte boundaries will be split into two transactions by the processor.

PCI Express* I/O Address Mapping

The processor can be programmed to direct non-memory (I/O) accesses to the PCI Express bus interface when processor initiated I/O cycle addresses are within the PCI Express I/O address range. This range is controlled using the I/O Base Address (IOBASE) and I/O Limit Address (IOLIMIT) registers in Device 1 Functions 0, 1, 2 configuration space.

Address decoding for this range is based on the following concept. The top 4 bits of the respective I/O Base and I/O Limit registers correspond to address bits A[15:12] of an I/O address. For the purpose of address decoding, the device assumes that the lower 12 address bits A[11:0] of the I/O base are zero and that address bits A[11:0] of the I/O limit address are FFFh. This forces the I/O address range alignment to a 4 KB boundary and produces a size granularity of 4 KB.

The processor positively decodes I/O accesses to PCI Express I/O address space as defined by the following equation:

```
I/O_Base_Address ≤ processor I/O Cycle Address ≤ I/
O_Limit_Address
```

The effective size of the range is programmed by the plug-and-play configuration software and it depends on the size of I/O space claimed by the PCI Express device.

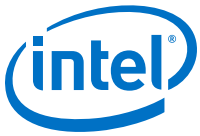

The processor also forwards accesses to the Legacy VGA I/O ranges according to the settings in the PEG configuration registers BCTRL (VGA Enable) and PCICMD (IOAE), unless a second adapter (monochrome) is present on the DMI Interface/PCI (or ISA). The presence of a second graphics adapter is determined by the MDAP configuration bit. When MDAP is set to 1, the processor will decode legacy monochrome I/O ranges and forward them to the DMI Interface. The I/O ranges decoded for the monochrome adapter are 3B4h, 3B5h, 3B8h, 3B9h, 3BAh, and 3BFh.

The PCICMD register can disable the routing of I/O cycles to PCI Express.

2.15 Direct Media Interface (DMI) Interface Decode Rules

All "SNOOP semantic" PCI Express* transactions are kept coherent with processor caches.

All "Snoop not required semantic" cycles must reference the main DRAM address range. PCI Express non-snoop initiated cycles are not snooped.

The processor accepts accesses from the DMI Interface to the following address ranges:

- All snoop memory read and write accesses to Main DRAM including PAM region (except stolen memory ranges, TSEG, A0000h – BFFFFh space)
- Write accesses to enabled VGA range, MBASE/MLIMIT, and PMBASE/PMLIMIT will be routed as peer cycles to the PCI Express interface.
- Write accesses above the top of usable DRAM and below 4 GB (not decoding to PCI Express or GMADR space) will be treated as master aborts.
- Read accesses above the top of usable DRAM and below 4 GB (not decoding to PCI Express) will be treated as unsupported requests.
- Reads and accesses above the TOUUD will be treated as unsupported requests on VC0/VCp.

DMI Interface memory read accesses that fall between TOLUD and 4 GB are considered invalid and will master abort. These invalid read accesses will be reassigned to address 000C_0000h and dispatch to DRAM. Reads will return unsupported request completion. Writes targeting PCI Express space will be treated as peer-to-peer cycles.

There is a known usage model for peer writes from DMI to PEG. A video capture card can be plugged into the PCH PCI bus. The video capture card can send video capture data (writes) directly into the frame buffer on an external graphics card (writes to the PEG port). As a result, peer writes from DMI to PEG must be supported.

I/O cycles and configuration cycles are not supported in the upstream direction. The result will be an unsupported request completion status.

DMI Accesses to the Processor that Cross Device Boundaries

The processor does not support transactions that cross device boundaries. This should not occur because PCI Express transactions are not allowed to cross a 4 KB boundary.

For reads, the processor will provide separate completion status for each naturally-aligned 64-byte block or, if chaining is enabled, each 128-byte block. If the starting address of a transaction hits a valid address, the portion of a request that hits that target device (PCI Express or DRAM) will complete normally.

If the starting transaction address hits an invalid address, the entire transaction will be remapped to address 000C_0000h and dispatched to DRAM. A single unsupported request completion will result.

Traffic Class (TC) / Virtual Channel (VC) Mapping Details

- VC0 (enabled by default)
  - Snoop port and Non-snoop Asynchronous transactions are supported.
  - Internal Graphics GMADR writes can occur. Unlike Front Side Bus (FSB) chipsets, these writes will NOT be snooped regardless of the snoop not required (SNR) bit.
  - Internal Graphics GMADR reads (unsupported).
  - Peer writes can occur. The SNR bit is ignored.
  - MSI can occur. These will route and be sent to the cores as Intlogical/ IntPhysical interrupts regardless of the SNR bit.
  - VLW messages can occur. These will route and be sent to the cores as VLW messages regardless of the SNR bit.
  - MCTP messages can occur. These are routed in a peer fashion.
- VCp (Optionally enabled):
  - Supports priority snoop traffic only. This VC is given higher priority at the snoop VC arbiter. Routed as an independent virtual channel and treated independently within the Cache module. VCp snoops are indicated as "high priority" in the snoop priority field. USB classic and USB2 traffic are expected to use this channel. **Note:** On prior chipsets, this was termed "snoop isochronous" traffic. "Snoop isochronous" is now termed "priority snoop" traffic.
  - SNR bit is ignored.
  - MSI on VCP is supported.
  - Peer read and write requests are not supported. Writes will route to address 000C_0000h with byte enables de-asserted, while reads will route to address 000C_0000h and an unsupported request completion.
  - Internal Graphics GMADR writes are NOT supported. These writes will route to address 000C_0000h with byte enables de-asserted.
  - Internal Graphics GMADR reads are not supported.
  - See DMI2 TC mapping for expected TC to VCp mapping. This has changed from DMI to DMI2.
- VC1 (Optionally enabled)
  - Supports non-snoop transactions only. (Used for isochronous traffic). The PCI Express* Egress port (PXPEPBAR) must also be programmed appropriately.
  - The snoop not required (SNR) bit must be set. Any transaction with the SNR bit not set will be treated as an unsupported request.
  - MSI and peer transactions are treated as unsupported requests.
  - No "pacer" arbitration or TWRR arbitration will occur. Never remaps to different port. (PCH takes care of Egress port remapping). The PCH meters TCm Intel ME accesses and Intel® High Definition Audio (Intel® HD Audio) TC1 access bandwidth.

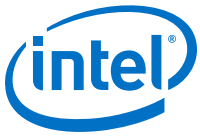

  - Internal Graphics GMADR writes and GMADR reads are not supported.
- VCm accesses
  - See the DMI2 specification for TC mapping to VCm. VCm access only map to Intel ME stolen DRAM. These transactions carry the direct physical DRAM address (no redirection or remapping of any kind will occur). This is how the PCH Intel ME accesses its dedicated DRAM stolen space.
  - DMI block will decode these transactions to ensure only Intel ME stolen memory is targeted, and abort otherwise.
  - VCm transactions will only route non-snoop.
  - VCm transactions will not go through VTd remap tables.
  - The remapbase/remaplimit registers to not apply to VCm transactions.

Figure 7. Example: DMI Upstream VC0 Memory Map

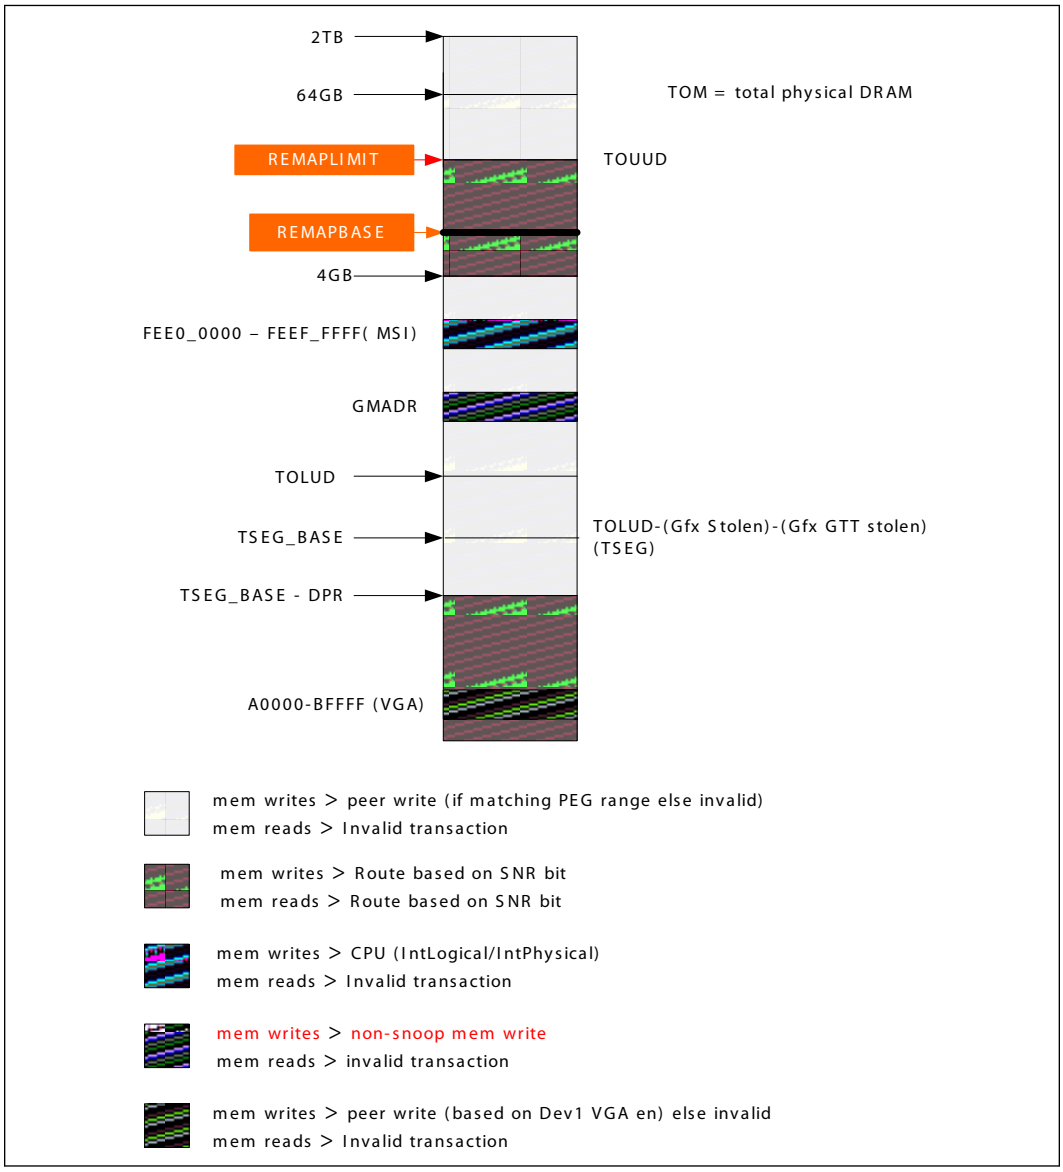

2.16 PCI Express* Interface Decode Rules

All "SNOOP semantic" PCI Express* transactions are kept coherent with processor caches. All "Snoop not required semantic" cycles must reference the direct DRAM address range. PCI Express non-snoop initiated cycles are not snooped. If a "Snoop not required semantic" cycle is outside of the address range mapped to system memory, then it will proceed as follows:

- Reads: Sent to DRAM address 000C_0000h (non-snooped) and will return "unsuccessful completion".

If PEG bus master enable is not set, all reads and writes are treated as unsupported requests.

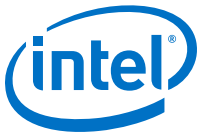

TC/VC Mapping Details

- VC0 (enabled by default)
  - — Snoop port and Non-snoop Asynchronous transactions are supported.
  - — Internal Graphics GMADR writes can occur. Unlike Front Side Bus (FSB) chipsets, these will NOT be snooped regardless of the snoop not required (SNR) bit.
  - — Internal Graphics GMADR reads (unsupported).
  - — Peer writes are only supported between PEG ports. PEG to DMI peer write accesses are NOT supported.
  - — MSI can occur. These will route to the cores (IntLogical/IntPhysical) regardless of the SNR bit.
- VC1 is not supported.
- VCm is not supported.

2.17 Legacy VGA and I/O Range Decode Rules

The legacy 128 KB VGA memory range 000A_0000h – 000B_FFFFh can be mapped to IGD (Device 2), PCI Express (Device 1 Functions), and/or to the DMI interface depending on the programming of the VGA steering bits. Priority for VGA mapping is constant in that the processor always decodes internally mapped devices first. Internal to the processor, decode precedence is always given to IGD. The processor always positively decodes internally mapped devices, namely the IGD. Subsequent decoding of regions mapped to either PCI Express port or the DMI Interface depends on the Legacy VGA configurations bits (VGA Enable and MDAP).

VGA range accesses will always be mapped as UC type memory.

Accesses to the VGA memory range are directed to IGD depend on the configuration. The configuration is specified by:

- Internal graphics controller in Device 2 is enabled (DEVEN.D2EN bit 4)
- Internal graphics VGA in Device 0 Function 0 is enabled through register GGC bit 1.
- IGD's memory accesses (PCICMD2 04h – 05h, MAE bit 1) in Device 2 configuration space are enabled.
- VGA compatibility memory accesses (VGA Miscellaneous Output register – MSR Register, bit 1) are enabled.
- Software sets the proper value for VGA Memory Map Mode register (VGA GR06 Register, bits 3:2). See the following table for translations.

Table 6. IGD Frame Buffer Accesses

Mem Access GR06(3:2)	A0000h - AFFFFh	B0000h - B7FFFh MDA	B8000h - BFFFFh
00	IGD	IGD	IGD
01	IGD	PCI Express bridge or DMI interface	PCI Express bridge or DMI interface
10	PCI Express bridge or DMI interface	IGD	PCI Express bridge or DMI interface
11	PCI Express bridge or DMI interface	PCI Express bridge or DMI interface	IGD

Note: Additional qualification within IGD comprehends internal MDA support. The VGA and MDA enabling bits detailed below control segments not mapped to IGD.

VGA I/O range is defined as addresses where A[15:0] are in the ranges 03B0h to 03BBh, and 03C0h to 03DFh. VGA I/O accesses are directed to IGD depends on the following configuration:

- Internal graphics controller in Device 2 is enabled through register DEVEN.D2EN bit 4.
- Internal graphics VGA in Device 0 Function 0 is enabled through register GGC bit 1.
- IGD's I/O accesses (PCICMD2 04 – 05h, IOAE bit 0) in Device 2 are enabled.
- VGA I/O decodes for IGD uses 16 address bits (15:0) there is no aliasing. This is different when compared to a bridge device (Device 1) that used only 10 address bits (A 9:0) for VGA I/O decode.
- VGA I/O input/output address select (VGA Miscellaneous Output register - MSR Register, bit 0) is used to select mapping of I/O access as defined in the following table.

Table 7. IGD VGA I/O Mapping

I/O Access MSRb0	3CX	3DX	3B0h – 3BBh	3BCh – 3BFh
0	IGD	PCI Express bridge or DMI interface	IGD	PCI Express bridge or DMI interface
1	IGD	IGD	PCI Express bridge or DMI interface	PCI Express bridge or DMI interface

Note: Additional qualification within IGD comprehends internal MDA support. The VGA and MDA enabling bits detailed below control ranges not mapped to IGD.

VGA Enable: Controls the routing of processor initiated transactions targeting VGA compatible I/O and memory address ranges. When this bit is set, the following processor accesses will be forwarded to the PCI Express:

- Memory accesses in the range 0A0000h to 0BFFFFh
- I/O addresses where A[9:0] are in the ranges 3B0h to 3BBh and 3C0h to 3DFh (including ISA address aliases – A[15:10] are not decoded)

When this bit is set to a "1":

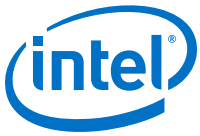

- Forwarding of these accesses issued by the processor is independent of the I/O address and memory address ranges defined by the previously defined base and limit registers.
- Forwarding of these accesses is also independent of the settings of the ISA Enable settings if this bit is "1".
- Accesses to I/O address range x3BCh – x3BFh are forwarded to the DMI Interface.

When this bit is set to a "0":

- Accesses to I/O address range x3BCh – x3BFh are treated like any other I/O accesses; the cycles are forwarded to PCI Express if the address is within IOBASE and IOLIMIT and ISA enable bit is not set. Otherwise, these accesses are forwarded to the DMI interface.
- VGA compatible memory and I/O range accesses are not forwarded to PCI Express but rather they are mapped to the DMI Interface, unless they are mapped to PCI Express using I/O and memory range registers defined above (IOBASE, IOLIMIT)

The following table shows the behavior for all combinations of MDA and VGA.

Table 8. VGA and MDA IO Transaction Mapping

VGA_en	MDAP	Range	Destination	Exceptions / Notes
0	0	VGA, MDA	DMI interface	
0	1	Illegal		Undefined behavior results
1	0	VGA	PCI Express	
1	1	VGA	PCI Express	
1	1	MDA	DMI interface	x3BCh – x3BEh will also go to DMI interface

The same registers control mapping of VGA I/O address ranges. The VGA I/O range is defined as addresses where A[9:0] are in the ranges 3B0h to 3BBh and 3C0h to 3DFh (inclusive of ISA address aliases – A[15:10] are not decoded). The function and interaction of these two bits is described below.

MDA Present (MDAP): This bit works with the VGA Enable bit in the BCTRL register of Device 1 to control the routing of processor-initiated transactions targeting MDA compatible I/O and memory address ranges. This bit should not be set when the VGA Enable bit is not set. If the VGA enable bit is set, accesses to I/O address range x3BCh – x3BFh are forwarded to the DMI Interface. If the VGA enable bit is not set, accesses to I/O address range x3BCh – x3BFh are treated just like any other I/O accesses; that is, the cycles are forwarded to PCI Express if the address is within IOBASE and IOLIMIT and the ISA enable bit is not set; otherwise, the accesses are forwarded to the DMI Interface. MDA resources are defined as the following:

Table 9. MDA Resources

Range Type	Address
Memory	0B0000h – 0B7FFFh
I/O	3B4h, 3B5h, 3B8h, 3B9h, 3BAh, 3BFh (Including ISA address aliases, A[15:10] are not used in decode)

Any I/O reference that includes the I/O locations listed above, or their aliases, will be forwarded to the DMI interface even if the reference includes I/O locations not listed above.

For I/O reads that are split into multiple DWord accesses, this decode applies to each DWord independently. For example, a read to x3B3h and x3B4h (quadword read to x3B0h with BE#=E7h) will result in a DWord read from PEG at 3B0h (BE#=Eh), and a DWord read from DMI at 3B4h (BE=7h). Since the processor will not issue I/O writes crossing the DWord boundary, this case does not exist for writes.

Summary of decode priority:

- Internal Graphics VGA, if enabled, gets:
  - 03C0h – 03CFh: always
  - 03B0h – 03BBh: if MSR[0]=0 (MSR is I/O register 03C2h)
  - 03D0h – 03DFh: if MSR[0]=1

Note: 03BCh – 03BFh never decodes to IGD; 3BCh – 3BEh are parallel port I/Os, and 3BFh is only used by true MDA devices.

- Else, if ISA Enable=1, DMI gets:
  - upper 768 bytes of each 1K block
- Else, IOBASE/IOLIMIT apply.

2.18 I/O Mapped Registers

The processor contains two registers that reside in the processor I/O address space - the Configuration Address (CONFIG_ADDRESS) Register and the Configuration Data (CONFIG_DATA) Register. The Configuration Address Register enables/disables the configuration space and determines what portion of configuration space is visible through the Configuration Data window.

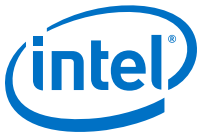

3.0 Host Device Configuration Registers

3.1 Host Bridge/DRAM Registers Summary

Offset	Register ID—Description	Default Value	Access
0	VID—Vendor Identification on page 48	8086h	RO
2	DID—Device Identification on page 48	0C00h	RO; RO_V
4	PCICMD—PCI Command on page 49	0006h	RO; RW
6	PCISTS—PCI Status on page 50	0090h	RO; RW1C
8	RID—Revision Identification on page 51	00h	RO
9	CC—Class Code on page 51	060000h	RO
E	HDR—Header Type on page 52	00h	RO
2C	SVID—Subsystem Vendor Identification on page 52	0000h	RW_O
2E	SID—Subsystem Identification on page 52	0000h	RW_O
34	CAPPTR—Capabilities Pointer on page 52	E0h	RO
40	PXPEPBAR—PCI Express Egress Port Base Address on page 52	0000000000000000h	RW
48	MCHBAR—Host Memory Mapped Register Range Base on page 53	0000000000000000h	RW
50	GGC—GMCH Graphics Control Register on page 54	0500h	RW_KL; RW_L
54	DEVEN—Device Enable on page 55	000000BFh	RO; RW_L; RW
58	PAVPC—Protected Audio Video Path Control on page 56	00000000h	RW_KL
5C	DPR—DMA Protected Range on page 57	00000000h	ROV; RW_L
60	PCIEXBAR—PCI Express Register Range Base Address on page 57	0000000000000000h	RW; RW_V
68	DMIBAR—Root Complex Register Range Base Address on page 58	0000000000000000h	RW
70	MESEG—Intel Management Engine Base Address Register on page 59	0000007FFFF00000h	RW_L
78	MESEG—Intel Management Engine Limit Address Register on page 59	0000000000000000h	RW_KL; RW_L
80	PAM0—Programmable Attribute Map 0 on page 60	00h	RW_KL; RW_L
81	PAM1—Programmable Attribute Map 1 on page 61	00h	RW_L
82	PAM2—Programmable Attribute Map 2 on page 62	00h	RW_L
83	PAM3—Programmable Attribute Map 3 on page 63	00h	RW_L
84	PAM4—Programmable Attribute Map 4 on page 64	00h	RW_L
85	PAM5—Programmable Attribute Map 5 on page 65	00h	RW_L
86	PAM6—Programmable Attribute Map 6 on page 66	00h	RW_L
87	LAC—Legacy Access Control on page 67	00h	RW

continued...

Offset	Register ID—Description	Default Value	Access
88	SMRAMC—System Management RAM Control on page 69	02h	RO; RW_L; RW_KL; RW_LV
90	REMAPBASE—Remap Base Address Register on page 70	0000007FFFF00000h	RW_KL; RW_L
98	REMAPLIMIT—Remap Limit Address Register on page 70	0000000000000000h	RW_KL; RW_L
A0	TOM—Top of Memory on page 71	0000007FFFF00000h	RW_KL; RW_L
A8	TOUUD—Top of Upper Usable DRAM on page 71	0000000000000000h	RW_KL; RW_L
B0	BDSM—Base Data of Stolen Memory on page 72	00000000h	RW_KL; RW_L
B4	BGSM—Base of GTT stolen Memory on page 73	00100000h	RW_KL; RW_L
B8	TSEGMB—TSEG Memory Base on page 73	00000000h	RW_KL; RW_L
BC	TOLUD—Top of Low Usable DRAM on page 73	00100000h	RW_KL; RW_L
C8	ERRSTS—Error Status on page 74	0000h	RW1CS
CA	ERRCMD—Error Command on page 75	0000h	RW
CC	SMICMD—SMI Command on page 76	0000h	RW
CE	SCICMD—SCI Command on page 76	0000h	RW
DC	SKPD—Scratchpad Data on page 77	00000000h	RW
E4	CAPID0—Capabilities A on page 77	00000000h	RO; RO_KFW
E8	CAPID0—Capabilities B on page 78	00000000h	RO

3.1.1 VID—Vendor Identification

This register combined with the Device Identification register uniquely identifies any PCI device.

B/D/F/Type: 0/0/0/CFG			Access: RO	
Size: 16	Default Value: 8086h		Address Offset: 0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:0	VID	Vendor Identification Number: PCI standard identification for Intel.	8086h	RO

3.1.2 DID—Device Identification

This register combined with the Vendor Identification register uniquely identifies any PCI device.

B/D/F/Type: 0/0/0/CFG			Access: RO; RO_V	
Size: 16	Default Value: 0C00h		Address Offset: 2h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:4	DID_MSB	Device Identification Number MSB: This is the upper part of device identification assigned to the processor.	0C0h	RO
3:2	DID_SKU	Device Identification Number SKU: This is the middle part of device identification assigned to the Processor.	0h	RO_V
1:0	DID_LSB	Device Identification Number LSB: This is the lower part of device identification assigned to the Processor.	0h	RO

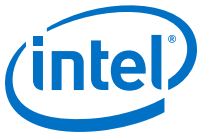

3.1.3 PCICMD—PCI Command

Since Device #0 does not physically reside on PCI_A many of the bits are not implemented.

B/D/F/Type: 0/0/0/CFG			Access: RO; RW	
Size: 16	**Default Value: 0006h**		**Address Offset: 4h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:10	RSVD	Reserved.	00h	RO
9	FB2B	Fast Back-to-Back Enable: This bit controls whether or not the master can do fast back-to-back write. Since device 0 is strictly a target this bit is not implemented and is hardwired to 0. Writes to this bit position have no effect.	0h	RO
8	SERRE	SERR Enable: This bit is a global enable bit for Device 0 SERR messaging. The processor communicates the SERR condition by sending an SERR message over DMI to the PCH. 1: The processor is enabled to generate SERR messages over DMI for specific Device 0 error conditions that are individually enabled in the ERRCMD and DMIUEMSK registers. The error status is reported in the ERRSTS, PCISTS, and DMIUEST registers. 0: The SERR message is not generated by the Host for Device 0. This bit only controls SERR messaging for Device 0. Other integrated devices have their own SERRE bits to control error reporting for error conditions occurring in each device. The control bits are used in a logical OR manner to enable the SERR DMI message mechanism. OPI N/A	0h	RW
7	ADSTEP	Address/Data Stepping Enable: Address/data stepping is not implemented in the processor, and this bit is hardwired to 0. Writes to this bit position have no effect.	0h	RO
6	PERRE	OPI - N/A Parity Error Enable: Controls whether or not the Master Data Parity Error bit in the PCI Status register can bet set. 0: Master Data Parity Error bit in PCI Status register can NOT be set. 1: Master Data Parity Error bit in PCI Status register CAN be set.	0h	RW
5	VGASNOOP	VGA Palette Snoop Enable: The processor does not implement this bit and it is hardwired to a 0. Writes to this bit position have no effect.	0h	RO
4	MWIE	Memory Write and Invalidate Enable: The processor will never issue memory write and invalidate commands. This bit is therefore hardwired to 0. Writes to this bit position will have no effect.	0h	RO
3	SCE	Special Cycle Enable: The processor does not implement this bit and it is hardwired to a 0. Writes to this bit position have no effect.	0h	RO

continued...

B/D/F/Type: 0/0/0/CFG			Access: RO; RW	
Size: 16	Default Value: 0006h		Address Offset: 4h	
Bit Range	Acronym	Description	Default	Access
2	BME	Bus Master Enable: The processor is always enabled as a master on the backbone. This bit is hardwired to a "1". Writes to this bit position have no effect.	1h	RO
1	MAE	Memory Access Enable: The processor always allows access to main memory, except when such access would violate security principles. Such exceptions are outside the scope of PCI control. This bit is not implemented and is hardwired to 1. Writes to this bit position have no effect.	1h	RO
0	IOAE	I/O Access Enable: This bit is not implemented in the processor and is hardwired to a 0. Writes to this bit position have no effect.	0h	RO

3.1.4 PCISTS—PCI Status

This status register reports the occurrence of error events on Device 0's PCI interface. Since Device 0 does not physically reside on PCI_A many of the bits are not implemented.

B/D/F/Type: 0/0/0/CFG			Access: RO; RW1C	
Size: 16	Default Value: 0090h		Address Offset: 6h	
Bit Range	Acronym	Description	Default	Access
15	DPE	Detected Parity Error: This bit is set when this Device recieves a Poisoned TLP.	0h	RW1C
14	SSE	Signaled System Error: This bit is set to 1 when Device 0 generates an SERR message over DMI for any enabled Device 0 error condition. Device 0 error conditions are enabled in the PCICMD, ERRCMD, and DMIUEMSK registers. Device 0 error flags are read/reset from the PCISTS, ERRSTS, or DMIUEST registers. Software clears this bit by writing a 1 to it.	0h	RW1C
13	RMAS	Received Master Abort Status: This bit is set when the processor generates a DMI request that receives an Unsupported Request completion packet. Software clears this bit by writing a 1 to it.	0h	RW1C
12	RTAS	Received Target Abort Status: This bit is set when the processor generates a DMI request that receives a Completer Abort completion packet. Software clears this bit by writing a 1 to it.	0h	RW1C
11	STAS	Signaled Target Abort Status: The processor will not generate a Target Abort DMI completion packet or Special Cycle. This bit is not implemented and is hardwired to a 0. Writes to this bit position have no effect.	0h	RO
10:9	DEVT	DEVSEL Timing: These bits are hardwired to "00". Writes to these bit positions have no affect. Device 0 does not physically connect to PCI_A. These bits are set to "00" (fast decode) so that optimum DEVSEL timing for PCI_A is not limited by the Host.	0h	RO
8	DPD	Master Data Parity Error Detected: This bit is set when DMI received a Poisoned completion from PCH. This bit can only be set when the Parity Error Enable bit in the PCI Command register is set.	0h	RW1C

continued...

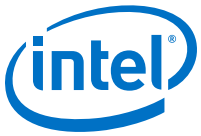

B/D/F/Type: 0/0/0/CFG			Access: RO; RW1C	
Size: 16	**Default Value: 0090h**		**Address Offset: 6h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7	FB2B	Fast Back-to-Back: This bit is hardwired to 1. Writes to these bit positions have no effect. Device 0 does not physically connect to PCI_A. This bit is set to 1 (indicating fast back-to-back capability) so that the optimum setting for PCI_A is not limited by the Host.	1h	RO
6	RSVD	Reserved.	0h	RO
5	MC66	66 MHz Capable: Does not apply to PCI Express. Must be hardwired to 0.	0h	RO
4	CLIST	Capability List: This bit is hardwired to 1 to indicate to the configuration software that this device/function implements a list of new capabilities. A list of new capabilities is accessed via register CAPPTR at configuration address offset 34h. Register CAPPTR contains an offset pointing to the start address within configuration space of this device where the Capability Identification register resides.	1h	RO
3:0	RSVD	Reserved.	0h	RO

3.1.5 RID—Revision Identification

This register contains the revision number of Device #0.
These bits are read only and writes to this register have no effect.

B/D/F/Type: 0/0/0/CFG			Access: RO	
Size: 8	**Default Value: 00h**		**Address Offset: 8h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:4	RID_MSB	Revision Identification Number MSB: Four MSB of RID	0h	RO
3:0	RID	Revision Identification Number: Four LSB of RID	0h	RO

3.1.6 CC—Class Code

This register identifies the basic function of the device, a more specific sub-class, and a register-specific programming interface.

B/D/F/Type: 0/0/0/CFG			Access: RO	
Size: 24	**Default Value: 060000h**		**Address Offset: 9h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
23:16	BCC	Base Class Code: This is an 8-bit value that indicates the base class code for the Host Bridge device. This code has the value 06h, indicating a Bridge device.	06h	RO
15:8	SUBCC	Sub-Class Code: This is an 8-bit value that indicates the category of Bridge into which the Host Bridge device falls. The code is 00h indicating a Host Bridge.	00h	RO
7:0	PI	Programming Interface: This is an 8-bit value that indicates the programming interface of this device. This value does not specify a particular register set layout and provides no practical use for this device.	00h	RO

3.1.7 HDR—Header Type

This register identifies the header layout of the configuration space. No physical register exists at this location.

B/D/F/Type: 0/0/0/CFG			Access: RO	
Size: 8	Default Value: 00h		Address Offset: Eh	
Bit Range	Acronym	Description	Default	Access
7:0	HDR	PCI Header: This field always returns 0 to indicate that the Host Bridge is a single function device with standard header layout. Reads and writes to this location have no effect.	00h	RO

3.1.8 SVID—Subsystem Vendor Identification

This value is used to identify the vendor of the subsystem.

B/D/F/Type: 0/0/0/CFG			Access: RW_O	
Size: 16	Default Value: 0000h		Address Offset: 2Ch	
Bit Range	Acronym	Description	Default	Access
15:0	SUBVID	Subsystem Vendor ID: This field should be programmed during boot-up to indicate the vendor of the system board. After it has been written once, it becomes read only.	0000h	RW_O

3.1.9 SID—Subsystem Identification

This value is used to identify a particular subsystem.

B/D/F/Type: 0/0/0/CFG			Access: RW_O	
Size: 16	Default Value: 0000h		Address Offset: 2Eh	
Bit Range	Acronym	Description	Default	Access
15:0	SUBID	Subsystem ID: This field should be programmed during BIOS initialization. After it has been written once, it becomes read only.	0000h	RW_O

3.1.10 CAPPTR—Capabilities Pointer

The CAPPTR provides the offset that is the pointer to the location of the first device capability in the capability list.

B/D/F/Type: 0/0/0/CFG			Access: RO	
Size: 8	Default Value: E0h		Address Offset: 34h	
Bit Range	Acronym	Description	Default	Access
7:0	CAPPTR	Capabilities Pointer: Pointer to the offset of the first capability ID register block. In this case the first capability is the product-specific Capability Identifier (CAPID0).	E0h	RO

3.1.11 PXPEPBAR—PCI Express Egress Port Base Address

This is the base address for the PCI Express Egress Port MMIO Configuration space. There is no physical memory within this 4KB window that can be addressed. The 4KB reserved by this register does not alias to any PCI 2.3 compliant memory mapped

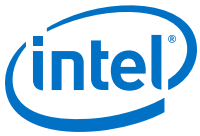

space. On reset, the EGRESS port MMIO configuration space is disabled and must be enabled by writing a 1 to PXPEPBAREN [Dev 0, offset 40h, bit 0].
All the bits in this register are locked in Intel TXT mode.

B/D/F/Type: 0/0/0/CFG			Access: RW	
Size: 64	Default Value: 0000000000000000h		Address Offset: 40h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:12	PXPEPBAR	This field corresponds to bits 38 to 12 of the base address PCI Express Egress Port MMIO configuration space. BIOS will program this register resulting in a base address for a 4KB block of contiguous memory address space. This register ensures that a naturally aligned 4KB space is allocated within the first 512GB of addressable memory space. System Software uses this base address to program the PCI Express Egress Port MMIO register set. All the bits in this register are locked in Intel TXT mode.	0000000h	RW
11:1	RSVD	Reserved.	000h	RO
0	PXPEPBAREN	0: PXPEPBAR is disabled and does not claim any memory 1: PXPEPBAR memory mapped accesses are claimed and decoded appropriately This register is locked by Intel TXT.	0h	RW

3.1.12 MCHBAR—Host Memory Mapped Register Range Base

This is the base address for the Host Memory Mapped Configuration space. There is no physical memory within this 32KB window that can be addressed. The 32KB reserved by this register does not alias to any PCI 2.3 compliant memory mapped space. On reset, the Host MMIO Memory Mapped Configuation space is disabled and must be enabled by writing a 1 to MCHBAREN [Dev 0, offset48h, bit 0].
All the bits in this register are locked in Intel TXT mode.
The register space contains memory control, initialization, timing, and buffer strength registers; clocking registers; and power and thermal management registers.

B/D/F/Type: 0/0/0/CFG			Access: RW	
Size: 64	Default Value: 0000000000000000h		Address Offset: 48h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:15	MCHBAR	This field corresponds to bits 38 to 15 of the base address Host Memory Mapped configuration space. BIOS will program this register resulting in a base address for a 32KB block of contiguous memory address space. This register ensures that a naturally aligned 32KB space is allocated within the first 512GB of addressable memory space. System Software uses this base address to program the Host Memory Mapped register set. All the bits in this register are locked in Intel TXT mode.	000000h	RW
14:1	RSVD	Reserved.	0000h	RO
0	MCHBAREN	0: MCHBAR is disabled and does not claim any memory 1: MCHBAR memory mapped accesses are claimed and decoded appropriately This register is locked in Intel TXT mode.	0h	RW

3.1.13 GGC—GMCH Graphics Control Register

All the bits in this register are Intel TXT lockable.

B/D/F/Type: 0/0/0/CFG			Access:	RW_KL; RW_L
Size: 16	Default Value: 0500h		Address Offset:	50h
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:8	GMS	This field is used to select the amount of Main Memory that is pre-allocated to support the Internal Graphics device in VGA (non-linear) and Native (linear) modes. The BIOS ensures that memory is pre-allocated only when Internal graphics is enabled. This register is also Intel TXT lockable. Hardware does not clear or set any of these bits automatically based on IGD being disabled/enabled. BIOS Requirement: BIOS must not set this field to 0h if IVD (bit 1 of this register) is 0. 0x00 = 0MB 0x01 = 32MB 0x02 = 64MB 0x03 = 96MB 0x04 = 128MB 0x05 = 160MB 0x06 = 192MB 0x07 = 224MB 0x08 = 256MB 0x09 = 288MB 0x0A = 320MB 0x0B = 352MB 0x0C = 384MB 0x0D = 416MB 0x0E = 448MB 0x0F = 480MB 0x10 = 512MB 0x20 = ...1024MB... 0x30 = ...1536MB... 0x3F = ...2016MB	05h	RW_L
7:6	GGMS	This field is used to select the amount of Main Memory that is pre-allocated to support the Internal Graphics Translation Table. The BIOS ensures that memory is pre-allocated only when Internal graphics is enabled. GSM is assumed to be a contiguous physical DRAM space with DSM, and BIOS needs to allocate a contiguous memory chunk. Hardware will derive the base of GSM from DSM only using the GSM size programmed in the register. Hardware functionality in case of programming this value to Reserved is not guaranteed. 0x0 = No Preallocated Memory 0x1 = 2MB of Preallocated Memory 0x2 = 4MB of Preallocated Memory 0x3 = 8MB of Preallocated Memory	0h	RW_L
5:3	RSVD	Reserved.	0h	RO

continued...

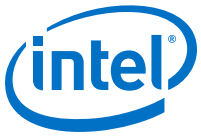

B/D/F/Type: 0/0/0/CFG			Access: RW_KL; RW_L	
Size: 16	**Default Value: 0500h**		**Address Offset: 50h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
2	VAMEN	Enables the use of the iGFX enbines for Versatile Acceleration. 1 - iGFX engines are in Versatile Acceleration Mode. Device 2 Class Code is 048000h. 0 - iGFX engines are in iGFX Mode. Device 2 Class Code is 030000h.	0h	RW_L
1	IVD	0: Enable. Device 2 (IGD) claims VGA memory and IO cycles, the Sub-Class Code within Device 2 Class Code register is 00. 1: Disable. Device 2 (IGD) does not claim VGA cycles (Mem and IO), and the Sub- Class Code field within Device 2 function 0 Class Code register is 80. BIOS Requirement: BIOS must not set this bit to 0 if the GMS field (bits 7:3 of this register) pre-allocates no memory. This bit MUST be set to 1 if Device 2 is disabled via a register (DEVEN[3] = 0). This register is locked by Intel TXT lock. 0 = Enable 1 = Disable	0h	RW_L
0	GGCLCK	When set to 1b, this bit will lock all bits in this register.	0h	RW_KL

3.1.14 DEVEN—Device Enable

Allows for enabling/disabling of PCI devices and functions that are within the processor package. The table below the bit definitions describes the behavior of all combinations of transactions to devices controlled by this register.
All the bits in this register are Intel TXT Lockable.

B/D/F/Type: 0/0/0/CFG			Access: RO; RW_L; RW	
Size: 32	**Default Value: 000000BFh**		**Address Offset: 54h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:15	RSVD	Reserved.	00000h	RO
14	D7EN	0: Bus 0 Device 7 is disabled and not visible. 1: Bus 0 Device 7 is enabled and visible. Non-production BIOS code should provide a setup option to enable Bus 0 Device 7. When enabled, Bus 0 Device 7 must be initialized in accordance to standard PCI device initialization procedures.	0h	RW
13:11	RSVD	Reserved.	0h	RO
10	D5EN	0: Bus 0 Device 5 is disabled and not visible. 1: Bus 0 Device 5 is enabled and visible. This bit will be set to 0b and remain 0b if Device 5 capability is disabled.	0h	RO
9:8	RSVD	Reserved.	0h	RO

continued...

B/D/F/Type: 0/0/0/CFG			Access: RO; RW_L; RW	
Size: 32	Default Value: 000000BFh		Address Offset: 54h	
Bit Range	Acronym	Description	Default	Access
7	D4EN	0: Bus 0 Device 4 is disabled and not visible. 1: Bus 0 Device 4 is enabled and visible. This bit will be set to 0b and remain 0b if Device 4 capability is disabled.	1h	RW_L
6	RSVD	Reserved.	0h	RO
5	D3EN	0: Bus 0 Device 3 is disabled and hidden 1: Bus 0 Device 3 is enabled and visible This bit will be set to 0b and remain 0b if Device 3 capability is disabled.	1h	RW_L
4	D2EN	0: Bus 0 Device 2 is disabled and hidden 1: Bus 0 Device 2 is enabled and visible This bit will be set to 0b and remain 0b if Device 2 capability is disabled.	1h	RW_L
3	D1F0EN	0: Bus 0 Device 1 Function 0 is disabled and hidden. 1: Bus 0 Device 1 Function 0 is enabled and visible. This bit will be set to 0b and remain 0b if PEG10 capability is disabled.	1h	RW_L
2	D1F1EN	0: Bus 0 Device 1 Function 1 is disabled and hidden. 1: Bus 0 Device 1 Function 1 is enabled and visible.	1h	RW_L
1	D1F2EN	0: Bus 0 Device 1 Function 2 is disabled and hidden. 1: Bus 0 Device 1 Function 2 is enabled and visible.	1h	RW_L
0	D0EN	Bus 0 Device 0 Function 0 may not be disabled and is therefore hardwired to 1.	1h	RO

3.1.15 PAVPC—Protected Audio Video Path Control

All the bits in this register are locked by Intel TXT. When locked the RW bits are RO.

B/D/F/Type: 0/0/0/CFG			Access: RW_KL	
Size: 32	Default Value: 00000000h		Address Offset: 58h	
Bit Range	Acronym	Description	Default	Access
31:3	RSVD	Reserved.	00000000h	RO
2	PAVPLCK	This bit locks all writeable contents in this register when set (including itself). Only a hardware reset can unlock the register again. This lock bit needs to be set only if PAVP is enabled (bit 1 of this register is asserted).	0h	RW_KL
1:0	RSVD	Reserved.	0h	RO

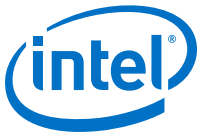

3.1.16 DPR—DMA Protected Range

DMA protected range register.

B/D/F/Type: 0/0/0/CFG			Access: ROV; RW_L	
Size: 32	**Default Value: 00000000h**		**Address Offset: 5Ch**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:3	RSVD	Reserved.	00000000h	RO
2	EPM	This field controls DMA accesses to the DMA Protected Range (DPR) region. 0: DPR is disabled 1: DPR is enabled. All DMA requests accessing DPR region are blocked. HW reports the status of DPR enable/disable through the PRS field in this register.	0h	RW_L
1	PRS	This field indicates the status of DPR. 0: DPR protection disabled 1: DPR protection enabled	0h	ROV
0	RSVD	Reserved.	0h	RO

3.1.17 PCIEXBAR—PCI Express Register Range Base Address

This is the base address for the PCI Express configuration space. This window of addresses contains the 4KB of configuration space for each PCI Express device that can potentially be part of the PCI Express Hierarchy associated with the Uncore. There is no actual physical memory within this window of up to 256MB that can be addressed. The actual size of this range is determined by a field in this register. Each PCI Express Hierarchy requires a PCI Express BASE register. The Uncore supports one PCI Express Hierarchy. The region reserved by this register does not alias to any PCI2.3 compliant memory mapped space. For example, the range reserved for MCHBAR is outside of PCIEXBAR space.

On reset, this register is disabled and must be enabled by writing a 1 to the enable field in this register. This base address shall be assigned on a boundary consistent with the number of buses (defined by the length field in this register), above TOLUD and still within 39-bit addressable memory space.

The PCI Express Base Address cannot be less than the maximum address written to the Top of physical memory register (TOLUD). Software must guarantee that these ranges do not overlap with known ranges located above TOLUD.

Software must ensure that the sum of the length of the enhanced configuration region + TOLUD + any other known ranges reserved above TOLUD is not greater than the 39-bit addessable limit of 512GB. In general, system implementation and the number of PCI/PCI Express/PCI-X buses supported in the hierarchy will dictate the length of the region.

All the bits in this register are locked in Intel TXT mode.

B/D/F/Type: 0/0/0/CFG			Access: RW; RW_V	
Size: 64	Default Value: 0000000000000000h		Address Offset: 60h	
Bit Range	Acronym	Description	Default	Access
63:39	RSVD	Reserved.	0000000h	RO
38:28	PCIEXBAR	This field corresponds to bits 38 to 28 of the base address for PCI Express enhanced configuration space. BIOS will program this register resulting in a base address for a contiguous memory address space. The size of the range is defined by bits [2:1] of this register. This Base address shall be assigned on a boundary consistent with the number of buses (defined by the Length field in this register) above TOLUD and still within the 39-bit addressable memory space. The address bits decoded depend on the length of the region defined by this register. This register is locked by Intel TXT. The address used to access the PCI Express configuration space for a specific device can be determined as follows: PCI Express Base Address + Bus Number * 1MB + Device Number * 32KB + Function Number * 4KB This address is the beginning of the 4KB space that contains both the PCI compatible configuration space and the PCI Express extended configuration space.	000h	RW
27	ADMSK128	This bit is either part of the PCI Express Base Address (RW) or part of the Address Mask (RO, read 0b), depending on the value of bits [2:1] in this register.	0h	RW_V
26	ADMSK64	This bit is either part of the PCI Express Base Address (RW) or part of the Address Mask (RO, read 0b), depending on the value of bits [2:1] in this register.	0h	RW_V
25:3	RSVD	Reserved.	000000h	RO
2:1	LENGTH	This field describes the length of this region. 00: 256MB (buses 0-255). Bits 38:28 are decoded in the PCI Express Base Address Field. 01: 128MB (buses 0-127). Bits 38:27 are decoded in the PCI Express Base Address Field. 10: 64MB (buses 0-63). Bits 38:26 are decoded in the PCI Express Base Address Field. 11: Reserved. Thsi register is locked by Intel TXT.	0h	RW
0	PCIEXBAREN	0: The PCIEXBAR register is disabled. Memory read and write transactions proceed s if there were no PCIEXBAR register. PCIEXBAR bits 38:26 are RW with no functionality behind them. 1: The PCIEXBAR register is enabled. Memory read and write transactions whose address bits 38:26 match PCIEXBAR will be translated to configuration reads and writes within the Uncore. These Translated cycles are routed as shown in the above table. This register is locked by Intel TXT.	0h	RW

3.1.18 DMIBAR—Root Complex Register Range Base Address

This is the base address for the Root Complex configuration space. This window of addresses contains the Root Complex Register set for the PCI Express Hierarchy associated with the Host Bridge. There is no physical memory within this 4KB window that can be addressed. The 4KB reserved by this register does not alias to any PCI 2.3

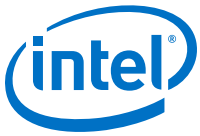

compliant memory mapped space. On reset, the Root Complex configuration space is disabled and must be enabled by writing a 1 to DMIBAREN [Dev 0, offset 68h, bit 0] All the bits in this register are locked in Intel TXT mode.

B/D/F/Type: 0/0/0/CFG			Access: RW	
Size: 64	Default Value: 0000000000000000h		Address Offset: 68h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:12	DMIBAR	This field corresponds to bits 38 to 12 of the base address DMI configuration space. BIOS will program this register resulting in a base address for a 4KB block of contiguous memory address space. This register ensures that a naturally aligned 4KB space is allocated within the first 512GB of addressable memory space. System Software uses this base address to program the DMI register set. All the Bits in this register are locked in Intel TXT mode.	0000000h	RW
11:1	RSVD	Reserved.	000h	RO
0	DMIBAREN	0: DMIBAR is disabled and does not claim any memory 1: DMIBAR memory mapped accesses are claimed and decoded appropriately This register is locked by Intel TXT.	0h	RW

3.1.19 MESEG—Intel Management Engine Base Address Register

This register determines the Base Address register of the memory range that is pre-allocated to the Intel Management Engine. Together with the MESEG_MASK register it controls the amount of memory allocated to the Intel ME.
This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be 0. Thus, the bottom of the defined memory address range will be aligned to a 1MB boundary.
This register is locked by Intel TXT.
NOTE: BIOS must program MESEG_BASE and MESEG_MASK so that Intel ME Stolen Memory is carved out from TOM.

B/D/F/Type: 0/0/0/CFG			Access: RW_L	
Size: 64	Default Value: 0000007FFFF00000h		Address Offset: 70h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:20	MEBASE	Corresponds to A[38:20] of the base address memory range that is allocated to the ME.	7FFFFh	RW_L
19:0	RSVD	Reserved.	00000h	RO

3.1.20 MESEG—Intel Management Engine Limit Address Register

This register determines the Mask Address register of the memory range that is pre-allocated to the Intel Management Engine. Together with the MESEG_BASE register it controls the amount of memory allocated to the Intel ME.
This register is locked by Intel TXT.
NOTE: BIOS must program MESEG_BASE and MESEG_MASK so that ME Stolen Memory is carved out from TOM.

B/D/F/Type: 0/0/0/CFG			Access: RW_KL; RW_L	
Size: 64	Default Value: 0000000000000000h		Address Offset: 78h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:20	MEMASK	This field indicates the bits that must match MEBASE in order to qualify as an Intel ME Memory Range access. For example, if the field is set to 7FFFFh, then ME Memory is 1MB in size. Another example is that if the field is set to 7FFFEh, then ME Memory is 2MB in size. Mask value should be such that once a bit is set to 1 all the more significant bit should be 1. It is not legal to set up mask with 0 and 1's interspersed. In other words, the size of ME Memory Range is limited to power of 2 times 1MB. MEBASE must be naturally aligned to the size of ME region.	00000h	RW_L
19:12	RSVD	Reserved.	00h	RO
11	ME_STLEN_EN	Indicates whether the Intel ME stolen Memory range is enabled or not.	0h	RW_L
10	MELCK	This field indicates whether all bits in the MESEG_BASE and MESEG_MASK registers are locked. When locked, updates to any field for these registers must be dropped.	0h	RW_KL
9:0	RSVD	Reserved.	000h	RO

3.1.21 PAM0—Programmable Attribute Map 0

This register controls the read, write and shadowing attributes of the BIOS range from F_0000h to F_FFFFh. The Uncore allows programmable memory attributes on 13 legacy memory segments of various sizes in the 768KB to 1MB address range. Seven Programmable Attribute Map (PAM) registers are used to support these features. Cacheability of these areas is controlled via the MTRR register in the core. Two bits are used to specify memory attributes for each memory segment. These bits apply to host accesses to the PAM areas. These attributes are:
RE - Read Enable. When RE=1, the host read accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when RE=0, the host read accesses are directed to DMI.
WE - Write Enable. When WE=1, the host write accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when WE=0, the host read accesses are directed to DMI.
The RE and WE attributes permit a memory segment to be Read Only, Write Only, Read/Write or Disabled. For example, if a memory segment has RE=1 and WE=0, the segment is Read Only.

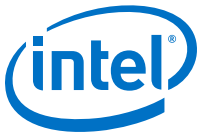

B/D/F/Type: 0/0/0/CFG			Access: RW_KL; RW_L	
Size: 8	**Default Value: 00h**		**Address Offset: 80h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:6	RSVD	Reserved.	0h	RO
5:4	HIENABLE	This field controls the steering of read and write cycles that address the BIOS area from 0F_0000h to 0F_FFFFh. 00: DRAM Disabled. All accesses are directed to DMI. 01: Read Only. All reads are sent to DRAM, all writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM, all reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.	0h	RW_L
3:1	RSVD	Reserved.	0h	RO
0	Lock	If this bit is set, all of the PAM* registers are locked (cannot be written)	0h	RW_KL

3.1.22 PAM1—Programmable Attribute Map 1

This register controls the read, write and shadowing attributes of the BIOS range from C_0000h to C_7FFFh. The Uncore allows programmable memory attributes on 13 legacy memory segments of various sizes in the 768KB to 1MB address range. Seven Programmable Attribute Map (PAM) registers are used to support these features. Cacheability of these areas is controlled via the MTRR register in the core. Two bits are used to specify memory attributes for each memory segment. These bits apply to host accesses to the PAM areas. These attributes are:
RE - Read Enable. When RE=1, the host read accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when RE=0, the host read accesses are directed to DMI.
WE - Write Enable. When WE=1, the host write accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when WE=0, the host read accesses are directed to DMI.
The RE and WE attributes permit a memory segment to be Read Only, Write Only, Read/Write or Disabled. For example, if a memory segment has RE=1 and WE=0, the segment is Read Only.

B/D/F/Type: 0/0/0/CFG			Access: RW_L	
Size: 8	**Default Value: 00h**		**Address Offset: 81h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:6	RSVD	Reserved.	0h	RO
5:4	HIENABLE	This field controls the steering of read and write cycles that address the BIOS area from 0C_4000h to 0C_7FFFh. 00: DRAM Disabled. All accesses are directed to DMI. 01: Read Only. All reads are sent to DRAM, all writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM, all reads are	0h	RW_L

continued...

B/D/F/Type: 0/0/0/CFG			Access: RW_L	
Size: 8	Default Value: 00h		Address Offset: 81h	
Bit Range	Acronym	Description	Default	Access
		serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.		
3:2	RSVD	Reserved.	0h	RO
1:0	LOENABLE	This field controls the steering of read and write cycles that address the BIOS area from 0C0000h to 0C3FFFh. 00: DRAM Disabled. All reads are sent to DRAM. All writes are forwarded to DMI. 01: Read Only. All reads are sent to DRAM. All writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM. All reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.	0h	RW_L

3.1.23 PAM2—Programmable Attribute Map 2

This register controls the read, write and shadowing attributes of the BIOS range from C_8000h to C_FFFFh. The Uncore allows programmable memory attributes on 13 legacy memory segments of various sizes in the 768KB to 1MB address range. Seven Programmable Attribute Map (PAM) registers are used to support these features. Cacheability of these areas is controlled via the MTRR register in the core. Two bits are used to specify memory attributes for each memory segment. These bits apply to host accesses to the PAM areas. These attributes are:
RE - Read Enable. When RE=1, the host read accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when RE=0, the host read accesses are directed to DMI.
WE - Write Enable. When WE=1, the host write accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when WE=0, the host read accesses are directed to DMI.
The RE and WE attributes permit a memory segment to be Read Only, Write Only, Read/Write or Disabled. For example, if a memory segment has RE=1 and WE=0, the segment is Read Only.

B/D/F/Type: 0/0/0/CFG			Access: RW_L	
Size: 8	Default Value: 00h		Address Offset: 82h	
Bit Range	Acronym	Description	Default	Access
7:6	RSVD	Reserved.	0h	RO
5:4	HIENABLE	This field controls the steering of read and write cycles that address the BIOS area from 0CC000h to 0CFFFFh. 00: DRAM Disabled. All accesses are directed to DMI. 01: Read Only. All reads are sent to DRAM, all writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM, all reads are	0h	RW_L

continued...

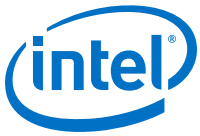

B/D/F/Type: 0/0/0/CFG			Access: RW_L	
Size: 8	**Default Value: 00h**		**Address Offset: 82h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.		
3:2	RSVD	Reserved.	0h	RO
1:0	LOENABLE	This field controls the steering of read and write cycles that address the BIOS area from 0C8000h to 0CBFFFh. 00: DRAM Disabled. All reads are sent to DRAM. All writes are forwarded to DMI. 01: Read Only. All reads are sent to DRAM. All writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM. All reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.	0h	RW_L

3.1.24 PAM3—Programmable Attribute Map 3

This register controls the read, write and shadowing attributes of the BIOS range from D0000h to D7FFFh. The Uncore allows programmable memory attributes on 13 legacy memory segments of various sizes in the 768KB to 1MB address range. Seven Programmable Attribute Map (PAM) registers are used to support these features. Cacheability of these areas is controlled via the MTRR register in the core.
Two bits are used to specify memory attributes for each memory segment. These bits apply to host accesses to the PAM areas. These attributes are:
RE - Read Enable. When RE=1, the host read accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when RE=0, the host read accesses are directed to DMI.
WE - Write Enable. When WE=1, the host write accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when WE=0, the host read accesses are directed to DMI.
The RE and WE attributes permit a memory segment to be Read Only, Write Only, Read/Write or Disabled. For example, if a memory segment has RE=1 and WE=0, the segment is Read Only.

B/D/F/Type: 0/0/0/CFG			Access: RW_L	
Size: 8	**Default Value: 00h**		**Address Offset: 83h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:6	RSVD	Reserved.	0h	RO
5:4	HIENABLE	This field controls the steering of read and write cycles that address the BIOS area from 0D4000h to 0D7FFFh. 00: DRAM Disabled. All accesses are directed to DMI. 01: Read Only. All reads are sent to DRAM, all writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM, all reads are	0h	RW_L

continued...

B/D/F/Type: 0/0/0/CFG			Access: RW_L	
Size: 8	Default Value: 00h		Address Offset: 83h	
Bit Range	Acronym	Description	Default	Access
		serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.		
3:2	RSVD	Reserved.	0h	RO
1:0	LOENABLE	This field controls the steering of read and write cycles that address the BIOS area from 0D0000h to 0D3FFFh. 00: DRAM Disabled. All reads are sent to DRAM. All writes are forwarded to DMI. 01: Read Only. All reads are sent to DRAM. All writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM. All reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.	0h	RW_L

3.1.25 PAM4—Programmable Attribute Map 4

This register controls the read, write and shadowing attributes of the BIOS range from D8000h to DFFFFh. The Uncore allows programmable memory attributes on 13 legacy memory segments of various sizes in the 768KB to 1MB address range. Seven Programmable Attribute Map (PAM) registers are used to support these features. Cacheability of these areas is controlled via the MTRR register in the core.
Two bits are used to specify memory attributes for each memory segment. These bits apply to host accesses to the PAM areas. These attributes are:
RE - Read Enable. When RE=1, the host read accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when RE=0, the host read accesses are directed to DMI.
WE - Write Enable. When WE=1, the host write accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when WE=0, the host read accesses are directed to DMI.
The RE and WE attributes permit a memory segment to be Read Only, Write Only, Read/Write or Disabled. For example, if a memory segment has RE=1 and WE=0, the segment is Read Only.

B/D/F/Type: 0/0/0/CFG			Access: RW_L	
Size: 8	Default Value: 00h		Address Offset: 84h	
Bit Range	Acronym	Description	Default	Access
7:6	RSVD	Reserved.	0h	RO
5:4	HIENABLE	This field controls the steering of read and write cycles that address the BIOS area from 0DC000h to 0DFFFFh. 00: DRAM Disabled. All accesses are directed to DMI. 01: Read Only. All reads are sent to DRAM, all writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM, all reads are	0h	RW_L

continued...

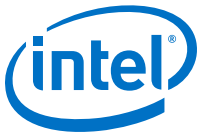

B/D/F/Type: 0/0/0/CFG			Access: RW_L	
Size: 8	**Default Value: 00h**		**Address Offset: 84h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.		
3:2	RSVD	Reserved.	0h	RO
1:0	LOENABLE	This field controls the steering of read and write cycles that address the BIOS area from 0D8000h to 0DBFFFh. 00: DRAM Disabled. All reads are sent to DRAM. All writes are forwarded to DMI. 01: Read Only. All reads are sent to DRAM. All writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM. All reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.	0h	RW_L

3.1.26 PAM5—Programmable Attribute Map 5

This register controls the read, write and shadowing attributes of the BIOS range from E_0000h to E_7FFFh. The Uncore allows programmable memory attributes on 13 legacy memory segments of various sizes in the 768KB to 1MB address range. Seven Programmable Attribute Map (PAM) registers are used to support these features. Cacheability of these areas is controlled via the MTRR register in the core. Two bits are used to specify memory attributes for each memory segment. These bits apply to host accesses to the PAM areas. These attributes are:
RE - Read Enable. When RE=1, the host read accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when RE=0, the host read accesses are directed to DMI.
WE - Write Enable. When WE=1, the host write accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when WE=0, the host read accesses are directed to DMI.
The RE and WE attributes permit a memory segment to be Read Only, Write Only, Read/Write or Disabled. For example, if a memory segment has RE=1 and WE=0, the segment is Read Only.

B/D/F/Type: 0/0/0/CFG			Access: RW_L	
Size: 8	**Default Value: 00h**		**Address Offset: 85h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:6	RSVD	Reserved.	0h	RO
5:4	HIENABLE	This field controls the steering of read and write cycles that address the BIOS area from 0E4000h to 0E7FFFh. 00: DRAM Disabled. All accesses are directed to DMI. 01: Read Only. All reads are sent to DRAM, all writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM, all reads are	0h	RW_L

continued...

B/D/F/Type: 0/0/0/CFG			Access: RW_L	
Size: 8	Default Value: 00h		Address Offset: 85h	
Bit Range	Acronym	Description	Default	Access
		serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.		
3:2	RSVD	Reserved.	0h	RO
1:0	LOENABLE	This field controls the steering of read and write cycles that address the BIOS area from 0E0000h to 0E3FFFh. 00: DRAM Disabled. All reads are sent to DRAM. All writes are forwarded to DMI. 01: Read Only. All reads are sent to DRAM. All writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM. All reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.	0h	RW_L

3.1.27 PAM6—Programmable Attribute Map 6

This register controls the read, write and shadowing attributes of the BIOS range from E_8000h to E_FFFFh. The Uncore allows programmable memory attributes on 13 legacy memory segments of various sizes in the 768KB to 1MB address range. Seven Programmable Attribute Map (PAM) registers are used to support these features. Cacheability of these areas is controlled via the MTRR register in the core. Two bits are used to specify memory attributes for each memory segment. These bits apply to host accesses to the PAM areas. These attributes are:
RE - Read Enable. When RE=1, the host read accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when RE=0, the host read accesses are directed to DMI.
WE - Write Enable. When WE=1, the host write accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when WE=0, the host read accesses are directed to DMI.
The RE and WE attributes permit a memory segment to be Read Only, Write Only, Read/Write or Disabled. For example, if a memory segment has RE=1 and WE=0, the segment is Read Only.

B/D/F/Type: 0/0/0/CFG			Access: RW_L	
Size: 8	Default Value: 00h		Address Offset: 86h	
Bit Range	Acronym	Description	Default	Access
7:6	RSVD	Reserved.	0h	RO
5:4	HIENABLE	This field controls the steering of read and write cycles that address the BIOS area from 0EC000h to 0EFFFFh. 00: DRAM Disabled. All accesses are directed to DMI. 01: Read Only. All reads are sent to DRAM, all writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM, all reads are	0h	RW_L

continued...

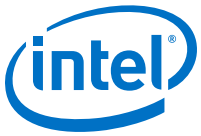

B/D/F/Type: 0/0/0/CFG			Access: RW_L	
Size: 8	**Default Value: 00h**		**Address Offset: 86h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.		
3:2	RSVD	Reserved.	0h	RO
1:0	LOENABLE	This field controls the steering of read and write cycles that address the BIOS area from 0E8000h to 0EBFFFh. 00: DRAM Disabled. All reads are sent to DRAM. All writes are forwarded to DMI. 01: Read Only. All reads are sent to DRAM. All writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM. All reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.	0h	RW_L

3.1.28 LAC—Legacy Access Control

This 8-bit register controls steering of MDA cycles and a fixed DRAM hole from 15-16MB.
There can only be at most one MDA device in the system.

B/D/F/Type: 0/0/0/CFG			Access: RW	
Size: 8	**Default Value: 00h**		**Address Offset: 87h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7	HEN	This field enables a memory hole in DRAM space. The DRAM that lies "behind" this space is not remapped. 0: No memory hole. 1: Memory hole from 15MB to 16MB. This bit is Intel TXT lockable.	0h	RW
6:3	RSVD	Reserved.	0h	RO
2	MDAP12	This bit works with the VGA Enable bits in the BCTRL register of Device 1 Function 2 to control the routing of processor initiated transactions targeting MDA compatible I/O and memory address ranges. This bit should not be set if device 1 function 2 VGA Enable bit is not set. If device 1 function 2 VGA enable bit is not set, then accesses to IO address range x3BCh-x3BFh remain on the backbone. If the VGA enable bit is set and MDA is not present, then accesses to IO address range x3BCh-x3BFh are forwarded to PCI Express through device 1 function 2 if the address is within the corresponding IOBASE and IOLIMIT, otherwise they remain on the backbone. MDA resources are defined as the following: Memory: 0B0000h - 0B7FFFh I/O: 3B4h, 3B5h, 3B8h, 3B9h, 3BAh, 3BFh, (including ISA address aliases, A[15:10] are not used in decode) Any I/O reference that includes the I/O locations listed above, or their aliases, will remain on the backbone even if the reference also includes I/O locations not listed above. The following table shows the behavior for all combinations of MDA and VGA:	0h	RW

continued...

B/D/F/Type: 0/0/0/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: 87h	
Bit Range	Acronym	Description	Default	Access
		VGAEN MDAP Description 0 0 All References to MDA and VGA space are not claimed by Device 1 Function 2. 0 1 Illegal combination 1 0 All VGA and MDA references are routed to PCI Express Graphics Attach device 1 function 2. 1 1 All VGA references are routed to PCI Express Graphics Attach device 1 function 2. MDA references are not claimed by device 1 function 2. VGA and MDA memory cycles can only be routed across PEG12 when MAE (PCICMD12[1]) is set. VGA and MDA I/O cycles can only be routed across PEG12 if IOAE (PCICMD12[0]) is set.		
1	MDAP11	This bit works with the VGA Enable bits in the BCTRL register of Device 1 Function 1 to control the routing of processor initiated transactions targeting MDA compatible I/O and memory address ranges. This bit should not be set if device 1 function 1 VGA Enable bit is not set. If device 1 function 1 VGA enable bit is not set, then accesses to IO address range x3BCh-x3BFh remain on the backbone. If the VGA enable bit is set and MDA is not present, then accesses to IO address range x3BCh-x3BFh are forwarded to PCI Express through device 1 function 1 if the address is within the corresponding IOBASE and IOLIMIT, otherwise they remain on the backbone. MDA resources are defined as the following: Memory: 0B0000h - 0B7FFFh I/O: 3B4h, 3B5h, 3B8h, 3B9h, 3BAh, 3BFh, (including ISA address aliases, A[15:10] are not used in decode) Any I/O reference that includes the I/O locations listed above, or their aliases, will remain on the backbone even if the reference also includes I/O locations not listed above. The following table shows the behavior for all combinations of MDA and VGA: VGAEN MDAP Description 0 0 All References to MDA and VGA space are not claimed by Device 1 Function 1. 0 1 Illegal combination 1 0 All VGA and MDA references are routed to PCI Express Graphics Attach device 1 function 1. 1 1 All VGA references are routed to PCI Express Graphics Attach device 1 function 1. MDA references are not claimed by device 1 function 1. VGA and MDA memory cycles can only be routed across PEG11 when MAE (PCICMD11[1]) is set. VGA and MDA I/O cycles can only be routed across PEG11 if IOAE (PCICMD11[0]) is set.	0h	RW
0	MDAP10	This bit works with the VGA Enable bits in the BCTRL register of Device 1 Function 0 to control the routing of processor initiated transactions targeting MDA compatible I/O and memory address ranges. This bit should not be set if device 1 function 0 VGA Enable bit is not set. If device 1 function 0 VGA enable bit is not set, then accesses to IO address range x3BCh-x3BFh remain on the backbone. If the VGA enable bit is set and MDA is not present, then accesses to IO address range x3BCh-x3BFh are forwarded to PCI Express through device 1 function 0 if the address is within the corresponding IOBASE and IOLIMIT, otherwise they remain on the backbone.	0h	RW

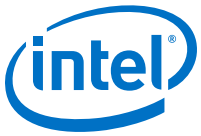

B/D/F/Type: 0/0/0/CFG			Access: RW	
Size: 8	**Default Value: 00h**		**Address Offset: 87h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		MDA resources are defined as the following: Memory: 0B0000h - 0B7FFFh I/O: 3B4h, 3B5h, 3B8h, 3B9h, 3BAh, 3BFh, (including ISA address aliases, A[15:10] are not used in decode) Any I/O reference that includes the I/O locations listed above, or their aliases, will remain on the backbone even if the reference also includes I/O locations not listed above. The following table shows the behavior for all combinations of MDA and VGA: VGAEN MDAP Description 0 0 All References to MDA and VGA space are not claimed by Device 1 Function 0. 0 1 Illegal combination 1 0 All VGA and MDA references are routed to PCI Express Graphics Attach device 1 function 0. 1 1 All VGA references are routed to PCI Express Graphics Attach device 1 function 0. MDA references are not claimed by device 1 function 0. VGA and MDA memory cycles can only be routed across PEG10 when MAE (PCICMD10[1]) is set. VGA and MDA I/O cycles can only be routed across PEG10 if IOAE (PCICMD10[0]) is set.		

3.1.29 SMRAMC—System Management RAM Control

The SMRAMC register controls how accesses to Compatible SMRAM spaces are treated. The Open, Close and Lock bits function only when G_SMRAME bit is set to 1. Also, the Open bit must be reset before the Lock bit is set.

B/D/F/Type: 0/0/0/CFG			Access: RO; RW_L; RW_KL; RW_LV	
Size: 8	**Default Value: 02h**		**Address Offset: 88h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7	RSVD	Reserved.	0h	RO
6	D_OPEN	When D_OPEN = 1 and D_LCK = 0, the SMM DRAM space is made visible even when SMM decode is not active. This is intended to help BIOS initialize SMM space. Software should ensure that D_OPEN = 1 and D_CLS = 1 are not set at the same time.	0h	RW_LV
5	D_CLS	When D_CLS = 1, SMM DRAM space is not accessible to data references, even if SMM decode is active. Code references may still access SMM DRAM space. This will allow SMM software to reference through SMM space to update the display even when SMM is mapped over the VGA range. Software should ensure that D_OPEN = 1 and D_CLS = 1 are not set at the same time.	0h	RW_L
4	D_LCK	When D_LCK=1, then D_OPEN is reset to 0 and all writeable fields in this register are locked (become RO). D_LCK can be set to 1 via a normal configuration space write but can only be cleared by a Full Reset. The combination of D_LCK and D_OPEN provide convenience with security. The BIOS can use the D_OPEN function to initialize SMM space and then use D_LCK to	0h	RW_KL

continued...

B/D/F/Type: 0/0/0/CFG			Access: RO; RW_L; RW_KL; RW_LV	
Size: 8	Default Value: 02h		Address Offset: 88h	
Bit Range	Acronym	Description	Default	Access
		"lock down" SMM space in the future so that no application software (or even BIOS itself) can violate the integrity of SMM space, even if the program has knowledge of the D_OPEN function.		
3	G_SMRAME	If set to '1', then Compatible SMRAM functions are enabled, providing 128KB of DRAM accessible at the A_0000h address while in SMM. Once D_LCK is set, this bit becomes RO.	0h	RW_L
2:0	C_BASE_SEG	This field indicates the location of SMM space. SMM DRAM is not remapped. It is simply made visible if the conditions are right to access SMM space, otherwise the access is forwarded to DMI. Only SMM space bewteen A_0000h and B_FFFFh is supported, so this field is hardwired to 010b.	2h	RO

3.1.30 REMAPBASE—Remap Base Address Register

B/D/F/Type: 0/0/0/CFG			Access: RW_KL; RW_L	
Size: 64	Default Value: 0000007FFFF00000h		Address Offset: 90h	
Bit Range	Acronym	Description	Default	Access
63:39	RSVD	Reserved.	0000000h	RO
38:20	REMAPBASE	The value in this register defines the lower boundary of the Remap window. The Remap window is inclusive of this address. In the decoder A[19:0] of the Remap Base Address are assumed to be 0's. Thus the bottom of the defined memory range will be aligned to a 1MB boundary. When the value in this register is greater than the value programmed into the Remap Limit register, the Remap window is disabled. These bits are Intel TXT lockable.	7FFFFh	RW_L
19:1	RSVD	Reserved.	00000h	RO
0	LOCK	This bit will lock all writeable settings in this register, including itself.	0h	RW_KL

3.1.31 REMAPLIMIT—Remap Limit Address Register

B/D/F/Type: 0/0/0/CFG			Access: RW_KL; RW_L	
Size: 64	Default Value: 0000000000000000h		Address Offset: 98h	
Bit Range	Acronym	Description	Default	Access
63:39	RSVD	Reserved.	0000000h	RO
38:20	REMAPLMT	The value in this register defines the upper boundary of the Remap window. The Remap window is inclusive of this address. In the decoder A[19:0] of the remap limit address are assumed to be F's. Thus the top of the defined range will be one byte less than a 1MB boundary.	00000h	RW_L

continued...

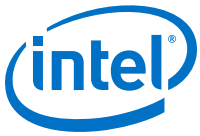

B/D/F/Type: 0/0/0/CFG			Access: RW_KL; RW_L	
Size: 64	**Default Value: 0000000000000000h**		**Address Offset: 98h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		When the value in this register is less than the value programmed into the Remap Base register, the Remap window is disabled. These Bits are Intel TXT lockable.		
19:1	RSVD	Reserved.	00000h	RO
0	LOCK	This bit will lock all writeable settings in this register, including itself.	0h	RW_KL

3.1.32 TOM—Top of Memory

This Register contains the size of physical memory. BIOS determines the memory size reported to the OS using this Register.

B/D/F/Type: 0/0/0/CFG			Access: RW_KL; RW_L	
Size: 64	**Default Value: 0000007FFFF00000h**		**Address Offset: A0h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:20	TOM	This register reflects the total amount of populated physical memory. This is NOT necessarily the highest main memory address (holes may exist in main memory address map due to addresses allocated for memory mapped IO). These bits correspond to address bits 38:20 (1MB granularity). Bits 19:0 are assumed to be 0. All the bits in this register are locked in Intel TXT mode.	7FFFFh	RW_L
19:1	RSVD	Reserved.	00000h	RO
0	LOCK	This bit will lock all writeable settings in this register, including itself.	0h	RW_KL

3.1.33 TOUUD—Top of Upper Usable DRAM

This 64 bit register defines the Top of Upper Usable DRAM.
Configuration software must set this value to TOM minus all ME stolen memory if reclaim is disabled. If reclaim is enabled, this value must be set to reclaim limit + 1byte, 1MB aligned, since reclaim limit is 1MB aligned. Address bits 19:0 are assumed to be 000_0000h for the purposes of address comparison. The Host interface positively decodes an address towards DRAM if the incoming address is less than the value programmed in this register and greater than or equal to 4GB.
BIOS Restriction: Minimum value for TOUUD is 4GB.
These bits are Intel TXT lockable.

B/D/F/Type: 0/0/0/CFG			Access: RW_KL; RW_L	
Size: 64	Default Value: 0000000000000000h		Address Offset: A8h	
Bit Range	Acronym	Description	Default	Access
63:39	RSVD	Reserved.	0000000h	RO
38:20	TOUUD	This register contains bits 38 to 20 of an address one byte above the maximum DRAM memory above 4G that is usable by the operating system. Configuration software must set this value to TOM minus all ME stolen memory if reclaim is disabled. If reclaim is enabled, this value must be set to reclaim limit 1MB aligned since reclaim limit + 1byte is 1MB aligned. Address bits 19:0 are assumed to be 000_0000h for the purposes of address comparison. The Host interface positively decodes an address towards DRAM if the incoming address is less than the value programmed in this register and greater than 4GB. All the bits in this register are locked in Intel TXT mode.	00000h	RW_L
19:1	RSVD	Reserved.	00000h	RO
0	LOCK	This bit will lock all writeable settings in this register, including itself.	0h	RW_KL

3.1.34 BDSM—Base Data of Stolen Memory

This register contains the base address of graphics data stolen DRAM memory. BIOS determines the base of graphics data stolen memory by subtracting the graphics data stolen memory size (PCI Device 0 offset 52 bits 7:4) from TOLUD (PCI Device 0 offset BC bits 31:20).

B/D/F/Type: 0/0/0/CFG			Access: RW_KL; RW_L	
Size: 32	Default Value: 00000000h		Address Offset: B0h	
Bit Range	Acronym	Description	Default	Access
31:20	BDSM	This register contains bits 31 to 20 of the base address of stolen DRAM memory. BIOS determines the base of graphics stolen memory by subtracting the graphics stolen memory size (PCI Device 0 offset 50 bits 15:8) from TOLUD (PCI Device 0 offset BC bits 31:20).	000h	RW_L
19:1	RSVD	Reserved.	00000h	RO
0	LOCK	This bit will lock all writeable settings in this register, including itself.	0h	RW_KL

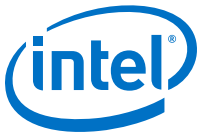

3.1.35 BGSM—Base of GTT stolen Memory

This register contains the base address of stolen DRAM memory for the GTT. BIOS determines the base of GTT stolen memory by subtracting the GTT graphics stolen memory size (PCI Device 0 offset 52 bits 9:8) from the Graphics Base of Data Stolen Memory (PCI Device 0 offset B0 bits 31:20).

B/D/F/Type: 0/0/0/CFG			Access: RW_KL; RW_L	
Size: 32	Default Value: 00100000h		Address Offset: B4h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:20	BGSM	This register contains the base address of stolen DRAM memory for the GTT. BIOS determines the base of GTT stolen memory by subtracting the GTT graphics stolen memory size (PCI Device 0 offset 50 bits 7:6) from the Graphics Base of Data Stolen Memory (PCI Device 0 offset B0 bits 31:20).	001h	RW_L
19:1	RSVD	Reserved.	00000h	RO
0	LOCK	This bit will lock all writeable settings in this register, including itself.	0h	RW_KL

3.1.36 TSEGMB—TSEG Memory Base

This register contains the base address of TSEG DRAM memory. BIOS determines the base of TSEG memory which must be at or below Graphics Base of GTT Stolen Memory (PCI Device 0 Offset B4 bits 31:20).
NOTE: BIOS must program TSEGMB to a 8MB naturally aligned boundary.

B/D/F/Type: 0/0/0/CFG			Access: RW_KL; RW_L	
Size: 32	Default Value: 00000000h		Address Offset: B8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:20	TSEGMB	This register contains the base address of TSEG DRAM memory. BIOS determines the base of TSEG memory which must be at or below Graphics Base of GTT Stolen Memory (PCI Device 0 Offset B4 bits 31:20). BIOS must program the value of TSEGMB to be the same as BGSM when TSEG is disabled.	000h	RW_L
19:1	RSVD	Reserved.	00000h	RO
0	LOCK	This bit will lock all writeable settings in this register, including itself.	0h	RW_KL

3.1.37 TOLUD—Top of Low Usable DRAM

This 32 bit register defines the Top of Low Usable DRAM. TSEG, GTT Graphics memory and Graphics Stolen Memory are within the DRAM space defined. From the top, the Host optionally claims 1 to 64MBs of DRAM for internal graphics if enabled, 1or 2MB of DRAM for GTT Graphics Stolen Memory (if enabled) and 1, 2, or 8 MB of DRAM for TSEG if enabled.

Programming Example:
C1DRB3 is set to 4GB
TSEG is enabled and TSEG size is set to 1MB
Internal Graphics is enabled, and Graphics Mode Select is set to 32MB

GTT Graphics Stolen Memory Size set to 2MB
BIOS knows the OS requires 1G of PCI space.
BIOS also knows the range from 0_FEC0_0000h to 0_FFFF_FFFFh is not usable by the system. This 20MB range at the very top of addressable memory space is lost to APIC and Intel TXT.
According to the above equation, TOLUD is originally calculated to: 4GB = 1_0000_0000h
The system memory requirements are: 4GB (max addressable space) - 1GB (pci space) - 35MB (lost memory) = 3GB - 35MB (minimum granularity) = 0_ECB0_0000h
Since 0_ECB0_0000h (PCI and other system requirements) is less than 1_0000_0000h, TOLUD should be programmed to ECBh.
These bits are Intel TXT lockable.

B/D/F/Type: 0/0/0/CFG			Access: RW_KL; RW_L	
Size: 32	**Default Value: 00100000h**		**Address Offset: BCh**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:20	TOLUD	This register contains bits 31 to 20 of an address one byte above the maximum DRAM memory below 4G that is usable by the operating system. Address bits 31 down to 20 programmed to 01h implies a minimum memory size of 1MB. Configuration software must set this value to the smaller of the following 2 choices: maximum amount memory in the system minus ME stolen memory plus one byte or the minimum address allocated for PCI memory. Address bits 19:0 are assumed to be 0_0000h for the purposes of address comparison. The Host interface positively decodes an address towards DRAM if the incoming address is less than the value programmed in this register. The Top of Low Usable DRAM is the lowest address above both Graphics Stolen memory and TSEG. BIOS determines the base of Graphics Stolen Memory by subtracting the Graphics Stolen Memory Size from TOLUD and further decrements by TSEG size to determine base of TSEG. All the Bits in this register are locked in Intel TXT mode. This register must be 1MB aligned when reclaim is enabled.	001h	RW_L
19:1	RSVD	Reserved.	00000h	RO
0	LOCK	This bit will lock all writeable settings in this register, including itself.	0h	RW_KL

3.1.38 ERRSTS—Error Status

This register is used to report various error conditions via the SERR DMI messaging mechanism. An SERR DMI message is generated on a zero to one transition of any of these flags (if enabled by the ERRCMD and PCICMD registers).
These bits are set regardless of whether or not the SERR is enabled and generated. After the error processing is complete, the error logging mechanism can be unlocked by clearing the appropriate status bit by software writing a '1' to it.

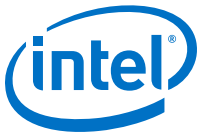

B/D/F/Type: 0/0/0/CFG			Access: RW1CS	
Size: 16	**Default Value: 0000h**		**Address Offset: C8h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:2	RSVD	Reserved.	0000h	RO
1	DMERR	If this bit is set to 1, a memory read data transfer had an uncorrectable multiple-bit error. When this bit is set, the column, row, bank, and rank that caused the error, and the error syndrome, are logged in the ECC Error Log register in the channel where the error occurred. Once this bit is set, the ECCERRLOGx fields are locked until the processor clears this bit by writing a 1. Software uses bits [1:0] to detect whether the logged error address is for a Single-bit or a Multiple-bit error.	0h	RW1CS
0	DSERR	If this bit is set to 1, a memory read data transfer had a single-bit correctable error and the corrected data was returned to the requesting agent. When this bit is set the column, row, bank, and rank where the error occurred and the syndrome of the error are logged in the ECC Error Log register in the channel where the error occurred. Once this bit is set the ECCERRLOGx fields are locked to further single-bit error updates until the processor clears this bit by writing a 1. A multiple bit error that occurs after this bit is set will overwrite the ECCERRLOGx fields with the multiple-bit error signature and the DMERR bit will also be set. A single bit error that occurs after a multibit error will set this bit but will not overwrite the other fields.	0h	RW1CS

3.1.39 ERRCMD—Error Command

This register controls the Host Bridge responses to various system errors. Since the Host Bridge does not have an SERR# signal, SERR messages are passed from the processor to the PCH over DMI.
When a bit in this register is set, a SERR message will be generated on DMI whenever the corresponding flag is set in the ERRSTS register. The actual generation of the SERR message is globally enabled for Device #0 via the PCI Command register.

B/D/F/Type: 0/0/0/CFG			Access: RW	
Size: 16	**Default Value: 0000h**		**Address Offset: CAh**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:2	RSVD	Reserved.	0000h	RO
1	DMERR	1: The Host Bridge generates an SERR message over DMI when it detects a multiple-bit error reported by the DRAM controller. 0: Reporting of this condition via SERR messaging is disabled. For systems not supporting ECC this bit must be disabled.	0h	RW
0	DSERR	1: The Host Bridge generates an SERR special cycle over DMI when the DRAM controller detects a single bit error. 0: Reporting of this condition via SERR messaging is	0h	RW

continued...

B/D/F/Type: 0/0/0/CFG			Access: RW	
Size: 16	Default Value: 0000h		Address Offset: CAh	
Bit Range	Acronym	Description	Default	Access
		disabled. For systems that do not support ECC this bit must be disabled.		

3.1.40 SMICMD—SMI Command

This register enables various errors to generate an SMI DMI special cycle. When an error flag is set in the ERRSTS register, it can generate an SERR, SMI, or SCI DMI special cycle when enabled in the ERRCMD, SMICMD, or SCICMD registers, respectively. Note that one and only one message type can be enabled.

B/D/F/Type: 0/0/0/CFG			Access: RW	
Size: 16	Default Value: 0000h		Address Offset: CCh	
Bit Range	Acronym	Description	Default	Access
15:2	RSVD	Reserved.	0000h	RO
1	DMESMI	1: The Host generates an SMI DMI message when it detects a multiple-bit error reported by the DRAM controller. 0: Reporting of this condition via SMI messaging is disabled. For systems not supporting ECC this bit must be disabled.	0h	RW
0	DSESMI	1: The Host generates an SMI DMI special cycle when the DRAM controller detects a single bit error. 0: Reporting of this condition via SMI messaging is disabled. For systems that do not support ECC this bit must be disabled.	0h	RW

3.1.41 SCICMD—SCI Command

This register enables various errors to generate an SMI DMI special cycle. When an error flag is set in the ERRSTS register, it can generate an SERR, SMI, or SCI DMI special cycle when enabled in the ERRCMD, SMICMD, or SCICMD registers, respectively. Note that one and only one message type can be enabled.

B/D/F/Type: 0/0/0/CFG			Access: RW	
Size: 16	Default Value: 0000h		Address Offset: CEh	
Bit Range	Acronym	Description	Default	Access
15:2	RSVD	Reserved.	0000h	RO
1	DMESCI	1: The Host generates an SCI DMI message when it detects a multiple-bit error reported by the DRAM controller. 0: Reporting of this condition via SCI messaging is disabled. For systems not supporting ECC this bit must be disabled.	0h	RW
0	DSESCI	1: The Host generates an SCI DMI special cycle when the DRAM controller detects a single bit error. 0: Reporting of this condition via SCI messaging is disabled. For systems that do not support ECC this bit must be disabled.	0h	RW

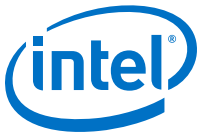

3.1.42 SKPD—Scratchpad Data

This register holds 32 writable bits with no functionality behind them. It is for the convenience of BIOS and graphics drivers.

B/D/F/Type: 0/0/0/CFG			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: DCh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:0	SKPD	1 DWORD of data storage.	00000000h	RW

3.1.43 CAPID0—Capabilities A

Control of bits in this register are only required for customer visible SKU differentiation.

B/D/F/Type: 0/0/0/CFG			Access: RO; RO_KFW	
Size: 32	Default Value: 00000000h		Address Offset: E4h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:26	RSVD	Reserved.	00h	RO
25	ECCDIS	0b ECC capable 1b Not ECC capable	0h	RO
24	RSVD	Reserved.	0h	RO
23	VTDD	0: Enable VTd 1: Disable VTd	0h	RO_KFW
22:15	RSVD	Reserved.	00h	RO
14	DDPCD	Allows Dual Channel operation but only supports 1 DIMM per channel. 0: 2 DIMMs per channel enabled 1: 2 DIMMs per channel disabled. This setting hardwires bits 2 and 3 of the rank population field for each channel to zero. (MCHBAR offset 260h, bits 22-23 for channel 0 and MCHBAR offset 660h, bits 22-23 for channel 1)	0h	RO
13	X2APIC_EN	Extended Interrupt Mode. 0b: Hardware does not support Extended APIC mode. 1b: Hardware supports Extended APIC mode.	0h	RO
12	PDCD	0: Capable of Dual Channels 1: Not Capable of Dual Channel - only single channel capable.	0h	RO
11:0	RSVD	Reserved.	000h	RO

3.1.44 CAPID0—Capabilities B

Control of bits in this register are only required for customer visible SKU differentiation.

B/D/F/Type: 0/0/0/CFG			Access: RO	
Size: 32	Default Value: 00000000h		Address Offset: E8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:29	RSVD	Reserved.	0h	RO
28	SMT	This setting indicates whether or not the processor is SMT capable.	0h	RO
27:25	CACHESZ	This setting indicates the supporting cache sizes.	0h	RO
24	RSVD	Reserved.	0h	RO
23:21	PLL_REF100_CFG	DDR3 Maximum Frequency Capability with 100 Memory. Maximum allowed memory frequency with 100 MHz ref clk. 0 - 100 MHz ref disabled 1 - upto DDR-1400 (7 x 200) 2 - upto DDR-1600 (8 x 200) 3 - upto DDR-1800 (8 x 200) 4 - upto DDR-2000 (10 x 200) 5 - upto DDR-2200 (11 x 200) 6 - upto DDR-2400 (12 x 200) 7 - no limit (but still limited by _DDR_FREQ200 to 3000)	0h	RO
20	PEGG3_DIS	0: Capable of running any of the Gen 3-compliant PEG controllers in Gen 3 mode (Devices 0/1/0, 0/1/1, 0/1/2) 1: Not capable of running any of the PEG controllers in Gen 3 mode	0h	RO
19	RSVD	Reserved.	0h	RO
18	ADDGFXEN	0 - Additive Graphics Disabled 1- Additive Graphics Enabled	0h	RO
17	ADDGFXCAP	0 - Capable of Additive Graphics 1 - Not capable of Additive Graphics	0h	RO
16:7	RSVD	Reserved.	000h	RO
6:4	DMFC	This field controls which values may be written to the Memory Frequency Select field 6:4 of the Clocking Configuration registers (MCHBAR Offset C00h). Any attempt to write an unsupported value will be ignored. 000: No limit (but still limited by %MAX_DDR_FREQ266 to 3200) 001: MC capable of up to DDR3 2667 010: MC capable of up to DDR3 2400 011: MC capable of up to DDR3 2133 100: MC capable of up to DDR3 1867 101: MC capable of up to DDR3 1600 110: MC capable of up to DDR3 1333 111: MC capable of up to DDR3 1067	0h	RO
3:0	RSVD	Reserved.	0h	RO

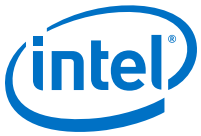

3.2 Integrated Graphics Device Registers Summary

Offset	Register ID—Description	Default Value	Access
0	VID2—Vendor Identification on page 80	8086h	RO
2	DID2—Device Identification on page 80	1602h	RO_V; RO
4	PCICMD—PCI Command on page 80	0000h	RW; RO
6	PCISTS2—PCI Status on page 81	0090h	RO_V; RO
8	RID2—Revision Identification on page 82	00h	RO
9	CC—Class Code on page 82	030000h	RO; RO_V
C	CLS—Cache Line Size on page 83	00h	RO
D	MLT2—Master Latency Timer on page 83	00h	RO
E	HDR2—Header Type on page 83	00h	RO
10	GTTMMADR—Graphics Translation Table, Memory Mapped Range Address on page 83	0000000000000004h	RO; RW
18	GMADR—Graphics Memory Range Address on page 84	000000000000000Ch	RO; RW_L; RW
20	IOBAR—I/O Base Address on page 85	00000001h	RO; RW
2C	SVID2—Subsystem Vendor Identification on page 85	0000h	RW_O
2E	SID2—Subsystem Identification on page 86	0000h	RW_O
30	ROMADR—Video BIOS ROM Base Address on page 86	00000000h	RO
34	CAPPOINT—Capabilities Pointer on page 86	90h	RO_V
3C	INTRLINE—Interrupt Line on page 86	00h	RW
3D	INTRPIN—Interrupt Pin on page 87	01h	RO
3E	MINGNT—Minimum Grant on page 87	00h	RO
3F	MAXLAT—Maximum Latency on page 87	00h	RO
44	CAPID0—Capabilities A on page 88	00000000h	RO_V
48	CAPID0—Capabilities B on page 88	00000000h	RO_V
54	DEVEN0—Device Enable on page 89	000000BFh	RO; RO_V
62	MSAC—Multi Size Aperture Control on page 90	01h	RW_KV; RW
90	MSI—Message Signaled Interrupts Capability ID on page 92	D005h	RO
92	MC—Message Control on page 92	0000h	RW; RO
94	MA—Message Address on page 93	00000000h	RO; RW
98	MD—Message Data on page 93	0000h	RW
A4	AFCIDNP—Advanced Features Capabilities Identifier and Next Pointer on page 94	0013h	RO
A8	AFCTL—Advanced Features Control on page 94	00h	RW1S
A9	AFSTS—Advanced Features Status on page 95	00h	RO

continued...

Offset	Register ID—Description	Default Value	Access
D0	PMCAPID—Power Management Capabilities ID on page 95	A401h	RO
D2	PMCAP—Power Management Capabilities on page 95	0022h	RO
D4	PMCS—Power Management Control/Status on page 96	0000h	RO_V; RO

3.2.1 VID2—Vendor Identification

This register combined with the Device Identification register uniquely identifies any PCI device.

B/D/F/Type: 0/2/0/CFG			Access: RO	
Size: 16	Default Value: 8086h		Address Offset: 0h	
Bit Range	Acronym	Description	Default	Access
15:0	VID	PCI standard identification for Intel.	8086h	RO

3.2.2 DID2—Device Identification

This register combined with the Vendor Identification register uniquely identifies any PCI device.
This is a 16 bit value assigned to Processor Graphics device.

B/D/F/Type: 0/2/0/CFG			Access: RO_V; RO	
Size: 16	Default Value: 1602h		Address Offset: 2h	
Bit Range	Acronym	Description	Default	Access
15:4	DID_MSB	This is the upper part of a 16 bit value assigned to the Graphics device.	160h	RO
3:2	DID_SKU	These are bits 3:2 of the 16 bit value assigned to the Processor Graphics device.	0h	RO_V
1:0	DID_LSB	This is the lower part of a 16 bit value assigned to the Processor Graphics device.	2h	RO_V

3.2.3 PCICMD—PCI Command

This 16-bit register provides basic control over the IGD's ability to respond to PCI cycles. The PCICMD Register in the IGD disables the IGD PCI compliant master accesses to main memory.

B/D/F/Type: 0/2/0/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: 4h	
Bit Range	Acronym	Description	Default	Access
15:11	RSVD	Reserved.	00h	RO
10	INTDIS	This bit disables the device from asserting INTx#. 0: Enable the assertion of this device's INTx# signal. 1: Disable the assertion of this device's INTx# signal. DO_INTx messages will not be sent to DMI.	0h	RW
9	FB2B	Not Implemented. Hardwired to 0.	0h	RO
8	SEN	Not Implemented. Hardwired to 0.	0h	RO

continued...

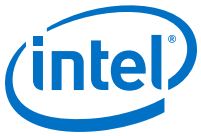

B/D/F/Type: 0/2/0/CFG			Access: RW; RO	
Size: 16	**Default Value: 0000h**		**Address Offset: 4h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7	WCC	Not Implemented. Hardwired to 0.	0h	RO
6	PER	Not Implemented. Hardwired to 0. Since the IGD belongs to the category of devices that does not corrupt programs or data in system memory or hard drives, the IGD ignores any parity error that it detects and continues with normal operation.	0h	RO
5	VPS	This bit is hardwired to 0 to disable snooping.	0h	RO
4	MWIE	Hardwired to 0. The IGD does not support memory write and invalidate commands.	0h	RO
3	SCE	This bit is hardwired to 0. The IGD ignores Special cycles.	0h	RO
2	BME	0: Disable IGD bus mastering. 1: Enable the IGD to function as a PCI compliant master.	0h	RW
1	MAE	This bit controls the IGD's response to memory space accesses. 0: Disable. 1: Enable.	0h	RW
0	IOAE	This bit controls the IGD's response to I/O space accesses. 0: Disable. 1: Enable.	0h	RW

3.2.4 PCISTS2—PCI Status

PCISTS is a 16-bit status register that reports the occurrence of a PCI compliant master abort and PCI compliant target abort.

PCISTS also indicates the DEVSEL# timing that has been set by the IGD.

B/D/F/Type: 0/2/0/CFG			Access: RO_V; RO	
Size: 16	**Default Value: 0090h**		**Address Offset: 6h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15	DPE	Since the IGD does not detect parity, this bit is always hardwired to 0.	0h	RO
14	SSE	The IGD never asserts SERR#, therefore this bit is hardwired to 0.	0h	RO
13	RMAS	The IGD never gets a Master Abort, therefore this bit is hardwired to 0.	0h	RO
12	RTAS	The IGD never gets a Target Abort, therefore this bit is hardwired to 0.	0h	RO
11	STAS	Hardwired to 0. The IGD does not use target abort semantics.	0h	RO
10:9	DEVT	N/A. These bits are hardwired to "00".	0h	RO
8	DPD	Since Parity Error Response is hardwired to disabled (and the IGD does not do any parity detection), this bit is hardwired to 0.	0h	RO
7	FB2B	Hardwired to 1. The IGD accepts fast back-to-back when the transactions are not to the same agent.	1h	RO

continued...

B/D/F/Type: 0/2/0/CFG			Access: RO_V; RO	
Size: 16	Default Value: 0090h		Address Offset: 6h	
Bit Range	Acronym	Description	Default	Access
6	UDF	Hardwired to 0.	0h	RO
5	C66	N/A - Hardwired to 0.	0h	RO
4	CLIST	This bit is set to 1 to indicate that the register at 34h provides an offset into the function's PCI Configuration Space containing a pointer to the location of the first item in the list.	1h	RO
3	INTSTS	This bit reflects the state of the interrupt in the device. Only when the Interrupt Disable bit in the command register is a 0 and this Interrupt Status bit is a 1, will the devices INTx# signal be asserted.	0h	RO_V
2:0	RSVD	Reserved.	0h	RO

3.2.5 RID2—Revision Identification

This register contains the revision number for Device #2 Functions 0. These bits are read only and writes to this register have no effect.

B/D/F/Type: 0/2/0/CFG			Access: RO	
Size: 8	Default Value: 00h		Address Offset: 8h	
Bit Range	Acronym	Description	Default	Access
7:4	RID_MSB	Four MSB of RID	0h	RO
3:0	RID	Four LSB of RID	0h	RO

3.2.6 CC—Class Code

This register contains the device programming interface information related to the Sub-Class Code and Base Class Code definition for the IGD. This register also contains the Base Class Code and the function sub-class in relation to the Base Class Code.

B/D/F/Type: 0/2/0/CFG			Access: RO; RO_V	
Size: 24	Default Value: 030000h		Address Offset: 9h	
Bit Range	Acronym	Description	Default	Access
23:16	BCC	This is an 8-bit value that indicates the base class code. When MGGC0[VAMEN] is 0 this code has the value 03h, indicating a Display Controller. When MGGC0[VAMEN] is 1 this code has the value 04h, indicating a Multimedia Device.	03h	RO_V
15:8	SUBCC	When MGGC0[VAMEN] is 0 this value will be determined based on Device 0 GGC register, GMS and IVD fields. 00h: VGA compatible 80h: Non VGA (GMS = "00h" or IVD = "1b") When MGGC0[VAMEN] is 1, this value is 80h, indicating other multimedia device.	00h	RO_V
7:0	PI	When MGGC0[VAMEN] is 0 this value is 00h, indicating a Display Controller. When MGGC0[VAMEN] is 1 this value is 00h, indicating a NOP.	00h	RO

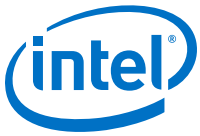

3.2.7 CLS—Cache Line Size

The IGD does not support this register as a PCI slave.

B/D/F/Type: 0/2/0/CFG			Access: RO	
Size: 8	Default Value: 00h		Address Offset: Ch	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:0	CLS	This field is hardwired to 0s. The IGD as a PCI compliant master does not use the Memory Write and Invalidate command and, in general, does not perform operations based on cache line size.	00h	RO

3.2.8 MLT2—Master Latency Timer

The IGD does not support the programmability of the master latency timer because it does not perform bursts.

B/D/F/Type: 0/2/0/CFG			Access: RO	
Size: 8	Default Value: 00h		Address Offset: Dh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:0	MLTCV	Hardwired to 0s.	00h	RO

3.2.9 HDR2—Header Type

This register contains the Header Type of the IGD.

B/D/F/Type: 0/2/0/CFG			Access: RO	
Size: 8	Default Value: 00h		Address Offset: Eh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7	MFUNC	Indicates if the device is a Multi-Function Device. The Value of this register is hardwired to 0, SNB graphics is a single function.	0h	RO
6:0	H	This is a 7-bit value that indicates the Header Code for the IGD. This code has the value 00h, indicating a type 0 configuration space format.	00h	RO

3.2.10 GTTMMADR—Graphics Translation Table, Memory Mapped Range Address

This register requests allocation for the combined Graphics Translation Table Modification Range and Memory Mapped Range. The range requires 16 MB combined for MMIO and Global GTT aperture, with 2MB of that used by MMIO and 8MB used by GTT. GTTADR will begin at (GTTMMADR + 8 MB) while the MMIO base address will be the same as GTTMMADR. The region between (GTTMMADR + 2MB) - (GTTMMADR + 8MB) is reserved.
For the Global GTT, this range is defined as a memory BAR in graphics device config space. It is an alias into which software is required to write Page Table Entry values (PTEs). Software may read PTE values from the global Graphics Translation Table (GTT). PTEs cannot be written directly into the global GTT memory area.
The device snoops writes to this region in order to invalidate any cached translations within the various TLB's implemented on-chip.
The allocation is for 16MB and the base address is defined by bits [38:24].

B/D/F/Type: 0/2/0/CFG			Access: RO; RW	
Size: 64	Default Value: 0000000000000004h		Address Offset: 10h	
Bit Range	Acronym	Description	Default	Access
63:39	RSVDRW	Must be set to 0 since addressing above 512GB is not supported.	0000000h	RW
38:24	MBA	Set by the OS, these bits correspond to address signals [38:24]. 16MB combined for MMIO and Global GTT table aperture (2MB for MMIO, 6MB reserved and 8 MB for GTT).	0000h	RW
23:4	ADM	Hardwired to 0s to indicate at least 16MB address range.	00000h	RO
3	PREFMEM	Hardwired to 0 to prevent prefetching.	0h	RO
2:1	MEMTYP	00 : To indicate 32 bit base address 01: Reserved 10 : To indicate 64 bit base address 11: Reserved	2h	RO
0	MIOS	Hardwired to 0 to indicate memory space.	0h	RO

3.2.11 GMADR—Graphics Memory Range Address

GMADR is the PCI aperture used by S/W to access tiled GFX surfaces in a linear fashion.

B/D/F/Type: 0/2/0/CFG			Access: RO; RW_L; RW	
Size: 64	Default Value: 000000000000000Ch		Address Offset: 18h	
Bit Range	Acronym	Description	Default	Access
63:39	RSVDRW	Must be set to 0 since addressing above 512GB is not supported.	0000000h	RW
38:32	MBA	Memory Base Address (MBA): Set by the OS, these bits correspond to address signals [38:32].	00h	RW
31	ADMSK4096	This Bit is either part of the Memory Base Address (RW) or part of the Address Mask (RO), depending on the value of MSAC[4:0]. See MSAC (Dev2, Func 0, offset 62h) for details.	0h	RW_L
30	ADMSK2048	This Bit is either part of the Memory Base Address (RW) or part of the Address Mask (RO), depending on the value of MSAC[4:0]. See MSAC (Dev2, Func 0, offset 62h) for details.	0h	RW_L
29	ADMSK1024	This Bit is either part of the Memory Base Address (RW) or part of the Address Mask (RO), depending on the value of MSAC[4:0]. See MSAC (Dev2, Func 0, offset 62h) for details.	0h	RW_L
28	ADMSK512	This Bit is either part of the Memory Base Address (RW) or part of the Address Mask (RO), depending on the value of MSAC[4:0]. See MSAC (Dev2, Func 0, offset 62h) for details.	0h	RW_L
27	ADMSK256	This bit is either part of the Memory Base Address (RW) or part of the Address Mask (RO), depending on the value of MSAC[4:0]. See MSAC (Dev 2, Func 0, offset 62h) for details.	0h	RW_L
26:4	ADM	Hardwired to 0s to indicate at least 128MB address range.	000000h	RO

continued...

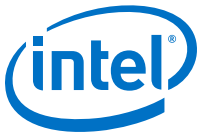

B/D/F/Type: 0/2/0/CFG			Access: RO; RW_L; RW	
Size: 64	**Default Value: 0000000000000000Ch**		**Address Offset: 18h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
3	PREFMEM	Hardwired to 1 to enable prefetching.	1h	RO
2:1	MEMTYP	Memory Type (MEMTYP): 00: indicate 32-bit address. 10: Indicate 64-bit address	2h	RO
0	MIOS	Hardwired to 0 to indicate memory space.	0h	RO

3.2.12 IOBAR—I/O Base Address

This register provides the Base offset of the I/O registers within Device #2. Bits 15:6 are programmable allowing the I/O Base to be located anywhere in 16bit I/O Address Space. Bits 2:1 are fixed and return zero; bit 0 is hardwired to a one indicating that 8 bytes of I/O space are decoded. Access to the 8Bs of IO space is allowed in PM state D0 when IO Enable (PCICMD bit 0) set. Access is disallowed in PM states D1-D3 or if IO Enable is clear or if Device #2 is turned off.
Note that access to this IO BAR is independent of VGA functionality within Device #2. If accesses to this IO bar are allowed, then all 8, 16 or 32 bit IO cycles from IA cores that fall within the 8B are claimed.

B/D/F/Type: 0/2/0/CFG			Access: RO; RW	
Size: 32	**Default Value: 00000001h**		**Address Offset: 20h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:6	IOBASE	Set by the OS, these bits correspond to address signals [15:6].	000h	RW
5:3	RSVD	Reserved.	0h	RO
2:1	MEMTYPE	Hardwired to 0s to indicate 32-bit address.	0h	RO
0	MIOS	Hardwired to "1" to indicate IO space.	1h	RO

3.2.13 SVID2—Subsystem Vendor Identification

This register is used to uniquely identify the subsystem where the PCI device resides.

B/D/F/Type: 0/2/0/CFG			Access: RW_O	
Size: 16	**Default Value: 0000h**		**Address Offset: 2Ch**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:0	SUBVID	This value is used to identify the vendor of the subsystem. This register should be programmed by BIOS during boot-up. Once written, this register becomes Read_Only. This register can only be cleared by a Reset.	0000h	RW_O

3.2.14 SID2—Subsystem Identification

This register is used to uniquely identify the subsystem where the PCI device resides.

B/D/F/Type: 0/2/0/CFG			Access: RW_O	
Size: 16	Default Value: 0000h		Address Offset: 2Eh	
Bit Range	Acronym	Description	Default	Access
15:0	SUBID	This value is used to identify a particular subsystem. This field should be programmed by BIOS during boot-up. Once written, this register becomes Read_Only. This register can only be cleared by a Reset.	0000h	RW_O

3.2.15 ROMADR—Video BIOS ROM Base Address

The IGD does not use a separate BIOS ROM, therefore this register is hardwired to 0s.

B/D/F/Type: 0/2/0/CFG			Access: RO	
Size: 32	Default Value: 00000000h		Address Offset: 30h	
Bit Range	Acronym	Description	Default	Access
31:18	RBA	Hardwired to 0's.	0000h	RO
17:11	ADMSK	Hardwired to 0s to indicate 256 KB address range.	00h	RO
10:1	RSVD	Reserved.	000h	RO
0	RBE	0: ROM not accessible.	0h	RO

3.2.16 CAPPOINT—Capabilities Pointer

This register points to a linked list of capabilities implemented by this device.

B/D/F/Type: 0/2/0/CFG			Access: RO_V	
Size: 8	Default Value: 90h		Address Offset: 34h	
Bit Range	Acronym	Description	Default	Access
7:0	CPV	This field contains an offset into the function's PCI Configuration Space for the first item in the New Capabilities Linked List, the MSI Capabilities ID registers at address 90h or the Power Management capability at D0h. This value is determined by the configuration in CAPL[0].	90h	RO_V

3.2.17 INTRLINE—Interrupt Line

This 8-bit register is used to communicate interrupt line routing information. It is read/write and must be implemented by the device. POST software will write the routing information into this register as it initializes and configures the system. The value in this register tells which input of the system interrupt controller(s) the device's interrupt pin is connected to. The device itself does not use this value, rather it is used by device drivers and operating systems to determine priority and vector information.

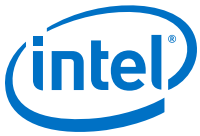

B/D/F/Type: 0/2/0/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: 3Ch	
Bit Range	Acronym	Description	Default	Access
7:0	INTCON	Used to communicate interrupt line routing information. POST software writes the routing information into this register as it initializes and configures the system. The value in this register indicates to which input of the system interrupt controller the device's interrupt pin is connected.	00h	RW

3.2.18 INTRPIN—Interrupt Pin

This register tells which interrupt pin the device uses. The Integrated Graphics Device uses INTA#.

B/D/F/Type: 0/2/0/CFG			Access: RO	
Size: 8	Default Value: 01h		Address Offset: 3Dh	
Bit Range	Acronym	Description	Default	Access
7:0	INTPIN	As a single function device, the IGD specifies INTA# as its interrupt pin. 01h: INTA#.	01h	RO

3.2.19 MINGNT—Minimum Grant

The Integrated Graphics Device has no requirement for the settings of Latency Timers.

B/D/F/Type: 0/2/0/CFG			Access: RO	
Size: 8	Default Value: 00h		Address Offset: 3Eh	
Bit Range	Acronym	Description	Default	Access
7:0	MGV	The IGD does not burst as a PCI compliant master.	00h	RO

3.2.20 MAXLAT—Maximum Latency

The Integrated Graphics Device has no requirement for the settings of Latency Timers.

B/D/F/Type: 0/2/0/CFG			Access: RO	
Size: 8	Default Value: 00h		Address Offset: 3Fh	
Bit Range	Acronym	Description	Default	Access
7:0	MLV	The IGD has no specific requirements for how often it needs to access the PCI bus.	00h	RO

3.2.21 CAPID0—Capabilities A

Control of bits in this register are only required for customer visible SKU differentiation.

B/D/F/Type: 0/2/0/CFG			Access: RO_V	
Size: 32	Default Value: 00000000h		Address Offset: 44h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:26	RSVD	Reserved.	00h	RO
25	ECCDIS	0b ECC capable 1b Not ECC capable	0h	RO_V
24	RSVD	Reserved.	0h	RO
23	VTDD	0: Enable VTd 1: Disable VTd	0h	RO_V
22:15	RSVD	Reserved.	00h	RO
14	DDPCD	Allows Dual Channel operation but only supports 1 DIMM per channel. 0: 2 DIMMs per channel enabled 1: 2 DIMMs per channel disabled. This setting hardwires bits 2 and 3 of the rank population field for each channel to zero. (MCHBAR offset 260h, bits 22-23 for channel 0 and MCHBAR offset 660h, bits 22-23 for channel 1)	0h	RO_V
13	X2APIC_EN	Extended Interrupt Mode. 0b: Hardware does not support Extended APIC mode. 1b: Hardware supports Extended APIC mode.	0h	RO_V
12	PDCD	0: Capable of Dual Channels 1: Not Capable of Dual Channel - only single channel capable.	0h	RO_V
11:0	RSVD	Reserved.	000h	RO

3.2.22 CAPID0—Capabilities B

Control of bits in this register are only required for customer visible SKU differentiation.

B/D/F/Type: 0/2/0/CFG			Access: RO_V	
Size: 32	Default Value: 00000000h		Address Offset: 48h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:29	RSVD	Reserved.	0h	RO
28	SMT	This setting indicates whether or not the processor is SMT capable.	0h	RO_V
27:25	CACHESZ	This setting indicates the supporting cache sizes.	0h	RO_V
24	RSVD	Reserved.	0h	RO
23:21	PLL_REF100_CFG	DDR3 Maximum Frequency Capability with 100 Memory. Maximum allowed memory frequency with 100 MHz ref clk. 0 - 100 MHz ref disabled	0h	RO_V

continued...

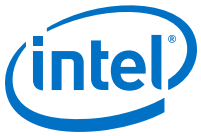

B/D/F/Type: 0/2/0/CFG			Access: RO_V	
Size: 32	**Default Value: 00000000h**		**Address Offset: 48h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		1 - upto DDR-1400 (7 x 200) 2 - upto DDR-1600 (8 x 200) 3 - upto DDR-1800 (8 x 200) 4 - upto DDR-2000 (10 x 200) 5 - upto DDR-2200 (11 x 200) 6 - upto DDR-2400 (12 x 200) 7 - no limit (but still limited by _DDR_FREQ200 to 3000)		
20	PEGG3_DIS	0: Capable of running any of the Gen 3-compliant PEG controllers in Gen 3 mode (Devices 0/1/0, 0/1/1, 0/1/2) 1: Not capable of running any of the PEG controllers in Gen 3 mode	0h	RO_V
19	RSVD	Reserved.	0h	RO
18	ADDGFXEN	0 - Additive Graphics Disabled 1- Additive Graphics Enabled	0h	RO_V
17	ADDGFXCAP	0 - Capable of Additive Graphics 1 - Not capable of Additive Graphics	0h	RO_V
16:7	RSVD	Reserved.	000h	RO
6:4	DMFC	This field controls which values may be written to the Memory Frequency Select field 6:4 of the Clocking Configuration registers (MCHBAR Offset C00h). Any attempt to write an unsupported value will be ignored. 000: No limit (but still limited by %MAX_DDR_FREQ266 to 3200) 001: MC capable of up to DDR3 2667 010: MC capable of up to DDR3 2400 011: MC capable of up to DDR3 2133 100: MC capable of up to DDR3 1867 101: MC capable of up to DDR3 1600 110: MC capable of up to DDR3 1333 111: MC capable of up to DDR3 1067	0h	RO_V
3:0	RSVD	Reserved.	0h	RO

3.2.23 DEVEN0—Device Enable

Allows for enabling/disabling of PCI devices and functions that are within the processor package. The table below the bit definitions describes the behavior of all combinations of transactions to devices controlled by this register.
All the bits in this register are Intel TXT Lockable.

B/D/F/Type: 0/2/0/CFG			Access: RO; RO_V	
Size: 32	**Default Value: 000000BFh**		**Address Offset: 54h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:15	RSVD	Reserved.	00000h	RO
14	D7EN	0: Bus 0 Device 7 is disabled and not visible. 1: Bus 0 Device 7 is enabled and visible. Non-production BIOS code should provide a setup option to enable Bus 0 Device 7. When enabled, Bus 0 Device 7 must be initialized in accordance to standard PCI device initialization procedures.	0h	RO_V

continued...

B/D/F/Type: 0/2/0/CFG			Access: RO; RO_V	
Size: 32	Default Value: 000000BFh		Address Offset: 54h	
Bit Range	Acronym	Description	Default	Access
13:11	RSVD	Reserved.	0h	RO
10	D5EN	0: Bus 0 Device 5 is disabled and not visible. 1: Bus 0 Device 5 is enabled and visible. This bit will be set to 0b and remain 0b if Device 5 capability is disabled.	0h	RO
9:8	RSVD	Reserved.	0h	RO
7	D4EN	0: Bus 0 Device 4 is disabled and not visible. 1: Bus 0 Device 4 is enabled and visible. This bit will be set to 0b and remain 0b if Device 4 capability is disabled.	1h	RO_V
6	RSVD	Reserved.	0h	RO
5	D3EN	0: Bus 0 Device 3 is disabled and hidden 1: Bus 0 Device 3 is enabled and visible This bit will be set to 0b and remain 0b if Device 3 capability is disabled.	1h	RO_V
4	D2EN	0: Bus 0 Device 2 is disabled and hidden 1: Bus 0 Device 2 is enabled and visible This bit will be set to 0b and remain 0b if Device 2 capability is disabled.	1h	RO_V
3	D1F0EN	0: Bus 0 Device 1 Function 0 is disabled and hidden. 1: Bus 0 Device 1 Function 0 is enabled and visible. This bit will be set to 0b and remain 0b if PEG10 capability is disabled.	1h	RO_V
2	D1F1EN	0: Bus 0 Device 1 Function 1 is disabled and hidden. 1: Bus 0 Device 1 Function 1 is enabled and visible.	1h	RO_V
1	D1F2EN	0: Bus 0 Device 1 Function 2 is disabled and hidden. 1: Bus 0 Device 1 Function 2 is enabled and visible.	1h	RO_V
0	D0EN	Bus 0 Device 0 Function 0 may not be disabled and is therefore hardwired to 1.	1h	RO

3.2.24 MSAC—Multi Size Aperture Control

This register determines the size of the graphics memory aperture in function 0 and in the trusted space. Only the system BIOS will write this register based on pre- boot address allocation efforts, but the graphics may read this register to determine the correct aperture size. System BIOS needs to save this value on boot so that it can reset it correctly during S3 resume.
This register is Intel TXT locked, becomes read-only when trusted environment is launched.

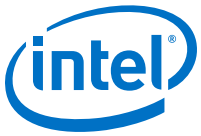

B/D/F/Type: 0/2/0/CFG			Access: RW_KV; RW	
Size: 8	**Default Value: 01h**		**Address Offset: 62h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:5	RSVDRW	Scratch Bits Only -- Have no physical effect on hardware	0h	RW
4	APSZ4	This field is used in conjuction with other APSZ* fields to determine the size of Aperture (GMADR) and affects certain bits of GMADR register. The description below is for all APSZ* fields 4:0 - 00000 = 128MB => GMADR.B[26:4] is hardwired to 0 00001 = 256MB => GMADR.B[27] = 0, RO 00010 = illegal (hardware will treat this as 00011) 00011 = 512MB => GMADR.B[28:27] = 0, RO 0100-00110 = illegal (hardware will treat this as 00111) 00111 = 1024MB => GMADR.B[29:27] = 0, RO 000-01110 = illegal (hardware will treat this as 01111) 01111 = 2048MB => GMADR.B[30:27] = 0, RO 10000-11110 = illegal (hardware will treat this as 11111) 11111 = 4096MB => GMADR.B[31:27] = 0, RO	0h	RW_KV
3	APSZ3	This field is used in conjuction with other APSZ* fields to determine the size of Aperture (GMADR) and affects certain bits of GMADR register. The description below is for all APSZ* fields 4:0 - 00000 = 128MB => GMADR.B[26:4] is hardwired to 0 00001 = 256MB => GMADR.B[27] = 0, RO 00010 = illegal (hardware will treat this as 00011) 00011 = 512MB => GMADR.B[28:27] = 0, RO 0100-00110 = illegal (hardware will treat this as 00111) 00111 = 1024MB => GMADR.B[29:27] = 0, RO 000-01110 = illegal (hardware will treat this as 01111) 01111 = 2048MB => GMADR.B[30:27] = 0, RO 10000-11110 = illegal (hardware will treat this as 11111) 11111 = 4096MB => GMADR.B[31:27] = 0, RO	0h	RW_KV
2	APSZ2	This field is used in conjuction with other APSZ* fields to determine the size of Aperture (GMADR) and affects certain bits of GMADR register. The description below is for all APSZ* fields 4:0 - 00000 = 128MB => GMADR.B[26:4] is hardwired to 0 00001 = 256MB => GMADR.B[27] = 0, RO 00010 = illegal (hardware will treat this as 00011) 00011 = 512MB => GMADR.B[28:27] = 0, RO 0100-00110 = illegal (hardware will treat this as 00111) 00111 = 1024MB => GMADR.B[29:27] = 0, RO 000-01110 = illegal (hardware will treat this as 01111) 01111 = 2048MB => GMADR.B[30:27] = 0, RO 10000-11110 = illegal (hardware will treat this as 11111) 11111 = 4096MB => GMADR.B[31:27] = 0, RO	0h	RW_KV
1	APSZ1	This field is used in conjuction with other APSZ* fields to determine the size of Aperture (GMADR) and affects certain bits of GMADR register. The description below is for all APSZ* fields 4:0 - 00000 = 128MB => GMADR.B[26:4] is hardwired to 0 00001 = 256MB => GMADR.B[27] = 0, RO 00010 = illegal (hardware will treat this as 00011) 00011 = 512MB => GMADR.B[28:27] = 0, RO 0100-00110 = illegal (hardware will treat this as 00111) 00111 = 1024MB => GMADR.B[29:27] = 0, RO 000-01110 = illegal (hardware will treat this as 01111)	0h	RW_KV

continued...

B/D/F/Type: 0/2/0/CFG			Access: RW_KV; RW	
Size: 8	Default Value: 01h		Address Offset: 62h	
Bit Range	Acronym	Description	Default	Access
		01111 = 2048MB => GMADR.B[30:27] = 0, RO 10000-11110 = illegal (hardware will treat this as 11111) 11111 = 4096MB => GMADR.B[31:27] = 0, RO		
0	APSZ0	This field is used in conjuction with other APSZ* fields to determine the size of Aperture (GMADR) and affects certain bits of GMADR register. The description below is for all APSZ* fields 4:0 - 00000 = 128MB => GMADR.B[26:4] is hardwired to 0 00001 = 256MB => GMADR.B[27] = 0, RO 00010 = illegal (hardware will treat this as 00011) 00011 = 512MB => GMADR.B[28:27] = 0, RO 0100-00110 = illegal (hardware will treat this as 00111) 00111 = 1024MB => GMADR.B[29:27] = 0, RO 000-01110 = illegal (hardware will treat this as 01111) 01111 = 2048MB => GMADR.B[30:27] = 0, RO 10000-11110 = illegal (hardware will treat this as 11111) 11111 = 4096MB => GMADR.B[31:27] = 0, RO	1h	RW_KV

3.2.25 MSI—Message Signaled Interrupts Capability ID

When a device supports MSI it can generate an interrupt request to the processor by writing a predefined data item (a message) to a predefined memory address. The reporting of the existence of this capability can be disabled by setting MSICH (CAPL[0] @ 7Fh). In that case walking this linked list will skip this capability and instead go directly to the PCI PM capability.

B/D/F/Type: 0/2/0/CFG			Access: RO	
Size: 16	Default Value: D005h		Address Offset: 90h	
Bit Range	Acronym	Description	Default	Access
15:8	POINTNEXT	This contains a pointer to the next item in the capabilities list which is the Power Management capability.	D0h	RO
7:0	CAPID	Value of 05h identifies this linked list item (capability structure) as being for MSI registers.	05h	RO

3.2.26 MC—Message Control

Message Signaled Interrupt control register. System software can modify bits in this register, but the device is prohibited from doing so. If the device writes the same message multiple times, only one of those messages is guaranteed to be serviced. If all of them must be serviced, the device must not generate the same message again until the driver services the earlier one.

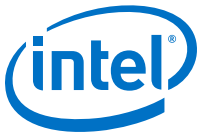

B/D/F/Type: 0/2/0/CFG			Access: RW; RO	
Size: 16	**Default Value: 0000h**		**Address Offset: 92h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:8	RSVD	Reserved.	00h	RO
7	CAP64B	Hardwired to 0 to indicate that the function does not implement the upper 32 bits of the Message address register and is incapable of generating a 64-bit memory address.	0h	RO
6:4	MME	System software programs this field to indicate the actual number of messages allocated to this device. This number will be equal to or less than the number actually requested. The encoding is the same as for the MMC field below.	0h	RW
3:1	MMC	System Software reads this field to determine the number of messages being requested by this device. Value: Number of requests 000: 1 All of the following are reserved in this implementation 001: 2 010: 4 011: 8 100: 16 101: 32 110: Reserved 111: Reserved	0h	RO
0	MSIEN	Controls the ability of this device to generate MSIs.	0h	RW

3.2.27 MA—Message Address

This register contains the Message Address for MSIs sent by the device.

B/D/F/Type: 0/2/0/CFG			Access: RO; RW	
Size: 32	**Default Value: 00000000h**		**Address Offset: 94h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:2	MESSADD	Used by system software to assign an MSI address to the device. The device handles an MSI by writing the padded contents of the MD register to this address.	00000000h	RW
1:0	FDWORD	Hardwired to 0 so that addresses assigned by system software are always aligned on a DWORD address boundary.	0h	RO

3.2.28 MD—Message Data

This register contains the Message Data for MSIs sent by the device.

B/D/F/Type: 0/2/0/CFG			Access: RW	
Size: 16	**Default Value: 0000h**		**Address Offset: 98h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:0	MESSDATA	Base message data pattern assigned by system software and used to handle an MSI from the device. When the device must generate an interrupt request, it	0000h	RW

B/D/F/Type: 0/2/0/CFG			Access: RW	
Size: 16	Default Value: 0000h		Address Offset: 98h	
Bit Range	Acronym	Description	Default	Access
		writes a 32-bit value to the memory address specified in the MA register. The upper 16 bits are always set to 0. The lower 16 bits are supplied by this register.		

3.2.29 AFCIDNP—Advanced Features Capabilities Identifier and Next Pointer

When this capability is linked into the list, the second function of the Internal Graphics Device can be reset independently of the first function.

B/D/F/Type: 0/2/0/CFG			Access: RO	
Size: 16	Default Value: 0013h		Address Offset: A4h	
Bit Range	Acronym	Description	Default	Access
15:8	NEXT_PTR	This contains a pointer to next item in capabilities list. This is the final capability in the list and must be set to 00h.	00h	RO
7:0	CAP_ID	A value of 13h identifies that this PCI Function is capable of Advanced Features.	13h	RO

3.2.30 AFCTL—Advanced Features Control

See Conventional PCI 3.0 Specification ECN for Advanced Capabilities, July 27th, 2006

B/D/F/Type: 0/2/0/CFG			Access: RW1S	
Size: 8	Default Value: 00h		Address Offset: A8h	
Bit Range	Acronym	Description	Default	Access
7:1	RSVD	Reserved.	00h	RO
0	INIT_FLR	A write of 1b initiates Function Level Reset (FLR). FLR requirements are defined in the PCI Express Base Specification. Registers and state information that do not apply to conventional PCI are exempt from the FLR requirements given there. Once written 1, FLR will be initiated. During FLR, a read will return 1`s since device 2 reads abort. Once FLR completes, hardware will clear the bit to 0.	0h	RW1S

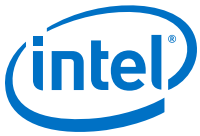

3.2.31 AFSTS—Advanced Features Status

See Conventional PCI 3.0 Specification ECN for Advanced Capabilities, July 27th, 2006

B/D/F/Type: 0/2/0/CFG			Access: RO	
Size: 8	Default Value: 00h		Address Offset: A9h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:1	RSVD	Reserved.	00h	RO
0	TP	1: The Function has issued one or more non-posted transactions which have not been completed, including non-posted transactions that a target has terminated with Retry. 0: All non-posted transactions have been completed.	0h	RO

3.2.32 PMCAPID—Power Management Capabilities ID

This register contains the PCI Power Management Capability ID and the next capability pointer.

B/D/F/Type: 0/2/0/CFG			Access: RO	
Size: 16	Default Value: A401h		Address Offset: D0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:8	NEXT_PTR	This contains a pointer to the next item in the capabilities list.	A4h	RO
7:0	CAP_ID	SIG defines this ID is 01h for power management.	01h	RO

3.2.33 PMCAP—Power Management Capabilities

This register provides information on the capabilities of the function related to powermanagement.

B/D/F/Type: 0/2/0/CFG			Access: RO	
Size: 16	Default Value: 0022h		Address Offset: D2h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:11	PMES	This field indicates the power states in which the IGD may assert PME#. Hardwired to 0 to indicate that the IGD does not assert the PME# signal.	00h	RO
10	D2	The D2 power management state is not supported. This bit is hardwired to 0.	0h	RO
9	D1	Hardwired to 0 to indicate that the D1 power management state is not supported.	0h	RO
8:6	RSVD	Reserved.	0h	RO
5	DSI	Hardwired to 1 to indicate that special initialization of the IGD is required before generic class device driver is to use it.	1h	RO

continued...

B/D/F/Type: 0/2/0/CFG			Access: RO	
Size: 16	Default Value: 0022h		Address Offset: D2h	
Bit Range	Acronym	Description	Default	Access
4	RSVD	Reserved.	0h	RO
3	PMECLK	Hardwired to 0 to indicate IGD does not support PME# generation.	0h	RO
2:0	VER	Hardwired to 010b to indicate that there are 4 bytes of power management registers implemented and that this device complies with revision 1.1 of the PCI Power Management Interface Specification.	2h	RO

3.2.34 PMCS—Power Management Control/Status

B/D/F/Type: 0/2/0/CFG			Access: RO_V; RO	
Size: 16	Default Value: 0000h		Address Offset: D4h	
Bit Range	Acronym	Description	Default	Access
15	PMESTS	This bit is 0 to indicate that IGD does not support PME# generation from D3 (cold).	0h	RO
14:13	DSCALE	The IGD does not support data register. This bit always returns 00 when read, write operations have no effect.	0h	RO
12:9	DSEL	The IGD does not support data register. This bit always returns 0h when read, write operations have no effect.	0h	RO
8	PMEEN	This bit is 0 to indicate that PME# assertion from D3 (cold) is disabled.	0h	RO
7:2	RSVD	Reserved.	00h	RO
1:0	PWRSTAT	This field indicates the current power state of the IGD and can be used to set the IGD into a new power state. If software attempts to write an unsupported state to this field, write operation must complete normally on the bus, but the data is discarded and no state change occurs. On a transition from D3 to D0 the graphics controller is optionally reset to initial values. Bits[1:0] Power state 00: D0 Default 01: D1 Not Supported 10: D2 Not Supported 11: D3	0h	RO_V

3.3 Audio Controller Registers Summary

Offset	Register ID—Description	Default Value	Access
0	VID—Vendor Identification on page 97	8086h	RO
2	DID—Device ID on page 98	0C0Ch	RO
4	PCICMD—PCI Command on page 98	0000h	RO; RW_V
6	STS—PCI Status on page 98	0010h	RO_V; RO
8	RID—Revision Identification on page 99	00h	RO
9	PI—Programming Interface on page 99	00h	RO

continued...

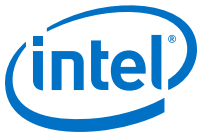

Offset	Register ID—Description	Default Value	Access
A	SCC—Sub Class Code on page 99	03h	RO
B	BCC—Base Class Code on page 100	04h	RO
C	CLS—Cache Line Size on page 100	00h	RW_V
10	HDALBAR—IntelÂ® HD Audio Base Lower Address on page 100	00000004h	RO; RW_V
14	HDAHBAR—IntelÂ® HD Audio Base Upper Address on page 100	00000000h	RW_V; RW
2C	SVID—Subsystem Vendor ID on page 101	00000000h	RW_O
34	CAPPTR—Capability Pointer on page 101	50h	RO
3C	INTLN—Interrupt Line on page 101	00h	RW_V
3D	INTPN—Interrupt Pin on page 101	01h	RO
44	CAPID0—Capabilities A on page 102	00000000h	RO_V
48	CAPID0—Capabilities B on page 102	00000000h	RO_V
4C	DEVEN—Device Enable on page 103	000000BFh	RO; RO_V
50	PID—PCI Power Management Capability ID on page 104	00026001h	RO; RW_O
54	PMCS—Power Management Control And Status on page 105	00000000h	RO_V; RO
60	MID—MSI Capability ID on page 105	7005h	RO
62	MMC—MSI Message Control on page 106	0000h	RW_V; RO
64	MMA—MSI Message Lower Address on page 106	00000000h	RO; RW_V
68	MMD—MSI Message Data on page 106	0000h	RW_V
70	PXID—PCI Express Capability ID on page 106	0010h	RO
72	PXC—PCI Express Capabilities on page 107	0091h	RO
74	DEVCAP—Device Capabilities on page 107	10000FC0h	RO
78	DEVC—Device Control on page 108	0800h	RW_V; RO; RW; RW1S
7A	DEVS—Device Status on page 108	0000h	RO; RO_V

3.3.1 VID—Vendor Identification

This register combined with the Device Identification register uniquely identify any PCI device.

B/D/F/Type: 0/3/0/CFG			Access: RO	
Size: 16	Default Value: 8086h		Address Offset: 0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:0	VID	Indicates that Intel is the vendor.	8086h	RO

3.3.2 DID—Device ID

This register combined with the Vendor Identification register uniquely identifies any PCI device.

B/D/F/Type: 0/3/0/CFG			Access: RO	
Size: 16	Default Value: 0C0Ch		Address Offset: 2h	
Bit Range	Acronym	Description	Default	Access
15:0	DID	Indicates the device number assigned.	0C0Ch	RO

3.3.3 PCICMD—PCI Command

B/D/F/Type: 0/3/0/CFG			Access: RO; RW_V	
Size: 16	Default Value: 0000h		Address Offset: 4h	
Bit Range	Acronym	Description	Default	Access
15:11	RSVD	Reserved.	00h	RO
10	INTDIS	Enables the device to assert an INTx#. When set, the Intel(r) HD Audio controller's INTx# signal will be de-asserted. When cleared, the INTx# signal may be asserted. Note that this bit does not affect the generation of MSI's.	0h	RW_V
9	FB2B	Not implemented. Hardwired to 0.	0h	RO
8	SEN	Functionality not implemented.	0h	RW_V
7	WCC	Not implemented. Hardwired to 0.	0h	RO
6	PER	Functionality not implemented.	0h	RW_V
5	VPS	Not implemented. Hardwired to 0.	0h	RO
4	MWIE	Not implemented. Hardwired to 0.	0h	RO
3	SCE	Not implemented. Hardwired to 0.	0h	RO
2	BME	1 = Enable, 0 = Disable. Controls standard PCI Express bus mastering capabilities for Memory and IO, reads and writes. Note that this also controls MSI generation since MSI are essentially Memory writes.	0h	RW_V
1	MAE	When set, enables memory space accesses to the Intel dHD Audio controller.	0h	RW_V
0	IOAE	The Intel dHD Audio controller does not implement IO Space, therefore this bit is hardwired to 0.	0h	RO

3.3.4 STS—PCI Status

B/D/F/Type: 0/3/0/CFG			Access: RO_V; RO	
Size: 16	Default Value: 0010h		Address Offset: 6h	
Bit Range	Acronym	Description	Default	Access
15	DPE	Not implemented. Hardwired to 0.	0h	RO
14	SERRS	Not implemented. Hardwired to 0.	0h	RO
13	RMA	Hardwired to '0' as master aborts are not tracked.	0h	RO

continued...

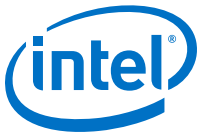

B/D/F/Type: 0/3/0/CFG			Access: RO_V; RO	
Size: 16	Default Value: 0010h		Address Offset: 6h	
Bit Range	Acronym	Description	Default	Access
12	RTA	Not implemented. Hardwired to 0.	0h	RO
11	STA	Not implemented. Hardwired to 0.	0h	RO
10:9	DEVT	Does not apply. Hardwired to 0.	0h	RO
8	MDPE	Not implemented. Hardwired to 0.	0h	RO
7	FBC	Does not apply. Hardwired to 0.	0h	RO
6	RSVD	Reserved.	0h	RO
5	C66	Does not apply. Hardwired to 0.	0h	RO
4	CLIST	Indicates that the controller contains a capabilities pointer list. The first item is pointed to by looking at configuration offset 34h.	1h	RO
3	IS	Reflects the state of the INTx# signal at the input of the enable/disable circuit. This bit is a 1 when the INTx# is asserted. This bit is a 0 after the interrupt is cleared (independent of the state of the Interrupt Disable bit in the command register). Note that this bit is not set by an MSI.	0h	RO_V
2:0	RSVD	Reserved.	0h	RO

3.3.5 RID—Revision Identification

B/D/F/Type: 0/3/0/CFG			Access: RO	
Size: 8	Default Value: 00h		Address Offset: 8h	
Bit Range	Acronym	Description	Default	Access
7:4	RID_MSB	Four MSB of RID	0h	RO
3:0	RID	Four LSB of RID	0h	RO

3.3.6 PI—Programming Interface

B/D/F/Type: 0/3/0/CFG			Access: RO	
Size: 8	Default Value: 00h		Address Offset: 9h	
Bit Range	Acronym	Description	Default	Access
7:0	PI	Value assigned to the dHD Audio controller.	00h	RO

3.3.7 SCC—Sub Class Code

B/D/F/Type: 0/3/0/CFG			Access: RO	
Size: 8	Default Value: 03h		Address Offset: Ah	
Bit Range	Acronym	Description	Default	Access
7:0	SCC	This indicates the device is an IntelÂ® HD Audio audio device, in the context of a multimedia device.	03h	RO

3.3.8 BCC—Base Class Code

B/D/F/Type: 0/3/0/CFG			Access: RO	
Size: 8	Default Value: 04h		Address Offset: Bh	
Bit Range	Acronym	Description	Default	Access
7:0	BCC	This register indicates that the function implements a multimedia device.	04h	RO

3.3.9 CLS—Cache Line Size

B/D/F/Type: 0/3/0/CFG			Access: RW_V	
Size: 8	Default Value: 00h		Address Offset: Ch	
Bit Range	Acronym	Description	Default	Access
7:0	CLS	Does not apply to PCI Express. PCI Express spec requires this to be implemented as a RW register but has no functional impact.	00h	RW_V

3.3.10 HDALBAR—IntelÂ® HD Audio Base Lower Address

B/D/F/Type: 0/3/0/CFG			Access: RO; RW_V	
Size: 32	Default Value: 00000004h		Address Offset: 10h	
Bit Range	Acronym	Description	Default	Access
31:14	MBA	Base address for the IntelÂ® HD Audio controllers memory mapped configuration registers. 16 Kbytes are requested by hardwiring bits 13:4 to 0s.	00000h	RW_V
13:4	ADM	Should be harwired to 0.	000h	RO
3	PREFMEM	Indicates that this BAR is NOT pre-fetchable.	0h	RO
2:1	MEMTYP	Indicates that this BAR can be located anywhere in 64-bit address space.	2h	RO
0	MIOS	Indicates that this BAR is located in memory space.	0h	RO

3.3.11 HDAHBAR—IntelÂ® HD Audio Base Upper Address

B/D/F/Type: 0/3/0/CFG			Access: RW_V; RW	
Size: 32	Default Value: 00000000h		Address Offset: 14h	
Bit Range	Acronym	Description	Default	Access
31:7	UMBA_RES	Reserved 25 bits of the Upper Base address for the Intel(r) HD Audio controller's memory mapped configuration registers.	0000000h	RW
6:0	UMBA	These are the lower 7 bits of the upper base address which can be programmed	00h	RW_V

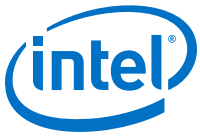

3.3.12 SVID—Subsystem Vendor ID

This value is used to identify the vendor of the subsystem.

B/D/F/Type: 0/3/0/CFG			Access: RW_O	
Size: 32	Default Value: 00000000h		Address Offset: 2Ch	
Bit Range	Acronym	Description	Default	Access
31:16	SID	Subsystem ID	0000h	RW_O
15:0	SVID	Subsystem Vendor ID	0000h	RW_O

3.3.13 CAPPTR—Capability Pointer

B/D/F/Type: 0/3/0/CFG			Access: RO	
Size: 8	Default Value: 50h		Address Offset: 34h	
Bit Range	Acronym	Description	Default	Access
7:0	CP	Indicates that the first capability pointer offset is offset 50h (Power Management Capability).	50h	RO

3.3.14 INTLN—Interrupt Line

B/D/F/Type: 0/3/0/CFG			Access: RW_V	
Size: 8	Default Value: 00h		Address Offset: 3Ch	
Bit Range	Acronym	Description	Default	Access
7:0	INTLN	Hardware does not use this field directly. It is used to communicate to software the interrupt line that the interrupt pin is connected to. This register is not affected by FLR.	00h	RW_V

3.3.15 INTPN—Interrupt Pin

B/D/F/Type: 0/3/0/CFG			Access: RO	
Size: 8	Default Value: 01h		Address Offset: 3Dh	
Bit Range	Acronym	Description	Default	Access
7:4	RSVD	Reserved.	0h	RO
3:0	IP	Interrupt Pin	1h	RO

3.3.16 CAPID0—Capabilities A

Control of bits in this register are only required for customer visible SKU differentiation.

B/D/F/Type: 0/3/0/CFG			Access: RO_V	
Size: 32	Default Value: 00000000h		Address Offset: 44h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:26	RSVD	Reserved.	00h	RO
25	ECCDIS	0b ECC capable 1b Not ECC capable	0h	RO_V
24	RSVD	Reserved.	0h	RO
23	VTDD	0: Enable VTd 1: Disable VTd	0h	RO_V
22:15	RSVD	Reserved.	00h	RO
14	DDPCD	Allows Dual Channel operation but only supports 1 DIMM per channel. 0: 2 DIMMs per channel enabled 1: 2 DIMMs per channel disabled. This setting hardwires bits 2 and 3 of the rank population field for each channel to zero. (MCHBAR offset 260h, bits 22-23 for channel 0 and MCHBAR offset 660h, bits 22-23 for channel 1)	0h	RO_V
13	X2APIC_EN	Extended Interrupt Mode. 0b: Hardware does not support Extended APIC mode. 1b: Hardware supports Extended APIC mode.	0h	RO_V
12	PDCD	0: Capable of Dual Channels 1: Not Capable of Dual Channel - only single channel capable.	0h	RO_V
11:0	RSVD	Reserved.	000h	RO

3.3.17 CAPID0—Capabilities B

Control of bits in this register are only required for customer visible SKU differentiation.

B/D/F/Type: 0/3/0/CFG			Access: RO_V	
Size: 32	Default Value: 00000000h		Address Offset: 48h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:29	RSVD	Reserved.	0h	RO
28	SMT	This setting indicates whether or not the processor is SMT capable.	0h	RO_V
27:25	CACHESZ	This setting indicates the supporting cache sizes.	0h	RO_V
24	RSVD	Reserved.	0h	RO
23:21	PLL_REF100_CFG	DDR3 Maximum Frequency Capability with 100 Memory. Maximum allowed memory frequency with 100 MHz ref clk. 0 - 100 MHz ref disabled	0h	RO_V

continued...

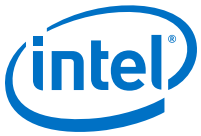

B/D/F/Type: 0/3/0/CFG			Access: RO_V	
Size: 32	**Default Value: 00000000h**		**Address Offset: 48h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		1 - upto DDR-1400 (7 x 200) 2 - upto DDR-1600 (8 x 200) 3 - upto DDR-1800 (8 x 200) 4 - upto DDR-2000 (10 x 200) 5 - upto DDR-2200 (11 x 200) 6 - upto DDR-2400 (12 x 200) 7 - no limit (but still limited by _DDR_FREQ200 to 3000)		
20	PEGG3_DIS	0: Capable of running any of the Gen 3-compliant PEG controllers in Gen 3 mode (Devices 0/1/0, 0/1/1, 0/1/2) 1: Not capable of running any of the PEG controllers in Gen 3 mode	0h	RO_V
19	RSVD	Reserved.	0h	RO
18	ADDGFXEN	0 - Additive Graphics Disabled 1- Additive Graphics Enabled	0h	RO_V
17	ADDGFXCAP	0 - Capable of Additive Graphics 1 - Not capable of Additive Graphics	0h	RO_V
16:7	RSVD	Reserved.	000h	RO
6:4	DMFC	This field controls which values may be written to the Memory Frequency Select field 6:4 of the Clocking Configuration registers (MCHBAR Offset C00h). Any attempt to write an unsupported value will be ignored. 000: No limit (but still limited by %MAX_DDR_FREQ266 to 3200) 001: MC capable of up to DDR3 2667 010: MC capable of up to DDR3 2400 011: MC capable of up to DDR3 2133 100: MC capable of up to DDR3 1867 101: MC capable of up to DDR3 1600 110: MC capable of up to DDR3 1333 111: MC capable of up to DDR3 1067	0h	RO_V
3:0	RSVD	Reserved.	0h	RO

3.3.18 DEVEN—Device Enable

Allows for enabling/disabling of PCI devices and functions that are within the processor package. The table below the bit definitions describes the behavior of all combinations of transactions to devices controlled by this register.
All the bits in this register are Intel TXT Lockable.

B/D/F/Type: 0/3/0/CFG			Access: RO; RO_V	
Size: 32	**Default Value: 000000BFh**		**Address Offset: 4Ch**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:15	RSVD	Reserved.	00000h	RO
14	D7EN	0: Bus 0 Device 7 is disabled and not visible. 1: Bus 0 Device 7 is enabled and visible. Non-production BIOS code should provide a setup option to enable Bus 0 Device 7. When enabled, Bus 0 Device 7 must be initialized in accordance to standard PCI device initialization procedures.	0h	RO_V

continued...

B/D/F/Type: 0/3/0/CFG			Access: RO; RO_V	
Size: 32	Default Value: 000000BFh		Address Offset: 4Ch	
Bit Range	Acronym	Description	Default	Access
13:11	RSVD	Reserved.	0h	RO
10	D5EN	0: Bus 0 Device 5 is disabled and not visible. 1: Bus 0 Device 5 is enabled and visible. This bit will be set to 0b and remain 0b if Device 5 capability is disabled.	0h	RO
9:8	RSVD	Reserved.	0h	RO
7	D4EN	0: Bus 0 Device 4 is disabled and not visible. 1: Bus 0 Device 4 is enabled and visible. This bit will be set to 0b and remain 0b if Device 4 capability is disabled.	1h	RO_V
6	RSVD	Reserved.	0h	RO
5	D3EN	0: Bus 0 Device 3 is disabled and hidden 1: Bus 0 Device 3 is enabled and visible This bit will be set to 0b and remain 0b if Device 3 capability is disabled.	1h	RO_V
4	D2EN	0: Bus 0 Device 2 is disabled and hidden 1: Bus 0 Device 2 is enabled and visible This bit will be set to 0b and remain 0b if Device 2 capability is disabled.	1h	RO_V
3	D1F0EN	0: Bus 0 Device 1 Function 0 is disabled and hidden. 1: Bus 0 Device 1 Function 0 is enabled and visible. This bit will be set to 0b and remain 0b if PEG10 capability is disabled.	1h	RO_V
2	D1F1EN	0: Bus 0 Device 1 Function 1 is disabled and hidden. 1: Bus 0 Device 1 Function 1 is enabled and visible.	1h	RO_V
1	D1F2EN	0: Bus 0 Device 1 Function 2 is disabled and hidden. 1: Bus 0 Device 1 Function 2 is enabled and visible.	1h	RO_V
0	D0EN	Bus 0 Device 0 Function 0 may not be disabled and is therefore hardwired to 1.	1h	RO

3.3.19 PID—PCI Power Management Capability ID

B/D/F/Type: 0/3/0/CFG			Access: RO; RW_O	
Size: 32	Default Value: 00026001h		Address Offset: 50h	
Bit Range	Acronym	Description	Default	Access
31:27	PMES	PME# cannot be generated.	00h	RO
26	D2S	The D2 state is not supported.	0h	RO
25	D1S	The D1 state is not supported.	0h	RO
24:22	AC	Not Supported.	0h	RO
21	DSI	Indicates that no device-specificinitialization is required.	0h	RO
20	RSVD	Reserved.	0h	RO
19	PMEC	Does not apply. Hardwired to 0.	0h	RO

continued...

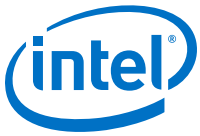

B/D/F/Type: 0/3/0/CFG			Access: RO; RW_O	
Size: 32	**Default Value: 00026001h**		**Address Offset: 50h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
18:16	VS	Indicates support for Revision 1.1 of the PCI PowerManagement Specification.	2h	RO
15:8	NEXT	Points to the next capability structure (MSI).	60h	RW_O
7:0	CAP	Indicates that this pointer is a PCI power management capability	01h	RO

3.3.20 PMCS—Power Management Control And Status

B/D/F/Type: 0/3/0/CFG			Access: RO_V; RO	
Size: 32	**Default Value: 00000000h**		**Address Offset: 54h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15	PMESTS	PME# Cannot be generated.	0h	RO
14:9	RSVD	Reserved.	00h	RO
8	PMEEN	Cannot be generated.	0h	RO
7:2	RSVD	Reserved.	00h	RO
1:0	PWRSTAT	This field is used both to determine the current power state of the Intel HD Audio controller and to set a new power state. The values are:00 - D0 state 11 - D3HOT state If software attempts to write a value of 10b or 01b in to this field, the writeoperation must complete normally; however, the data is discarded and nostate change occurs. When in the D3HOT states, the Intel HD Audio controller's configuration spaceis available, but the I/O and memory spaces are not. Additionally, interrupts are blocked.When software changes this value from the D3 HOT state to the D0 state, an internal warm (soft) reset is generated, and software must re-initialize the function.	0h	RO_V

3.3.21 MID—MSI Capability ID

B/D/F/Type: 0/3/0/CFG			Access: RO	
Size: 16	**Default Value: 7005h**		**Address Offset: 60h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:8	NEXT	Points to the PCI Express* capability structure.	70h	RO
7:0	CAP	Indicates that this pointer is a MSI capability	05h	RO

3.3.22 MMC—MSI Message Control

B/D/F/Type: 0/3/0/CFG			Access: RW_V; RO	
Size: 16	Default Value: 0000h		Address Offset: 62h	
Bit Range	Acronym	Description	Default	Access
15:8	RSVD	Reserved.	00h	RO
7	ADD64	Indicates the ability to generate a 64-bit message address	0h	RO
6:4	MME	Normally this is a RW register. However, since only 1 message is supported, these bits are hardwired to 000 = 1 message.	0h	RO
3:1	MMC	Hardwired to 0 indicating request for 1 message	0h	RO
0	ME	If set to 1 an MSI will be generated instead of an INTx#signal. If set to 0, an MSI may not be generated.	0h	RW_V

3.3.23 MMA—MSI Message Lower Address

B/D/F/Type: 0/3/0/CFG			Access: RO; RW_V	
Size: 32	Default Value: 00000000h		Address Offset: 64h	
Bit Range	Acronym	Description	Default	Access
31:2	MA	Address used for MSI Message.	00000000h	RW_V
1:0	RESRV	Does not apply. Hardwired to 0.	0h	RO

3.3.24 MMD—MSI Message Data

B/D/F/Type: 0/3/0/CFG			Access: RW_V	
Size: 16	Default Value: 0000h		Address Offset: 68h	
Bit Range	Acronym	Description	Default	Access
15:0	MD	Data used for MSI Message.	0000h	RW_V

3.3.25 PXID—PCI Express Capability ID

B/D/F/Type: 0/3/0/CFG			Access: RO	
Size: 16	Default Value: 0010h		Address Offset: 70h	
Bit Range	Acronym	Description	Default	Access
15:8	NEXT	Indicates that this is the last capability structure in the list.	00h	RO
7:0	CAP	Indicates that this pointer is a PCI Express capability structure.	10h	RO

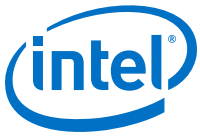

3.3.26 PXC—PCI Express Capabilities

B/D/F/Type: 0/3/0/CFG			Access: RO	
Size: 16	Default Value: 0091h		Address Offset: 72h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:14	RSVD	Reserved.	0h	RO
13:9	IMN	Hardwired to 0.	00h	RO
8	SI	Hardwired to 0.	0h	RO
7:4	DPT	Indicates that this is a Root Complex IntegratedEndpoint Device.	9h	RO
3:0	CV	Indicates version #1 PCI Express capability	1h	RO

3.3.27 DEVCAP—Device Capabilities

B/D/F/Type: 0/3/0/CFG			Access: RO	
Size: 32	Default Value: 10000FC0h		Address Offset: 74h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:29	RSVD	Reserved.	0h	RO
28	FLR	A indicates that the Display HD Audio controller supports the Function Level Reset capability.	1h	RO
27:26	SPLS	Hardwired to 0.	0h	RO
25:18	SPLV	Hardwired to 0.	00h	RO
17:15	RSVD	Reserved.	0h	RO
14	PIP	Hardwired to 0.	0h	RO
13	AIP	Hardwired to 0.	0h	RO
12	ABP	Hardwired to 0.	0h	RO
11:9	L1CAP	Endpoint L1 Acceptable Latency	7h	RO
8:6	L0SCAP	Endpoint L0s Acceptable Latency	7h	RO
5	ETCAP	Indicates 5 bit tag supported.	0h	RO
4:3	PFCAP	Indicates phantom functions not supported.	0h	RO
2:0	MPCAP	Indicates 128B maximum payload size capability.	0h	RO

3.3.28 DEVC—Device Control

B/D/F/Type: 0/3/0/CFG			Access: RW_V; RO; RW; RW1S	
Size: 16	Default Value: 0800h		Address Offset: 78h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15	IFLR	Used to initiate FLR transition. A write of initiates FLR transition. Since hardware must not respond to any cycles until FLR completion, the read value by software from this bit is .	0h	RW1S
14:12	MRRS	Hardwired to 000 enabling 128 B maximum read request size.	0h	RO
11	NSNPEN	When set to 1 the IntelÂ® HD Audio controller is permitted to set the No Snoop bit in the Requester Attributes of a bus master transaction. In this case VC0, VCp, or VC1 may be used for isochronous transfers. When set to 0 the Intel HD Audio controller will not set the No Snoop bit. In the case isochronous transfers will not use VC1(VCi) even if it is enabled since VC1 is never snooped. Isochronous transfers will use either VCp or VC0. This bit is not affected by D3HOT to D0 reset or FLR.	1h	RW
10	AUXPEN	Hardwired to 0 indicating Intel HD Audio device does not draw AUX power.	0h	RO
9	PFEN	Hardwired to 0 disabling phantom functions.	0h	RO
8	ETEN	Hardwired to 0 enabling 5-bit tag.	0h	RO
7:5	MAXPAY	Hardwired to 000 indicating 128 B.	0h	RO
4	ROEN	Hardwired to 0 disabling relaxed ordering.	0h	RO
3	URREN	Functionality not implemented. This bit is RW to pass PCIe compliance testing.	0h	RW_V
2	FEREN	Functionality not implemented. This bit is RW to pass PCIe compliance testing.	0h	RW_V
1	NFEREN	Functionality not implemented. This bit is RW to pass PCIe compliance testing.	0h	RW_V
0	CEREN	Functionality not implemented. This bit is RW to pass PCIe compliance testing.	0h	RW_V

3.3.29 DEVS—Device Status

B/D/F/Type: 0/3/0/CFG			Access: RO; RO_V	
Size: 16	Default Value: 0000h		Address Offset: 7Ah	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:6	RSVD	Reserved.	000h	RO
5	TXP	A 1 indicates that the Intel HD Audio controller has issued Non-Posted requests which have not been completed. A 0 indicates that Completions for all Non-Posted Requests have been received.	0h	RO_V
4	AUXDET	Hardwired to 1 indicating the device is connected to Suspend power.	0h	RO

continued...

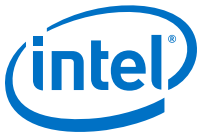

B/D/F/Type: 0/3/0/CFG			Access: RO; RO_V	
Size: 16	**Default Value: 0000h**		**Address Offset: 7Ah**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
3	URDET	Not implemented. Hardwired to 0.	0h	RO
2	FEDET	Not implemented. Hardwired to 0.	0h	RO
1	NFEDET	Not implemented. Hardwired to 0.	0h	RO
0	CEDET	Not implemented. Hardwired to 0.	0h	RO

4.0 PCI Express* (PCIe*) Registers

4.1 PCI Express Controller (x16) Registers Summary

Offset	Register ID—Description	Default Value	Access
0	VID—Vendor Identification on page 111	8086h	RO
2	DID—Device Identification on page 112	0C01h	RO
4	PCICMD—PCI Command on page 112	0000h	RW; RO
6	PCISTS—PCI Status on page 113	0010h	RO_V; RO; RW1C
8	RID—Revision Identification on page 115	00h	RO
9	CC—Class Code on page 116	060400h	RO
C	CL—Cache Line Size on page 116	00h	RW
E	HDR—Header Type on page 116	81h	RO
18	PBUSN—Primary Bus Number on page 117	00h	RO
19	SBUSN—Secondary Bus Number on page 117	00h	RW
1A	SUBUSN—Subordinate Bus Number on page 117	00h	RW
1C	IOBASE—I/O Base Address on page 117	F0h	RW
1D	IOLIMIT—I/O Limit Address on page 118	00h	RW
1E	SSTS—Secondary Status on page 118	0000h	RO; RW1C
20	MBASE—Memory Base Address on page 119	FFF0h	RW
22	MLIMIT—Memory Limit Address on page 120	0000h	RW
24	PMBASE—Prefetchable Memory Base Address on page 120	FFF1h	RO; RW
26	PMLIMIT—Prefetchable Memory Limit Address on page 121	0001h	RO; RW
28	PMBASEU—Prefetchable Memory Base Address Upper on page 121	00000000h	RW
2C	PMLIMITU—Prefetchable Memory Limit Address Upper on page 122	00000000h	RW
34	CAPPTR—Capabilities Pointer on page 122	88h	RO
3C	INTRLINE—Interrupt Line on page 123	00h	RW
3D	INTRPIN—Interrupt Pin on page 123	01h	RW_O; RO
3E	BCTRL—Bridge Control on page 124	0000h	RW; RO
80	PM—Power Management Capabilities on page 125	C8039001h	RO; RO_V
84	PM—Power Management Control/Status on page 126	00000008h	RW; RO
88	SS—Subsystem ID and Vendor ID Capabilities on page 127	0000800Dh	RO
8C	SS—Subsystem ID and Subsystem Vendor ID on page 128	00008086h	RW_O

continued...

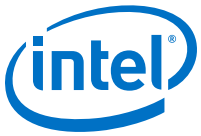

Offset	Register ID—Description	Default Value	Access
90	MSI—Message Signaled Interrupts Capability ID on page 128	A005h	RO
92	MC—Message Control on page 128	0000h	RW; RO
94	MA—Message Address on page 129	00000000h	RO; RW
98	MD—Message Data on page 129	0000h	RW
A0	PEG—PCI Express Capability List on page 130	0010h	RO
A2	PEG—PCI Express Capabilities on page 130	0142h	RO; RW_O
A4	DCAP—Device Capabilities on page 130	00008001h	RW_O; RO
A8	DCTL—Device Control on page 131	0020h	RW; RO
AA	DSTS—Device Status on page 132	0000h	RW1C; RO
B0	LCTL—Link Control on page 133	0000h	RW; RO; RW_V
B2	LSTS—Link Status on page 135	1001h	RO_V; RO; RW1C
B4	SLOTCAP—Slot Capabilities on page 136	00040000h	RO; RW_O
B8	SLOTCTL—Slot Control on page 138	0000h	RO
BA	SLOTSTS—Slot Status on page 140	0000h	RO; RW1C; RO_V
BC	RCTL—Root Control on page 141	00000000h	RW; RO; RWS_V; RO_V
C0	RSTS—Root Status on page 143	00000000h	RO_V; RW1C; RO
C4	DCAP2—Device Capabilites 2 on page 143	00000B80h	RO; RW_O
C8	DCTL2—Device Control 2 on page 145	0000h	RW; RO; RW_V
D0	LCTL2—Link Control 2 on page 146	0003h	RWS; RWS_V
D2	LSTS2—Link Status 2 on page 149	0000h	RO_V; RW1C
104	PVCCAP1—Port VC Capability Register 1 on page 150	00000000h	RO
108	PVCCAP2—Port VC Capability Register 2 on page 150	00000000h	RO
10C	PVCCTL—Port VC Control on page 150	0000h	RO; RW
110	VC0RCAP—VC0 Resource Capability on page 151	00000001h	RO
114	VC0RCTL—VC0 Resource Control on page 151	800000FFh	RO; RW
11A	VC0RSTS—VC0 Resource Status on page 152	0002h	RO_V

4.1.1 VID—Vendor Identification

This register combined with the Device Identification register uniquely identify any PCI device.

B/D/F/Type: 0/1/0/CFG			Access: RO	
Size: 16	Default Value: 8086h		Address Offset: 0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:0	VID	Vendor Identification: PCI standard identification for Intel.	8086h	RO

4.1.2 DID—Device Identification

This register combined with the Vendor Identification register uniquely identifies any PCI device.

B/D/F/Type: 0/1/0/CFG			Access: RO	
Size: 16	Default Value: 0C01h		Address Offset: 2h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:0	DID_MSB	Device Identification Number MSB: Identifier assigned to the Processor root port (virtual PCI-to-PCI bridge, PCI Express Graphics port).	0C01h	RO

4.1.3 PCICMD—PCI Command

B/D/F/Type: 0/1/0/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: 4h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:11	RSVD	Reserved.	00h	RO
10	INTAAD	INTA Assertion Disable: 0: This device is permitted to generate INTA interrupt messages. 1: This device is prevented from generating interrupt messages. Any INTA emulation interrupts already asserted must be de-asserted when this bit is set. Only affects interrupts generated by the device (PCI INTA from a PME or Hot Plug event) controlled by this command register. It does not affect upstream MSIs, upstream PCI INTA-INTD assert and deassert messages.	0h	RW
9	FB2B	Fast Back-to-Back Enable: Not Applicable or Implemented. Hardwired to 0.	0h	RO
8	SERRE	SERR# Message Enable: Controls the root port's SERR# messaging. The processor communicates the SERR# condition by sending an SERR message to the PCH. This bit, when set, enables reporting of non-fatal and fatal errors detected by the device to the Root Complex. Note that errors are reported if enabled either through this bit or through the PCI-Express specific bits in the Device Control Register. In addtion, for Type 1 configuration space header devices, this bit, when set, enables transmission by the primary interface of ERR_NONFATAL and ERR_FATAL error messages forwarded from the secondary interface. This bit does not affect the transmission of forwarded ERR_COR messages. 0: The SERR message is generated by the root port only under conditions enabled individually through the Device Control Register. 1: The root port is enabled to generate SERR messages which will be sent to the PCH for specific root port error conditions generated/detected or received on the secondary side of the virtual PCI to PCI bridge. The status of SERRs generated is reported in the PCISTS register.	0h	RW
7	RSVD	Reserved.	0h	RO
6	PERRE	Parity Error Response Enable: Controls whether or not the Master Data Parity Error bit in the PCI Status register can bet set. 0: Master Data Parity Error bit in PCI Status register	0h	RW

continued...

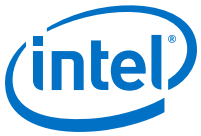

B/D/F/Type: 0/1/0/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: 4h	
Bit Range	Acronym	Description	Default	Access
		can NOT be set. 1: Master Data Parity Error bit in PCI Status register CAN be set.		
5	VGAPS	VGA Palette Snoop: Not Applicable or Implemented. Hardwired to 0.	0h	RO
4	MWIE	Memory Write and Invalidate Enable: Not Applicable or Implemented. Hardwired to 0.	0h	RO
3	SCE	Special Cycle Enable: Not Applicable or Implemented. Hardwired to 0.	0h	RO
2	BME	Bus Master Enable: Bus Master Enable (BME): Controls the ability of the PEG port to forward Memory Read/Write Requests in the upstream direction. 0: This device is prevented from making memory requests to its primary bus. Note that according to PCI Specification, as MSI interrupt messages are in-band memory writes, disabling the bus master enable bit prevents this device from generating MSI interrupt messages or passing them from its secondary bus to its primary bus. Upstream memory writes/reads, peer writes/reads, and MSIs will all be treated as illegal cycles. Writes are aborted. Reads are aborted and will return Unsupported Request status (or Master abort) in its completion packet. 1: This device is allowed to issue requests to its primary bus. Completions for previously issued memory read requests on the primary bus will be issued when the data is available. This bit does not affect forwarding of Completions from the primary interface to the secondary interface.	0h	RW
1	MAE	Memory Access Enable: 0: All of device's memory space is disabled. 1: Enable the Memory and Pre-fetchable memory address ranges defined in the MBASE, MLIMIT, PMBASE, and PMLIMIT registers.	0h	RW
0	IOAE	IO Access Enable: 0: All of devic's I/O space is disabled. 1: Enable the I/O address range defined in the IOBASE, and IOLIMIT registers.	0h	RW

4.1.4 PCISTS—PCI Status

This register reports the occurrence of error conditions associated with primary side of the "virtual" Host-PCI Express bridge embedded within the Root port.

B/D/F/Type: 0/1/0/CFG			Access: RO_V; RO; RW1C	
Size: 16	Default Value: 0010h		Address Offset: 6h	
Bit Range	Acronym	Description	Default	Access
15	DPE	Detected Parity Error: This bit is Set by a Function whenever it receives a Poisoned TLP, regardless of the state the Parity Error Response bit in the Command register. On a Function with a Type 1	0h	RW1C

continued...

B/D/F/Type: 0/1/0/CFG			Access: RO_V; RO; RW1C	
Size: 16	Default Value: 0010h		Address Offset: 6h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		Configuration header, the bit is Set when the Poisoned TLP is received by its Primary Side. Default value of this bit is 0b. This bit will be set only for completions of requests enountering ECC error in DRAM. Poisoned Peer 2 peer posted forwarded will not set this bit. They are reported at the receiveing port.		
14	SSE	Signaled System Error: This bit is set when this Device sends an SERR due to detecting an ERR_FATAL or ERR_NONFATAL condition and the SERR Enable bit in the Command register is '1'. Both received (if enabled by BCTRL1[1]) and internally detected error messages do not affect this field.	0h	RW1C
13	RMAS	Received Master Abort Status: This bit is Set when a Requester receives a Completion with Unsupported Request Completion Status. On a Function with a Type 1 Configuration header, the bit is Set when the Unsupported Request is received by its Primary Side. Not applicable. We do not have UR on primary interface	0h	RO
12	RTAS	Received Target Abort Status: This bit is Set when a Requester receives a Completion with Completer Abort Completion Status. On a Function with a Type 1 Configuration header, the bit is Set when the Completer Abort is received by its Primary Side. Default value of this bit is 0b. Not Applicable or Implemented. Hardwired to 0. The concept of a Completer abort does not exist on primary side of this device.	0h	RO
11	STAS	Signaled Target Abort Status: This bit is Set when a Function completes a Posted or Non- Posted Request as a Completer Abort error. This applies to a Function with a Type 1 Configuration header when the Completer Abort was generated by its Primary Side. Default value of this bit is 0b. Not Applicable or Implemented. Hardwired to 0. The concept of a target abort does not exist on primary side of this device.	0h	RO
10:9	DEVT	DEVSELB Timing: This device is not the subtractively decoded device on bus 0. This bit field is therefore hardwired to 00 to indicate that the device uses the fastest possible decode. Does not apply to PCI Express and must be hardwired to 00b.	0h	RO
8	PMDPE	Master Data Parity Error: This bit is Set by a Requester (Primary Side for Type 1 Configuration Space header Function) if the Parity Error Response bit in the Command register is 1b and either of the following two conditions occurs:	0h	RW1C

continued...

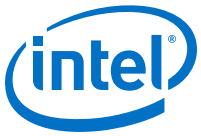

B/D/F/Type: 0/1/0/CFG			Access: RO_V; RO; RW1C	
Size: 16	Default Value: 0010h		Address Offset: 6h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		Requester receives a Completion marked poisoned Requester poisons a write Request If the Parity Error Response bit is 0b, this bit is never Set. Default value of this bit is 0b. This bit will be set only for completions of requests enountering ECC error in DRAM. Poisoned Peer 2 peer posted forwarded will not set this bit. They are reported at the receiveing port.		
7	FB2B	Fast Back-to-Back: Not Applicable or Implemented. Hardwired to 0.	0h	RO
6	RSVD	Reserved.	0h	RO
5	CAP66	66/60MHz capability: Not Applicable or Implemented. Hardwired to 0.	0h	RO
4	CAPL	Capabilities List: Indicates that a capabilities list is present. Hardwired to 1.	1h	RO
3	INTAS	INTx Status: Indicates that an interrupt message is pending internally to the device. Only PME and Hot Plug sources feed into this status bit (not PCI INTA-INTD assert and deassert messages). The INTA Assertion Disable bit, PCICMD1[10], has no effect on this bit. Note that INTA emulation interrupts received across the link are not reflected in this bit.	0h	RO_V
2:0	RSVD	Reserved.	0h	RO

4.1.5 RID—Revision Identification

This register contains the revision number of Device #1.
These bits are read only and writes to this register have no effect.

B/D/F/Type: 0/1/0/CFG			Access: RO	
Size: 8	Default Value: 00h		Address Offset: 8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:4	RID_MSB	Revision Identification Number MSB: This is an 8-bit value that indicates the revision identification number for the root port.	0h	RO
3:0	RID	Revision Identification Number: This is an 8-bit value that indicates the revision identification number for the root port.	0h	RO

4.1.6 CC—Class Code

This register identifies the basic function of the device, a more specific sub-class, and a register- specific programming interface.

B/D/F/Type: 0/1/0/CFG			Access: RO	
Size: 24	Default Value: 060400h		Address Offset: 9h	
Bit Range	Acronym	Description	Default	Access
23:16	BCC	Base Class Code: Indicates the base class code for this device. This code has the value 06h, indicating a Bridge device.	06h	RO
15:8	SUBCC	Sub-Class Code: Indicates the sub-class code for this device. The code is 04h indicating a PCI to PCI Bridge.	04h	RO
7:0	PI	Programming Interface: Indicates the programming interface of this device. This value does not specify a particular register set layout and provides no practical use for this device.	00h	RO

4.1.7 CL—Cache Line Size

B/D/F/Type: 0/1/0/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: Ch	
Bit Range	Acronym	Description	Default	Access
7:0	CLS	Cache Line Size: Implemented by PCI Express devices as a read-write field for legacy compatibility purposes but has no impact on any PCI Express device functionality.	00h	RW

4.1.8 HDR—Header Type

This register identifies the header layout of the configuration space. No physical register exists at this location.

B/D/F/Type: 0/1/0/CFG			Access: RO	
Size: 8	Default Value: 81h		Address Offset: Eh	
Bit Range	Acronym	Description	Default	Access
7:0	HDR	Header Type Register: Device #1 returns 81 to indicate that this is a multi function device with bridge header layout. Device #6 returns 01 to indicate that this is a single function device with bridge header layout.	81h	RO

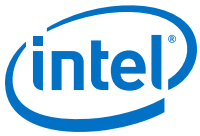

4.1.9 PBUSN—Primary Bus Number

This register identifies that this "virtual" Host-PCI Express bridge is connected to PCI bus #0.

B/D/F/Type: 0/1/0/CFG			Access: RO	
Size: 8	Default Value: 00h		Address Offset: 18h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:0	BUSN	Primary Bus Number: Configuration software typically programs this field with the number of the bus on the primary side of the bridge. Since the processor root port is an internal device and its primary bus is always 0, these bits are read only and are hardwired to 0.	00h	RO

4.1.10 SBUSN—Secondary Bus Number

This register identifies the bus number assigned to the second bus side of the "virtual" bridge i.e. to PCI Express. This number is programmed by the PCI configuration software to allow mapping of configuration cycles to PCI Express.

B/D/F/Type: 0/1/0/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: 19h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:0	BUSN	Secondary Bus Number: This field is programmed by configuration software with the bus number assigned to PCI Express.	00h	RW

4.1.11 SUBUSN—Subordinate Bus Number

This register identifies the subordinate bus (if any) that resides at the level below PCI Express. This number is programmed by the PCI configuration software to allow mapping of configuration cycles to PCI Express.

B/D/F/Type: 0/1/0/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: 1Ah	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:0	BUSN	Subordinate Bus Number: This register is programmed by configuration software with the number of the highest subordinate bus that lies behind the Processor root port bridge. When only a single PCI device resides on the PCI Express segment, this register will contain the same value as the SBUSN1 register.	00h	RW

4.1.12 IOBASE—I/O Base Address

This register controls the processor to PCI Express-G I/O access routing based on the following formula:

IO_BASE=< address =<IO_LIMIT

Only upper 4 bits are programmable. For the purpose of address decode address bits A[11:0] are treated as 0. Thus the bottom of the defined I/O address range will be aligned to a 4KB boundary.

B/D/F/Type: 0/1/0/CFG			Access: RW	
Size: 8	Default Value: F0h		Address Offset: 1Ch	
Bit Range	Acronym	Description	Default	Access
7:4	IOBASE	I/O Address Base: Corresponds to A[15:12] of the I/O addresses passed by the root port to PCI Express-G.	Fh	RW
3:0	RSVD	Reserved.	0h	RO

4.1.13 IOLIMIT—I/O Limit Address

This register controls the processor to PCI Express-G I/O access routing based on the following formula:

IO_BASE=< address =<IO_LIMIT

Only upper 4 bits are programmable. For the purpose of address decode address bits A[11:0] are assumed to be FFFh. Thus, the top of the defined I/O address range will be at the top of a 4KB aligned address block.

B/D/F/Type: 0/1/0/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: 1Dh	
Bit Range	Acronym	Description	Default	Access
7:4	IOLIMIT	I/O Address Limit: Corresponds to A[15:12] of the I/O address limit of the root port. Devices between this upper limit and IOBASE1 will be passed to the PCI Express hierarchy associated with this device.	0h	RW
3:0	RSVD	Reserved.	0h	RO

4.1.14 SSTS—Secondary Status

SSTS is a 16-bit status register that reports the occurrence of error conditions associated with secondary side (i.e. PCI Express-G side) of the "virtual" PCI-PCI bridge embedded within the Processor.

B/D/F/Type: 0/1/0/CFG			Access: RO; RW1C	
Size: 16	Default Value: 0000h		Address Offset: 1Eh	
Bit Range	Acronym	Description	Default	Access
15	DPE	Detected Parity Error: This bit is set by the Secondary Side for a Type 1 Configuration Space header device whenever it receives a Poisoned TLP, regardless of the state of the Parity Error Response Enable bit in the Bridge Control Register.	0h	RW1C
14	RSE	Received System Error: This bit is set when the Secondary Side for a Type 1 configuration space header device receives an ERR_FATAL or ERR_NONFATAL.	0h	RW1C
13	RMA	Received Master Abort: This bit is set when the Secondary Side for Type 1 Configuration Space Header Device (for requests initiated by the Type 1 Header Device itself) receives a Completion with Unsupported Request Completion Status.	0h	RW1C

continued...

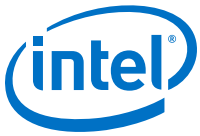

B/D/F/Type: 0/1/0/CFG			Access: RO; RW1C	
Size: 16	**Default Value: 0000h**		**Address Offset: 1Eh**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
12	RTA	Received Target Abort: This bit is set when the Secondary Side for Type 1 Configuration Space Header Device (for requests initiated by the Type 1 Header Device itself) receives a Completion with Completer Abort Completion Status.	0h	RW1C
11	STA	Signaled Target Abort: Not Applicable or Implemented. Hardwired to 0. The processor does not generate Target Aborts (The root port will never complete a request using the Completer Abort Completion status). UR detected inside the processor (such as in MC) will be reported in primary side status.	0h	RO
10:9	DEVT	DEVSELB Timing: Not Applicable or Implemented. Hardwired to 0.	0h	RO
8	SMDPE	Master Data Parity Error: When set indicates that the processor received across the link (upstream) a Read Data Completion Poisoned TLP (EP=1). This bit can only be set when the Parity Error Enable bit in the Bridge Control register is set.	0h	RW1C
7	FB2B	Fast Back-to-Back: Not Applicable or Implemented. Hardwired to 0.	0h	RO
6	RSVD	Reserved.	0h	RO
5	CAP66	66/60 MHz capability: Not Applicable or Implemented. Hardwired to 0.	0h	RO
4:0	RSVD	Reserved.	00h	RO

4.1.15 MBASE—Memory Base Address

This register controls the processor to PCI Express non-prefetchable memory access routing based on the following formula:

MEMORY_BASE=< address =<MEMORY_LIMIT

The upper 12 bits of the register are read/write and correspond to the upper 12 address bits A[31:20] of the 32 bit address. The bottom 4 bits of this register are read-only and return zeroes when read. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be 0. Thus, the bottom of the defined memory address range will be aligned to a 1MB boundary.

B/D/F/Type: 0/1/0/CFG			Access: RW	
Size: 16	**Default Value: FFF0h**		**Address Offset: 20h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:4	MBASE	Memory Address Base: Corresponds to A[31:20] of the lower limit of the memory range that will be passed to PCI Express.	FFFh	RW
3:0	RSVD	Reserved.	0h	RO

4.1.16 MLIMIT—Memory Limit Address

This register controls the processor to PCI Express non-prefetchable memory access routing based on the following formula:
MEMORY_BASE=< address =<MEMORY_LIMIT
The upper 12 bits of the register are read/write and correspond to the upper 12 address bits A[31:20] of the 32 bit address. The bottom 4 bits of this register are read-only and return zeroes when read. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be FFFFFh. Thus, the top of the defined memory address range will be at the top of a 1MB aligned memory block. NOTE: Memory range covered by MBASE and MLIMIT registers are used to map non-prefetchable PCI Express address ranges (typically where control/status memory-mapped I/O data structures of the graphics controller will reside) and PMBASE and PMLIMIT are used to map prefetchable address ranges (typically graphics local memory). This segregation allows application of USWC space attribute to be performed in a true plug-and-play manner to the prefetchable address range for improved Processor-PCI Express memory access performance.
Note also that configuration software is responsible for programming all address range registers (prefetchable, non-prefetchable) with the values that provide exclusive address ranges i.e. prevent overlap with each other and/or with the ranges covered with the main memory. There is no provision in the processor hardware to enforce prevention of overlap and operations of the system in the case of overlap are not guaranteed.

B/D/F/Type: 0/1/0/CFG			Access: RW	
Size: 16	Default Value: 0000h		Address Offset: 22h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:4	MLIMIT	Memory Address Limit: Corresponds to A[31:20] of the upper limit of the address range passed to PCI Express.	000h	RW
3:0	RSVD	Reserved.	0h	RO

4.1.17 PMBASE—Prefetchable Memory Base Address

This register in conjunction with the corresponding Upper Base Address register controls the processor to PCI Express prefetchable memory access routing based on the following formula:
PREFETCHABLE_MEMORY_BASE =< address =< PREFETCHABLE_MEMORY_LIMIT
The upper 12 bits of this register are read/write and correspond to address bits A[31:20] of the 40-bit address. The lower 8 bits of the Upper Base Address register are read/write and correspond to address bits A[39:32] of the 40-bit address. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be 0. Thus, the bottom of the defined memory address range will be aligned to a 1MB boundary.

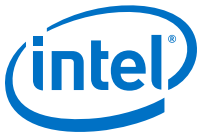

B/D/F/Type: 0/1/0/CFG			Access: RO; RW	
Size: 16	**Default Value: FFF1h**		**Address Offset: 24h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:4	PMBASE	Prefetchable Memory Base Address: Corresponds to A[31:20] of the lower limit of the memory range that will be passed to PCI Express.	FFFh	RW
3:0	AS64	64-bit Address Support: Indicates that the upper 32 bits of the prefetchable memory region base address are contained in the Prefetchable Memory base Upper Address register at 28h.	1h	RO

4.1.18 PMLIMIT—Prefetchable Memory Limit Address

This register in conjunction with the corresponding Upper Limit Address register controls the processor to PCI Express prefetchable memory access routing based on the following formula:

PREFETCHABLE_MEMORY_BASE =< address =< PREFETCHABLE_MEMORY_LIMIT

The upper 12 bits of this register are read/write and correspond to address bits A[31:20] of the 40-bit address. The lower 8 bits of the Upper Limit Address register are read/write and correspond to address bits A[39:32] of the 40-bit address. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be FFFFFh. Thus, the top of the defined memory address range will be at the top of a 1MB aligned memory block. Note that prefetchable memory range is supported to allow segregation by the configuration software between the memory ranges that must be defined as UC and the ones that can be designated as a USWC (i.e. prefetchable) from the processor perspective.

B/D/F/Type: 0/1/0/CFG			Access: RO; RW	
Size: 16	**Default Value: 0001h**		**Address Offset: 26h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:4	PMLIMIT	Prefetchable Memory Address Limit: Corresponds to A[31:20] of the upper limit of the address range passed to PCI Express.	000h	RW
3:0	AS64B	64-bit Address Support: Indicates that the upper 32 bits of the prefetchable memory region limit address are contained in the Prefetchable Memory Base Limit Address register at 2Ch	1h	RO

4.1.19 PMBASEU—Prefetchable Memory Base Address Upper

The functionality associated with this register is present in the PEG design implementation.

This register in conjunction with the corresponding Upper Base Address register controls the processor to PCI Express prefetchable memory access routing based on the following formula:

PREFETCHABLE_MEMORY_BASE =< address =< PREFETCHABLE_MEMORY_LIMIT

The upper 12 bits of this register are read/write and correspond to address bits A[31:20] of the 39-bit address. The lower 7 bits of the Upper Base Address register are read/write and correspond to address bits A[38:32] of the 39-bit address. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be 0. Thus, the bottom of the defined memory address range will be aligned to a 1MB boundary.

B/D/F/Type: 0/1/0/CFG			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 28h	
Bit Range	Acronym	Description	Default	Access
31:0	PMBASEU	Prefetchable Memory Base Address: Corresponds to A[63:32] of the lower limit of the prefetchable memory range that will be passed to PCI Express.	00000000h	RW

4.1.20 PMLIMITU—Prefetchable Memory Limit Address Upper

The functionality associated with this register is present in the PEG design implementation.

This register in conjunction with the corresponding Upper Limit Address register controls the processor to PCI Express prefetchable memory access routing based on the following formula:

PREFETCHABLE_MEMORY_BASE =< address =< PREFETCHABLE_MEMORY_LIMIT

The upper 12 bits of this register are read/write and correspond to address bits A[31:20] of the 39-bit address. The lower 7 bits of the Upper Limit Address register are read/write and correspond to address bits A[39:32] of the 39-bit address. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be FFFFFh. Thus, the top of the defined memory address range will be at the top of a 1MB aligned memory block.

Note that prefetchable memory range is supported to allow segregation by the configuration software between the memory ranges that must be defined as UC and the ones that can be designated as a USWC (i.e. prefetchable) from the processor perspective.

B/D/F/Type: 0/1/0/CFG			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 2Ch	
Bit Range	Acronym	Description	Default	Access
31:0	PMLIMITU	Prefetchable Memory Address Limit: Corresponds to A[63:32] of the upper limit of the prefetchable Memory range that will be passed to PCI Express.	00000000h	RW

4.1.21 CAPPTR—Capabilities Pointer

The capabilities pointer provides the address offset to the location of the first entry in this device's linked list of capabilities.

B/D/F/Type: 0/1/0/CFG			Access: RO	
Size: 8	Default Value: 88h		Address Offset: 34h	
Bit Range	Acronym	Description	Default	Access
7:0	CAPPTR1	First Capability: The first capability in the list is the Subsystem ID and Subsystem Vendor ID Capability.	88h	RO

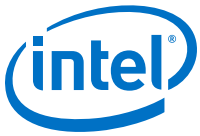

4.1.22 INTRLINE—Interrupt Line

This register contains interrupt line routing information. The device itself does not use this value, rather it is used by device drivers and operating systems to determine priority and vector information.

B/D/F/Type: 0/1/0/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: 3Ch	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:0	INTCON	Interrupt Connection: Used to communicate interrupt line routing information. BIOS Requirement: POST software writes the routing information into this register as it initializes and configures the system. The value indicates to which input of the system interrupt controller this device's interrupt pin is connected.	00h	RW

4.1.23 INTRPIN—Interrupt Pin

This register specifies which interrupt pin this device uses.

B/D/F/Type: 0/1/0/CFG			Access: RW_O; RO	
Size: 8	Default Value: 01h		Address Offset: 3Dh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:3	INTPINH	Interrupt Pin High:	00h	RO
2:0	INTPIN	Interrupt Pin: As a multifunction device, the PCI Express device may specify any INTx (x=A,B,C,D) as its interrupt pin. The Interrupt Pin register tells which interrupt pin the device (or device function) uses. A value of 1 corresponds to INTA# (Default) A value of 2 corresponds to INTB# A value of 3 corresponds to INTC# A value of 4 corresponds to INTD# Devices (or device functions) that do not use an interrupt pin must put a 0 in this register. The values 05h through FFh are reserved. This register is write once. BIOS must set this register to select the INTx to be used by this root port.	1h	RW_O

4.1.24 BCTRL—Bridge Control

This register provides extensions to the PCICMD register that are specific to PCI-PCI bridges. The BCTRL provides additional control for the secondary interface (i.e. PCI Express) as well as some bits that affect the overall behavior of the "virtual" Host-PCI Express bridge embedded within the processor, e.g. VGA compatible address ranges mapping.

B/D/F/Type: 0/1/0/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: 3Eh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:12	RSVD	Reserved.	0h	RO
11	DTSERRE	Discard Timer SERR# Enable: Not Applicable or Implemented. Hardwired to 0.	0h	RO
10	DTSTS	Discard Timer Status: Not Applicable or Implemented. Hardwired to 0.	0h	RO
9	SDT	Secondary Discard Timer: Not Applicable or Implemented. Hardwired to 0.	0h	RO
8	PDT	Primary Discard Timer: Not Applicable or Implemented. Hardwired to 0.	0h	RO
7	FB2BEN	Fast Back-to-Back Enable: Not Applicable or Implemented. Hardwired to 0.	0h	RO
6	SRESET	Secondary Bus Reset: Setting this bit triggers a hot reset on the corresponding PCI Express Port. This will force the LTSSM to transition to the Hot Reset state (via Recovery) from L0, L0s, or L1 states.	0h	RW
5	MAMODE	Master Abort Mode: Does not apply to PCI Express. Hardwired to 0.	0h	RO
4	VGA16D	VGA 16-bit Decode: Enables the PCI-to-PCI bridge to provide 16-bit decoding of VGA I/O address precluding the decoding of alias addresses every 1 KB. This bit only has meaning if bit 3 (VGA Enable) of this register is also set to 1, enabling VGA I/O decoding and forwarding by the bridge. 0: Execute 10-bit address decodes on VGA I/O accesses. 1: Execute 16-bit address decodes on VGA I/O accesses.	0h	RW
3	VGAEN	VGA Enable: Controls the routing of processor initiated transactions targeting VGA compatible I/O and memory address ranges. See the VGAEN/MDAP table in device 0, offset 97h[0].	0h	RW
2	ISAEN	ISA Enable: Needed to exclude legacy resource decode to route ISA resources to legacy decode path. Modifies the response by the root port to an I/O access issued by the processor that target ISA I/O addresses. This applies only to I/O addresses that are enabled by the IOBASE and IOLIMIT registers. 0: All addresses defined by the IOBASE and IOLIMIT for processor I/O transactions will be mapped to PCI Express. 1: The root port will not forward to PCI Express any	0h	RW

continued...

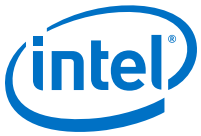

B/D/F/Type: 0/1/0/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: 3Eh	
Bit Range	Acronym	Description	Default	Access
		I/O transactions addressing the last 768 bytes in each 1KB block even if the addresses are within the range defined by the IOBASE and IOLIMIT registers.		
1	SERREN	SERR Enable: 0: No forwarding of error messages from secondary side to primary side that could result in an SERR. 1: ERR_COR, ERR_NONFATAL, and ERR_FATAL messages result in SERR message when individually enabled by the Root Control register.	0h	RW
0	PEREN	Parity Error Response Enable: Controls whether or not the Master Data Parity Error bit in the Secondary Status register is set when the root port receives across the link (upstream) a Read Data Completion Poisoned TLP 0: Master Data Parity Error bit in Secondary Status register can NOT be set. 1: Master Data Parity Error bit in Secondary Status register CAN be set.	0h	RW

4.1.25 PM—Power Management Capabilities

B/D/F/Type: 0/1/0/CFG			Access: RO; RO_V	
Size: 32	Default Value: C8039001h		Address Offset: 80h	
Bit Range	Acronym	Description	Default	Access
31:27	PMES	PME Support: This field indicates the power states in which this device may indicate PME wake via PCI Express messaging. D0, D3hot & D3cold. This device is not required to do anything to support D3hot & D3cold, it simply must report that those states are supported. Refer to the PCI Power Management 1.1 specification for encoding explanation and other power management details.	19h	RO
26	D2PSS	D2 Power State Support: Hardwired to 0 to indicate that the D2 power management state is NOT supported.	0h	RO
25	D1PSS	D1 Power State Support: Hardwired to 0 to indicate that the D1 power management state is NOT supported.	0h	RO
24:22	AUXC	Auxiliary Current: Hardwired to 0 to indicate that there are no 3.3Vaux auxiliary current requirements.	0h	RO
21	DSI	Device Specific Initialization: Hardwired to 0 to indicate that special initialization of this device is NOT required before generic class device driver is to use it.	0h	RO
20	APS	Auxiliary Power Source: Hardwired to 0.	0h	RO
19	PMECLK	PME Clock: Hardwired to 0 to indicate this device does NOT support PME# generation.	0h	RO
18:16	PCIPMCV	PCI PM CAP Version: Version - A value of 011b indicates that this function complies with revision 1.2 of the PCI Power Management Interface Specification. --Was Previously Hardwired to 02h to indicate there are 4 bytes of power	3h	RO

continued...

B/D/F/Type: 0/1/0/CFG			Access: RO; RO_V	
Size: 32	Default Value: C8039001h		Address Offset: 80h	
Bit Range	Acronym	Description	Default	Access
		management registers implemented and that this device complies with revision 1.1 of the PCI Power Management Interface Specification.		
15:8	PNC	Pointer to Next Capability: This contains a pointer to the next item in the capabilities list. If MSICH (CAPL[0] @ 7Fh) is 0, then the next item in the capabilities list is the Message Signaled Interrupts (MSI) capability at 90h. If MSICH (CAPL[0] @ 7Fh) is 1, then the next item in the capabilities list is the PCI Express capability at A0h.	90h	RO_V
7:0	CID	Capability ID: Value of 01h identifies this linked list item (capability structure) as being for PCI Power Management registers.	01h	RO

4.1.26 PM—Power Management Control/Status

B/D/F/Type: 0/1/0/CFG			Access: RW; RO	
Size: 32	Default Value: 00000008h		Address Offset: 84h	
Bit Range	Acronym	Description	Default	Access
31:16	RSVD	Reserved.	0000h	RO
15	PMESTS	PME Status: Indicates that this device does not support PMEB generation from D3cold.	0h	RO
14:13	DSCALE	Data Scale: Indicates that this device does not support the power management data register.	0h	RO
12:9	DSEL	Data Select: Indicates that this device does not support the power management data register.	0h	RO
8	PMEE	PME Enable: Indicates that this device does not generate PMEB assertion from any D-state. 0: PMEB generation not possible from any D State 1: PMEB generation enabled from any D State The setting of this bit has no effect on hardware. See PM_CAP[15:11]	0h	RW
7:4	RSVD	Reserved.	0h	RO
3	NSR	No Soft Reset: No Soft Reset. When set to 1 this bit indicates that the device is transitioning from D3hot to D0 because the power state commands do not perform a internal reset. Config context is preserved. Upon transition no additonal operating sys intervention is required to preserve configuration context beyond writing the power state bits. When clear the devices do not perform an internal reset upon transitioning from D3hot to D0 via software control of the power state bits. Regardless of this bit the devices that transition from a	1h	RO

continued...

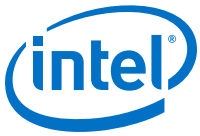

B/D/F/Type: 0/1/0/CFG			Access: RW; RO	
Size: 32	**Default Value: 00000008h**		**Address Offset: 84h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		D3hot to D0 by a system or bus segment reset will return to the device state D0 uniintialized with only PME context preserved if PME is supported and enabled.		
2	RSVD	Reserved.	0h	RO
1:0	PS	Power State: Indicates the current power state of this device and can be used to set the device into a new power state. If software attempts to write an unsupported state to this field, write operation must complete normally on the bus, but the data is discarded and no state change occurs. 00: D0 01: D1 (Not supported in this device.) 10: D2 (Not supported in this device.) 11: D3 Support of D3cold does not require any special action. While in the D3hot state, this device can only act as the target of PCI configuration transactions (for power management control). This device also cannot generate interrupts or respond to MMR cycles in the D3 state. The device must return to the D0 state in order to be fully-functional. When the Power State is other than D0, the bridge will Master Abort (i.e. not claim) any downstream cycles (with exception of type 0 config cycles). Consequently, these unclaimed cycles will go down DMI and come back up as Unsupported Requests, which the processor logs as Master Aborts in Device 0 PCISTS[13] There is no additional hardware functionality required to support these Power States.	0h	RW

4.1.27 SS—Subsystem ID and Vendor ID Capabilities

This capability is used to uniquely identify the subsystem where the PCI device resides. Because this device is an integrated part of the system and not an add-in device, it is anticipated that this capability will never be used. However, it is necessary because Microsoft will test for its presence.

B/D/F/Type: 0/1/0/CFG			Access: RO	
Size: 32	**Default Value: 0000800Dh**		**Address Offset: 88h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:8	PNC	Pointer to Next Capability: This contains a pointer to the next item in the capabilities list which is the PCI Power Management capability.	80h	RO
7:0	CID	Capability ID: Value of 0Dh identifies this linked list item (capability structure) as being for SSID/SSVID registers in a PCI-to-PCI Bridge.	0Dh	RO

4.1.28 SS—Subsystem ID and Subsystem Vendor ID

System BIOS can be used as the mechanism for loading the SSID/SVID values. These values must be preserved through power management transitions and a hardware reset.

B/D/F/Type: 0/1/0/CFG			Access: RW_O	
Size: 32	Default Value: 00008086h		Address Offset: 8Ch	
Bit Range	Acronym	Description	Default	Access
31:16	SSID	Subsystem ID: Identifies the particular subsystem and is assigned by the vendor.	0000h	RW_O
15:0	SSVID	Subsystem Vendor ID: Identifies the manufacturer of the subsystem and is the same as the vendor ID which is assigned by the PCI Special Interest Group.	8086h	RW_O

4.1.29 MSI—Message Signaled Interrupts Capability ID

When a device supports MSI it can generate an interrupt request to the processor by writing a predefined data item (a message) to a predefined memory address.

The reporting of the existence of this capability can be disabled by setting MSICH (CAPL[0] @ 7Fh). In that case walking this linked list will skip this capability and instead go directly from the PCI PM capability to the PCI Express capability.

B/D/F/Type: 0/1/0/CFG			Access: RO	
Size: 16	Default Value: A005h		Address Offset: 90h	
Bit Range	Acronym	Description	Default	Access
15:8	PNC	Pointer to Next Capability: This contains a pointer to the next item in the capabilities list which is the PCI Express capability.	A0h	RO
7:0	CID	Capability ID: Value of 05h identifies this linked list item (capability structure) as being for MSI registers.	05h	RO

4.1.30 MC—Message Control

System software can modify bits in this register, but the device is prohibited from doing so.

If the device writes the same message multiple times, only one of those messages is guaranteed to be serviced. If all of them must be serviced, the device must not generate the same message again until the driver services the earlier one.

B/D/F/Type: 0/1/0/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: 92h	
Bit Range	Acronym	Description	Default	Access
15:8	RSVD	Reserved.	00h	RO
7	B64AC	64-bit Address Capable: Hardwired to 0 to indicate that the function does not implement the upper 32 bits of the Message Address register and is incapable of generating a 64-bit memory address. This may need to change in future implementations when addressable system memory exceeds the 32b/4GB limit.	0h	RO

continued...

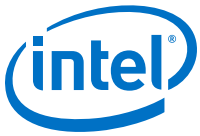

B/D/F/Type: 0/1/0/CFG			Access: RW; RO	
Size: 16	**Default Value: 0000h**		**Address Offset: 92h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
6:4	MME	Multiple Message Enable: System software programs this field to indicate the actual number of messages allocated to this device. This number will be equal to or less than the number actually requested. The encoding is the same as for the MMC field below.	0h	RW
3:1	MMC	Multiple Message Capable: System software reads this field to determine the number of messages being requested by this device. Value: Number of Messages Requested 000: 1 All of the following are reserved in this implementation: 001: 2 010: 4 011: 8 100: 16 101: 32 110: Reserved 111: Reserved	0h	RO
0	MSIEN	MSI Enable: Controls the ability of this device to generate MSIs. 0: MSI will not be generated. 1: MSI will be generated when we receive PME messages. INTA will not be generated and INTA Status (PCISTS1[3]) will not be set.	0h	RW

4.1.31 MA—Message Address

B/D/F/Type: 0/1/0/CFG			Access: RO; RW	
Size: 32	**Default Value: 00000000h**		**Address Offset: 94h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:2	MA	Message Address: Used by system software to assign an MSI address to the device. The device handles an MSI by writing the padded contents of the MD register to this address.	00000000h	RW
1:0	FDWA	Force DWord Align: Hardwired to 0 so that addresses assigned by system software are always aligned on a dword address boundary.	0h	RO

4.1.32 MD—Message Data

B/D/F/Type: 0/1/0/CFG			Access: RW	
Size: 16	**Default Value: 0000h**		**Address Offset: 98h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:0	MD	Message Data: Base message data pattern assigned by system software and used to handle an MSI from the device. When the device must generate an interrupt request, it writes a 32-bit value to the memory address specified in the MA register. The upper 16 bits are always set to 0. The lower 16 bits are supplied by this register.	0000h	RW

4.1.33 PEG—PCI Express Capability List

Enumerates the PCI Express capability structure.

B/D/F/Type: 0/1/0/CFG			Access: RO	
Size: 16	Default Value: 0010h		Address Offset: A0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:8	PNC	Pointer to Next Capability: This value terminates the capabilities list. The Virtual Channel capability and any other PCI Express specific capabilities that are reported via this mechanism are in a separate capabilities list located entirely within PCI Express Extended Configuration Space.	00h	RO
7:0	CID	Capability ID: Identifies this linked list item (capability structure) as being for PCI Express registers.	10h	RO

4.1.34 PEG—PCI Express Capabilities

Indicates PCI Express device capabilities.

B/D/F/Type: 0/1/0/CFG			Access: RO; RW_O	
Size: 16	Default Value: 0142h		Address Offset: A2h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:14	RSVD	Reserved.	0h	RO
13:9	IMN	Interrupt Message Number: Not Applicable or Implemented. Hardwired to 0.	00h	RO
8	SI	Slot Implemented: 0: The PCI Express Link associated with this port is connected to an integrated component or is disabled. 1: The PCI Express Link associated with this port is connected to a slot. BIOS Requirement: This field must be initialized appropriately if a slot connection is not implemented.	1h	RW_O
7:4	DPT	Device/Port Type: Hardwired to 4h to indicate root port of PCI Express Root Complex.	4h	RO
3:0	PCIECV	PCI Express Capability Version: PCI Express Capability Version (PCIECV): Hardwired to 2h to indicate compliance to the PCI Express Capabilities Register Expansion ECN.	2h	RO

4.1.35 DCAP—Device Capabilities

Indicates PCI Express device capabilities.

B/D/F/Type: 0/1/0/CFG			Access: RW_O; RO	
Size: 32	Default Value: 00008001h		Address Offset: A4h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15	RBER	Role Based Error Reporting (RBER): Indicates that this device implements the functionality defined in the Error Reporting ECN as required by the PCI Express 1.1 spec.	1h	RO
14:6	RSVD	Reserved.	000h	RO

continued...

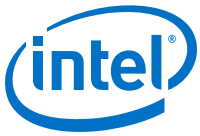

B/D/F/Type: 0/1/0/CFG			Access: RW_O; RO	
Size: 32	**Default Value: 00008001h**		**Address Offset: A4h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
5	ETFS	Extended Tag Field Supported: Hardwired to indicate support for 5-bit Tags as a Requestor.	0h	RO
4:3	PFS	Phantom Functions Supported: Not Applicable or Implemented. Hardwired to 0.	0h	RO
2:0	MPS	Max Payload Size: Default indicates 256B max supported payload for Transaction Layer Packets (TLP) for x16 PEG only. x8 and x4 PEG are limited to 128B support.	1h	RW_O

4.1.36 DCTL—Device Control

Provides control for PCI Express device specific capabilities.
The error reporting enable bits are in reference to errors detected by this device, not error messages received across the link. The reporting of error messages (ERR_CORR, ERR_NONFATAL, ERR_FATAL) received by Root Port is controlled exclusively by Root Port Command Register.

B/D/F/Type: 0/1/0/CFG			Access: RW; RO	
Size: 16	**Default Value: 0020h**		**Address Offset: A8h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15	RSVD	Reserved.	0h	RO
14:12	MRRS	Reserved for Max Read Request Size:	0h	RO
11	NSE	Reserved for Enable No Snoop:	0h	RO
10:8	RSVD	Reserved.	0h	RO
7:5	MPS	000b 128B max supported payload for Transaction Layer Packets (TLP). 001b 256B max supported payload for Transaction Layer Packets (TLP). All other encodings are reserved. As a receiver, the Device must handle TLPs as large as the set value. As transmitter, the Device must not generate TLPs exceeding the set value. NOTE: A payload size of 256B is only supported for Device 1 Function 0.	1h	RW
4	ROE	Reserved for Enable Relaxed Ordering:	0h	RO
3	URRE	Unsupported Request Reporting Enable: Unsupported Request Reporting Enable (URRE): When set, allows signaling ERR_NONFATAL, ERR_FATAL, or ERR_CORR to the Root Control register when detecting an unmasked Unsupported Request (UR). An ERR_CORR is signaled when an unmasked Advisory Non-Fatal UR is received. An ERR_FATAL or ERR_NONFATAL is sent to the Root Control register when an uncorrectable non-Advisory UR is received with the severity bit set in the Uncorrectable Error Severity register.	0h	RW

continued...

B/D/F/Type: 0/1/0/CFG			Access: RW; RO	
Size: 16	Default Value: 0020h		Address Offset: A8h	
Bit Range	Acronym	Description	Default	Access
2	FERE	Fatal Error Reporting Enable: Fatal Error Reporting Enable (FERE): When set, enables signaling of ERR_FATAL to the Root Control register due to internally detected errors or error messages received across the link. Other bits also control the full scope of related error reporting.	0h	RW
1	NERE	Non-Fatal Error Reporting Enable: Non-Fatal Error Reporting Enable (NERE): When set, enables signaling of ERR_NONFATAL to the Rool Control register due to internally detected errors or error messages received across the link. Other bits also control the full scope of related error reporting.	0h	RW
0	CERE	Correctable Error Reporting Enable: Correctable Error Reporting Enable (CERE): When set, enables signaling of ERR_CORR to the Root Control register due to internally detected errors or error messages received across the link. Other bits also control the full scope of related error reporting.	0h	RW

4.1.37 DSTS—Device Status

Reflects status corresponding to controls in the Device Control register. The error reporting bits are in reference to errors detected by this device, not errors messages received across the link.

B/D/F/Type: 0/1/0/CFG			Access: RW1C; RO	
Size: 16	Default Value: 0000h		Address Offset: AAh	
Bit Range	Acronym	Description	Default	Access
15:6	RSVD	Reserved.	000h	RO
5	TP	Transactions Pending: 0: All pending transactions (including completions for any outstanding non-posted requests on any used virtual channel) have been completed. 1: Indicates that the device has transaction(s) pending (including completions for any outstanding non-posted requests for all used Traffic Classes). Not Applicable or Implemented. Hardwired to 0.	0h	RO
4	RSVD	Reserved.	0h	RO
3	URD	Unsupported Request Detected: When set this bit indicates that the Device received an Unsupported Request. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control Register. Additionally, the Non-Fatal Error Detected bit or the Fatal Error Detected bit is set according to the setting of the Unsupported Request Error Severity bit. In production systems setting the Fatal Error Detected bit is not an option as support for AER will not be reported.	0h	RW1C
2	FED	Fatal Error Detected: When set this bit indicates that fatal error(s) were detected. Errors are logged in this register regardless of whether error reporting is enabled or not in	0h	RW1C

continued...

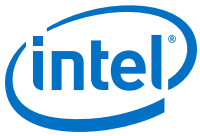

B/D/F/Type: 0/1/0/CFG			Access: RW1C; RO	
Size: 16	Default Value: 0000h		Address Offset: AAh	
Bit Range	Acronym	Description	Default	Access
		the Device Control register. When Advanced Error Handling is enabled, errors are logged in this register regardless of the settings of the uncorrectable error mask register.		
1	NFED	Non-Fatal Error Detected: When set this bit indicates that non-fatal error(s) were detected. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. When Advanced Error Handling is enabled, errors are logged in this register regardless of the settings of the uncorrectable error mask register.	0h	RW1C
0	CED	Correctable Error Detected: When set this bit indicates that correctable error(s) were detected. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. When Advanced Error Handling is enabled, errors are logged in this register regardless of the settings of the correctable error mask register.	0h	RW1C

4.1.38 LCTL—Link Control

Allows control of PCI Express link.

B/D/F/Type: 0/1/0/CFG			Access: RW; RO; RW_V	
Size: 16	Default Value: 0000h		Address Offset: B0h	
Bit Range	Acronym	Description	Default	Access
15:12	RSVD	Reserved.	0h	RO
11	LABIE	Link Autonomous Bandwidth Interrupt Enable: Link Autonomous Bandwidth Interrupt Enable - When Set, this bit enables the generation of an interrupt to indicate that the Link Autonomous Bandwidth Status bit has been Set. This bit is not applicable and is reserved for Endpoint devices, PCI Express to PCI/PCI-X bridges, and Upstream Ports of Switches. Devices that do not implement the Link Bandwidth Notification capability must hardwire this bit to 0b.	0h	RW
10	LBMIE	Link Bandwidth Management Interrupt Enable: Link Bandwidth Management Interrupt Enable - When Set, this bit enables the generation of an interrupt to indicate that the Link Bandwidth Management Status bit has been Set. This bit is not applicable and is reserved for Endpoint devices, PCI Express to PCI/PCI-X bridges, and Upstream Ports of Switches.	0h	RW
9	HAWD	Hardware Autonomous Width Disable: Hardware Autonomous Width Disable - When Set, this bit disables hardware from changing the Link width for reasons other than attempting to correct unreliable Link operation by reducing Link width. Devices that do not implement the ability autonomously to change Link width are permitted to hardwire this bit to 0b.	0h	RW

continued...

B/D/F/Type: 0/1/0/CFG			Access: RW; RO; RW_V	
Size: 16	Default Value: 0000h		Address Offset: B0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
8	ECPM	Enable Clock Power Management: Applicable only for form factors that support a "Clock Request" (CLKREQ#) mechanism, this enable functions as follows 0b - Clock power management is disabled and device must hold CLKREQ# signal low 1b - When this bit is set to 1 the device is permitted to use CLKREQ# signal to power manage link clock according to protocol defined in appropriate form factor specification. Default value of this field is 0b. Components that do not support Clock Power Management (as indicated by a 0b value in the Clock Power Management bit of the Link Capabilities Register) must hardwire this bit to 0b.	0h	RO
7	ES	Extended Synch: Extended synch 0: Standard Fast Training Sequence (FTS). 1: Forces the transmission of additional ordered sets when exiting the L0s state and when in the Recovery state. This mode provides external devices (e.g., logic analyzers) monitoring the Link time to achieve bit and symbol lock before the link enters L0 and resumes communication. This is a test mode only and may cause other undesired side effects such as buffer overflows or underruns.	0h	RW
6	CCC	Common Clock Configuration: 0: Indicates that this component and the component at the opposite end of this Link are operating with asynchronous reference clock. 1: Indicates that this component and the component at the opposite end of this Link are operating with a distributed common reference clock. The state of this bit affects the L0s Exit Latency reported in LCAP[14:12] and the N_FTS value advertised during link training. See PEGL0SLAT at offset 22Ch.	0h	RW
5	RL	Retrain Link: 0b Normal operation. 1b Full Link retraining is initiated by directing the Physical Layer LTSSM from L0, L0s, or L1 states to the Recovery state. This bit always returns 0 when read. This bit is cleared automatically (no need to write a 0). Intel Reserved Text: PEG10: This bit is set implicitly if CAPID0_B[SPEGFX1] or CAPID0_B[DPEGFX1] are set to 1b, and in addition BCTRL[VGAEN] is set to 1b. PEG11: This bit is set implicitly if CAPID0_B[DPEGFX1] is set to 1b and in addition BCTRL[VGAEN] is set to 1b. PEG12: This bit is set implicitly if CAPID0_B[DPEGFX1] is set to 1b and in addition BCTRL[VGAEN] is set to 1b.	0h	RW_V
4	LD	Link Disable: 0: Normal operation 1: Link is disabled. Forces the LTSSM to transition to the Disabled state (via Recovery) from L0, L0s, or L1 states. Link retraining happens automatically on 0 to 1 transition, just like when coming out of reset. Writes to this bit are immediately reflected in the value read from the bit, regardless of actual Link state.	0h	RW

continued...

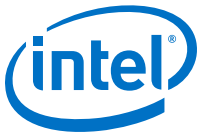

B/D/F/Type: 0/1/0/CFG			Access: RW; RO; RW_V	
Size: 16	**Default Value: 0000h**		**Address Offset: B0h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
3	RCB	Read Completion Boundary: Hardwired to 0 to indicate 64 byte.	0h	RO
2	RSVD	Reserved.	0h	RO
1:0	ASPM	Active State PM: Controls the level of active state power management supported on the given link. 00: Disabled 01: L0s Entry Supported 10: L1 Entry Supported 11: L0s and L1 Entry Supported	0h	RW

4.1.39 LSTS—Link Status

Indicates PCI Express link status.

B/D/F/Type: 0/1/0/CFG			Access: RO_V; RO; RW1C	
Size: 16	**Default Value: 1001h**		**Address Offset: B2h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15	LABWS	Link Autonomous Bandwidth Status: This bit is set to 1b by hardware to indicate that hardware has autonomously changed link speed or width, without the port transitioning through DL_Down status, for reasons other than to attempt to correct unreliable link operation. This bit must be set if the Physical Layer reports a speed or width change was initiated by the downstream component that was indicated as an autonomous change.	0h	RW1C
14	LBWMS	Link Bandwidth Management Status: This bit is set to 1b by hardware to indicate that either of the following has occurred without the port transitioning through DL_Down status: A link retraining initiated by a write of 1b to the Retrain Link bit has completed. Note: This bit is Set following any write of 1b to the Retrain Link bit, including when the Link is in the process of retraining for some other reason. Hardware has autonomously changed link speed or width to attempt to correct unreliable link operation, either through an LTSSM timeout or a higher level process This bit must be set if the Physical Layer reports a speed or width change was initiated by the downstream component that was not indicated as an autonomous change.	0h	RW1C
13	DLLLA	Data Link Layer Link Active (Optional): This bit indicates the status of the Data Link Control and Management State Machine. It returns a 1b to indicate the DL_Active state, 0b otherwise. This bit must be implemented if the corresponding Data Link Layer Active Capability bit is implemented. Otherwise, this bit must be hardwired to 0b.	0h	RO_V

continued...

B/D/F/Type: 0/1/0/CFG			Access: RO_V; RO; RW1C	
Size: 16	Default Value: 1001h		Address Offset: B2h	
Bit Range	Acronym	Description	Default	Access
12	SCC	Slot Clock Configuration: 0: The device uses an independent clock irrespective of the presence of a reference on the connector. 1: The device uses the same physical reference clock that the platform provides on the connector.	1h	RO
11	LTRN	Link Training: Indicates that the Physical Layer LTSSM is in the Configuration or Recovery state, or that 1b was written to the Retrain Link bit but Link training has not yet begun. Hardware clears this bit when the LTSSM exits the Configuration/Recovery state once Link training is complete.	0h	RO_V
10	RSVD	Reserved.	0h	RO
9:4	NLW	Negotiated Link Width: Indicates negotiated link width. This field is valid only when the link is in the L0, L0s, or L1 states (after link width negotiation is successfully completed). 00h: Reserved 01h: X1 02h: X2 04h: X4 08h: X8 10h: X16 All other encodings are reserved.	00h	RO_V
3:0	CLS	Current Link Speed: This field indicates the negotiated Link speed of the given PCI Express Link. The encoding is the binary value of the bit location in the Supported Link Speeds Vector (in the Link Capabilities 2 register) that corresponds to the current Link speed. For example, a value of 0010b in this field indicates that the current Link speed is that corresponding to bit 2 in the Supported Link Speeds Vector, which is 5.0 GT/s. The value in this field is undefined when the Link is not up.	1h	RO_V

4.1.40 SLOTCAP—Slot Capabilities

PCI Express Slot related registers allow for the support of Hot Plug.

B/D/F/Type: 0/1/0/CFG			Access: RO; RW_O	
Size: 32	Default Value: 00040000h		Address Offset: B4h	
Bit Range	Acronym	Description	Default	Access
31:19	PSN	Physical Slot Number: Indicates the physical slot number attached to this Port. BIOS Requirement: This field must be initialized by BIOS to a value that assigns a slot number that is globally unique within the chassis.	0000h	RW_O
18	NCCS	No Command Completed Support: When set to 1b, this bit indicates that this slot does not generate software notification when an issued command is completed by the Hot-Plug Controller. This bit is only permitted to be set to 1b if the hotplug capable port is able to accept writes to all fields of the Slot Control register without delay between successive writes.	1h	RO

continued...

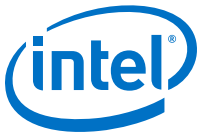

B/D/F/Type: 0/1/0/CFG			Access: RO; RW_O	
Size: 32	**Default Value: 00040000h**		**Address Offset: B4h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
17	EIP	Reserved for Electromechanical Interlock Present: When set to 1b, this bit indicates that an Electromechanical Interlock is implemented on the chassis for this slot.	0h	RO
16:15	SPLS	Slot Power Limit Scale: Specifies the scale used for the Slot Power Limit Value. 00: 1.0x 01: 0.1x 10: 0.01x 11: 0.001x If this field is written, the link sends a Set_Slot_Power_Limit message.	0h	RW_O
14:7	SPLV	Slot Power Limit Value: In combination with the Slot Power Limit Scale value, specifies the upper limit on power supplied by slot. Power limit (in Watts) is calculated by multiplying the value in this field by the value in the Slot Power Limit Scale field. If this field is written, the link sends a Set_Slot_Power_Limit message.	00h	RW_O
6	HPC	Reserved for Hot-plug Capable: When set to 1b, this bit indicates thta this slot is capable of supporting hot-plug operations.	0h	RO
5	HPS	Reserved for Hot-plug Surprise: When set to 1b, this bit indicates that an adapter present in this slot might be removed from the system without any prior notification. This is a form factor specific capability. this bit is an indication to the operating system to allow for such removal without impacting continued software operation.	0h	RO
4	PIP	Reserved for Power Indicator Present: When set to 1b, this bit indicates that a Power Indicator is electrically controlled by the chassis for this slot.	0h	RO
3	AIP	Reserved for Attention Indicator Present: When set to 1b, this bit indicates that an Attention Indicator is electrically controlled by the chassis.	0h	RO
2	MSP	Reserved for MRL Sensor Present: When set to 1b, this bit indicates that an MRL Sensor is implemented on the chassis for this slot.	0h	RO
1	PCP	Reserved for Power Controller Present: When set to 1b, this bit indicates that a software programmable Power Controller is implemented for this slot/adapter (depending on form factor).	0h	RO
0	ABP	Reserved for Attention Button Present: When set to 1b, this bit indicates that an Attention Button for this slot is electrically controlled by the chassis.	0h	RO

4.1.41 SLOTCTL—Slot Control

PCI Express Slot related registers allow for the support of Hot Plug.

B/D/F/Type: 0/1/0/CFG			Access: RO	
Size: 16	Default Value: 0000h		Address Offset: B8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:13	RSVD	Reserved.	0h	RO
12	DLLSCE	Reserved for Data Link Layer State Changed Enable: Reserved for Data Link Layer State Changed Enable (DLLSCE): If the Data Link Layer Link Active capability is implemented, when set to 1b, this field enables software notification when Data Link Layer Link Active field is changed. If the Data Link Layer Link Active capability is not implemented, this bit is permitted to be read-only with a value of 0b.	0h	RO
11	EIC	Reserved for Electromechanical Interlock Control: If an Electromechanical Interlock is implemented, a write of 1b to this field causes the state of the interlock to toggle. A write of 0b to this field has no effect. A read to this register always returns a 0.	0h	RO
10	PCC	Reserved for Power Controller Control: If a Power Controller is implemented, this field when written sets the power state of the slot per the defined encodings. Reads of this field must reflect the value from the latest write, even if the corresponding hotplug command is not complete, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined. Depending on the form factor, the power is turned on/off either to the slot or within the adapter. Note that in some cases the power controller may autonomously remove slot power or not respond to a power-up request based on a detected fault condition, independent of the Power Controller Control setting. The defined encodings are: 0b Power On 1b Power Off If the Power Controller Implemented field in the Slot Capabilities register is set to 0b, then writes to this field have no effect and the read value of this field is undefined.	0h	RO
9:8	PIC	Reserved Power Indicator Control: Reserved Power Indicator Control (PIC): If a Power Indicator is implemented, writes to this field set the Power Indicator to the written state. Reads of this field must reflect the value from the latest write, even if the corresponding hot-plug command is not complete, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined. 00: Reserved 01: On 10: Blink 11: Off If the Power Indicator Present bit in the Slot Capabilities register is 0b, this field is permitted to be read-only with a value of 00b.	0h	RO
7:6	AIC	Reserved for Attention Indicator Control: Reserved for Attention Indicator Control (AIC): If an Attention Indicator is implemented, writes to this	0h	RO

continued...

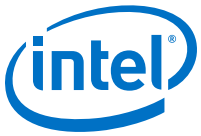

B/D/F/Type: 0/1/0/CFG			Access: RO	
Size: 16	**Default Value: 0000h**		**Address Offset: B8h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		field set the Attention Indicator to the written state. Reads of this field must reflect the value from the latest write, even if the corresponding hot-plug command is not complete, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined. If the indicator is electrically controlled by chassis, the indicator is controlled directly by the downstream port through implementation specific mechanisms. 00: Reserved 01: On 10: Blink 11: Off If the Attention Indicator Present bit in the Slot Capabilities register is 0b, this field is permitted to be readonly with a value of 00b.		
5	HPIE	Reserved for Hot-plug Interrupt Enable: When set to 1b, this bit enables generation of an interrupt on enabled hot-plug events Default value of this field is 0b. If the Hot Plug Capable field in the Slot Capabilities register is set to 0b, this bit is permitted to be read-only with a value of 0b.	0h	RO
4	CCI	Reserved for Command Completed Interrupt Enable: If Command Completed notification is supported (as indicated by No Command Completed Support field of Slot Capabilities Register), when set to 1b, this bit enables software notification when a hot-plug command is completed by the Hot-Plug Controller. Default value of this field is 0b. If Command Completed notification is not supported, this bit must be hardwired to 0b.	0h	RO
3	PDCE	Presence Detect Changed Enable: When set to 1b, this bit enables software notification on a presence detect changed event.	0h	RO
2	MSCE	Reserved for MRL Sensor Changed Enable: When set to 1b, this bit enables software notification on a MRL sensor changed event. Default value of this field is 0b. If the MRL Sensor Present field in the Slot Capabilities register is set to 0b, this bit is permitted to be read-only with a value of 0b.	0h	RO
1	PFDE	Reserved for Power Fault Detected Enable: When set to 1b, this bit enables software notification on a power fault event. Default value of this field is 0b. If Power Fault detection is not supported, this bit is permitted to be read-only with a value of 0b	0h	RO
0	ABPE	Reserved for Attention Button Pressed Enable: When set to 1b, this bit enables software notification on an attention button pressed event.	0h	RO

4.1.42 SLOTSTS—Slot Status

PCI Express Slot related registers.

B/D/F/Type: 0/1/0/CFG			Access:	RO; RW1C; RO_V
Size: 16	Default Value: 0000h		Address Offset:	BAh
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:9	RSVD	Reserved.	00h	RO
8	DLLSC	Reserved for Data Link Layer State Changed: This bit is set when the value reported in the Data Link Layer Link Active field of the Link Status register is changed. In response to a Data Link Layer State Changed event, software must read the Data Link Layer Link Active field of the Link Status register to determine if the link is active before initiating configuration cycles to the hot plugged device.	0h	RO
7	EIS	Reserved for Electromechanical Interlock Status: If an Electromechanical Interlock is implemented, this bit indicates the current status of the Electromechanical Interlock. Defined encodings are: 0b Electromechanical Interlock Disengaged 1b Electromechanical Interlock Engaged	0h	RO
6	PDS	Presence Detect State: --In band presence detect state: 0b: Slot Empty 1b: Card present in slot This bit indicates the presence of an adapter in the slot, reflected by the logical "OR" of the Physical Layer in-band presence detect mechanism and, if present, any out-of-band presence detect mechanism defined for the slot's corresponding form factor. Note that the in-band presence detect mechanism requires that power be applied to an adapter for its presence to be detected. Consequently, form factors that require a power controller for hot-plug must implement a physical pin presence detect mechanism. Defined encodings are: 0b Slot Empty 1b Card Present in slot This register must be implemented on all Downstream Ports that implement slots. For Downstream Ports not connected to slots (where the Slot Implemented bit of the PCI Express Capabilities Register is 0b), this bit must return 1b.	0h	RO_V
5	MSS	Reserved for MRL Sensor State: This register reports the status of the MRL sensor if it is implemented. Defined encodings are: 0b MRL Closed 1b MRL Open	0h	RO
4	CC	Reserved for Command Completed: If Command Completed notification is supported (as indicated by No Command Completed Support field of Slot Capabilities Register), this bit is set when a hot-plug command has completed and the Hot-Plug Controller is ready to accept a subsequent command. The Command Completed status bit is set as an indication to host software that the Hot-Plug Controller has processed the previous command and is ready to receive the next command; it provides no guarantee that the action corresponding to the command	0h	RO

continued...

B/D/F/Type: 0/1/0/CFG			Access: RO; RW1C; RO_V	
Size: 16	Default Value: 0000h		Address Offset: BAh	
Bit Range	Acronym	Description	Default	Access
		is complete. If Command Completed notification is not supported, this bit must be hardwired to 0b.		
3	PDC	Presence Detect Changed: --A pulse indication that the inband presence detect state has changed This bit is set when the value reported in Presence Detect State is changed.	0h	RW1C
2	MSC	Reserved for MRL Sensor Changed: If an MRL sensor is implemented, this bit is set when a MRL Sensor state change is detected. If an MRL sensor is not implemented, this bit must not be set.	0h	RO
1	PFD	Reserved for Power Fault Detected: If a Power Controller that supports power fault detection is implemented, this bit is set when the Power Controller detects a power fault at this slot. Note that, depending on hardware capability, it is possible that a power fault can be detected at any time, independent of the Power Controller Control setting or the occupancy of the slot. If power fault detection is not supported, this bit must not be set.	0h	RO
0	ABP	Reserved for Attention Button Pressed: If an Attention Button is implemented, this bit is set when the attention button is pressed. If an Attention Button is not supported, this bit must not be set.	0h	RO

4.1.43 RCTL—Root Control

Allows control of PCI Express Root Complex specific parameters. The system error control bits in this register determine if corresponding SERRs are generated when our device detects an error (reported in this device's Device Status register) or when an error message is received across the link. Reporting of SERR as controlled by these bits takes precedence over the SERR Enable in the PCI Command Register.

B/D/F/Type: 0/1/0/CFG			Access: RW; RO; RWS_V; RO_V	
Size: 32	Default Value: 00000000h		Address Offset: BCh	
Bit Range	Acronym	Description	Default	Access
31:19	RSVD	Reserved.	0000h	RO
18	INSWL2	this bit indicates whether the current state of the LTSSM is L2, and it got there through the SWL2 flow. This bit is part of the RTD3 flow ,and may be used as indication for initiate power savings measures.	0h	RO_V
17:9	RSVD	Reserved.	000h	RO
8	SWL2DIS	SWL2 disable bit, when set writing to bits RCTL.L23ER or RCTL.L22DT wont have any affect over H/W	0h	RWS_V
7	RSVD	Reserved.	0h	RO

continued...

B/D/F/Type: 0/1/0/CFG			Access: RW; RO; RWS_V; RO_V	
Size: 32	Default Value: 00000000h		Address Offset: BCh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
6	L22DT	L23 to Detect Transition: When set by software, the LTSSM moves from L2 to Detect state. This bit is cleared by hardware once the LTSSM is in detect (and in general, while not in L2 state)	0h	RWS_V
5	L23ER	L23_Rdy Entry Request: When set by software, the corresponding PCIe root port will initiate the sequence to put the link into L2/L3 Ready state. PME_Turn_Off message will be sent and the corresponding PME_TO_Ack response will be returned by the device. Once the link enters L2, this bit will be cleared by hardware.	0h	RWS_V
4	CSVE	Reserved for CRS Software Visibility Enable: This bit, when set, enables the Root Port to return Configuration Request Retry Status (CRS) Completion Status to software. Root Ports that do not implement this capability must hardwire this bit to 0b.	0h	RO
3	PMEIE	PME Interrupt Enable: 0: No interrupts are generated as a result of receiving PME messages. 1: Enables interrupt generation upon receipt of a PME message as reflected in the PME Status bit of the Root Status Register. A PME interrupt is also generated if the PME Status bit of the Root Status Register is set when this bit is set from a cleared state. If the bit change from 1 to 0 and interrupt is pending than interrupt is deasserted	0h	RW
2	SEFEE	System Error on Fatal Error Enable: Controls the Root Complex's response to fatal errors. 0: No SERR generated on receipt of fatal error. 1: Indicates that an SERR should be generated if a fatal error is reported by any of the devices in the hierarchy associated with this Root Port, or by the Root Port itself.	0h	RW
1	SENFUEE	System Error on Non-Fatal Uncorrectable Error Enable: Controls the Root Complex's response to non-fatal errors. 0: No SERR generated on receipt of non-fatal error. 1: Indicates that an SERR should be generated if a non-fatal error is reported by any of the devices in the hierarchy associated with this Root Port, or by the Root Port itself.	0h	RW
0	SECEE	System Error on Correctable Error Enable: Controls the Root Complex's response to correctable errors. 0: No SERR generated on receipt of correctable error. 1: Indicates that an SERR should be generated if a correctable error is reported by any of the devices in the hierarchy associated with this Root Port, or by the Root Port itself.	0h	RW

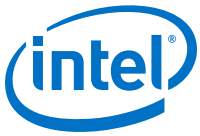

4.1.44 RSTS—Root Status

Provides information about PCI Express Root Complex specific parameters.

B/D/F/Type: 0/1/0/CFG			Access: RO_V; RW1C; RO	
Size: 32	**Default Value: 00000000h**		**Address Offset: C0h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:18	RSVD	Reserved.	0000h	RO
17	PMEP	PME Pending: Indicates that another PME is pending when the PME Status bit is set. When the PME Status bit is cleared by software; the PME is delivered by hardware by setting the PME Status bit again and updating the Requestor ID appropriately. The PME pending bit is cleared by hardware if no more PMEs are pending. Intel Reserved text: In order to simplify design and validation, if a 2nd PME arrives, before the 1st is processed, we will drop it. The initiating device must re-initiate the PME, as described in the spec so there should be no functional impact by taking this approach. This means this bit is hardwired to '0' and will never be set to '1'.	0h	RO
16	PMES	PME Status: Indicates that PME was asserted by the requestor ID indicated in the PME Requestor ID field. Subsequent PMEs are kept pending until the status register is cleared by writing a 1 to this field. An interrupt is asserted If PMEIE is asserted and PMES is changing from 0 to 1 An interrupt is deasserted If PMEIE is asserted and PMES is changing from 1 to 0 An Assert_PMEGPE is sent upstream If PMEGPEE in PEG Legacy cControl register (PEGLC) is asserted and PMES is changing from 0 to 1 An Deassert_PMEGPE is sent upstream If PMEGPEE in PEG Legacy cControl register (PEGLC) is asserted and PMES is changing from 1 to 0 An interrupt is deasserted If PMEIE is asserted and PMES is changing from 1 to 0 Intel Reserved text: In order to simplify design and validation, if a 2nd PME arrives, before the 1st is processed, we will drop it. The initiating device must re-initiate the PME, as described in the spec so there should be no functional impact by taking this approach.	0h	RW1C
15:0	PMERID	PME Requestor ID: Indicates the PCI requestor ID of the last PME requestor.	0000h	RO_V

4.1.45 DCAP2—Device Capabilites 2

B/D/F/Type: 0/1/0/CFG			Access: RO; RW_O	
Size: 32	**Default Value: 00000B80h**		**Address Offset: C4h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:20	RSVD	Reserved.	000h	RO
19:18	OBFF_SUPPORTED	OBFF Supported 00b OBFF Not Supported 01b OBFF supported using Message signaling only	0h	RW_O

continued...

B/D/F/Type: 0/1/0/CFG			Access: RO; RW_O	
Size: 32	Default Value: 00000B80h		Address Offset: C4h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		10b OBFF supported using WAKE# signaling only 11b OBFF supported using WAKE# and Message signaling The value reported in this field must indicate support for WAKE# signaling only if: - for a Downstream Port, driving the WAKE# signal for OBFF is supported and the connector or component connected Downstream is known to receive that same WAKE# signal - for an Upstream Port, receiving the WAKE# signal for OBFF is supported and, if the component is on an add-in-card, that the component is connected to the WAKE# signal on the connector. Root Ports, Switch Ports, and Endpoints are permitted to implement this capability. For a multi-Function device associated with an Upstream Port, each Function must report the same value for this field. For Bridges and Ports that do not implement this capability, this field must be hardwired to 00b.		
17:12	RSVD	Reserved.	00h	RO
11	LTRS	Latency Tolerance and BW reporting Mechanism Supported: A value of 1b indicates support for the optional Latency Tolerance & Bandwidth Requirement Reporting (LTBWR) mechanism capability. Root Ports, Switches and Endpoints are permitted to implement this capability. For Switches that implement LTBWR, this bit must be set only at the upstream port. For a multi-Function device, each Function must report the same value for this bit. For Bridges, Downstream Ports, and components that do not implement this capability, this bit must be hardwired to 0b.	1h	RO
10	RSVD	Reserved.	0h	RO
9	ATOMIC128SUP	128-bit CAS atomic operation completion support. This bit must be set to 1b if the Function supports this optional capability.	1h	RO
8	ATOMIC64SUP	64-bit atomic operation completion support. Includes FetchAdd, Swap, and CAS AtomicOps. This bit must be set to 1b if the Function supports this optional capability.	1h	RO
7	ATOMIC32SUP	32-bit atomic operation completion support. Includes FetchAdd, Swap, and CAS AtomicOps. This bit must be set to 1b if the Function supports this optional capability.	1h	RO
6	ATOMIC_OP_ROUTING_SUPPORT	Atomic Operation Routing Supported. If set then then atomic operations are supported.	0h	RO

continued...

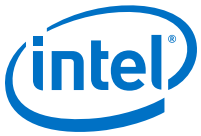

B/D/F/Type: 0/1/0/CFG			Access: RO; RW_O	
Size: 32	Default Value: 00000B80h		Address Offset: C4h	
Bit Range	Acronym	Description	Default	Access
5	ARIFS	ARI Forwarding Supported: Applicable only to Switch Downstream Ports and Root Ports; must be 0b for other Function types. This bit must be set to 1b if a Switch Downstream Port or Root Port supports this optional capability.	0h	RW_O
4	CTODS	Completion Timeout Disabled Supported: A value of 1b indicates support for the Completion Timeout Disable mechanism. The Completion Timeout Disable mechanism is required for Endpoints that issue Requests on their own behalf and PCI Express to PCI/PCI-X Bridges that take ownership of Requests issued on PCI Express. This mechanism is optional for Root Ports. The Root port does nopt support completion timeout disable	0h	RO
3:0	CTOR	Completion Timer Ranges Supported: device Function support for the optional Completion Timeout programmability mechanism. This mechanism allows system software to modify the Completion Timeout value. This field is applicable only to Root Ports, Endpoints that issue Requests on their own behalf, and PCI Express to PCI/PCI-X Bridges that take ownership of Requests issued on PCI Express. For all other Functions this field is reserved and must be hardwired to 0000b. 0000b Completion Timeout programming not supported - the Function must implement a timeout value in the range 50 us to 50 ms.	0h	RO

4.1.46 DCTL2—Device Control 2

B/D/F/Type: 0/1/0/CFG			Access: RW; RO; RW_V	
Size: 16	Default Value: 0000h		Address Offset: C8h	
Bit Range	Acronym	Description	Default	Access
15	RSVD	Reserved.	0h	RO
14:13	OBFFEN		0h	RW
12:11	RSVD	Reserved.	0h	RO
10	LTREN	Latency Tolerance Reporting Mechansim Enable: When Set to 1b, this bit enables the Latency Tolerance & Reporting (LTR) mechanism. This bit is required for all Functions that support the LTR Capability. For a Multi-Function device associated with an upstream port of a device that implements LTBWR, the bit in Function 0 is of type RW, and only Function 0 controls the components Link behavior. In all other Functions of that device, this bit is of type RsvdP. Components that do not implement LTR are permitted to hardwire this bit to 0b. Default value of this bit is 0b.	0h	RW_V

continued...

B/D/F/Type: 0/1/0/CFG			Access: RW; RO; RW_V	
Size: 16	Default Value: 0000h		Address Offset: C8h	
Bit Range	Acronym	Description	Default	Access
		This bit is cleared when the port goes to DL_down state. HW ignores the value of this bit.		
9:7	RSVD	Reserved.	0h	RO
6	ATOMIC_OP_REQUESTER_EN	AtomicOp Requester Enable Applicable only to Endpoints and Root Ports; must be hardwired to 0b for other Function types. The Function is allowed to initiate AtomicOp Requests only if this bit and the Bus Master Enable bit in the Command register are both Set. This bit is required to be RW if the Endpoint or Root Port is capable of initiating AtomicOp Requests, but otherwise is permitted to be hardwired to 0b. This bit does not serve as a capability bit. This bit is permitted to be RW even if no AtomicOp Requester capabilities are supported by the Endpoint or Root Port.	0h	RO
5	ARIFEN	ARI Forward Enable: When set, the Downstream Port disables its traditional Device Number field being 0 enforcement when turning a Type 1 Configuration Request into a Type 0 Configuration Request, permitting access to Extended Functions in an ARI Device immediately below the Port. Default value of this bit is 0b. Must be hardwired to 0b if the ARI Forwarding Supported bit is 0b.	0h	RW
4:0	RSVD	Reserved.	00h	RO

4.1.47 LCTL2—Link Control 2

B/D/F/Type: 0/1/0/CFG			Access: RWS; RWS_V	
Size: 16	Default Value: 0003h		Address Offset: D0h	
Bit Range	Acronym	Description	Default	Access
15:12	ComplianceDeemphasis	Compliance De-emphasis: For 8 GT/s Data Rate: This field sets the Transmitter Preset level in Polling.Compliance state if the entry occurred due to the Enter Compliance bit being 1b. This bit sets the de-emphasis level in Polling.Compliance state if the entry occurred due to the Enter Compliance bit being 1b. Defined encodings are: 0001b -3.5 dB 0000b -6 dB When the Link is operating at 2.5 GT/s, the setting of this bit has	0h	RWS

continued...

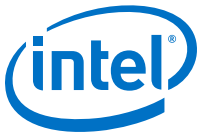

B/D/F/Type: 0/1/0/CFG			Access: RWS; RWS_V	
Size: 16	**Default Value: 0003h**		**Address Offset: D0h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		no effect. Components that support only 2.5 GT/s speed are permitted to hardwire this bit to 0b. For a Multi-Function device associated with an Upstream Port, the bit in Function 0 is of type RWS, and only Function 0 controls the component's Link behavior. In all other Functions of that device, this bit is of type RsvdP. The default value of this bit is 0000b. This bit is intended for debug, compliance testing purposes. System firmware and software is allowed to modify this bit only during debug or compliance testing.		
11	compsos	Compliance SOS: When set to 1b, the LTSSM is required to send SKP Ordered Sets periodically in between the (modified) compliance patterns. For a Multi-Function device associated with an Upstream Port, the bit in Function 0 is of type RWS, and only Function 0 controls the component's Link behavior. In all other Functions of that device, this bit is of type RsvdP. The default value of this bit is 0b. This bit is applicable when the Link is operating at 2.5 GT/s or 5 GT/s data rates only. Components that support only the 2.5 GT/s speed are permitted to hardwire this field to 0b.	0h	RWS
10	entermodcompliance	Enter Modified Compliance: When this bit is set to 1b, the device transmits modified compliance pattern if the LTSSM enters Polling.Compliance state. Components that support only the 2.5GT/s speed are permitted to hardwire this bit to 0b. Default value of this field is 0b.	0h	RWS
9:7	txmargin	Transmit Margin: This field controls the value of the non-deemphasized voltage level at the Transmitter pins. This field is reset to 000b on entry to the LTSSM Polling.Configuration substate (see Chapter 4 for details of how the transmitter voltage level is determined in various states). Encodings: 000: Normal operating range 001: 800-1200 mV for full swing and 400-700 mV for half-swing 010 - (n-1): Values must be monotonic with a non-zero slope. The value of n must be greater than 3 and less than 7. At least two of these must be below the normal operating range n : 200-400 mV for full-swing and 100-200 mV for half-swing n -111: reserved Default value is 000b. Components that support only the 2.5GT/s speed are	0h	RWS_V

continued...

B/D/F/Type: 0/1/0/CFG			Access: RWS; RWS_V	
Size: 16	Default Value: 0003h		Address Offset: D0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		permitted to hardwire this bit to 0b. When operating in 5GT/s mode with full swing, the deemphasis ratio must be maintained within +/- 1dB from the spec defined operational value (either -3.5 or -6 dB).		
6	selectabledeemphasis	Selectable De-emphasis: When the Link is operating at 5GT/s speed, selects the level of de-emphasis. Encodings: 1b -3.5 dB 0b -6 dB Default value is implementation specific, unless a specific value is required for a selected form factor or platform. When the Link is operating at 2.5GT/s speed, the setting of this bit has no effect. Components that support only the 2.5GT/s speed are permitted to hardwire this bit to 0b.	0h	RWS
5	HASD	Hardware Autonomous Speed Disable: When set to 1b this bit disables hardware from changing the link speed for reasons other than attempting to correct unreliable link operation by reducing link speed.	0h	RWS
4	EC	Enter Compliance: Software is permitted to force a link to enter Compliance mode at the speed indicated in the Target Link Speed field by setting this bit to 1b in both components on a link and then initiating a hot reset on the link.	0h	RWS
3:0	TLS	Target Link Speed: For Downstream Ports, this field sets an upper limit on Link operational speed by restricting the values advertised by the Upstream component in its training sequences. The encoding is the binary value of the bit in the Supported Link Speeds Vector (in the Link Capabilities 2 register) that corresponds to the desired target Link speed. All other encodings are reserved. For example, 5.0 GT/s corresponds to bit 2 in the Supported Link Speeds Vector, so the encoding for a 5.0 GT/s target Link speed in this field is 0010b. If a value is written to this field that does not correspond to a supported speed (as indicated by the Max Link Speed Vector), the result is undefined. The default value of this field is the highest Link speed supported by the component (as reported in the Max Link Speed field of the Link Capabilities register) unless the corresponding platform/form factor requires a different default value. For both Upstream and Downstream Ports, this field is used to set the target compliance mode speed when software is using the Enter Compliance bit to force a Link into compliance mode. For a Multi-Function device associated with an Upstream Port, the field in Function 0 is of type RWS, and only Function 0	3h	RWS

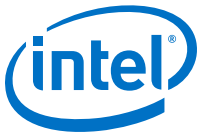

B/D/F/Type: 0/1/0/CFG			Access: RWS; RWS_V	
Size: 16	**Default Value: 0003h**		**Address Offset: D0h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		controls the components Link behavior. In all other Functions of that device, this field is of type RsvdP.		

4.1.48 LSTS2—Link Status 2

B/D/F/Type: 0/1/0/CFG			Access: RO_V; RW1C	
Size: 16	**Default Value: 0000h**		**Address Offset: D2h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:6	RSVD	Reserved.	000h	RO
5	LNKEQREQ	This bit is Set by hardware to request the Link equalization process to be performed on the Link.	0h	RW1C
4	EQPH3SUCC	Equalization Phase 3 Successful When set to 1b, this bit indicates that Phase 3 of the Transmitter Equalization procedure has successfully completed.	0h	RO_V
3	EQPH2SUCC	Equalization Phase 2 Successful When set to 1b, this bit indicates that Phase 2 of the Transmitter Equalization procedure has successfully completed.	0h	RO_V
2	EQPH1SUCC	Equalization Phase 1 Successful When set to 1b, this bit indicates that Phase 1 of the Transmitter Equalization procedure has successfully completed.	0h	RO_V
1	EQCOMPLETE	Equalization Complete When set to 1b, this bit indicates that the Transmitter Equalization procedure has completed.	0h	RO_V
0	CURDELVL	Current De-emphasis Level: Current De-emphasis Level - When the Link is operating at 5 GT/s speed, this reflects the level of de-emphasis. Encodings: 1b -3.5 dB 0b -6 dB When the Link is operating at 2.5 GT/s speed, this bit is 0b.	0h	RO_V

4.1.49 PVCCAP1—Port VC Capability Register 1

Describes the configuration of PCI Express Virtual Channels associated with this port.

B/D/F/Type: 0/1/0/CFG			Access: RO	
Size: 32	Default Value: 00000000h		Address Offset: 104h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:7	RSVD	Reserved.	0000000h	RO
6:4	LPEVCC	Low Priority Extended VC Count: Indicates the number of (extended) Virtual Channels in addition to the default VC belonging to the low-priority VC (LPVC) group that has the lowest priority with respect to other VC resources in a strict-priority VC Arbitration. The value of 0 in this field implies strict VC arbitration.	0h	RO
3	RSVD	Reserved.	0h	RO
2:0	EVCC	Extended VC Count: Indicates the number of (extended) Virtual Channels in addition to the default VC supported by the device.	0h	RO

4.1.50 PVCCAP2—Port VC Capability Register 2

Describes the configuration of PCI Express Virtual Channels associated with this port.

B/D/F/Type: 0/1/0/CFG			Access: RO	
Size: 32	Default Value: 00000000h		Address Offset: 108h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:24	VCATO	VC Arbitration Table Offset: Indicates the location of the VC Arbitration Table. This field contains the zero-based offset of the table in DQWORDS (16 bytes) from the base address of the Virtual Channel Capability Structure. A value of 0 indicates that the table is not present (due to fixed VC priority).	00h	RO
23:8	RSVD	Reserved.	0000h	RO
7:0	VCAC	Reserved for VC Arbitration Capability:	00h	RO

4.1.51 PVCCTL—Port VC Control

B/D/F/Type: 0/1/0/CFG			Access: RO; RW	
Size: 16	Default Value: 0000h		Address Offset: 10Ch	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:4	RSVD	Reserved.	000h	RO
3:1	VCAS	VC Arbitration Select: This field will be programmed by software to the only possible value as indicated in the VC Arbitration Capability field. Since there is no other VC supported than the default, this field is reserved.	0h	RW
0	VCARB	Reserved for Load VC Arbitration Table: Used for software to update the VC Arbitration Table when VC arbitration uses the VC Arbitration Table. As a VC Arbitration Table is never used by this component this field will never be used.	0h	RO

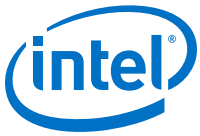

4.1.52 VC0RCAP—VC0 Resource Capability

B/D/F/Type: 0/1/0/CFG			Access: RO	
Size: 32	**Default Value: 00000001h**		**Address Offset: 110h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:24	PATO	Reserved for Port Arbitration Table Offset:	00h	RO
23	RSVD	Reserved.	0h	RO
22:16	MTS	Reserved for Maximum Time Slots:	00h	RO
15	RSNPT	Reject Snoop Transactions: Reject Snoop Transactions (RSNPT): 0: Transactions with or without the No Snoop bit set within the TLP header are allowed on this VC. 1: When Set, any transaction for which the No Snoop attribute is applicable but is not Set within the TLP Header will be rejected as an Unsupported Request	0h	RO
14:8	RSVD	Reserved.	00h	RO
7:0	PAC	Port Arbitration Capability: Port Arbitration Capability - Indicates types of Port Arbitration supported by the VC resource. This field is valid for all Switch Ports, Root Ports that support peer-to to-peer traffic, and RCRBs, but not for PCI Express Endpoint devices or Root Ports that do not support peer to peer traffic. Each bit location within this field corresponds to a Port Arbitration Capability defined below. When more than one bit in this field is Set, it indicates that the VC resource can be configured to provide different arbitration services. Software selects among these capabilities by writing to the Port Arbitration Select field (see below). Defined bit positions are: Bit 0 Non-configurable hardware-fixed arbitration scheme, e.g., Round Robin (RR) Bit 1 Weighted Round Robin (WRR) arbitration with 32 phases Bit 2 WRR arbitration with 64 phases Bit 3 WRR arbitration with 128 phases Bit 4 Time-based WRR with 128 phases Bit 5 WRR arbitration with 256 phases Bits 6-7 Reserved Processor only supported arbitration indicates "Non-configurable hardware-fixed arbitration scheme".	01h	RO

4.1.53 VC0RCTL—VC0 Resource Control

Controls the resources associated with PCI Express Virtual Channel 0.

B/D/F/Type: 0/1/0/CFG			Access: RO; RW	
Size: 32	**Default Value: 800000FFh**		**Address Offset: 114h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31	VC0E	VC0 Enable: For VC0 this is hardwired to 1 and read only as VC0 can never be disabled.	1h	RO
30:27	RSVD	Reserved.	0h	RO
26:24	VC0ID	VC0 ID: Assigns a VC ID to the VC resource. For VC0 this is hardwired to 0 and read only.	0h	RO

continued...

B/D/F/Type: 0/1/0/CFG			Access: RO; RW	
Size: 32	Default Value: 800000FFh		Address Offset: 114h	
Bit Range	Acronym	Description	Default	Access
23:20	RSVD	Reserved.	0h	RO
19:17	PAS	Port Arbitration Select: Port Arbitration Select - This field configures the VC resource to provide a particular Port Arbitration service. This field is valid for RCRBs, Root Ports that support peer to peer traffic, and Switch Ports, but not for PCI Express Endpoint devices or Root Ports that do not support peer to peer traffic. The permissible value of this field is a number corresponding to one of the asserted bits in the Port Arbitration Capability field of the VC resource. This field does not affect the root port behavior.	0h	RW
16	RSVD	Reserved.	0h	RO
15:8	TCHVC0M	TC High VC0 Map: Allow usage of high order TCs. BIOS should keep this field zeroed to allow usage of the reserved TC[3] for other purposes	00h	RW
7:1	TCVC0M	TC/VC0 Map: Indicates the TCs (Traffic Classes) that are mapped to the VC resource. Bit locations within this field correspond to TC values. For example, when bit 7 is set in this field, TC7 is mapped to this VC resource. When more than one bit in this field is set, it indicates that multiple TCs are mapped to the VC resource. In order to remove one or more TCs from the TC/VC Map of an enabled VC, software must ensure that no new or outstanding transactions with the TC labels are targeted at the given Link.	7Fh	RW
0	TC0VC0M	TC0/VC0 Map: Traffic Class 0 is always routed to VC0.	1h	RO

4.1.54 VC0RSTS—VC0 Resource Status

Reports the Virtual Channel specific status.

B/D/F/Type: 0/1/0/CFG			Access: RO_V	
Size: 16	Default Value: 0002h		Address Offset: 11Ah	
Bit Range	Acronym	Description	Default	Access
15:2	RSVD	Reserved.	0000h	RO
1	VC0NP	VC0 Negotiation Pending: 0: The VC negotiation is complete. 1: The VC resource is still in the process of negotiation (initialization or disabling). This bit indicates the status of the process of Flow Control initialization. It is set by default on Reset, as well as whenever the corresponding Virtual Channel is Disabled or the Link is in the DL_Down state. It is cleared when the link successfully exits the FC_INIT2 state. Before using a Virtual Channel, software must check whether the VC Negotiation Pending fields for that Virtual Channel are cleared in both Components on a Link.	1h	RO_V
0	RSVD	Reserved.	0h	RO

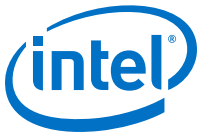

4.2 PCI Express Controller (x8) Registers Summary

Offset	Register ID—Description	Default Value	Access
0	VID—Vendor Identification on page 154	8086h	RO
2	DID—Device Identification on page 155	0C05h	RO
4	PCICMD—PCI Command on page 155	0000h	RW; RO
6	PCISTS—PCI Status on page 156	0010h	RO_V; RO; RW1C
8	RID—Revision Identification on page 158	00h	RO
9	CC—Class Code on page 159	060400h	RO
C	CL—Cache Line Size on page 159	00h	RW
E	HDR—Header Type on page 159	81h	RO
18	PBUSN—Primary Bus Number on page 160	00h	RO
19	SBUSN—Secondary Bus Number on page 160	00h	RW
1A	SUBUSN—Subordinate Bus Number on page 160	00h	RW
1C	IOBASE—I/O Base Address on page 160	F0h	RW
1D	IOLIMIT—I/O Limit Address on page 161	00h	RW
1E	SSTS—Secondary Status on page 161	0000h	RO; RW1C
20	MBASE—Memory Base Address on page 162	FFF0h	RW
22	MLIMIT—Memory Limit Address on page 163	0000h	RW
24	PMBASE—Prefetchable Memory Base Address on page 163	FFF1h	RO; RW
26	PMLIMIT—Prefetchable Memory Limit Address on page 164	0001h	RO; RW
28	PMBASEU—Prefetchable Memory Base Address Upper on page 164	00000000h	RW
2C	PMLIMITU—Prefetchable Memory Limit Address Upper on page 165	00000000h	RW
34	CAPPTR—Capabilities Pointer on page 165	88h	RO
3C	INTRLINE—Interrupt Line on page 166	00h	RW
3D	INTRPIN—Interrupt Pin on page 166	01h	RW_O; RO
3E	BCTRL—Bridge Control on page 167	0000h	RW; RO
80	PM—Power Management Capabilities on page 168	C8039001h	RO; RO_V
84	PM—Power Management Control/Status on page 169	00000008h	RW; RO
88	SS—Subsystem ID and Vendor ID Capabilities on page 170	0000800Dh	RO
8C	SS—Subsystem ID and Subsystem Vendor ID on page 171	00008086h	RW_O
90	MSI—Message Signaled Interrupts Capability ID on page 171	A005h	RO
92	MC—Message Control on page 171	0000h	RW; RO
94	MA—Message Address on page 172	00000000h	RO; RW
98	MD—Message Data on page 172	0000h	RW
A0	PEG—PCI Express Capability List on page 173	0010h	RO
A2	PEG—PCI Express Capabilities on page 173	0142h	RO; RW_O

continued...

Offset	Register ID—Description	Default Value	Access
A4	DCAP—Device Capabilities on page 173	00008000h	RW_O; RO
A8	DCTL—Device Control on page 174	0000h	RW; RO
AA	DSTS—Device Status on page 175	0000h	RW1C; RO
B0	LCTL—Link Control on page 176	0000h	RW; RO; RW_V
B2	LSTS—Link Status on page 178	1001h	RO_V; RO; RW1C
B4	SLOTCAP—Slot Capabilities on page 179	00040000h	RO; RW_O
B8	SLOTCTL—Slot Control on page 181	0000h	RO
BA	SLOTSTS—Slot Status on page 183	0000h	RO; RW1C; RO_V
BC	RCTL—Root Control on page 184	00000000h	RW; RO; RWS_V; RO_V
C0	RSTS—Root Status on page 186	00000000h	RO_V; RW1C; RO
C4	DCAP2—Device Capabilites 2 on page 186	00000B80h	RO; RW_O
C8	DCTL2—Device Control 2 on page 188	0000h	RW; RO; RW_V
D0	LCTL2—Link Control 2 on page 189	0003h	RWS; RWS_V
D2	LSTS2—Link Status 2 on page 192	0000h	RO_V; RW1C
104	PVCCAP1—Port VC Capability Register 1 on page 193	00000000h	RO
108	PVCCAP2—Port VC Capability Register 2 on page 193	00000000h	RO
10C	PVCCTL—Port VC Control on page 193	0000h	RO; RW
110	VC0RCAP—VC0 Resource Capability on page 194	00000001h	RO
114	VC0RCTL—VC0 Resource Control on page 194	800000FFh	RO; RW
11A	VC0RSTS—VC0 Resource Status on page 195	0002h	RO_V

4.2.1 VID—Vendor Identification

This register combined with the Device Identification register uniquely identify any PCI device.

B/D/F/Type: 0/1/1/CFG			Access: RO	
Size: 16	Default Value: 8086h		Address Offset: 0h	
Bit Range	Acronym	Description	Default	Access
15:0	VID	Vendor Identification: PCI standard identification for Intel.	8086h	RO

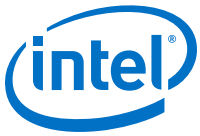

4.2.2 DID—Device Identification

This register combined with the Vendor Identification register uniquely identifies any PCI device.

B/D/F/Type: 0/1/1/CFG			Access: RO	
Size: 16	Default Value: 0C05h		Address Offset: 2h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:0	DID_MSB	Device Identification Number MSB: Identifier assigned to the Processor root port (virtual PCI-to-PCI bridge, PCI Express Graphics port).	0C05h	RO

4.2.3 PCICMD—PCI Command

B/D/F/Type: 0/1/1/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: 4h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:11	RSVD	Reserved.	00h	RO
10	INTAAD	INTA Assertion Disable: 0: This device is permitted to generate INTA interrupt messages. 1: This device is prevented from generating interrupt messages. Any INTA emulation interrupts already asserted must be de-asserted when this bit is set. Only affects interrupts generated by the device (PCI INTA from a PME or Hot Plug event) controlled by this command register. It does not affect upstream MSIs, upstream PCI INTA-INTD assert and deassert messages.	0h	RW
9	FB2B	Fast Back-to-Back Enable: Not Applicable or Implemented. Hardwired to 0.	0h	RO
8	SERRE	SERR# Message Enable: Controls the root port's SERR# messaging. The processor communicates the SERR# condition by sending an SERR message to the PCH. This bit, when set, enables reporting of non-fatal and fatal errors detected by the device to the Root Complex. Note that errors are reported if enabled either through this bit or through the PCI-Express specific bits in the Device Control Register. In addtion, for Type 1 configuration space header devices, this bit, when set, enables transmission by the primary interface of ERR_NONFATAL and ERR_FATAL error messages forwarded from the secondary interface. This bit does not affect the transmission of forwarded ERR_COR messages. 0: The SERR message is generated by the root port only under conditions enabled individually through the Device Control Register. 1: The root port is enabled to generate SERR messages which will be sent to the PCH for specific root port error conditions generated/detected or received on the secondary side of the virtual PCI to PCI bridge. The status of SERRs generated is reported in the PCISTS register.	0h	RW
7	RSVD	Reserved.	0h	RO
6	PERRE	Parity Error Response Enable: Controls whether or not the Master Data Parity Error bit in the PCI Status register can bet set. 0: Master Data Parity Error bit in PCI Status register	0h	RW

continued...

B/D/F/Type: 0/1/1/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: 4h	
Bit Range	Acronym	Description	Default	Access
		can NOT be set. 1: Master Data Parity Error bit in PCI Status register CAN be set.		
5	VGAPS	VGA Palette Snoop: Not Applicable or Implemented. Hardwired to 0.	0h	RO
4	MWIE	Memory Write and Invalidate Enable: Not Applicable or Implemented. Hardwired to 0.	0h	RO
3	SCE	Special Cycle Enable: Not Applicable or Implemented. Hardwired to 0.	0h	RO
2	BME	Bus Master Enable: Bus Master Enable (BME): Controls the ability of the PEG port to forward Memory Read/Write Requests in the upstream direction. 0: This device is prevented from making memory requests to its primary bus. Note that according to PCI Specification, as MSI interrupt messages are in-band memory writes, disabling the bus master enable bit prevents this device from generating MSI interrupt messages or passing them from its secondary bus to its primary bus. Upstream memory writes/reads, peer writes/reads, and MSIs will all be treated as illegal cycles. Writes are aborted. Reads are aborted and will return Unsupported Request status (or Master abort) in its completion packet. 1: This device is allowed to issue requests to its primary bus. Completions for previously issued memory read requests on the primary bus will be issued when the data is available. This bit does not affect forwarding of Completions from the primary interface to the secondary interface.	0h	RW
1	MAE	Memory Access Enable: 0: All of device's memory space is disabled. 1: Enable the Memory and Pre-fetchable memory address ranges defined in the MBASE, MLIMIT, PMBASE, and PMLIMIT registers.	0h	RW
0	IOAE	IO Access Enable: 0: All of devic's I/O space is disabled. 1: Enable the I/O address range defined in the IOBASE, and IOLIMIT registers.	0h	RW

4.2.4 PCISTS—PCI Status

This register reports the occurrence of error conditions associated with primary side of the "virtual" Host-PCI Express bridge embedded within the Root port.

B/D/F/Type: 0/1/1/CFG			Access: RO_V; RO; RW1C	
Size: 16	Default Value: 0010h		Address Offset: 6h	
Bit Range	Acronym	Description	Default	Access
15	DPE	Detected Parity Error: This bit is Set by a Function whenever it receives a Poisoned TLP, regardless of the state the Parity Error Response bit in the Command register. On a Function with a Type 1	0h	RW1C

continued...

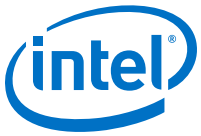

B/D/F/Type: 0/1/1/CFG			Access: RO_V; RO; RW1C	
Size: 16	**Default Value: 0010h**		**Address Offset: 6h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		Configuration header, the bit is Set when the Poisoned TLP is received by its Primary Side. Default value of this bit is 0b. This bit will be set only for completions of requests enountering ECC error in DRAM. Poisoned Peer 2 peer posted forwarded will not set this bit. They are reported at the receiveing port.		
14	SSE	Signaled System Error: This bit is set when this Device sends an SERR due to detecting an ERR_FATAL or ERR_NONFATAL condition and the SERR Enable bit in the Command register is '1'. Both received (if enabled by BCTRL1[1]) and internally detected error messages do not affect this field.	0h	RW1C
13	RMAS	Received Master Abort Status: This bit is Set when a Requester receives a Completion with Unsupported Request Completion Status. On a Function with a Type 1 Configuration header, the bit is Set when the Unsupported Request is received by its Primary Side. Not applicable. We do not have UR on primary interface	0h	RO
12	RTAS	Received Target Abort Status: This bit is Set when a Requester receives a Completion with Completer Abort Completion Status. On a Function with a Type 1 Configuration header, the bit is Set when the Completer Abort is received by its Primary Side. Default value of this bit is 0b. Not Applicable or Implemented. Hardwired to 0. The concept of a Completer abort does not exist on primary side of this device.	0h	RO
11	STAS	Signaled Target Abort Status: This bit is Set when a Function completes a Posted or Non- Posted Request as a Completer Abort error. This applies to a Function with a Type 1 Configuration header when the Completer Abort was generated by its Primary Side. Default value of this bit is 0b. Not Applicable or Implemented. Hardwired to 0. The concept of a target abort does not exist on primary side of this device.	0h	RO
10:9	DEVT	DEVSELB Timing: This device is not the subtractively decoded device on bus 0. This bit field is therefore hardwired to 00 to indicate that the device uses the fastest possible decode. Does not apply to PCI Express and must be hardwired to 00b.	0h	RO
8	PMDPE	Master Data Parity Error: This bit is Set by a Requester (Primary Side for Type 1 Configuration Space header Function) if the Parity Error Response bit in the Command register is 1b and either of the following two conditions occurs:	0h	RW1C

continued...

B/D/F/Type: 0/1/1/CFG			Access: RO_V; RO; RW1C	
Size: 16	Default Value: 0010h		Address Offset: 6h	
Bit Range	Acronym	Description	Default	Access
		Requester receives a Completion marked poisoned Requester poisons a write Request If the Parity Error Response bit is 0b, this bit is never Set. Default value of this bit is 0b. This bit will be set only for completions of requests enountering ECC error in DRAM. Poisoned Peer 2 peer posted forwarded will not set this bit. They are reported at the receiveing port.		
7	FB2B	Fast Back-to-Back: Not Applicable or Implemented. Hardwired to 0.	0h	RO
6	RSVD	Reserved.	0h	RO
5	CAP66	66/60MHz capability: Not Applicable or Implemented. Hardwired to 0.	0h	RO
4	CAPL	Capabilities List: Indicates that a capabilities list is present. Hardwired to 1.	1h	RO
3	INTAS	INTx Status: Indicates that an interrupt message is pending internally to the device. Only PME and Hot Plug sources feed into this status bit (not PCI INTA-INTD assert and deassert messages). The INTA Assertion Disable bit, PCICMD1[10], has no effect on this bit. Note that INTA emulation interrupts received across the link are not reflected in this bit.	0h	RO_V
2:0	RSVD	Reserved.	0h	RO

4.2.5 RID—Revision Identification

This register contains the revision number of Device #1.
These bits are read only and writes to this register have no effect.

B/D/F/Type: 0/1/1/CFG			Access: RO	
Size: 8	Default Value: 00h		Address Offset: 8h	
Bit Range	Acronym	Description	Default	Access
7:4	RID_MSB	Revision Identification Number MSB: This is an 8-bit value that indicates the revision identification number for the root port.	0h	RO
3:0	RID	Revision Identification Number: This is an 8-bit value that indicates the revision identification number for the root port.	0h	RO

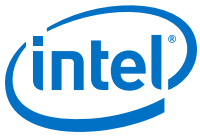

4.2.6 CC—Class Code

This register identifies the basic function of the device, a more specific sub-class, and a register- specific programming interface.

B/D/F/Type: 0/1/1/CFG			Access: RO	
Size: 24	Default Value: 060400h		Address Offset: 9h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
23:16	BCC	Base Class Code: Indicates the base class code for this device. This code has the value 06h, indicating a Bridge device.	06h	RO
15:8	SUBCC	Sub-Class Code: Indicates the sub-class code for this device. The code is 04h indicating a PCI to PCI Bridge.	04h	RO
7:0	PI	Programming Interface: Indicates the programming interface of this device. This value does not specify a particular register set layout and provides no practical use for this device.	00h	RO

4.2.7 CL—Cache Line Size

B/D/F/Type: 0/1/1/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: Ch	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:0	CLS	Cache Line Size: Implemented by PCI Express devices as a read-write field for legacy compatibility purposes but has no impact on any PCI Express device functionality.	00h	RW

4.2.8 HDR—Header Type

This register identifies the header layout of the configuration space. No physical register exists at this location.

B/D/F/Type: 0/1/1/CFG			Access: RO	
Size: 8	Default Value: 81h		Address Offset: Eh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:0	HDR	Header Type Register: Device #1 returns 81 to indicate that this is a multi function device with bridge header layout. Device #6 returns 01 to indicate that this is a single function device with bridge header layout.	81h	RO

4.2.9 PBUSN—Primary Bus Number

This register identifies that this "virtual" Host-PCI Express bridge is connected to PCI bus #0.

B/D/F/Type: 0/1/1/CFG			Access: RO	
Size: 8	Default Value: 00h		Address Offset: 18h	
Bit Range	Acronym	Description	Default	Access
7:0	BUSN	Primary Bus Number: Configuration software typically programs this field with the number of the bus on the primary side of the bridge. Since the processor root port is an internal device and its primary bus is always 0, these bits are read only and are hardwired to 0.	00h	RO

4.2.10 SBUSN—Secondary Bus Number

This register identifies the bus number assigned to the second bus side of the "virtual" bridge i.e. to PCI Express. This number is programmed by the PCI configuration software to allow mapping of configuration cycles to PCI Express.

B/D/F/Type: 0/1/1/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: 19h	
Bit Range	Acronym	Description	Default	Access
7:0	BUSN	Secondary Bus Number: This field is programmed by configuration software with the bus number assigned to PCI Express.	00h	RW

4.2.11 SUBUSN—Subordinate Bus Number

This register identifies the subordinate bus (if any) that resides at the level below PCI Express. This number is programmed by the PCI configuration software to allow mapping of configuration cycles to PCI Express.

B/D/F/Type: 0/1/1/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: 1Ah	
Bit Range	Acronym	Description	Default	Access
7:0	BUSN	Subordinate Bus Number: This register is programmed by configuration software with the number of the highest subordinate bus that lies behind the Processor root port bridge. When only a single PCI device resides on the PCI Express segment, this register will contain the same value as the SBUSN1 register.	00h	RW

4.2.12 IOBASE—I/O Base Address

This register controls the processor to PCI Express-G I/O access routing based on the following formula:

IO_BASE=< address =<IO_LIMIT

Only upper 4 bits are programmable. For the purpose of address decode address bits A[11:0] are treated as 0. Thus the bottom of the defined I/O address range will be aligned to a 4KB boundary.

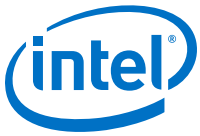

B/D/F/Type: 0/1/1/CFG			Access: RW	
Size: 8	Default Value: F0h		Address Offset: 1Ch	
Bit Range	Acronym	Description	Default	Access
7:4	IOBASE	I/O Address Base: Corresponds to A[15:12] of the I/O addresses passed by the root port to PCI Express-G.	Fh	RW
3:0	RSVD	Reserved.	0h	RO

4.2.13 IOLIMIT—I/O Limit Address

This register controls the processor to PCI Express-G I/O access routing based on the following formula:

IO_BASE=< address =<IO_LIMIT

Only upper 4 bits are programmable. For the purpose of address decode address bits A[11:0] are assumed to be FFFh. Thus, the top of the defined I/O address range will be at the top of a 4KB aligned address block.

B/D/F/Type: 0/1/1/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: 1Dh	
Bit Range	Acronym	Description	Default	Access
7:4	IOLIMIT	I/O Address Limit: Corresponds to A[15:12] of the I/O address limit of the root port. Devices between this upper limit and IOBASE1 will be passed to the PCI Express hierarchy associated with this device.	0h	RW
3:0	RSVD	Reserved.	0h	RO

4.2.14 SSTS—Secondary Status

SSTS is a 16-bit status register that reports the occurrence of error conditions associated with secondary side (i.e. PCI Express-G side) of the "virtual" PCI-PCI bridge embedded within the Processor.

B/D/F/Type: 0/1/1/CFG			Access: RO; RW1C	
Size: 16	Default Value: 0000h		Address Offset: 1Eh	
Bit Range	Acronym	Description	Default	Access
15	DPE	Detected Parity Error: This bit is set by the Secondary Side for a Type 1 Configuration Space header device whenever it receives a Poisoned TLP, regardless of the state of the Parity Error Response Enable bit in the Bridge Control Register.	0h	RW1C
14	RSE	Received System Error: This bit is set when the Secondary Side for a Type 1 configuration space header device receives an ERR_FATAL or ERR_NONFATAL.	0h	RW1C
13	RMA	Received Master Abort: This bit is set when the Secondary Side for Type 1 Configuration Space Header Device (for requests initiated by the Type 1 Header Device itself) receives a Completion with Unsupported Request Completion Status.	0h	RW1C

continued...

B/D/F/Type: 0/1/1/CFG			Access: RO; RW1C	
Size: 16	Default Value: 0000h		Address Offset: 1Eh	
Bit Range	Acronym	Description	Default	Access
12	RTA	Received Target Abort: This bit is set when the Secondary Side for Type 1 Configuration Space Header Device (for requests initiated by the Type 1 Header Device itself) receives a Completion with Completer Abort Completion Status.	0h	RW1C
11	STA	Signaled Target Abort: Not Applicable or Implemented. Hardwired to 0. The processor does not generate Target Aborts (The root port will never complete a request using the Completer Abort Completion status). UR detected inside the processor (such as in MC) will be reported in primary side status.	0h	RO
10:9	DEVT	DEVSELB Timing: Not Applicable or Implemented. Hardwired to 0.	0h	RO
8	SMDPE	Master Data Parity Error: When set indicates that the processor received across the link (upstream) a Read Data Completion Poisoned TLP (EP=1). This bit can only be set when the Parity Error Enable bit in the Bridge Control register is set.	0h	RW1C
7	FB2B	Fast Back-to-Back: Not Applicable or Implemented. Hardwired to 0.	0h	RO
6	RSVD	Reserved.	0h	RO
5	CAP66	66/60 MHz capability: Not Applicable or Implemented. Hardwired to 0.	0h	RO
4:0	RSVD	Reserved.	00h	RO

4.2.15 MBASE—Memory Base Address

This register controls the processor to PCI Express non-prefetchable memory access routing based on the following formula:

MEMORY_BASE=< address =<MEMORY_LIMIT

The upper 12 bits of the register are read/write and correspond to the upper 12 address bits A[31:20] of the 32 bit address. The bottom 4 bits of this register are read-only and return zeroes when read. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be 0. Thus, the bottom of the defined memory address range will be aligned to a 1MB boundary.

B/D/F/Type: 0/1/1/CFG			Access: RW	
Size: 16	Default Value: FFF0h		Address Offset: 20h	
Bit Range	Acronym	Description	Default	Access
15:4	MBASE	Memory Address Base: Corresponds to A[31:20] of the lower limit of the memory range that will be passed to PCI Express.	FFFh	RW
3:0	RSVD	Reserved.	0h	RO

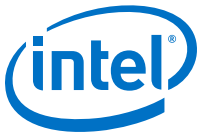

4.2.16 MLIMIT—Memory Limit Address

This register controls the processor to PCI Express non-prefetchable memory access routing based on the following formula:
MEMORY_BASE=< address =<MEMORY_LIMIT
The upper 12 bits of the register are read/write and correspond to the upper 12 address bits A[31:20] of the 32 bit address. The bottom 4 bits of this register are read-only and return zeroes when read. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be FFFFFh. Thus, the top of the defined memory address range will be at the top of a 1MB aligned memory block. NOTE: Memory range covered by MBASE and MLIMIT registers are used to map non-prefetchable PCI Express address ranges (typically where control/status memory-mapped I/O data structures of the graphics controller will reside) and PMBASE and PMLIMIT are used to map prefetchable address ranges (typically graphics local memory). This segregation allows application of USWC space attribute to be performed in a true plug-and-play manner to the prefetchable address range for improved Processor-PCI Express memory access performance.
Note also that configuration software is responsible for programming all address range registers (prefetchable, non-prefetchable) with the values that provide exclusive address ranges i.e. prevent overlap with each other and/or with the ranges covered with the main memory. There is no provision in the processor hardware to enforce prevention of overlap and operations of the system in the case of overlap are not guaranteed.

B/D/F/Type: 0/1/1/CFG			Access: RW	
Size: 16	Default Value: 0000h		Address Offset: 22h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:4	MLIMIT	Memory Address Limit: Corresponds to A[31:20] of the upper limit of the address range passed to PCI Express.	000h	RW
3:0	RSVD	Reserved.	0h	RO

4.2.17 PMBASE—Prefetchable Memory Base Address

This register in conjunction with the corresponding Upper Base Address register controls the processor to PCI Express prefetchable memory access routing based on the following formula:
PREFETCHABLE_MEMORY_BASE =< address =< PREFETCHABLE_MEMORY_LIMIT
The upper 12 bits of this register are read/write and correspond to address bits A[31:20] of the 40-bit address. The lower 8 bits of the Upper Base Address register are read/write and correspond to address bits A[39:32] of the 40-bit address. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be 0. Thus, the bottom of the defined memory address range will be aligned to a 1MB boundary.

B/D/F/Type: 0/1/1/CFG			Access: RO; RW	
Size: 16	Default Value: FFF1h		Address Offset: 24h	
Bit Range	Acronym	Description	Default	Access
15:4	PMBASE	Prefetchable Memory Base Address: Corresponds to A[31:20] of the lower limit of the memory range that will be passed to PCI Express.	FFFh	RW
3:0	AS64	64-bit Address Support: Indicates that the upper 32 bits of the prefetchable memory region base address are contained in the Prefetchable Memory base Upper Address register at 28h.	1h	RO

4.2.18 PMLIMIT—Prefetchable Memory Limit Address

This register in conjunction with the corresponding Upper Limit Address register controls the processor to PCI Express prefetchable memory access routing based on the following formula:

PREFETCHABLE_MEMORY_BASE =< address =< PREFETCHABLE_MEMORY_LIMIT

The upper 12 bits of this register are read/write and correspond to address bits A[31:20] of the 40-bit address. The lower 8 bits of the Upper Limit Address register are read/write and correspond to address bits A[39:32] of the 40-bit address. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be FFFFFh. Thus, the top of the defined memory address range will be at the top of a 1MB aligned memory block. Note that prefetchable memory range is supported to allow segregation by the configuration software between the memory ranges that must be defined as UC and the ones that can be designated as a USWC (i.e. prefetchable) from the processor perspective.

B/D/F/Type: 0/1/1/CFG			Access: RO; RW	
Size: 16	Default Value: 0001h		Address Offset: 26h	
Bit Range	Acronym	Description	Default	Access
15:4	PMLIMIT	Prefetchable Memory Address Limit: Corresponds to A[31:20] of the upper limit of the address range passed to PCI Express.	000h	RW
3:0	AS64B	64-bit Address Support: Indicates that the upper 32 bits of the prefetchable memory region limit address are contained in the Prefetchable Memory Base Limit Address register at 2Ch	1h	RO

4.2.19 PMBASEU—Prefetchable Memory Base Address Upper

The functionality associated with this register is present in the PEG design implementation.

This register in conjunction with the corresponding Upper Base Address register controls the processor to PCI Express prefetchable memory access routing based on the following formula:

PREFETCHABLE_MEMORY_BASE =< address =< PREFETCHABLE_MEMORY_LIMIT

The upper 12 bits of this register are read/write and correspond to address bits A[31:20] of the 39-bit address. The lower 7 bits of the Upper Base Address register are read/write and correspond to address bits A[38:32] of the 39-bit address. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be 0. Thus, the bottom of the defined memory address range will be aligned to a 1MB boundary.

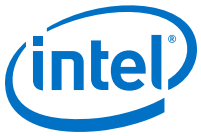

B/D/F/Type: 0/1/1/CFG			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 28h	
Bit Range	Acronym	Description	Default	Access
31:0	PMBASEU	Prefetchable Memory Base Address: Corresponds to A[63:32] of the lower limit of the prefetchable memory range that will be passed to PCI Express.	00000000h	RW

4.2.20 PMLIMITU—Prefetchable Memory Limit Address Upper

The functionality associated with this register is present in the PEG design implementation.

This register in conjunction with the corresponding Upper Limit Address register controls the processor to PCI Express prefetchable memory access routing based on the following formula:

PREFETCHABLE_MEMORY_BASE =< address =< PREFETCHABLE_MEMORY_LIMIT

The upper 12 bits of this register are read/write and correspond to address bits A[31:20] of the 39-bit address. The lower 7 bits of the Upper Limit Address register are read/write and correspond to address bits A[39:32] of the 39-bit address. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be FFFFFh. Thus, the top of the defined memory address range will be at the top of a 1MB aligned memory block.

Note that prefetchable memory range is supported to allow segregation by the configuration software between the memory ranges that must be defined as UC and the ones that can be designated as a USWC (i.e. prefetchable) from the processor perspective.

B/D/F/Type: 0/1/1/CFG			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 2Ch	
Bit Range	Acronym	Description	Default	Access
31:0	PMLIMITU	Prefetchable Memory Address Limit: Corresponds to A[63:32] of the upper limit of the prefetchable Memory range that will be passed to PCI Express.	00000000h	RW

4.2.21 CAPPTR—Capabilities Pointer

The capabilities pointer provides the address offset to the location of the first entry in this device's linked list of capabilities.

B/D/F/Type: 0/1/1/CFG			Access: RO	
Size: 8	Default Value: 88h		Address Offset: 34h	
Bit Range	Acronym	Description	Default	Access
7:0	CAPPTR1	First Capability: The first capability in the list is the Subsystem ID and Subsystem Vendor ID Capability.	88h	RO

4.2.22 INTRLINE—Interrupt Line

This register contains interrupt line routing information. The device itself does not use this value, rather it is used by device drivers and operating systems to determine priority and vector information.

B/D/F/Type: 0/1/1/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: 3Ch	
Bit Range	Acronym	Description	Default	Access
7:0	INTCON	Interrupt Connection: Used to communicate interrupt line routing information. BIOS Requirement: POST software writes the routing information into this register as it initializes and configures the system. The value indicates to which input of the system interrupt controller this device's interrupt pin is connected.	00h	RW

4.2.23 INTRPIN—Interrupt Pin

This register specifies which interrupt pin this device uses.

B/D/F/Type: 0/1/1/CFG			Access: RW_O; RO	
Size: 8	Default Value: 01h		Address Offset: 3Dh	
Bit Range	Acronym	Description	Default	Access
7:3	INTPINH	Interrupt Pin High:	00h	RO
2:0	INTPIN	Interrupt Pin: As a multifunction device, the PCI Express device may specify any INTx (x=A,B,C,D) as its interrupt pin. The Interrupt Pin register tells which interrupt pin the device (or device function) uses. A value of 1 corresponds to INTA# (Default) A value of 2 corresponds to INTB# A value of 3 corresponds to INTC# A value of 4 corresponds to INTD# Devices (or device functions) that do not use an interrupt pin must put a 0 in this register. The values 05h through FFh are reserved. This register is write once. BIOS must set this register to select the INTx to be used by this root port.	1h	RW_O

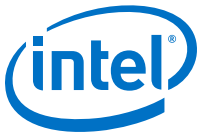

4.2.24 BCTRL—Bridge Control

This register provides extensions to the PCICMD register that are specific to PCI-PCI bridges. The BCTRL provides additional control for the secondary interface (i.e. PCI Express) as well as some bits that affect the overall behavior of the "virtual" Host-PCI Express bridge embedded within the processor, e.g. VGA compatible address ranges mapping.

B/D/F/Type: 0/1/1/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: 3Eh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:12	RSVD	Reserved.	0h	RO
11	DTSERRE	Discard Timer SERR# Enable: Not Applicable or Implemented. Hardwired to 0.	0h	RO
10	DTSTS	Discard Timer Status: Not Applicable or Implemented. Hardwired to 0.	0h	RO
9	SDT	Secondary Discard Timer: Not Applicable or Implemented. Hardwired to 0.	0h	RO
8	PDT	Primary Discard Timer: Not Applicable or Implemented. Hardwired to 0.	0h	RO
7	FB2BEN	Fast Back-to-Back Enable: Not Applicable or Implemented. Hardwired to 0.	0h	RO
6	SRESET	Secondary Bus Reset: Setting this bit triggers a hot reset on the corresponding PCI Express Port. This will force the LTSSM to transition to the Hot Reset state (via Recovery) from L0, L0s, or L1 states.	0h	RW
5	MAMODE	Master Abort Mode: Does not apply to PCI Express. Hardwired to 0.	0h	RO
4	VGA16D	VGA 16-bit Decode: Enables the PCI-to-PCI bridge to provide 16-bit decoding of VGA I/O address precluding the decoding of alias addresses every 1 KB. This bit only has meaning if bit 3 (VGA Enable) of this register is also set to 1, enabling VGA I/O decoding and forwarding by the bridge. 0: Execute 10-bit address decodes on VGA I/O accesses. 1: Execute 16-bit address decodes on VGA I/O accesses.	0h	RW
3	VGAEN	VGA Enable: Controls the routing of processor initiated transactions targeting VGA compatible I/O and memory address ranges. See the VGAEN/MDAP table in device 0, offset 97h[0].	0h	RW
2	ISAEN	ISA Enable: Needed to exclude legacy resource decode to route ISA resources to legacy decode path. Modifies the response by the root port to an I/O access issued by the processor that target ISA I/O addresses. This applies only to I/O addresses that are enabled by the IOBASE and IOLIMIT registers. 0: All addresses defined by the IOBASE and IOLIMIT for processor I/O transactions will be mapped to PCI Express. 1: The root port will not forward to PCI Express any	0h	RW

continued...

B/D/F/Type: 0/1/1/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: 3Eh	
Bit Range	Acronym	Description	Default	Access
		I/O transactions addressing the last 768 bytes in each 1KB block even if the addresses are within the range defined by the IOBASE and IOLIMIT registers.		
1	SERREN	SERR Enable: 0: No forwarding of error messages from secondary side to primary side that could result in an SERR. 1: ERR_COR, ERR_NONFATAL, and ERR_FATAL messages result in SERR message when individually enabled by the Root Control register.	0h	RW
0	PEREN	Parity Error Response Enable: Controls whether or not the Master Data Parity Error bit in the Secondary Status register is set when the root port receives across the link (upstream) a Read Data Completion Poisoned TLP 0: Master Data Parity Error bit in Secondary Status register can NOT be set. 1: Master Data Parity Error bit in Secondary Status register CAN be set.	0h	RW

4.2.25 PM—Power Management Capabilities

B/D/F/Type: 0/1/1/CFG			Access: RO; RO_V	
Size: 32	Default Value: C8039001h		Address Offset: 80h	
Bit Range	Acronym	Description	Default	Access
31:27	PMES	PME Support: This field indicates the power states in which this device may indicate PME wake via PCI Express messaging. D0, D3hot & D3cold. This device is not required to do anything to support D3hot & D3cold, it simply must report that those states are supported. Refer to the PCI Power Management 1.1 specification for encoding explanation and other power management details.	19h	RO
26	D2PSS	D2 Power State Support: Hardwired to 0 to indicate that the D2 power management state is NOT supported.	0h	RO
25	D1PSS	D1 Power State Support: Hardwired to 0 to indicate that the D1 power management state is NOT supported.	0h	RO
24:22	AUXC	Auxiliary Current: Hardwired to 0 to indicate that there are no 3.3Vaux auxiliary current requirements.	0h	RO
21	DSI	Device Specific Initialization: Hardwired to 0 to indicate that special initialization of this device is NOT required before generic class device driver is to use it.	0h	RO
20	APS	Auxiliary Power Source: Hardwired to 0.	0h	RO
19	PMECLK	PME Clock: Hardwired to 0 to indicate this device does NOT support PME# generation.	0h	RO
18:16	PCIPMCV	PCI PM CAP Version: Version - A value of 011b indicates that this function complies with revision 1.2 of the PCI Power Management Interface Specification. --Was Previously Hardwired to 02h to indicate there are 4 bytes of power	3h	RO

continued...

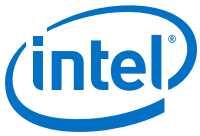

B/D/F/Type: 0/1/1/CFG			Access: RO; RO_V	
Size: 32	**Default Value: C8039001h**		**Address Offset: 80h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		management registers implemented and that this device complies with revision 1.1 of the PCI Power Management Interface Specification.		
15:8	PNC	Pointer to Next Capability: This contains a pointer to the next item in the capabilities list. If MSICH (CAPL[0] @ 7Fh) is 0, then the next item in the capabilities list is the Message Signaled Interrupts (MSI) capability at 90h. If MSICH (CAPL[0] @ 7Fh) is 1, then the next item in the capabilities list is the PCI Express capability at A0h.	90h	RO_V
7:0	CID	Capability ID: Value of 01h identifies this linked list item (capability structure) as being for PCI Power Management registers.	01h	RO

4.2.26 PM—Power Management Control/Status

B/D/F/Type: 0/1/1/CFG			Access: RW; RO	
Size: 32	**Default Value: 00000008h**		**Address Offset: 84h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15	PMESTS	PME Status: Indicates that this device does not support PMEB generation from D3cold.	0h	RO
14:13	DSCALE	Data Scale: Indicates that this device does not support the power management data register.	0h	RO
12:9	DSEL	Data Select: Indicates that this device does not support the power management data register.	0h	RO
8	PMEE	PME Enable: Indicates that this device does not generate PMEB assertion from any D-state. 0: PMEB generation not possible from any D State 1: PMEB generation enabled from any D State The setting of this bit has no effect on hardware. See PM_CAP[15:11]	0h	RW
7:4	RSVD	Reserved.	0h	RO
3	NSR	No Soft Reset: No Soft Reset. When set to 1 this bit indicates that the device is transitioning from D3hot to D0 because the power state commands do not perform a internal reset. Config context is preserved. Upon transition no additonal operating sys intervention is required to preserve configuration context beyond writing the power state bits. When clear the devices do not perform an internal reset upon transitioning from D3hot to D0 via software control of the power state bits. Regardless of this bit the devices that transition from a	1h	RO

continued...

B/D/F/Type: 0/1/1/CFG			Access: RW; RO	
Size: 32	Default Value: 00000008h		Address Offset: 84h	
Bit Range	Acronym	Description	Default	Access
		D3hot to D0 by a system or bus segment reset will return to the device state D0 uniintialized with only PME context preserved if PME is supported and enabled.		
2	RSVD	Reserved.	0h	RO
1:0	PS	Power State: Indicates the current power state of this device and can be used to set the device into a new power state. If software attempts to write an unsupported state to this field, write operation must complete normally on the bus, but the data is discarded and no state change occurs. 00: D0 01: D1 (Not supported in this device.) 10: D2 (Not supported in this device.) 11: D3 Support of D3cold does not require any special action. While in the D3hot state, this device can only act as the target of PCI configuration transactions (for power management control). This device also cannot generate interrupts or respond to MMR cycles in the D3 state. The device must return to the D0 state in order to be fully-functional. When the Power State is other than D0, the bridge will Master Abort (i.e. not claim) any downstream cycles (with exception of type 0 config cycles). Consequently, these unclaimed cycles will go down DMI and come back up as Unsupported Requests, which the processor logs as Master Aborts in Device 0 PCISTS[13] There is no additional hardware functionality required to support these Power States.	0h	RW

4.2.27 SS—Subsystem ID and Vendor ID Capabilities

This capability is used to uniquely identify the subsystem where the PCI device resides. Because this device is an integrated part of the system and not an add-in device, it is anticipated that this capability will never be used. However, it is necessary because Microsoft will test for its presence.

B/D/F/Type: 0/1/1/CFG			Access: RO	
Size: 32	Default Value: 0000800Dh		Address Offset: 88h	
Bit Range	Acronym	Description	Default	Access
31:16	RSVD	Reserved.	0000h	RO
15:8	PNC	Pointer to Next Capability: This contains a pointer to the next item in the capabilities list which is the PCI Power Management capability.	80h	RO
7:0	CID	Capability ID: Value of 0Dh identifies this linked list item (capability structure) as being for SSID/SSVID registers in a PCI-to-PCI Bridge.	0Dh	RO

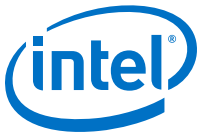

4.2.28 SS—Subsystem ID and Subsystem Vendor ID

System BIOS can be used as the mechanism for loading the SSID/SVID values. These values must be preserved through power management transitions and a hardware reset.

B/D/F/Type: 0/1/1/CFG			Access: RW_O	
Size: 32	Default Value: 00008086h		Address Offset: 8Ch	
Bit Range	Acronym	Description	Default	Access
31:16	SSID	Subsystem ID: Identifies the particular subsystem and is assigned by the vendor.	0000h	RW_O
15:0	SSVID	Subsystem Vendor ID: Identifies the manufacturer of the subsystem and is the same as the vendor ID which is assigned by the PCI Special Interest Group.	8086h	RW_O

4.2.29 MSI—Message Signaled Interrupts Capability ID

When a device supports MSI it can generate an interrupt request to the processor by writing a predefined data item (a message) to a predefined memory address.

The reporting of the existence of this capability can be disabled by setting MSICH (CAPL[0] @ 7Fh). In that case walking this linked list will skip this capability and instead go directly from the PCI PM capability to the PCI Express capability.

B/D/F/Type: 0/1/1/CFG			Access: RO	
Size: 16	Default Value: A005h		Address Offset: 90h	
Bit Range	Acronym	Description	Default	Access
15:8	PNC	Pointer to Next Capability: This contains a pointer to the next item in the capabilities list which is the PCI Express capability.	A0h	RO
7:0	CID	Capability ID: Value of 05h identifies this linked list item (capability structure) as being for MSI registers.	05h	RO

4.2.30 MC—Message Control

System software can modify bits in this register, but the device is prohibited from doing so.

If the device writes the same message multiple times, only one of those messages is guaranteed to be serviced. If all of them must be serviced, the device must not generate the same message again until the driver services the earlier one.

B/D/F/Type: 0/1/1/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: 92h	
Bit Range	Acronym	Description	Default	Access
15:8	RSVD	Reserved.	00h	RO
7	B64AC	64-bit Address Capable: Hardwired to 0 to indicate that the function does not implement the upper 32 bits of the Message Address register and is incapable of generating a 64-bit memory address. This may need to change in future implementations when addressable system memory exceeds the 32b/4GB limit.	0h	RO

continued...

B/D/F/Type: 0/1/1/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: 92h	
Bit Range	Acronym	Description	Default	Access
6:4	MME	Multiple Message Enable: System software programs this field to indicate the actual number of messages allocated to this device. This number will be equal to or less than the number actually requested. The encoding is the same as for the MMC field below.	0h	RW
3:1	MMC	Multiple Message Capable: System software reads this field to determine the number of messages being requested by this device. Value: Number of Messages Requested 000: 1 All of the following are reserved in this implementation: 001: 2 010: 4 011: 8 100: 16 101: 32 110: Reserved 111: Reserved	0h	RO
0	MSIEN	MSI Enable: Controls the ability of this device to generate MSIs. 0: MSI will not be generated. 1: MSI will be generated when we receive PME messages. INTA will not be generated and INTA Status (PCISTS1[3]) will not be set.	0h	RW

4.2.31 MA—Message Address

B/D/F/Type: 0/1/1/CFG			Access: RO; RW	
Size: 32	Default Value: 00000000h		Address Offset: 94h	
Bit Range	Acronym	Description	Default	Access
31:2	MA	Message Address: Used by system software to assign an MSI address to the device. The device handles an MSI by writing the padded contents of the MD register to this address.	00000000h	RW
1:0	FDWA	Force DWord Align: Hardwired to 0 so that addresses assigned by system software are always aligned on a dword address boundary.	0h	RO

4.2.32 MD—Message Data

B/D/F/Type: 0/1/1/CFG			Access: RW	
Size: 16	Default Value: 0000h		Address Offset: 98h	
Bit Range	Acronym	Description	Default	Access
15:0	MD	Message Data: Base message data pattern assigned by system software and used to handle an MSI from the device. When the device must generate an interrupt request, it writes a 32-bit value to the memory address specified in the MA register. The upper 16 bits are always set to 0. The lower 16 bits are supplied by this register.	0000h	RW

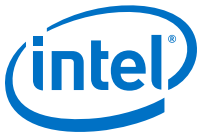

4.2.33 PEG—PCI Express Capability List

Enumerates the PCI Express capability structure.

B/D/F/Type: 0/1/1/CFG			Access: RO	
Size: 16	Default Value: 0010h		Address Offset: A0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:8	PNC	Pointer to Next Capability: This value terminates the capabilities list. The Virtual Channel capability and any other PCI Express specific capabilities that are reported via this mechanism are in a separate capabilities list located entirely within PCI Express Extended Configuration Space.	00h	RO
7:0	CID	Capability ID: Identifies this linked list item (capability structure) as being for PCI Express registers.	10h	RO

4.2.34 PEG—PCI Express Capabilities

Indicates PCI Express device capabilities.

B/D/F/Type: 0/1/1/CFG			Access: RO; RW_O	
Size: 16	Default Value: 0142h		Address Offset: A2h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:14	RSVD	Reserved.	0h	RO
13:9	IMN	Interrupt Message Number: Not Applicable or Implemented. Hardwired to 0.	00h	RO
8	SI	Slot Implemented: 0: The PCI Express Link associated with this port is connected to an integrated component or is disabled. 1: The PCI Express Link associated with this port is connected to a slot. BIOS Requirement: This field must be initialized appropriately if a slot connection is not implemented.	1h	RW_O
7:4	DPT	Device/Port Type: Hardwired to 4h to indicate root port of PCI Express Root Complex.	4h	RO
3:0	PCIECV	PCI Express Capability Version: PCI Express Capability Version (PCIECV): Hardwired to 2h to indicate compliance to the PCI Express Capabilities Register Expansion ECN.	2h	RO

4.2.35 DCAP—Device Capabilities

Indicates PCI Express device capabilities.

B/D/F/Type: 0/1/1/CFG			Access: RW_O; RO	
Size: 32	Default Value: 00008000h		Address Offset: A4h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15	RBER	Role Based Error Reporting (RBER): Indicates that this device implements the functionality defined in the Error Reporting ECN as required by the PCI Express 1.1 spec.	1h	RO
14:6	RSVD	Reserved.	000h	RO

continued...

B/D/F/Type: 0/1/1/CFG			Access: RW_O; RO	
Size: 32	Default Value: 00008000h		Address Offset: A4h	
Bit Range	Acronym	Description	Default	Access
5	ETFS	Extended Tag Field Supported: Hardwired to indicate support for 5-bit Tags as a Requestor.	0h	RO
4:3	PFS	Phantom Functions Supported: Not Applicable or Implemented. Hardwired to 0.	0h	RO
2:0	MPS	Max Payload Size: Default indicates 256B max supported payload for Transaction Layer Packets (TLP) for x16 PEG only. x8 and x4 PEG are limited to 128B support.	0h	RW_O

4.2.36 DCTL—Device Control

Provides control for PCI Express device specific capabilities.
The error reporting enable bits are in reference to errors detected by this device, not error messages received across the link. The reporting of error messages (ERR_CORR, ERR_NONFATAL, ERR_FATAL) received by Root Port is controlled exclusively by Root Port Command Register.

B/D/F/Type: 0/1/1/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: A8h	
Bit Range	Acronym	Description	Default	Access
15	RSVD	Reserved.	0h	RO
14:12	MRRS	Reserved for Max Read Request Size:	0h	RO
11	NSE	Reserved for Enable No Snoop:	0h	RO
10:8	RSVD	Reserved.	0h	RO
7:5	MPS	000b 128B max supported payload for Transaction Layer Packets (TLP). 001b 256B max supported payload for Transaction Layer Packets (TLP). All other encodings are reserved. As a receiver, the Device must handle TLPs as large as the set value. As transmitter, the Device must not generate TLPs exceeding the set value. NOTE: A payload size of 256B is only supported for Device 1 Function 0.	0h	RW
4	ROE	Reserved for Enable Relaxed Ordering:	0h	RO
3	URRE	Unsupported Request Reporting Enable: Unsupported Request Reporting Enable (URRE): When set, allows signaling ERR_NONFATAL, ERR_FATAL, or ERR_CORR to the Root Control register when detecting an unmasked Unsupported Request (UR). An ERR_CORR is signaled when an unmasked Advisory Non-Fatal UR is received. An ERR_FATAL or ERR_NONFATAL is sent to the Root Control register when an uncorrectable non-Advisory UR is received with the severity bit set in the Uncorrectable Error Severity register.	0h	RW

continued...

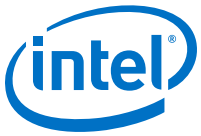

B/D/F/Type: 0/1/1/CFG			Access: RW; RO	
Size: 16	**Default Value: 0000h**		**Address Offset: A8h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
2	FERE	Fatal Error Reporting Enable: Fatal Error Reporting Enable (FERE): When set, enables signaling of ERR_FATAL to the Root Control register due to internally detected errors or error messages received across the link. Other bits also control the full scope of related error reporting.	0h	RW
1	NERE	Non-Fatal Error Reporting Enable: Non-Fatal Error Reporting Enable (NERE): When set, enables signaling of ERR_NONFATAL to the Rool Control register due to internally detected errors or error messages received across the link. Other bits also control the full scope of related error reporting.	0h	RW
0	CERE	Correctable Error Reporting Enable: Correctable Error Reporting Enable (CERE): When set, enables signaling of ERR_CORR to the Root Control register due to internally detected errors or error messages received across the link. Other bits also control the full scope of related error reporting.	0h	RW

4.2.37 DSTS—Device Status

Reflects status corresponding to controls in the Device Control register. The error reporting bits are in reference to errors detected by this device, not errors messages received across the link.

B/D/F/Type: 0/1/1/CFG			Access: RW1C; RO	
Size: 16	**Default Value: 0000h**		**Address Offset: AAh**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:6	RSVD	Reserved.	000h	RO
5	TP	Transactions Pending: 0: All pending transactions (including completions for any outstanding non-posted requests on any used virtual channel) have been completed. 1: Indicates that the device has transaction(s) pending (including completions for any outstanding non-posted requests for all used Traffic Classes). Not Applicable or Implemented. Hardwired to 0.	0h	RO
4	RSVD	Reserved.	0h	RO
3	URD	Unsupported Request Detected: When set this bit indicates that the Device received an Unsupported Request. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control Register. Additionally, the Non-Fatal Error Detected bit or the Fatal Error Detected bit is set according to the setting of the Unsupported Request Error Severity bit. In production systems setting the Fatal Error Detected bit is not an option as support for AER will not be reported.	0h	RW1C
2	FED	Fatal Error Detected: When set this bit indicates that fatal error(s) were detected. Errors are logged in this register regardless of whether error reporting is enabled or not in	0h	RW1C

continued...

B/D/F/Type: 0/1/1/CFG			Access: RW1C; RO	
Size: 16	Default Value: 0000h		Address Offset: AAh	
Bit Range	Acronym	Description	Default	Access
		the Device Control register. When Advanced Error Handling is enabled, errors are logged in this register regardless of the settings of the uncorrectable error mask register.		
1	NFED	Non-Fatal Error Detected: When set this bit indicates that non-fatal error(s) were detected. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. When Advanced Error Handling is enabled, errors are logged in this register regardless of the settings of the uncorrectable error mask register.	0h	RW1C
0	CED	Correctable Error Detected: When set this bit indicates that correctable error(s) were detected. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. When Advanced Error Handling is enabled, errors are logged in this register regardless of the settings of the correctable error mask register.	0h	RW1C

4.2.38 LCTL—Link Control

Allows control of PCI Express link.

B/D/F/Type: 0/1/1/CFG			Access: RW; RO; RW_V	
Size: 16	Default Value: 0000h		Address Offset: B0h	
Bit Range	Acronym	Description	Default	Access
15:12	RSVD	Reserved.	0h	RO
11	LABIE	Link Autonomous Bandwidth Interrupt Enable: Link Autonomous Bandwidth Interrupt Enable - When Set, this bit enables the generation of an interrupt to indicate that the Link Autonomous Bandwidth Status bit has been Set. This bit is not applicable and is reserved for Endpoint devices, PCI Express to PCI/PCI-X bridges, and Upstream Ports of Switches. Devices that do not implement the Link Bandwidth Notification capability must hardwire this bit to 0b.	0h	RW
10	LBMIE	Link Bandwidth Management Interrupt Enable: Link Bandwidth Management Interrupt Enable - When Set, this bit enables the generation of an interrupt to indicate that the Link Bandwidth Management Status bit has been Set. This bit is not applicable and is reserved for Endpoint devices, PCI Express to PCI/PCI-X bridges, and Upstream Ports of Switches.	0h	RW
9	HAWD	Hardware Autonomous Width Disable: Hardware Autonomous Width Disable - When Set, this bit disables hardware from changing the Link width for reasons other than attempting to correct unreliable Link operation by reducing Link width. Devices that do not implement the ability autonomously to change Link width are permitted to hardwire this bit to 0b.	0h	RW

continued...

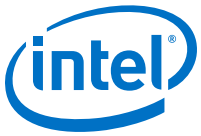

B/D/F/Type: 0/1/1/CFG			Access: RW; RO; RW_V	
Size: 16	**Default Value: 0000h**		**Address Offset: B0h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
8	ECPM	Enable Clock Power Management: Applicable only for form factors that support a "Clock Request" (CLKREQ#) mechanism, this enable functions as follows 0b - Clock power management is disabled and device must hold CLKREQ# signal low 1b - When this bit is set to 1 the device is permitted to use CLKREQ# signal to power manage link clock according to protocol defined in appropriate form factor specification. Default value of this field is 0b. Components that do not support Clock Power Management (as indicated by a 0b value in the Clock Power Management bit of the Link Capabilities Register) must hardwire this bit to 0b.	0h	RO
7	ES	Extended Synch: Extended synch 0: Standard Fast Training Sequence (FTS). 1: Forces the transmission of additional ordered sets when exiting the L0s state and when in the Recovery state. This mode provides external devices (e.g., logic analyzers) monitoring the Link time to achieve bit and symbol lock before the link enters L0 and resumes communication. This is a test mode only and may cause other undesired side effects such as buffer overflows or underruns.	0h	RW
6	CCC	Common Clock Configuration: 0: Indicates that this component and the component at the opposite end of this Link are operating with asynchronous reference clock. 1: Indicates that this component and the component at the opposite end of this Link are operating with a distributed common reference clock. The state of this bit affects the L0s Exit Latency reported in LCAP[14:12] and the N_FTS value advertised during link training. See PEGL0SLAT at offset 22Ch.	0h	RW
5	RL	Retrain Link: 0b Normal operation. 1b Full Link retraining is initiated by directing the Physical Layer LTSSM from L0, L0s, or L1 states to the Recovery state. This bit always returns 0 when read. This bit is cleared automatically (no need to write a 0). Intel Reserved Text: PEG10: This bit is set implicitly if CAPID0_B[SPEGFX1] or CAPID0_B[DPEGFX1] are set to 1b, and in addition BCTRL[VGAEN] is set to 1b. PEG11: This bit is set implicitly if CAPID0_B[DPEGFX1] is set to 1b and in addition BCTRL[VGAEN] is set to 1b. PEG12: This bit is set implicitly if CAPID0_B[DPEGFX1] is set to 1b and in addition BCTRL[VGAEN] is set to 1b.	0h	RW_V
4	LD	Link Disable: 0: Normal operation 1: Link is disabled. Forces the LTSSM to transition to the Disabled state (via Recovery) from L0, L0s, or L1 states. Link retraining happens automatically on 0 to 1 transition, just like when coming out of reset. Writes to this bit are immediately reflected in the value read from the bit, regardless of actual Link state.	0h	RW

continued...

B/D/F/Type: 0/1/1/CFG			Access: RW; RO; RW_V	
Size: 16	Default Value: 0000h		Address Offset: B0h	
Bit Range	Acronym	Description	Default	Access
3	RCB	Read Completion Boundary: Hardwired to 0 to indicate 64 byte.	0h	RO
2	RSVD	Reserved.	0h	RO
1:0	ASPM	Active State PM: Controls the level of active state power management supported on the given link. 00: Disabled 01: L0s Entry Supported 10: L1 Entry Supported 11: L0s and L1 Entry Supported	0h	RW

4.2.39 LSTS—Link Status

Indicates PCI Express link status.

B/D/F/Type: 0/1/1/CFG			Access: RO_V; RO; RW1C	
Size: 16	Default Value: 1001h		Address Offset: B2h	
Bit Range	Acronym	Description	Default	Access
15	LABWS	Link Autonomous Bandwidth Status: This bit is set to 1b by hardware to indicate that hardware has autonomously changed link speed or width, without the port transitioning through DL_Down status, for reasons other than to attempt to correct unreliable link operation. This bit must be set if the Physical Layer reports a speed or width change was initiated by the downstream component that was indicated as an autonomous change.	0h	RW1C
14	LBWMS	Link Bandwidth Management Status: This bit is set to 1b by hardware to indicate that either of the following has occurred without the port transitioning through DL_Down status: A link retraining initiated by a write of 1b to the Retrain Link bit has completed. Note: This bit is Set following any write of 1b to the Retrain Link bit, including when the Link is in the process of retraining for some other reason. Hardware has autonomously changed link speed or width to attempt to correct unreliable link operation, either through an LTSSM timeout or a higher level process This bit must be set if the Physical Layer reports a speed or width change was initiated by the downstream component that was not indicated as an autonomous change.	0h	RW1C
13	DLLLA	Data Link Layer Link Active (Optional): This bit indicates the status of the Data Link Control and Management State Machine. It returns a 1b to indicate the DL_Active state, 0b otherwise. This bit must be implemented if the corresponding Data Link Layer Active Capability bit is implemented. Otherwise, this bit must be hardwired to 0b.	0h	RO_V

continued...

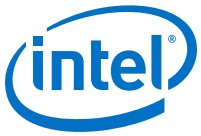

B/D/F/Type: 0/1/1/CFG			Access: RO_V; RO; RW1C	
Size: 16	**Default Value: 1001h**		**Address Offset: B2h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
12	SCC	Slot Clock Configuration: 0: The device uses an independent clock irrespective of the presence of a reference on the connector. 1: The device uses the same physical reference clock that the platform provides on the connector.	1h	RO
11	LTRN	Link Training: Indicates that the Physical Layer LTSSM is in the Configuration or Recovery state, or that 1b was written to the Retrain Link bit but Link training has not yet begun. Hardware clears this bit when the LTSSM exits the Configuration/Recovery state once Link training is complete.	0h	RO_V
10	RSVD	Reserved.	0h	RO
9:4	NLW	Negotiated Link Width: Indicates negotiated link width. This field is valid only when the link is in the L0, L0s, or L1 states (after link width negotiation is successfully completed). 00h: Reserved 01h: X1 02h: X2 04h: X4 08h: X8 10h: X16 All other encodings are reserved.	00h	RO_V
3:0	CLS	Current Link Speed: This field indicates the negotiated Link speed of the given PCI Express Link. The encoding is the binary value of the bit location in the Supported Link Speeds Vector (in the Link Capabilities 2 register) that corresponds to the current Link speed. For example, a value of 0010b in this field indicates that the current Link speed is that corresponding to bit 2 in the Supported Link Speeds Vector, which is 5.0 GT/s. The value in this field is undefined when the Link is not up.	1h	RO_V

4.2.40 SLOTCAP—Slot Capabilities

PCI Express Slot related registers allow for the support of Hot Plug.

B/D/F/Type: 0/1/1/CFG			Access: RO; RW_O	
Size: 32	**Default Value: 00040000h**		**Address Offset: B4h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:19	PSN	Physical Slot Number: Indicates the physical slot number attached to this Port. BIOS Requirement: This field must be initialized by BIOS to a value that assigns a slot number that is globally unique within the chassis.	0000h	RW_O
18	NCCS	No Command Completed Support: When set to 1b, this bit indicates that this slot does not generate software notification when an issued command is completed by the Hot-Plug Controller. This bit is only permitted to be set to 1b if the hotplug capable port is able to accept writes to all fields of the Slot Control register without delay between successive writes.	1h	RO

continued...

B/D/F/Type: 0/1/1/CFG			Access: RO; RW_O	
Size: 32	Default Value: 00040000h		Address Offset: B4h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
17	EIP	Reserved for Electromechanical Interlock Present: When set to 1b, this bit indicates that an Electromechanical Interlock is implemented on the chassis for this slot.	0h	RO
16:15	SPLS	Slot Power Limit Scale: Specifies the scale used for the Slot Power Limit Value. 00: 1.0x 01: 0.1x 10: 0.01x 11: 0.001x If this field is written, the link sends a Set_Slot_Power_Limit message.	0h	RW_O
14:7	SPLV	Slot Power Limit Value: In combination with the Slot Power Limit Scale value, specifies the upper limit on power supplied by slot. Power limit (in Watts) is calculated by multiplying the value in this field by the value in the Slot Power Limit Scale field. If this field is written, the link sends a Set_Slot_Power_Limit message.	00h	RW_O
6	HPC	Reserved for Hot-plug Capable: When set to 1b, this bit indicates thta this slot is capable of supporting hot-plug operations.	0h	RO
5	HPS	Reserved for Hot-plug Surprise: When set to 1b, this bit indicates that an adapter present in this slot might be removed from the system without any prior notification. This is a form factor specific capability. this bit is an indication to the operating system to allow for such removal without impacting continued software operation.	0h	RO
4	PIP	Reserved for Power Indicator Present: When set to 1b, this bit indicates that a Power Indicator is electrically controlled by the chassis for this slot.	0h	RO
3	AIP	Reserved for Attention Indicator Present: When set to 1b, this bit indicates that an Attention Indicator is electrically controlled by the chassis.	0h	RO
2	MSP	Reserved for MRL Sensor Present: When set to 1b, this bit indicates that an MRL Sensor is implemented on the chassis for this slot.	0h	RO
1	PCP	Reserved for Power Controller Present: When set to 1b, this bit indicates that a software programmable Power Controller is implemented for this slot/adapter (depending on form factor).	0h	RO
0	ABP	Reserved for Attention Button Present: When set to 1b, this bit indicates that an Attention Button for this slot is electrically controlled by the chassis.	0h	RO

4.2.41 SLOTCTL—Slot Control

PCI Express Slot related registers allow for the support of Hot Plug.

B/D/F/Type: 0/1/1/CFG			Access: RO	
Size: 16	Default Value: 0000h		Address Offset: B8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:13	RSVD	Reserved.	0h	RO
12	DLLSCE	Reserved for Data Link Layer State Changed Enable: Reserved for Data Link Layer State Changed Enable (DLLSCE): If the Data Link Layer Link Active capability is implemented, when set to 1b, this field enables software notification when Data Link Layer Link Active field is changed. If the Data Link Layer Link Active capability is not implemented, this bit is permitted to be read-only with a value of 0b.	0h	RO
11	EIC	Reserved for Electromechanical Interlock Control: If an Electromechanical Interlock is implemented, a write of 1b to this field causes the state of the interlock to toggle. A write of 0b to this field has no effect. A read to this register always returns a 0.	0h	RO
10	PCC	Reserved for Power Controller Control: If a Power Controller is implemented, this field when written sets the power state of the slot per the defined encodings. Reads of this field must reflect the value from the latest write, even if the corresponding hotplug command is not complete, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined. Depending on the form factor, the power is turned on/off either to the slot or within the adapter. Note that in some cases the power controller may autonomously remove slot power or not respond to a power-up request based on a detected fault condition, independent of the Power Controller Control setting. The defined encodings are: 0b Power On 1b Power Off If the Power Controller Implemented field in the Slot Capabilities register is set to 0b, then writes to this field have no effect and the read value of this field is undefined.	0h	RO
9:8	PIC	Reserved Power Indicator Control: Reserved Power Indicator Control (PIC): If a Power Indicator is implemented, writes to this field set the Power Indicator to the written state. Reads of this field must reflect the value from the latest write, even if the corresponding hot-plug command is not complete, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined. 00: Reserved 01: On 10: Blink 11: Off If the Power Indicator Present bit in the Slot Capabilities register is 0b, this field is permitted to be read-only with a value of 00b.	0h	RO
7:6	AIC	Reserved for Attention Indicator Control: Reserved for Attention Indicator Control (AIC): If an Attention Indicator is implemented, writes to this	0h	RO

continued...

B/D/F/Type: 0/1/1/CFG			Access: RO	
Size: 16	Default Value: 0000h		Address Offset: B8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		field set the Attention Indicator to the written state. Reads of this field must reflect the value from the latest write, even if the corresponding hot-plug command is not complete, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined. If the indicator is electrically controlled by chassis, the indicator is controlled directly by the downstream port through implementation specific mechanisms. 00: Reserved 01: On 10: Blink 11: Off If the Attention Indicator Present bit in the Slot Capabilities register is 0b, this field is permitted to be readonly with a value of 00b.		
5	HPIE	Reserved for Hot-plug Interrupt Enable: When set to 1b, this bit enables generation of an interrupt on enabled hot-plug events Default value of this field is 0b. If the Hot Plug Capable field in the Slot Capabilities register is set to 0b, this bit is permitted to be read-only with a value of 0b.	0h	RO
4	CCI	Reserved for Command Completed Interrupt Enable: If Command Completed notification is supported (as indicated by No Command Completed Support field of Slot Capabilities Register), when set to 1b, this bit enables software notification when a hot-plug command is completed by the Hot-Plug Controller. Default value of this field is 0b. If Command Completed notification is not supported, this bit must be hardwired to 0b.	0h	RO
3	PDCE	Presence Detect Changed Enable: When set to 1b, this bit enables software notification on a presence detect changed event.	0h	RO
2	MSCE	Reserved for MRL Sensor Changed Enable: When set to 1b, this bit enables software notification on a MRL sensor changed event. Default value of this field is 0b. If the MRL Sensor Present field in the Slot Capabilities register is set to 0b, this bit is permitted to be read-only with a value of 0b.	0h	RO
1	PFDE	Reserved for Power Fault Detected Enable: When set to 1b, this bit enables software notification on a power fault event. Default value of this field is 0b. If Power Fault detection is not supported, this bit is permitted to be read-only with a value of 0b	0h	RO
0	ABPE	Reserved for Attention Button Pressed Enable: When set to 1b, this bit enables software notification on an attention button pressed event.	0h	RO

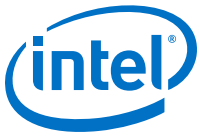

4.2.42 SLOTSTS—Slot Status

PCI Express Slot related registers.

B/D/F/Type: 0/1/1/CFG			Access: RO; RW1C; RO_V	
Size: 16	**Default Value: 0000h**		**Address Offset: BAh**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:9	RSVD	Reserved.	00h	RO
8	DLLSC	Reserved for Data Link Layer State Changed: This bit is set when the value reported in the Data Link Layer Link Active field of the Link Status register is changed. In response to a Data Link Layer State Changed event, software must read the Data Link Layer Link Active field of the Link Status register to determine if the link is active before initiating configuration cycles to the hot plugged device.	0h	RO
7	EIS	Reserved for Electromechanical Interlock Status: If an Electromechanical Interlock is implemented, this bit indicates the current status of the Electromechanical Interlock. Defined encodings are: 0b Electromechanical Interlock Disengaged 1b Electromechanical Interlock Engaged	0h	RO
6	PDS	Presence Detect State: --In band presence detect state: 0b: Slot Empty 1b: Card present in slot This bit indicates the presence of an adapter in the slot, reflected by the logical "OR" of the Physical Layer in-band presence detect mechanism and, if present, any out-of-band presence detect mechanism defined for the slot's corresponding form factor. Note that the in-band presence detect mechanism requires that power be applied to an adapter for its presence to be detected. Consequently, form factors that require a power controller for hot-plug must implement a physical pin presence detect mechanism. Defined encodings are: 0b Slot Empty 1b Card Present in slot This register must be implemented on all Downstream Ports that implement slots. For Downstream Ports not connected to slots (where the Slot Implemented bit of the PCI Express Capabilities Register is 0b), this bit must return 1b.	0h	RO_V
5	MSS	Reserved for MRL Sensor State: This register reports the status of the MRL sensor if it is implemented. Defined encodings are: 0b MRL Closed 1b MRL Open	0h	RO
4	CC	Reserved for Command Completed: If Command Completed notification is supported (as indicated by No Command Completed Support field of Slot Capabilities Register), this bit is set when a hot-plug command has completed and the Hot-Plug Controller is ready to accept a subsequent command. The Command Completed status bit is set as an indication to host software that the Hot-Plug Controller has processed the previous command and is ready to receive the next command; it provides no guarantee that the action corresponding to the command	0h	RO

continued...

B/D/F/Type: 0/1/1/CFG			Access: RO; RW1C; RO_V	
Size: 16	Default Value: 0000h		Address Offset: BAh	
Bit Range	Acronym	Description	Default	Access
		is complete. If Command Completed notification is not supported, this bit must be hardwired to 0b.		
3	PDC	Presence Detect Changed: --A pulse indication that the inband presence detect state has changed This bit is set when the value reported in Presence Detect State is changed.	0h	RW1C
2	MSC	Reserved for MRL Sensor Changed: If an MRL sensor is implemented, this bit is set when a MRL Sensor state change is detected. If an MRL sensor is not implemented, this bit must not be set.	0h	RO
1	PFD	Reserved for Power Fault Detected: If a Power Controller that supports power fault detection is implemented, this bit is set when the Power Controller detects a power fault at this slot. Note that, depending on hardware capability, it is possible that a power fault can be detected at any time, independent of the Power Controller Control setting or the occupancy of the slot. If power fault detection is not supported, this bit must not be set.	0h	RO
0	ABP	Reserved for Attention Button Pressed: If an Attention Button is implemented, this bit is set when the attention button is pressed. If an Attention Button is not supported, this bit must not be set.	0h	RO

4.2.43 RCTL—Root Control

Allows control of PCI Express Root Complex specific parameters. The system error control bits in this register determine if corresponding SERRs are generated when our device detects an error (reported in this device's Device Status register) or when an error message is received across the link. Reporting of SERR as controlled by these bits takes precedence over the SERR Enable in the PCI Command Register.

B/D/F/Type: 0/1/1/CFG			Access: RW; RO; RWS_V; RO_V	
Size: 32	Default Value: 00000000h		Address Offset: BCh	
Bit Range	Acronym	Description	Default	Access
31:19	RSVD	Reserved.	0000h	RO
18	INSWL2	this bit indicates whether the current state of the LTSSM is L2, and it got there through the SWL2 flow. This bit is part of the RTD3 flow ,and may be used as indication for initiate power savings measures.	0h	RO_V
17:9	RSVD	Reserved.	000h	RO
8	SWL2DIS	SWL2 disable bit, when set writing to bits RCTL.L23ER or RCTL.L22DT wont have any affect over H/W	0h	RWS_V
7	RSVD	Reserved.	0h	RO

continued...

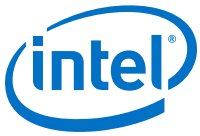

B/D/F/Type: 0/1/1/CFG			Access: RW; RO; RWS_V; RO_V	
Size: 32	**Default Value: 00000000h**		**Address Offset: BCh**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
6	L22DT	L23 to Detect Transition: When set by software, the LTSSM moves from L2 to Detect state. This bit is cleared by hardware once the LTSSM is in detect (and in general, while not in L2 state)	0h	RWS_V
5	L23ER	L23_Rdy Entry Request: When set by software, the corresponding PCIe root port will initiate the sequence to put the link into L2/L3 Ready state. PME_Turn_Off message will be sent and the corresponding PME_TO_Ack response will be returned by the device. Once the link enters L2, this bit will be cleared by hardware.	0h	RWS_V
4	CSVE	Reserved for CRS Software Visibility Enable: This bit, when set, enables the Root Port to return Configuration Request Retry Status (CRS) Completion Status to software. Root Ports that do not implement this capability must hardwire this bit to 0b.	0h	RO
3	PMEIE	PME Interrupt Enable: 0: No interrupts are generated as a result of receiving PME messages. 1: Enables interrupt generation upon receipt of a PME message as reflected in the PME Status bit of the Root Status Register. A PME interrupt is also generated if the PME Status bit of the Root Status Register is set when this bit is set from a cleared state. If the bit change from 1 to 0 and interrupt is pending than interrupt is deasserted	0h	RW
2	SEFEE	System Error on Fatal Error Enable: Controls the Root Complex's response to fatal errors. 0: No SERR generated on receipt of fatal error. 1: Indicates that an SERR should be generated if a fatal error is reported by any of the devices in the hierarchy associated with this Root Port, or by the Root Port itself.	0h	RW
1	SENFUEE	System Error on Non-Fatal Uncorrectable Error Enable: Controls the Root Complex's response to non-fatal errors. 0: No SERR generated on receipt of non-fatal error. 1: Indicates that an SERR should be generated if a non-fatal error is reported by any of the devices in the hierarchy associated with this Root Port, or by the Root Port itself.	0h	RW
0	SECEE	System Error on Correctable Error Enable: Controls the Root Complex's response to correctable errors. 0: No SERR generated on receipt of correctable error. 1: Indicates that an SERR should be generated if a correctable error is reported by any of the devices in the hierarchy associated with this Root Port, or by the Root Port itself.	0h	RW

4.2.44 RSTS—Root Status

Provides information about PCI Express Root Complex specific parameters.

B/D/F/Type: 0/1/1/CFG			Access:	RO_V; RW1C; RO
Size: 32	Default Value: 00000000h		Address Offset:	C0h
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:18	RSVD	Reserved.	0000h	RO
17	PMEP	PME Pending: Indicates that another PME is pending when the PME Status bit is set. When the PME Status bit is cleared by software; the PME is delivered by hardware by setting the PME Status bit again and updating the Requestor ID appropriately. The PME pending bit is cleared by hardware if no more PMEs are pending. Intel Reserved text: In order to simplify design and validation, if a 2nd PME arrives, before the 1st is processed, we will drop it. The initiating device must re-initiate the PME, as described in the spec so there should be no functional impact by taking this approach. This means this bit is hardwired to '0' and will never be set to '1'.	0h	RO
16	PMES	PME Status: Indicates that PME was asserted by the requestor ID indicated in the PME Requestor ID field. Subsequent PMEs are kept pending until the status register is cleared by writing a 1 to this field. An interrupt is asserted If PMEIE is asserted and PMES is changing from 0 to 1 An interrupt is deasserted If PMEIE is asserted and PMES is changing from 1 to 0 An Assert_PMEGPE is sent upstream If PMEGPEE in PEG Legacy cControl register (PEGLC) is asserted and PMES is changing from 0 to 1 An Deassert_PMEGPE is sent upstream If PMEGPEE in PEG Legacy cControl register (PEGLC) is asserted and PMES is changing from 1 to 0 An interrupt is deasserted If PMEIE is asserted and PMES is changing from 1 to 0 Intel Reserved text: In order to simplify design and validation, if a 2nd PME arrives, before the 1st is processed, we will drop it. The initiating device must re-initiate the PME, as described in the spec so there should be no functional impact by taking this approach.	0h	RW1C
15:0	PMERID	PME Requestor ID: Indicates the PCI requestor ID of the last PME requestor.	0000h	RO_V

4.2.45 DCAP2—Device Capabilites 2

B/D/F/Type: 0/1/1/CFG			Access:	RO; RW_O
Size: 32	Default Value: 00000B80h		Address Offset:	C4h
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:20	RSVD	Reserved.	000h	RO
19:18	OBFF_SUPPORTED	OBFF Supported 00b OBFF Not Supported 01b OBFF supported using Message signaling only	0h	RW_O

continued...

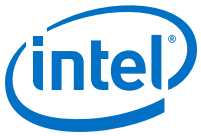

B/D/F/Type: 0/1/1/CFG			Access: RO; RW_O	
Size: 32	**Default Value: 00000B80h**		**Address Offset: C4h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		10b OBFF supported using WAKE# signaling only 11b OBFF supported using WAKE# and Message signaling The value reported in this field must indicate support for WAKE# signaling only if: - for a Downstream Port, driving the WAKE# signal for OBFF is supported and the connector or component connected Downstream is known to receive that same WAKE# signal - for an Upstream Port, receiving the WAKE# signal for OBFF is supported and, if the component is on an add-in-card, that the component is connected to the WAKE# signal on the connector. Root Ports, Switch Ports, and Endpoints are permitted to implement this capability. For a multi-Function device associated with an Upstream Port, each Function must report the same value for this field. For Bridges and Ports that do not implement this capability, this field must be hardwired to 00b.		
17:12	RSVD	Reserved.	00h	RO
11	LTRS	Latency Tolerance and BW reporting Mechanism Supported: A value of 1b indicates support for the optional Latency Tolerance & Bandwidth Requirement Reporting (LTBWR) mechanism capability. Root Ports, Switches and Endpoints are permitted to implement this capability. For Switches that implement LTBWR, this bit must be set only at the upstream port. For a multi-Function device, each Function must report the same value for this bit. For Bridges, Downstream Ports, and components that do not implement this capability, this bit must be hardwired to 0b.	1h	RO
10	RSVD	Reserved.	0h	RO
9	ATOMIC128SUP	128-bit CAS atomic operation completion support. This bit must be set to 1b if the Function supports this optional capability.	1h	RO
8	ATOMIC64SUP	64-bit atomic operation completion support. Includes FetchAdd, Swap, and CAS AtomicOps. This bit must be set to 1b if the Function supports this optional capability.	1h	RO
7	ATOMIC32SUP	32-bit atomic operation completion support. Includes FetchAdd, Swap, and CAS AtomicOps. This bit must be set to 1b if the Function supports this optional capability.	1h	RO
6	ATOMIC_OP_ROUTING_SUPPORT	Atomic Operation Routing Supported. If set then then atomic operations are supported.	0h	RO

continued...

B/D/F/Type: 0/1/1/CFG			Access: RO; RW_O	
Size: 32	Default Value: 00000B80h		Address Offset: C4h	
Bit Range	Acronym	Description	Default	Access
5	ARIFS	ARI Forwarding Supported: Applicable only to Switch Downstream Ports and Root Ports; must be 0b for other Function types. This bit must be set to 1b if a Switch Downstream Port or Root Port supports this optional capability.	0h	RW_O
4	CTODS	Completion Timeout Disabled Supported: A value of 1b indicates support for the Completion Timeout Disable mechanism. The Completion Timeout Disable mechanism is required for Endpoints that issue Requests on their own behalf and PCI Express to PCI/PCI-X Bridges that take ownership of Requests issued on PCI Express. This mechanism is optional for Root Ports. The Root port does nopt support completion timeout disable	0h	RO
3:0	CTOR	Completion Timer Ranges Supported: device Function support for the optional Completion Timeout programmability mechanism. This mechanism allows system software to modify the Completion Timeout value. This field is applicable only to Root Ports, Endpoints that issue Requests on their own behalf, and PCI Express to PCI/PCI-X Bridges that take ownership of Requests issued on PCI Express. For all other Functions this field is reserved and must be hardwired to 0000b. 0000b Completion Timeout programming not supported - the Function must implement a timeout value in the range 50 us to 50 ms.	0h	RO

4.2.46 DCTL2—Device Control 2

B/D/F/Type: 0/1/1/CFG			Access: RW; RO; RW_V	
Size: 16	Default Value: 0000h		Address Offset: C8h	
Bit Range	Acronym	Description	Default	Access
15	RSVD	Reserved.	0h	RO
14:13	OBFFEN		0h	RW
12:11	RSVD	Reserved.	0h	RO
10	LTREN	Latency Tolerance Reporting Mechansim Enable: When Set to 1b, this bit enables the Latency Tolerance & Reporting (LTR) mechanism. This bit is required for all Functions that support the LTR Capability. For a Multi-Function device associated with an upstream port of a device that implements LTBWR, the bit in Function 0 is of type RW, and only Function 0 controls the components Link behavior. In all other Functions of that device, this bit is of type RsvdP. Components that do not implement LTR are permitted to hardwire this bit to 0b. Default value of this bit is 0b.	0h	RW_V

continued...

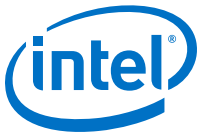

B/D/F/Type: 0/1/1/CFG			Access: RW; RO; RW_V	
Size: 16	Default Value: 0000h		Address Offset: C8h	
Bit Range	Acronym	Description	Default	Access
		This bit is cleared when the port goes to DL_down state. HW ignores the value of this bit.		
9:7	RSVD	Reserved.	0h	RO
6	ATOMIC_OP_REQUESTER_EN	AtomicOp Requester Enable Applicable only to Endpoints and Root Ports; must be hardwired to 0b for other Function types. The Function is allowed to initiate AtomicOp Requests only if this bit and the Bus Master Enable bit in the Command register are both Set. This bit is required to be RW if the Endpoint or Root Port is capable of initiating AtomicOp Requests, but otherwise is permitted to be hardwired to 0b. This bit does not serve as a capability bit. This bit is permitted to be RW even if no AtomicOp Requester capabilities are supported by the Endpoint or Root Port.	0h	RO
5	ARIFEN	ARI Forward Enable: When set, the Downstream Port disables its traditional Device Number field being 0 enforcement when turning a Type 1 Configuration Request into a Type 0 Configuration Request, permitting access to Extended Functions in an ARI Device immediately below the Port. Default value of this bit is 0b. Must be hardwired to 0b if the ARI Forwarding Supported bit is 0b.	0h	RW
4:0	RSVD	Reserved.	00h	RO

4.2.47 LCTL2—Link Control 2

B/D/F/Type: 0/1/1/CFG			Access: RWS; RWS_V	
Size: 16	Default Value: 0003h		Address Offset: D0h	
Bit Range	Acronym	Description	Default	Access
15:12	ComplianceDeemphasis	Compliance De-emphasis: For 8 GT/s Data Rate: This field sets the Transmitter Preset level in Polling.Compliance state if the entry occurred due to the Enter Compliance bit being 1b. This bit sets the de-emphasis level in Polling.Compliance state if the entry occurred due to the Enter Compliance bit being 1b. Defined encodings are: 0001b -3.5 dB 0000b -6 dB When the Link is operating at 2.5 GT/s, the setting of this bit has	0h	RWS

continued...

B/D/F/Type: 0/1/1/CFG			Access: RWS; RWS_V	
Size: 16	Default Value: 0003h		Address Offset: D0h	
Bit Range	Acronym	Description	Default	Access
		no effect. Components that support only 2.5 GT/s speed are permitted to hardwire this bit to 0b. For a Multi-Function device associated with an Upstream Port, the bit in Function 0 is of type RWS, and only Function 0 controls the component's Link behavior. In all other Functions of that device, this bit is of type RsvdP. The default value of this bit is 0000b. This bit is intended for debug, compliance testing purposes. System firmware and software is allowed to modify this bit only during debug or compliance testing.		
11	compsos	Compliance SOS: When set to 1b, the LTSSM is required to send SKP Ordered Sets periodically in between the (modified) compliance patterns. For a Multi-Function device associated with an Upstream Port, the bit in Function 0 is of type RWS, and only Function 0 controls the component's Link behavior. In all other Functions of that device, this bit is of type RsvdP. The default value of this bit is 0b. This bit is applicable when the Link is operating at 2.5 GT/s or 5 GT/s data rates only. Components that support only the 2.5 GT/s speed are permitted to hardwire this field to 0b.	0h	RWS
10	entermodcompliance	Enter Modified Compliance: When this bit is set to 1b, the device transmits modified compliance pattern if the LTSSM enters Polling.Compliance state. Components that support only the 2.5GT/s speed are permitted to hardwire this bit to 0b. Default value of this field is 0b.	0h	RWS
9:7	txmargin	Transmit Margin: This field controls the value of the non-deemphasized voltage level at the Transmitter pins. This field is reset to 000b on entry to the LTSSM Polling.Configuration substate (see Chapter 4 for details of how the transmitter voltage level is determined in various states). Encodings: 000: Normal operating range 001: 800-1200 mV for full swing and 400-700 mV for half-swing 010 - (n-1): Values must be monotonic with a non-zero slope. The value of n must be greater than 3 and less than 7. At least two of these must be below the normal operating range n : 200-400 mV for full-swing and 100-200 mV for half-swing n -111: reserved Default value is 000b. Components that support only the 2.5GT/s speed are	0h	RWS_V

continued...

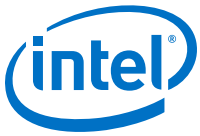

B/D/F/Type: 0/1/1/CFG			Access: RWS; RWS_V	
Size: 16	**Default Value: 0003h**		**Address Offset: D0h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		permitted to hardwire this bit to 0b. When operating in 5GT/s mode with full swing, the deemphasis ratio must be maintained within +/- 1dB from the spec defined operational value (either -3.5 or -6 dB).		
6	selectabledeemphasis	Selectable De-emphasis: When the Link is operating at 5GT/s speed, selects the level of de-emphasis. Encodings: 1b -3.5 dB 0b -6 dB Default value is implementation specific, unless a specific value is required for a selected form factor or platform. When the Link is operating at 2.5GT/s speed, the setting of this bit has no effect. Components that support only the 2.5GT/s speed are permitted to hardwire this bit to 0b.	0h	RWS
5	HASD	Hardware Autonomous Speed Disable: When set to 1b this bit disables hardware from changing the link speed for reasons other than attempting to correct unreliable link operation by reducing link speed.	0h	RWS
4	EC	Enter Compliance: Software is permitted to force a link to enter Compliance mode at the speed indicated in the Target Link Speed field by setting this bit to 1b in both components on a link and then initiating a hot reset on the link.	0h	RWS
3:0	TLS	Target Link Speed: For Downstream Ports, this field sets an upper limit on Link operational speed by restricting the values advertised by the Upstream component in its training sequences. The encoding is the binary value of the bit in the Supported Link Speeds Vector (in the Link Capabilities 2 register) that corresponds to the desired target Link speed. All other encodings are reserved. For example, 5.0 GT/s corresponds to bit 2 in the Supported Link Speeds Vector, so the encoding for a 5.0 GT/s target Link speed in this field is 0010b. If a value is written to this field that does not correspond to a supported speed (as indicated by the Max Link Speed Vector), the result is undefined. The default value of this field is the highest Link speed supported by the component (as reported in the Max Link Speed field of the Link Capabilities register) unless the corresponding platform/form factor requires a different default value. For both Upstream and Downstream Ports, this field is used to set the target compliance mode speed when software is using the Enter Compliance bit to force a Link into compliance mode. For a Multi-Function device associated with an Upstream Port, the field in Function 0 is of type RWS, and only Function 0	3h	RWS

B/D/F/Type: 0/1/1/CFG			Access: RWS; RWS_V	
Size: 16	Default Value: 0003h		Address Offset: D0h	
Bit Range	Acronym	Description	Default	Access
		controls the components Link behavior. In all other Functions of that device, this field is of type RsvdP.		

4.2.48 LSTS2—Link Status 2

B/D/F/Type: 0/1/1/CFG			Access: RO_V; RW1C	
Size: 16	Default Value: 0000h		Address Offset: D2h	
Bit Range	Acronym	Description	Default	Access
15:6	RSVD	Reserved.	000h	RO
5	LNKEQREQ	This bit is Set by hardware to request the Link equalization process to be performed on the Link.	0h	RW1C
4	EQPH3SUCC	Equalization Phase 3 Successful When set to 1b, this bit indicates that Phase 3 of the Transmitter Equalization procedure has successfully completed.	0h	RO_V
3	EQPH2SUCC	Equalization Phase 2 Successful When set to 1b, this bit indicates that Phase 2 of the Transmitter Equalization procedure has successfully completed.	0h	RO_V
2	EQPH1SUCC	Equalization Phase 1 Successful When set to 1b, this bit indicates that Phase 1 of the Transmitter Equalization procedure has successfully completed.	0h	RO_V
1	EQCOMPLETE	Equalization Complete When set to 1b, this bit indicates that the Transmitter Equalization procedure has completed.	0h	RO_V
0	CURDELVL	Current De-emphasis Level: Current De-emphasis Level - When the Link is operating at 5 GT/s speed, this reflects the level of de-emphasis. Encodings: 1b -3.5 dB 0b -6 dB When the Link is operating at 2.5 GT/s speed, this bit is 0b.	0h	RO_V

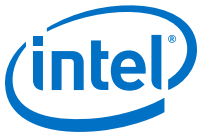

4.2.49 PVCCAP1—Port VC Capability Register 1

Describes the configuration of PCI Express Virtual Channels associated with this port.

B/D/F/Type: 0/1/1/CFG			Access: RO	
Size: 32	Default Value: 00000000h		Address Offset: 104h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:7	RSVD	Reserved.	0000000h	RO
6:4	LPEVCC	Low Priority Extended VC Count: Indicates the number of (extended) Virtual Channels in addition to the default VC belonging to the low-priority VC (LPVC) group that has the lowest priority with respect to other VC resources in a strict-priority VC Arbitration. The value of 0 in this field implies strict VC arbitration.	0h	RO
3	RSVD	Reserved.	0h	RO
2:0	EVCC	Extended VC Count: Indicates the number of (extended) Virtual Channels in addition to the default VC supported by the device.	0h	RO

4.2.50 PVCCAP2—Port VC Capability Register 2

Describes the configuration of PCI Express Virtual Channels associated with this port.

B/D/F/Type: 0/1/1/CFG			Access: RO	
Size: 32	Default Value: 00000000h		Address Offset: 108h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:24	VCATO	VC Arbitration Table Offset: Indicates the location of the VC Arbitration Table. This field contains the zero-based offset of the table in DQWORDS (16 bytes) from the base address of the Virtual Channel Capability Structure. A value of 0 indicates that the table is not present (due to fixed VC priority).	00h	RO
23:8	RSVD	Reserved.	0000h	RO
7:0	VCAC	Reserved for VC Arbitration Capability:	00h	RO

4.2.51 PVCCTL—Port VC Control

B/D/F/Type: 0/1/1/CFG			Access: RO; RW	
Size: 16	Default Value: 0000h		Address Offset: 10Ch	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:4	RSVD	Reserved.	000h	RO
3:1	VCAS	VC Arbitration Select: This field will be programmed by software to the only possible value as indicated in the VC Arbitration Capability field. Since there is no other VC supported than the default, this field is reserved.	0h	RW
0	VCARB	Reserved for Load VC Arbitration Table: Used for software to update the VC Arbitration Table when VC arbitration uses the VC Arbitration Table. As a VC Arbitration Table is never used by this component this field will never be used.	0h	RO

4.2.52 VC0RCAP—VC0 Resource Capability

B/D/F/Type: 0/1/1/CFG			Access: RO	
Size: 32	Default Value: 00000001h		Address Offset: 110h	
Bit Range	Acronym	Description	Default	Access
31:24	PATO	Reserved for Port Arbitration Table Offset:	00h	RO
23	RSVD	Reserved.	0h	RO
22:16	MTS	Reserved for Maximum Time Slots:	00h	RO
15	RSNPT	Reject Snoop Transactions: Reject Snoop Transactions (RSNPT): 0: Transactions with or without the No Snoop bit set within the TLP header are allowed on this VC. 1: When Set, any transaction for which the No Snoop attribute is applicable but is not Set within the TLP Header will be rejected as an Unsupported Request	0h	RO
14:8	RSVD	Reserved.	00h	RO
7:0	PAC	Port Arbitration Capability: Port Arbitration Capability - Indicates types of Port Arbitration supported by the VC resource. This field is valid for all Switch Ports, Root Ports that support peer-to to-peer traffic, and RCRBs, but not for PCI Express Endpoint devices or Root Ports that do not support peer to peer traffic. Each bit location within this field corresponds to a Port Arbitration Capability defined below. When more than one bit in this field is Set, it indicates that the VC resource can be configured to provide different arbitration services. Software selects among these capabilities by writing to the Port Arbitration Select field (see below). Defined bit positions are: Bit 0 Non-configurable hardware-fixed arbitration scheme, e.g., Round Robin (RR) Bit 1 Weighted Round Robin (WRR) arbitration with 32 phases Bit 2 WRR arbitration with 64 phases Bit 3 WRR arbitration with 128 phases Bit 4 Time-based WRR with 128 phases Bit 5 WRR arbitration with 256 phases Bits 6-7 Reserved Processor only supported arbitration indicates "Non-configurable hardware-fixed arbitration scheme".	01h	RO

4.2.53 VC0RCTL—VC0 Resource Control

Controls the resources associated with PCI Express Virtual Channel 0.

B/D/F/Type: 0/1/1/CFG			Access: RO; RW	
Size: 32	Default Value: 800000FFh		Address Offset: 114h	
Bit Range	Acronym	Description	Default	Access
31	VC0E	VC0 Enable: For VC0 this is hardwired to 1 and read only as VC0 can never be disabled.	1h	RO
30:27	RSVD	Reserved.	0h	RO
26:24	VC0ID	VC0 ID: Assigns a VC ID to the VC resource. For VC0 this is hardwired to 0 and read only.	0h	RO

continued...

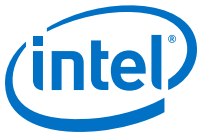

B/D/F/Type: 0/1/1/CFG			Access: RO; RW	
Size: 32	**Default Value: 800000FFh**		**Address Offset: 114h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
23:20	RSVD	Reserved.	0h	RO
19:17	PAS	Port Arbitration Select: Port Arbitration Select - This field configures the VC resource to provide a particular Port Arbitration service. This field is valid for RCRBs, Root Ports that support peer to peer traffic, and Switch Ports, but not for PCI Express Endpoint devices or Root Ports that do not support peer to peer traffic. The permissible value of this field is a number corresponding to one of the asserted bits in the Port Arbitration Capability field of the VC resource. This field does not affect the root port behavior.	0h	RW
16	RSVD	Reserved.	0h	RO
15:8	TCHVC0M	TC High VC0 Map: Allow usage of high order TCs. BIOS should keep this field zeroed to allow usage of the reserved TC[3] for other purposes	00h	RW
7:1	TCVC0M	TC/VC0 Map: Indicates the TCs (Traffic Classes) that are mapped to the VC resource. Bit locations within this field correspond to TC values. For example, when bit 7 is set in this field, TC7 is mapped to this VC resource. When more than one bit in this field is set, it indicates that multiple TCs are mapped to the VC resource. In order to remove one or more TCs from the TC/VC Map of an enabled VC, software must ensure that no new or outstanding transactions with the TC labels are targeted at the given Link.	7Fh	RW
0	TC0VC0M	TC0/VC0 Map: Traffic Class 0 is always routed to VC0.	1h	RO

4.2.54 VC0RSTS—VC0 Resource Status

Reports the Virtual Channel specific status.

B/D/F/Type: 0/1/1/CFG			Access: RO_V	
Size: 16	**Default Value: 0002h**		**Address Offset: 11Ah**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:2	RSVD	Reserved.	0000h	RO
1	VC0NP	VC0 Negotiation Pending: 0: The VC negotiation is complete. 1: The VC resource is still in the process of negotiation (initialization or disabling). This bit indicates the status of the process of Flow Control initialization. It is set by default on Reset, as well as whenever the corresponding Virtual Channel is Disabled or the Link is in the DL_Down state. It is cleared when the link successfully exits the FC_INIT2 state. Before using a Virtual Channel, software must check whether the VC Negotiation Pending fields for that Virtual Channel are cleared in both Components on a Link.	1h	RO_V
0	RSVD	Reserved.	0h	RO

4.3 PCI Express Controller (x4) Registers Summary

Offset	Register ID—Description	Default Value	Access
0	VID—Vendor Identification on page 197	8086h	RO
2	DID—Device Identification on page 198	0C09h	RO
4	PCICMD—PCI Command on page 198	0000h	RW; RO
6	PCISTS—PCI Status on page 199	0010h	RO_V; RO; RW1C
8	RID—Revision Identification on page 201	00h	RO
9	CC—Class Code on page 202	060400h	RO
C	CL—Cache Line Size on page 202	00h	RW
E	HDR—Header Type on page 202	81h	RO
18	PBUSN—Primary Bus Number on page 203	00h	RO
19	SBUSN—Secondary Bus Number on page 203	00h	RW
1A	SUBUSN—Subordinate Bus Number on page 203	00h	RW
1C	IOBASE—I/O Base Address on page 203	F0h	RW
1D	IOLIMIT—I/O Limit Address on page 204	00h	RW
1E	SSTS—Secondary Status on page 204	0000h	RO; RW1C
20	MBASE—Memory Base Address on page 205	FFF0h	RW
22	MLIMIT—Memory Limit Address on page 206	0000h	RW
24	PMBASE—Prefetchable Memory Base Address on page 206	FFF1h	RO; RW
26	PMLIMIT—Prefetchable Memory Limit Address on page 207	0001h	RO; RW
28	PMBASEU—Prefetchable Memory Base Address Upper on page 207	00000000h	RW
2C	PMLIMITU—Prefetchable Memory Limit Address Upper on page 208	00000000h	RW
34	CAPPTR—Capabilities Pointer on page 208	88h	RO
3C	INTRLINE—Interrupt Line on page 209	00h	RW
3D	INTRPIN—Interrupt Pin on page 209	01h	RW_O; RO
3E	BCTRL—Bridge Control on page 210	0000h	RW; RO
80	PM—Power Management Capabilities on page 211	C8039001h	RO; RO_V
84	PM—Power Management Control/Status on page 212	00000008h	RW; RO
88	SS—Subsystem ID and Vendor ID Capabilities on page 213	0000800Dh	RO
8C	SS—Subsystem ID and Subsystem Vendor ID on page 214	00008086h	RW_O
90	MSI—Message Signaled Interrupts Capability ID on page 214	A005h	RO
92	MC—Message Control on page 214	0000h	RW; RO
94	MA—Message Address on page 215	00000000h	RO; RW
98	MD—Message Data on page 215	0000h	RW
A0	PEG—PCI Express Capability List on page 216	0010h	RO
A2	PEG—PCI Express Capabilities on page 216	0142h	RO; RW_O

continued...

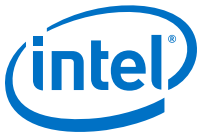

Offset	Register ID—Description	Default Value	Access
A4	DCAP—Device Capabilities on page 216	00008000h	RW_O; RO
A8	DCTL—Device Control on page 217	0000h	RW; RO
AA	DSTS—Device Status on page 218	0000h	RW1C; RO
B0	LCTL—Link Control on page 219	0000h	RW; RO; RW_V
B2	LSTS—Link Status on page 221	1001h	RO_V; RO; RW1C
B4	SLOTCAP—Slot Capabilities on page 222	00040000h	RO; RW_O
B8	SLOTCTL—Slot Control on page 224	0000h	RO
BA	SLOTSTS—Slot Status on page 226	0000h	RO; RW1C; RO_V
BC	RCTL—Root Control on page 227	00000000h	RW; RO; RWS_V; RO_V
C0	RSTS—Root Status on page 229	00000000h	RO_V; RW1C; RO
C4	DCAP2—Device Capabilites 2 on page 229	00000B80h	RO; RW_O
C8	DCTL2—Device Control 2 on page 231	0000h	RW; RO; RW_V
D0	LCTL2—Link Control 2 on page 232	0003h	RWS; RWS_V
D2	LSTS2—Link Status 2 on page 235	0000h	RO_V; RW1C
104	PVCCAP1—Port VC Capability Register 1 on page 236	00000000h	RO
108	PVCCAP2—Port VC Capability Register 2 on page 236	00000000h	RO
10C	PVCCTL—Port VC Control on page 236	0000h	RO; RW
110	VC0RCAP—VC0 Resource Capability on page 237	00000001h	RO
114	VC0RCTL—VC0 Resource Control on page 237	800000FFh	RO; RW
11A	VC0RSTS—VC0 Resource Status on page 238	0002h	RO_V

4.3.1 VID—Vendor Identification

This register combined with the Device Identification register uniquely identify any PCI device.

B/D/F/Type: 0/1/2/CFG			Access: RO	
Size: 16	Default Value: 8086h		Address Offset: 0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:0	VID	Vendor Identification: PCI standard identification for Intel.	8086h	RO

4.3.2 DID—Device Identification

This register combined with the Vendor Identification register uniquely identifies any PCI device.

B/D/F/Type: 0/1/2/CFG			Access: RO	
Size: 16	Default Value: 0C09h		Address Offset: 2h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:0	DID_MSB	Device Identification Number MSB: Identifier assigned to the Processor root port (virtual PCI-to-PCI bridge, PCI Express Graphics port).	0C09h	RO

4.3.3 PCICMD—PCI Command

B/D/F/Type: 0/1/2/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: 4h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:11	RSVD	Reserved.	00h	RO
10	INTAAD	INTA Assertion Disable: 0: This device is permitted to generate INTA interrupt messages. 1: This device is prevented from generating interrupt messages. Any INTA emulation interrupts already asserted must be de-asserted when this bit is set. Only affects interrupts generated by the device (PCI INTA from a PME or Hot Plug event) controlled by this command register. It does not affect upstream MSIs, upstream PCI INTA-INTD assert and deassert messages.	0h	RW
9	FB2B	Fast Back-to-Back Enable: Not Applicable or Implemented. Hardwired to 0.	0h	RO
8	SERRE	SERR# Message Enable: Controls the root port's SERR# messaging. The processor communicates the SERR# condition by sending an SERR message to the PCH. This bit, when set, enables reporting of non-fatal and fatal errors detected by the device to the Root Complex. Note that errors are reported if enabled either through this bit or through the PCI-Express specific bits in the Device Control Register. In addtion, for Type 1 configuration space header devices, this bit, when set, enables transmission by the primary interface of ERR_NONFATAL and ERR_FATAL error messages forwarded from the secondary interface. This bit does not affect the transmission of forwarded ERR_COR messages. 0: The SERR message is generated by the root port only under conditions enabled individually through the Device Control Register. 1: The root port is enabled to generate SERR messages which will be sent to the PCH for specific root port error conditions generated/detected or received on the secondary side of the virtual PCI to PCI bridge. The status of SERRs generated is reported in the PCISTS register.	0h	RW
7	RSVD	Reserved.	0h	RO
6	PERRE	Parity Error Response Enable: Controls whether or not the Master Data Parity Error bit in the PCI Status register can bet set. 0: Master Data Parity Error bit in PCI Status register	0h	RW

continued...

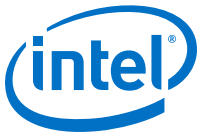

B/D/F/Type: 0/1/2/CFG			Access: RW; RO	
Size: 16	**Default Value: 0000h**		**Address Offset: 4h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		can NOT be set. 1: Master Data Parity Error bit in PCI Status register CAN be set.		
5	VGAPS	VGA Palette Snoop: Not Applicable or Implemented. Hardwired to 0.	0h	RO
4	MWIE	Memory Write and Invalidate Enable: Not Applicable or Implemented. Hardwired to 0.	0h	RO
3	SCE	Special Cycle Enable: Not Applicable or Implemented. Hardwired to 0.	0h	RO
2	BME	Bus Master Enable: Bus Master Enable (BME): Controls the ability of the PEG port to forward Memory Read/Write Requests in the upstream direction. 0: This device is prevented from making memory requests to its primary bus. Note that according to PCI Specification, as MSI interrupt messages are in-band memory writes, disabling the bus master enable bit prevents this device from generating MSI interrupt messages or passing them from its secondary bus to its primary bus. Upstream memory writes/reads, peer writes/reads, and MSIs will all be treated as illegal cycles. Writes are aborted. Reads are aborted and will return Unsupported Request status (or Master abort) in its completion packet. 1: This device is allowed to issue requests to its primary bus. Completions for previously issued memory read requests on the primary bus will be issued when the data is available. This bit does not affect forwarding of Completions from the primary interface to the secondary interface.	0h	RW
1	MAE	Memory Access Enable: 0: All of device's memory space is disabled. 1: Enable the Memory and Pre-fetchable memory address ranges defined in the MBASE, MLIMIT, PMBASE, and PMLIMIT registers.	0h	RW
0	IOAE	IO Access Enable: 0: All of devic's I/O space is disabled. 1: Enable the I/O address range defined in the IOBASE, and IOLIMIT registers.	0h	RW

4.3.4 PCISTS—PCI Status

This register reports the occurrence of error conditions associated with primary side of the "virtual" Host-PCI Express bridge embedded within the Root port.

B/D/F/Type: 0/1/2/CFG			Access: RO_V; RO; RW1C	
Size: 16	**Default Value: 0010h**		**Address Offset: 6h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15	DPE	Detected Parity Error: This bit is Set by a Function whenever it receives a Poisoned TLP, regardless of the state the Parity Error Response bit in the Command register. On a Function with a Type 1	0h	RW1C

continued...

B/D/F/Type: 0/1/2/CFG			Access: RO_V; RO; RW1C	
Size: 16	Default Value: 0010h		Address Offset: 6h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		Configuration header, the bit is Set when the Poisoned TLP is received by its Primary Side. Default value of this bit is 0b. This bit will be set only for completions of requests enountering ECC error in DRAM. Poisoned Peer 2 peer posted forwarded will not set this bit. They are reported at the receiveing port.		
14	SSE	Signaled System Error: This bit is set when this Device sends an SERR due to detecting an ERR_FATAL or ERR_NONFATAL condition and the SERR Enable bit in the Command register is '1'. Both received (if enabled by BCTRL1[1]) and internally detected error messages do not affect this field.	0h	RW1C
13	RMAS	Received Master Abort Status: This bit is Set when a Requester receives a Completion with Unsupported Request Completion Status. On a Function with a Type 1 Configuration header, the bit is Set when the Unsupported Request is received by its Primary Side. Not applicable. We do not have UR on primary interface	0h	RO
12	RTAS	Received Target Abort Status: This bit is Set when a Requester receives a Completion with Completer Abort Completion Status. On a Function with a Type 1 Configuration header, the bit is Set when the Completer Abort is received by its Primary Side. Default value of this bit is 0b. Not Applicable or Implemented. Hardwired to 0. The concept of a Completer abort does not exist on primary side of this device.	0h	RO
11	STAS	Signaled Target Abort Status: This bit is Set when a Function completes a Posted or Non- Posted Request as a Completer Abort error. This applies to a Function with a Type 1 Configuration header when the Completer Abort was generated by its Primary Side. Default value of this bit is 0b. Not Applicable or Implemented. Hardwired to 0. The concept of a target abort does not exist on primary side of this device.	0h	RO
10:9	DEVT	DEVSELB Timing: This device is not the subtractively decoded device on bus 0. This bit field is therefore hardwired to 00 to indicate that the device uses the fastest possible decode. Does not apply to PCI Express and must be hardwired to 00b.	0h	RO
8	PMDPE	Master Data Parity Error: This bit is Set by a Requester (Primary Side for Type 1 Configuration Space header Function) if the Parity Error Response bit in the Command register is 1b and either of the following two conditions occurs:	0h	RW1C

continued...

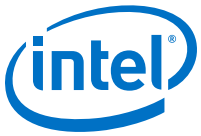

B/D/F/Type: 0/1/2/CFG			Access: RO_V; RO; RW1C	
Size: 16	Default Value: 0010h		Address Offset: 6h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		Requester receives a Completion marked poisoned Requester poisons a write Request If the Parity Error Response bit is 0b, this bit is never Set. Default value of this bit is 0b. This bit will be set only for completions of requests enountering ECC error in DRAM. Poisoned Peer 2 peer posted forwarded will not set this bit. They are reported at the receiveing port.		
7	FB2B	Fast Back-to-Back: Not Applicable or Implemented. Hardwired to 0.	0h	RO
6	RSVD	Reserved.	0h	RO
5	CAP66	66/60MHz capability: Not Applicable or Implemented. Hardwired to 0.	0h	RO
4	CAPL	Capabilities List: Indicates that a capabilities list is present. Hardwired to 1.	1h	RO
3	INTAS	INTx Status: Indicates that an interrupt message is pending internally to the device. Only PME and Hot Plug sources feed into this status bit (not PCI INTA-INTD assert and deassert messages). The INTA Assertion Disable bit, PCICMD1[10], has no effect on this bit. Note that INTA emulation interrupts received across the link are not reflected in this bit.	0h	RO_V
2:0	RSVD	Reserved.	0h	RO

4.3.5 RID—Revision Identification

This register contains the revision number of Device #1.
These bits are read only and writes to this register have no effect.

B/D/F/Type: 0/1/2/CFG			Access: RO	
Size: 8	Default Value: 00h		Address Offset: 8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:4	RID_MSB	Revision Identification Number MSB: This is an 8-bit value that indicates the revision identification number for the root port.	0h	RO
3:0	RID	Revision Identification Number: This is an 8-bit value that indicates the revision identification number for the root port.	0h	RO

4.3.6 CC—Class Code

This register identifies the basic function of the device, a more specific sub-class, and a register- specific programming interface.

B/D/F/Type: 0/1/2/CFG			Access: RO	
Size: 24	Default Value: 060400h		Address Offset: 9h	
Bit Range	Acronym	Description	Default	Access
23:16	BCC	Base Class Code: Indicates the base class code for this device. This code has the value 06h, indicating a Bridge device.	06h	RO
15:8	SUBCC	Sub-Class Code: Indicates the sub-class code for this device. The code is 04h indicating a PCI to PCI Bridge.	04h	RO
7:0	PI	Programming Interface: Indicates the programming interface of this device. This value does not specify a particular register set layout and provides no practical use for this device.	00h	RO

4.3.7 CL—Cache Line Size

B/D/F/Type: 0/1/2/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: Ch	
Bit Range	Acronym	Description	Default	Access
7:0	CLS	Cache Line Size: Implemented by PCI Express devices as a read-write field for legacy compatibility purposes but has no impact on any PCI Express device functionality.	00h	RW

4.3.8 HDR—Header Type

This register identifies the header layout of the configuration space. No physical register exists at this location.

B/D/F/Type: 0/1/2/CFG			Access: RO	
Size: 8	Default Value: 81h		Address Offset: Eh	
Bit Range	Acronym	Description	Default	Access
7:0	HDR	Header Type Register: Device #1 returns 81 to indicate that this is a multi function device with bridge header layout. Device #6 returns 01 to indicate that this is a single function device with bridge header layout.	81h	RO

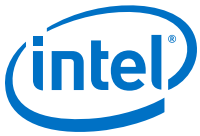

4.3.9 PBUSN—Primary Bus Number

This register identifies that this "virtual" Host-PCI Express bridge is connected to PCI bus #0.

B/D/F/Type: 0/1/2/CFG			Access: RO	
Size: 8	Default Value: 00h		Address Offset: 18h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:0	BUSN	Primary Bus Number: Configuration software typically programs this field with the number of the bus on the primary side of the bridge. Since the processor root port is an internal device and its primary bus is always 0, these bits are read only and are hardwired to 0.	00h	RO

4.3.10 SBUSN—Secondary Bus Number

This register identifies the bus number assigned to the second bus side of the "virtual" bridge i.e. to PCI Express. This number is programmed by the PCI configuration software to allow mapping of configuration cycles to PCI Express.

B/D/F/Type: 0/1/2/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: 19h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:0	BUSN	Secondary Bus Number: This field is programmed by configuration software with the bus number assigned to PCI Express.	00h	RW

4.3.11 SUBUSN—Subordinate Bus Number

This register identifies the subordinate bus (if any) that resides at the level below PCI Express. This number is programmed by the PCI configuration software to allow mapping of configuration cycles to PCI Express.

B/D/F/Type: 0/1/2/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: 1Ah	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:0	BUSN	Subordinate Bus Number: This register is programmed by configuration software with the number of the highest subordinate bus that lies behind the Processor root port bridge. When only a single PCI device resides on the PCI Express segment, this register will contain the same value as the SBUSN1 register.	00h	RW

4.3.12 IOBASE—I/O Base Address

This register controls the processor to PCI Express-G I/O access routing based on the following formula:

IO_BASE=< address =<IO_LIMIT

Only upper 4 bits are programmable. For the purpose of address decode address bits A[11:0] are treated as 0. Thus the bottom of the defined I/O address range will be aligned to a 4KB boundary.

B/D/F/Type: 0/1/2/CFG			Access: RW	
Size: 8	Default Value: F0h		Address Offset: 1Ch	
Bit Range	Acronym	Description	Default	Access
7:4	IOBASE	I/O Address Base: Corresponds to A[15:12] of the I/O addresses passed by the root port to PCI Express-G.	Fh	RW
3:0	RSVD	Reserved.	0h	RO

4.3.13 IOLIMIT—I/O Limit Address

This register controls the processor to PCI Express-G I/O access routing based on the following formula:

IO_BASE=< address =<IO_LIMIT

Only upper 4 bits are programmable. For the purpose of address decode address bits A[11:0] are assumed to be FFFh. Thus, the top of the defined I/O address range will be at the top of a 4KB aligned address block.

B/D/F/Type: 0/1/2/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: 1Dh	
Bit Range	Acronym	Description	Default	Access
7:4	IOLIMIT	I/O Address Limit: Corresponds to A[15:12] of the I/O address limit of the root port. Devices between this upper limit and IOBASE1 will be passed to the PCI Express hierarchy associated with this device.	0h	RW
3:0	RSVD	Reserved.	0h	RO

4.3.14 SSTS—Secondary Status

SSTS is a 16-bit status register that reports the occurrence of error conditions associated with secondary side (i.e. PCI Express-G side) of the "virtual" PCI-PCI bridge embedded within the Processor.

B/D/F/Type: 0/1/2/CFG			Access: RO; RW1C	
Size: 16	Default Value: 0000h		Address Offset: 1Eh	
Bit Range	Acronym	Description	Default	Access
15	DPE	Detected Parity Error: This bit is set by the Secondary Side for a Type 1 Configuration Space header device whenever it receives a Poisoned TLP, regardless of the state of the Parity Error Response Enable bit in the Bridge Control Register.	0h	RW1C
14	RSE	Received System Error: This bit is set when the Secondary Side for a Type 1 configuration space header device receives an ERR_FATAL or ERR_NONFATAL.	0h	RW1C
13	RMA	Received Master Abort: This bit is set when the Secondary Side for Type 1 Configuration Space Header Device (for requests initiated by the Type 1 Header Device itself) receives a Completion with Unsupported Request Completion Status.	0h	RW1C

continued...

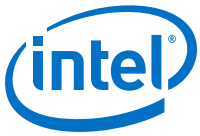

B/D/F/Type: 0/1/2/CFG			Access: RO; RW1C	
Size: 16	**Default Value: 0000h**		**Address Offset: 1Eh**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
12	RTA	Received Target Abort: This bit is set when the Secondary Side for Type 1 Configuration Space Header Device (for requests initiated by the Type 1 Header Device itself) receives a Completion with Completer Abort Completion Status.	0h	RW1C
11	STA	Signaled Target Abort: Not Applicable or Implemented. Hardwired to 0. The processor does not generate Target Aborts (The root port will never complete a request using the Completer Abort Completion status). UR detected inside the processor (such as in MC) will be reported in primary side status.	0h	RO
10:9	DEVT	DEVSELB Timing: Not Applicable or Implemented. Hardwired to 0.	0h	RO
8	SMDPE	Master Data Parity Error: When set indicates that the processor received across the link (upstream) a Read Data Completion Poisoned TLP (EP=1). This bit can only be set when the Parity Error Enable bit in the Bridge Control register is set.	0h	RW1C
7	FB2B	Fast Back-to-Back: Not Applicable or Implemented. Hardwired to 0.	0h	RO
6	RSVD	Reserved.	0h	RO
5	CAP66	66/60 MHz capability: Not Applicable or Implemented. Hardwired to 0.	0h	RO
4:0	RSVD	Reserved.	00h	RO

4.3.15 MBASE—Memory Base Address

This register controls the processor to PCI Express non-prefetchable memory access routing based on the following formula:

MEMORY_BASE=< address =<MEMORY_LIMIT

The upper 12 bits of the register are read/write and correspond to the upper 12 address bits A[31:20] of the 32 bit address. The bottom 4 bits of this register are read-only and return zeroes when read. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be 0. Thus, the bottom of the defined memory address range will be aligned to a 1MB boundary.

B/D/F/Type: 0/1/2/CFG			Access: RW	
Size: 16	**Default Value: FFF0h**		**Address Offset: 20h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:4	MBASE	Memory Address Base: Corresponds to A[31:20] of the lower limit of the memory range that will be passed to PCI Express.	FFFh	RW
3:0	RSVD	Reserved.	0h	RO

4.3.16 MLIMIT—Memory Limit Address

This register controls the processor to PCI Express non-prefetchable memory access routing based on the following formula:
MEMORY_BASE=< address =<MEMORY_LIMIT
The upper 12 bits of the register are read/write and correspond to the upper 12 address bits A[31:20] of the 32 bit address. The bottom 4 bits of this register are read-only and return zeroes when read. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be FFFFFh. Thus, the top of the defined memory address range will be at the top of a 1MB aligned memory block. NOTE: Memory range covered by MBASE and MLIMIT registers are used to map non-prefetchable PCI Express address ranges (typically where control/status memory-mapped I/O data structures of the graphics controller will reside) and PMBASE and PMLIMIT are used to map prefetchable address ranges (typically graphics local memory). This segregation allows application of USWC space attribute to be performed in a true plug-and-play manner to the prefetchable address range for improved Processor-PCI Express memory access performance.
Note also that configuration software is responsible for programming all address range registers (prefetchable, non-prefetchable) with the values that provide exclusive address ranges i.e. prevent overlap with each other and/or with the ranges covered with the main memory. There is no provision in the processor hardware to enforce prevention of overlap and operations of the system in the case of overlap are not guaranteed.

B/D/F/Type: 0/1/2/CFG			Access: RW	
Size: 16	Default Value: 0000h		Address Offset: 22h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:4	MLIMIT	Memory Address Limit: Corresponds to A[31:20] of the upper limit of the address range passed to PCI Express.	000h	RW
3:0	RSVD	Reserved.	0h	RO

4.3.17 PMBASE—Prefetchable Memory Base Address

This register in conjunction with the corresponding Upper Base Address register controls the processor to PCI Express prefetchable memory access routing based on the following formula:
PREFETCHABLE_MEMORY_BASE =< address =< PREFETCHABLE_MEMORY_LIMIT
The upper 12 bits of this register are read/write and correspond to address bits A[31:20] of the 40-bit address. The lower 8 bits of the Upper Base Address register are read/write and correspond to address bits A[39:32] of the 40-bit address. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be 0. Thus, the bottom of the defined memory address range will be aligned to a 1MB boundary.

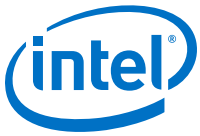

B/D/F/Type: 0/1/2/CFG			Access: RO; RW	
Size: 16	**Default Value: FFF1h**		**Address Offset: 24h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:4	PMBASE	Prefetchable Memory Base Address: Corresponds to A[31:20] of the lower limit of the memory range that will be passed to PCI Express.	FFFh	RW
3:0	AS64	64-bit Address Support: Indicates that the upper 32 bits of the prefetchable memory region base address are contained in the Prefetchable Memory base Upper Address register at 28h.	1h	RO

4.3.18 PMLIMIT—Prefetchable Memory Limit Address

This register in conjunction with the corresponding Upper Limit Address register controls the processor to PCI Express prefetchable memory access routing based on the following formula:

PREFETCHABLE_MEMORY_BASE =< address =< PREFETCHABLE_MEMORY_LIMIT

The upper 12 bits of this register are read/write and correspond to address bits A[31:20] of the 40-bit address. The lower 8 bits of the Upper Limit Address register are read/write and correspond to address bits A[39:32] of the 40-bit address. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be FFFFFh. Thus, the top of the defined memory address range will be at the top of a 1MB aligned memory block. Note that prefetchable memory range is supported to allow segregation by the configuration software between the memory ranges that must be defined as UC and the ones that can be designated as a USWC (i.e. prefetchable) from the processor perspective.

B/D/F/Type: 0/1/2/CFG			Access: RO; RW	
Size: 16	**Default Value: 0001h**		**Address Offset: 26h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:4	PMLIMIT	Prefetchable Memory Address Limit: Corresponds to A[31:20] of the upper limit of the address range passed to PCI Express.	000h	RW
3:0	AS64B	64-bit Address Support: Indicates that the upper 32 bits of the prefetchable memory region limit address are contained in the Prefetchable Memory Base Limit Address register at 2Ch	1h	RO

4.3.19 PMBASEU—Prefetchable Memory Base Address Upper

The functionality associated with this register is present in the PEG design implementation.

This register in conjunction with the corresponding Upper Base Address register controls the processor to PCI Express prefetchable memory access routing based on the following formula:

PREFETCHABLE_MEMORY_BASE =< address =< PREFETCHABLE_MEMORY_LIMIT

The upper 12 bits of this register are read/write and correspond to address bits A[31:20] of the 39-bit address. The lower 7 bits of the Upper Base Address register are read/write and correspond to address bits A[38:32] of the 39-bit address. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be 0. Thus, the bottom of the defined memory address range will be aligned to a 1MB boundary.

B/D/F/Type: 0/1/2/CFG			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 28h	
Bit Range	Acronym	Description	Default	Access
31:0	PMBASEU	Prefetchable Memory Base Address: Corresponds to A[63:32] of the lower limit of the prefetchable memory range that will be passed to PCI Express.	00000000h	RW

4.3.20 PMLIMITU—Prefetchable Memory Limit Address Upper

The functionality associated with this register is present in the PEG design implementation.

This register in conjunction with the corresponding Upper Limit Address register controls the processor to PCI Express prefetchable memory access routing based on the following formula:

PREFETCHABLE_MEMORY_BASE =< address =< PREFETCHABLE_MEMORY_LIMIT

The upper 12 bits of this register are read/write and correspond to address bits A[31:20] of the 39-bit address. The lower 7 bits of the Upper Limit Address register are read/write and correspond to address bits A[39:32] of the 39-bit address. This register must be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be FFFFFh. Thus, the top of the defined memory address range will be at the top of a 1MB aligned memory block.

Note that prefetchable memory range is supported to allow segregation by the configuration software between the memory ranges that must be defined as UC and the ones that can be designated as a USWC (i.e. prefetchable) from the processor perspective.

B/D/F/Type: 0/1/2/CFG			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 2Ch	
Bit Range	Acronym	Description	Default	Access
31:0	PMLIMITU	Prefetchable Memory Address Limit: Corresponds to A[63:32] of the upper limit of the prefetchable Memory range that will be passed to PCI Express.	00000000h	RW

4.3.21 CAPPTR—Capabilities Pointer

The capabilities pointer provides the address offset to the location of the first entry in this device's linked list of capabilities.

B/D/F/Type: 0/1/2/CFG			Access: RO	
Size: 8	Default Value: 88h		Address Offset: 34h	
Bit Range	Acronym	Description	Default	Access
7:0	CAPPTR1	First Capability: The first capability in the list is the Subsystem ID and Subsystem Vendor ID Capability.	88h	RO

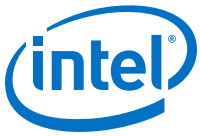

4.3.22 INTRLINE—Interrupt Line

This register contains interrupt line routing information. The device itself does not use this value, rather it is used by device drivers and operating systems to determine priority and vector information.

B/D/F/Type: 0/1/2/CFG			Access: RW	
Size: 8	Default Value: 00h		Address Offset: 3Ch	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:0	INTCON	Interrupt Connection: Used to communicate interrupt line routing information. BIOS Requirement: POST software writes the routing information into this register as it initializes and configures the system. The value indicates to which input of the system interrupt controller this device's interrupt pin is connected.	00h	RW

4.3.23 INTRPIN—Interrupt Pin

This register specifies which interrupt pin this device uses.

B/D/F/Type: 0/1/2/CFG			Access: RW_O; RO	
Size: 8	Default Value: 01h		Address Offset: 3Dh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7:3	INTPINH	Interrupt Pin High:	00h	RO
2:0	INTPIN	Interrupt Pin: As a multifunction device, the PCI Express device may specify any INTx (x=A,B,C,D) as its interrupt pin. The Interrupt Pin register tells which interrupt pin the device (or device function) uses. A value of 1 corresponds to INTA# (Default) A value of 2 corresponds to INTB# A value of 3 corresponds to INTC# A value of 4 corresponds to INTD# Devices (or device functions) that do not use an interrupt pin must put a 0 in this register. The values 05h through FFh are reserved. This register is write once. BIOS must set this register to select the INTx to be used by this root port.	1h	RW_O

4.3.24 BCTRL—Bridge Control

This register provides extensions to the PCICMD register that are specific to PCI-PCI bridges. The BCTRL provides additional control for the secondary interface (i.e. PCI Express) as well as some bits that affect the overall behavior of the "virtual" Host-PCI Express bridge embedded within the processor, e.g. VGA compatible address ranges mapping.

B/D/F/Type: 0/1/2/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: 3Eh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:12	RSVD	Reserved.	0h	RO
11	DTSERRE	Discard Timer SERR# Enable: Not Applicable or Implemented. Hardwired to 0.	0h	RO
10	DTSTS	Discard Timer Status: Not Applicable or Implemented. Hardwired to 0.	0h	RO
9	SDT	Secondary Discard Timer: Not Applicable or Implemented. Hardwired to 0.	0h	RO
8	PDT	Primary Discard Timer: Not Applicable or Implemented. Hardwired to 0.	0h	RO
7	FB2BEN	Fast Back-to-Back Enable: Not Applicable or Implemented. Hardwired to 0.	0h	RO
6	SRESET	Secondary Bus Reset: Setting this bit triggers a hot reset on the corresponding PCI Express Port. This will force the LTSSM to transition to the Hot Reset state (via Recovery) from L0, L0s, or L1 states.	0h	RW
5	MAMODE	Master Abort Mode: Does not apply to PCI Express. Hardwired to 0.	0h	RO
4	VGA16D	VGA 16-bit Decode: Enables the PCI-to-PCI bridge to provide 16-bit decoding of VGA I/O address precluding the decoding of alias addresses every 1 KB. This bit only has meaning if bit 3 (VGA Enable) of this register is also set to 1, enabling VGA I/O decoding and forwarding by the bridge. 0: Execute 10-bit address decodes on VGA I/O accesses. 1: Execute 16-bit address decodes on VGA I/O accesses.	0h	RW
3	VGAEN	VGA Enable: Controls the routing of processor initiated transactions targeting VGA compatible I/O and memory address ranges. See the VGAEN/MDAP table in device 0, offset 97h[0].	0h	RW
2	ISAEN	ISA Enable: Needed to exclude legacy resource decode to route ISA resources to legacy decode path. Modifies the response by the root port to an I/O access issued by the processor that target ISA I/O addresses. This applies only to I/O addresses that are enabled by the IOBASE and IOLIMIT registers. 0: All addresses defined by the IOBASE and IOLIMIT for processor I/O transactions will be mapped to PCI Express. 1: The root port will not forward to PCI Express any	0h	RW

continued...

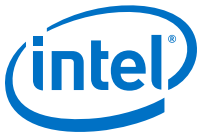

B/D/F/Type: 0/1/2/CFG			Access: RW; RO	
Size: 16	**Default Value: 0000h**		**Address Offset: 3Eh**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		I/O transactions addressing the last 768 bytes in each 1KB block even if the addresses are within the range defined by the IOBASE and IOLIMIT registers.		
1	SERREN	SERR Enable: 0: No forwarding of error messages from secondary side to primary side that could result in an SERR. 1: ERR_COR, ERR_NONFATAL, and ERR_FATAL messages result in SERR message when individually enabled by the Root Control register.	0h	RW
0	PEREN	Parity Error Response Enable: Controls whether or not the Master Data Parity Error bit in the Secondary Status register is set when the root port receives across the link (upstream) a Read Data Completion Poisoned TLP 0: Master Data Parity Error bit in Secondary Status register can NOT be set. 1: Master Data Parity Error bit in Secondary Status register CAN be set.	0h	RW

4.3.25 PM—Power Management Capabilities

B/D/F/Type: 0/1/2/CFG			Access: RO; RO_V	
Size: 32	**Default Value: C8039001h**		**Address Offset: 80h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:27	PMES	PME Support: This field indicates the power states in which this device may indicate PME wake via PCI Express messaging. D0, D3hot & D3cold. This device is not required to do anything to support D3hot & D3cold, it simply must report that those states are supported. Refer to the PCI Power Management 1.1 specification for encoding explanation and other power management details.	19h	RO
26	D2PSS	D2 Power State Support: Hardwired to 0 to indicate that the D2 power management state is NOT supported.	0h	RO
25	D1PSS	D1 Power State Support: Hardwired to 0 to indicate that the D1 power management state is NOT supported.	0h	RO
24:22	AUXC	Auxiliary Current: Hardwired to 0 to indicate that there are no 3.3Vaux auxiliary current requirements.	0h	RO
21	DSI	Device Specific Initialization: Hardwired to 0 to indicate that special initialization of this device is NOT required before generic class device driver is to use it.	0h	RO
20	APS	Auxiliary Power Source: Hardwired to 0.	0h	RO
19	PMECLK	PME Clock: Hardwired to 0 to indicate this device does NOT support PME# generation.	0h	RO
18:16	PCIPMCV	PCI PM CAP Version: Version - A value of 011b indicates that this function complies with revision 1.2 of the PCI Power Management Interface Specification. --Was Previously Hardwired to 02h to indicate there are 4 bytes of power	3h	RO

continued...

B/D/F/Type: 0/1/2/CFG			Access: RO; RO_V	
Size: 32	Default Value: C8039001h		Address Offset: 80h	
Bit Range	Acronym	Description	Default	Access
		management registers implemented and that this device complies with revision 1.1 of the PCI Power Management Interface Specification.		
15:8	PNC	Pointer to Next Capability: This contains a pointer to the next item in the capabilities list. If MSICH (CAPL[0] @ 7Fh) is 0, then the next item in the capabilities list is the Message Signaled Interrupts (MSI) capability at 90h. If MSICH (CAPL[0] @ 7Fh) is 1, then the next item in the capabilities list is the PCI Express capability at A0h.	90h	RO_V
7:0	CID	Capability ID: Value of 01h identifies this linked list item (capability structure) as being for PCI Power Management registers.	01h	RO

4.3.26 PM—Power Management Control/Status

B/D/F/Type: 0/1/2/CFG			Access: RW; RO	
Size: 32	Default Value: 00000008h		Address Offset: 84h	
Bit Range	Acronym	Description	Default	Access
31:16	RSVD	Reserved.	0000h	RO
15	PMESTS	PME Status: Indicates that this device does not support PMEB generation from D3cold.	0h	RO
14:13	DSCALE	Data Scale: Indicates that this device does not support the power management data register.	0h	RO
12:9	DSEL	Data Select: Indicates that this device does not support the power management data register.	0h	RO
8	PMEE	PME Enable: Indicates that this device does not generate PMEB assertion from any D-state. 0: PMEB generation not possible from any D State 1: PMEB generation enabled from any D State The setting of this bit has no effect on hardware. See PM_CAP[15:11]	0h	RW
7:4	RSVD	Reserved.	0h	RO
3	NSR	No Soft Reset: No Soft Reset. When set to 1 this bit indicates that the device is transitioning from D3hot to D0 because the power state commands do not perform a internal reset. Config context is preserved. Upon transition no additonal operating sys intervention is required to preserve configuration context beyond writing the power state bits. When clear the devices do not perform an internal reset upon transitioning from D3hot to D0 via software control of the power state bits. Regardless of this bit the devices that transition from a	1h	RO

continued...

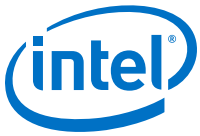

B/D/F/Type: 0/1/2/CFG			Access: RW; RO	
Size: 32	**Default Value: 00000008h**		**Address Offset: 84h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		D3hot to D0 by a system or bus segment reset will return to the device state D0 uniintialized with only PME context preserved if PME is supported and enabled.		
2	RSVD	Reserved.	0h	RO
1:0	PS	Power State: Indicates the current power state of this device and can be used to set the device into a new power state. If software attempts to write an unsupported state to this field, write operation must complete normally on the bus, but the data is discarded and no state change occurs. 00: D0 01: D1 (Not supported in this device.) 10: D2 (Not supported in this device.) 11: D3 Support of D3cold does not require any special action. While in the D3hot state, this device can only act as the target of PCI configuration transactions (for power management control). This device also cannot generate interrupts or respond to MMR cycles in the D3 state. The device must return to the D0 state in order to be fully-functional. When the Power State is other than D0, the bridge will Master Abort (i.e. not claim) any downstream cycles (with exception of type 0 config cycles). Consequently, these unclaimed cycles will go down DMI and come back up as Unsupported Requests, which the processor logs as Master Aborts in Device 0 PCISTS[13] There is no additional hardware functionality required to support these Power States.	0h	RW

4.3.27 SS—Subsystem ID and Vendor ID Capabilities

This capability is used to uniquely identify the subsystem where the PCI device resides. Because this device is an integrated part of the system and not an add-in device, it is anticipated that this capability will never be used. However, it is necessary because Microsoft will test for its presence.

B/D/F/Type: 0/1/2/CFG			Access: RO	
Size: 32	**Default Value: 0000800Dh**		**Address Offset: 88h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:8	PNC	Pointer to Next Capability: This contains a pointer to the next item in the capabilities list which is the PCI Power Management capability.	80h	RO
7:0	CID	Capability ID: Value of 0Dh identifies this linked list item (capability structure) as being for SSID/SSVID registers in a PCI-to-PCI Bridge.	0Dh	RO

4.3.28 SS—Subsystem ID and Subsystem Vendor ID

System BIOS can be used as the mechanism for loading the SSID/SVID values. These values must be preserved through power management transitions and a hardware reset.

B/D/F/Type: 0/1/2/CFG			Access: RW_O	
Size: 32	Default Value: 00008086h		Address Offset: 8Ch	
Bit Range	Acronym	Description	Default	Access
31:16	SSID	Subsystem ID: Identifies the particular subsystem and is assigned by the vendor.	0000h	RW_O
15:0	SSVID	Subsystem Vendor ID: Identifies the manufacturer of the subsystem and is the same as the vendor ID which is assigned by the PCI Special Interest Group.	8086h	RW_O

4.3.29 MSI—Message Signaled Interrupts Capability ID

When a device supports MSI it can generate an interrupt request to the processor by writing a predefined data item (a message) to a predefined memory address.

The reporting of the existence of this capability can be disabled by setting MSICH (CAPL[0] @ 7Fh). In that case walking this linked list will skip this capability and instead go directly from the PCI PM capability to the PCI Express capability.

B/D/F/Type: 0/1/2/CFG			Access: RO	
Size: 16	Default Value: A005h		Address Offset: 90h	
Bit Range	Acronym	Description	Default	Access
15:8	PNC	Pointer to Next Capability: This contains a pointer to the next item in the capabilities list which is the PCI Express capability.	A0h	RO
7:0	CID	Capability ID: Value of 05h identifies this linked list item (capability structure) as being for MSI registers.	05h	RO

4.3.30 MC—Message Control

System software can modify bits in this register, but the device is prohibited from doing so.

If the device writes the same message multiple times, only one of those messages is guaranteed to be serviced. If all of them must be serviced, the device must not generate the same message again until the driver services the earlier one.

B/D/F/Type: 0/1/2/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: 92h	
Bit Range	Acronym	Description	Default	Access
15:8	RSVD	Reserved.	00h	RO
7	B64AC	64-bit Address Capable: Hardwired to 0 to indicate that the function does not implement the upper 32 bits of the Message Address register and is incapable of generating a 64-bit memory address. This may need to change in future implementations when addressable system memory exceeds the 32b/4GB limit.	0h	RO

continued...

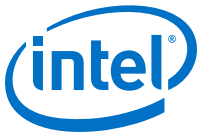

B/D/F/Type: 0/1/2/CFG			Access: RW; RO	
Size: 16	**Default Value: 0000h**		**Address Offset: 92h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
6:4	MME	Multiple Message Enable: System software programs this field to indicate the actual number of messages allocated to this device. This number will be equal to or less than the number actually requested. The encoding is the same as for the MMC field below.	0h	RW
3:1	MMC	Multiple Message Capable: System software reads this field to determine the number of messages being requested by this device. Value: Number of Messages Requested 000: 1 All of the following are reserved in this implementation: 001: 2 010: 4 011: 8 100: 16 101: 32 110: Reserved 111: Reserved	0h	RO
0	MSIEN	MSI Enable: Controls the ability of this device to generate MSIs. 0: MSI will not be generated. 1: MSI will be generated when we receive PME messages. INTA will not be generated and INTA Status (PCISTS1[3]) will not be set.	0h	RW

4.3.31 MA—Message Address

B/D/F/Type: 0/1/2/CFG			Access: RO; RW	
Size: 32	**Default Value: 00000000h**		**Address Offset: 94h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:2	MA	Message Address: Used by system software to assign an MSI address to the device. The device handles an MSI by writing the padded contents of the MD register to this address.	00000000h	RW
1:0	FDWA	Force DWord Align: Hardwired to 0 so that addresses assigned by system software are always aligned on a dword address boundary.	0h	RO

4.3.32 MD—Message Data

B/D/F/Type: 0/1/2/CFG			Access: RW	
Size: 16	**Default Value: 0000h**		**Address Offset: 98h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:0	MD	Message Data: Base message data pattern assigned by system software and used to handle an MSI from the device. When the device must generate an interrupt request, it writes a 32-bit value to the memory address specified in the MA register. The upper 16 bits are always set to 0. The lower 16 bits are supplied by this register.	0000h	RW

4.3.33 PEG—PCI Express Capability List

Enumerates the PCI Express capability structure.

B/D/F/Type: 0/1/2/CFG			Access: RO	
Size: 16	Default Value: 0010h		Address Offset: A0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:8	PNC	Pointer to Next Capability: This value terminates the capabilities list. The Virtual Channel capability and any other PCI Express specific capabilities that are reported via this mechanism are in a separate capabilities list located entirely within PCI Express Extended Configuration Space.	00h	RO
7:0	CID	Capability ID: Identifies this linked list item (capability structure) as being for PCI Express registers.	10h	RO

4.3.34 PEG—PCI Express Capabilities

Indicates PCI Express device capabilities.

B/D/F/Type: 0/1/2/CFG			Access: RO; RW_O	
Size: 16	Default Value: 0142h		Address Offset: A2h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:14	RSVD	Reserved.	0h	RO
13:9	IMN	Interrupt Message Number: Not Applicable or Implemented. Hardwired to 0.	00h	RO
8	SI	Slot Implemented: 0: The PCI Express Link associated with this port is connected to an integrated component or is disabled. 1: The PCI Express Link associated with this port is connected to a slot. BIOS Requirement: This field must be initialized appropriately if a slot connection is not implemented.	1h	RW_O
7:4	DPT	Device/Port Type: Hardwired to 4h to indicate root port of PCI Express Root Complex.	4h	RO
3:0	PCIECV	PCI Express Capability Version: PCI Express Capability Version (PCIECV): Hardwired to 2h to indicate compliance to the PCI Express Capabilities Register Expansion ECN.	2h	RO

4.3.35 DCAP—Device Capabilities

Indicates PCI Express device capabilities.

B/D/F/Type: 0/1/2/CFG			Access: RW_O; RO	
Size: 32	Default Value: 00008000h		Address Offset: A4h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15	RBER	Role Based Error Reporting (RBER): Indicates that this device implements the functionality defined in the Error Reporting ECN as required by the PCI Express 1.1 spec.	1h	RO
14:6	RSVD	Reserved.	000h	RO

continued...

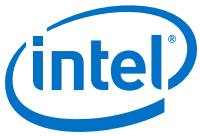

B/D/F/Type: 0/1/2/CFG			Access: RW_O; RO	
Size: 32	**Default Value: 00008000h**		**Address Offset: A4h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
5	ETFS	Extended Tag Field Supported: Hardwired to indicate support for 5-bit Tags as a Requestor.	0h	RO
4:3	PFS	Phantom Functions Supported: Not Applicable or Implemented. Hardwired to 0.	0h	RO
2:0	MPS	Max Payload Size: Default indicates 256B max supported payload for Transaction Layer Packets (TLP) for x16 PEG only. x8 and x4 PEG are limited to 128B support.	0h	RW_O

4.3.36 DCTL—Device Control

Provides control for PCI Express device specific capabilities.
The error reporting enable bits are in reference to errors detected by this device, not error messages received across the link. The reporting of error messages (ERR_CORR, ERR_NONFATAL, ERR_FATAL) received by Root Port is controlled exclusively by Root Port Command Register.

B/D/F/Type: 0/1/2/CFG			Access: RW; RO	
Size: 16	**Default Value: 0000h**		**Address Offset: A8h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15	RSVD	Reserved.	0h	RO
14:12	MRRS	Reserved for Max Read Request Size:	0h	RO
11	NSE	Reserved for Enable No Snoop:	0h	RO
10:8	RSVD	Reserved.	0h	RO
7:5	MPS	000b 128B max supported payload for Transaction Layer Packets (TLP). 001b 256B max supported payload for Transaction Layer Packets (TLP). All other encodings are reserved. As a receiver, the Device must handle TLPs as large as the set value. As transmitter, the Device must not generate TLPs exceeding the set value. NOTE: A payload size of 256B is only supported for Device 1 Function 0.	0h	RW
4	ROE	Reserved for Enable Relaxed Ordering:	0h	RO
3	URRE	Unsupported Request Reporting Enable: Unsupported Request Reporting Enable (URRE): When set, allows signaling ERR_NONFATAL, ERR_FATAL, or ERR_CORR to the Root Control register when detecting an unmasked Unsupported Request (UR). An ERR_CORR is signaled when an unmasked Advisory Non-Fatal UR is received. An ERR_FATAL or ERR_NONFATAL is sent to the Root Control register when an uncorrectable non-Advisory UR is received with the severity bit set in the Uncorrectable Error Severity register.	0h	RW

continued...

B/D/F/Type: 0/1/2/CFG			Access: RW; RO	
Size: 16	Default Value: 0000h		Address Offset: A8h	
Bit Range	Acronym	Description	Default	Access
2	FERE	Fatal Error Reporting Enable: Fatal Error Reporting Enable (FERE): When set, enables signaling of ERR_FATAL to the Root Control register due to internally detected errors or error messages received across the link. Other bits also control the full scope of related error reporting.	0h	RW
1	NERE	Non-Fatal Error Reporting Enable: Non-Fatal Error Reporting Enable (NERE): When set, enables signaling of ERR_NONFATAL to the Rool Control register due to internally detected errors or error messages received across the link. Other bits also control the full scope of related error reporting.	0h	RW
0	CERE	Correctable Error Reporting Enable: Correctable Error Reporting Enable (CERE): When set, enables signaling of ERR_CORR to the Root Control register due to internally detected errors or error messages received across the link. Other bits also control the full scope of related error reporting.	0h	RW

4.3.37 DSTS—Device Status

Reflects status corresponding to controls in the Device Control register. The error reporting bits are in reference to errors detected by this device, not errors messages received across the link.

B/D/F/Type: 0/1/2/CFG			Access: RW1C; RO	
Size: 16	Default Value: 0000h		Address Offset: AAh	
Bit Range	Acronym	Description	Default	Access
15:6	RSVD	Reserved.	000h	RO
5	TP	Transactions Pending: 0: All pending transactions (including completions for any outstanding non-posted requests on any used virtual channel) have been completed. 1: Indicates that the device has transaction(s) pending (including completions for any outstanding non-posted requests for all used Traffic Classes). Not Applicable or Implemented. Hardwired to 0.	0h	RO
4	RSVD	Reserved.	0h	RO
3	URD	Unsupported Request Detected: When set this bit indicates that the Device received an Unsupported Request. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control Register. Additionally, the Non-Fatal Error Detected bit or the Fatal Error Detected bit is set according to the setting of the Unsupported Request Error Severity bit. In production systems setting the Fatal Error Detected bit is not an option as support for AER will not be reported.	0h	RW1C
2	FED	Fatal Error Detected: When set this bit indicates that fatal error(s) were detected. Errors are logged in this register regardless of whether error reporting is enabled or not in	0h	RW1C

continued...

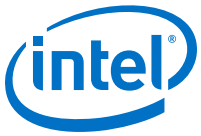

B/D/F/Type: 0/1/2/CFG			Access: RW1C; RO	
Size: 16	**Default Value: 0000h**		**Address Offset: AAh**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		the Device Control register. When Advanced Error Handling is enabled, errors are logged in this register regardless of the settings of the uncorrectable error mask register.		
1	NFED	Non-Fatal Error Detected: When set this bit indicates that non-fatal error(s) were detected. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. When Advanced Error Handling is enabled, errors are logged in this register regardless of the settings of the uncorrectable error mask register.	0h	RW1C
0	CED	Correctable Error Detected: When set this bit indicates that correctable error(s) were detected. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. When Advanced Error Handling is enabled, errors are logged in this register regardless of the settings of the correctable error mask register.	0h	RW1C

4.3.38 LCTL—Link Control

Allows control of PCI Express link.

B/D/F/Type: 0/1/2/CFG			Access: RW; RO; RW_V	
Size: 16	**Default Value: 0000h**		**Address Offset: B0h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:12	RSVD	Reserved.	0h	RO
11	LABIE	Link Autonomous Bandwidth Interrupt Enable: Link Autonomous Bandwidth Interrupt Enable - When Set, this bit enables the generation of an interrupt to indicate that the Link Autonomous Bandwidth Status bit has been Set. This bit is not applicable and is reserved for Endpoint devices, PCI Express to PCI/PCI-X bridges, and Upstream Ports of Switches. Devices that do not implement the Link Bandwidth Notification capability must hardwire this bit to 0b.	0h	RW
10	LBMIE	Link Bandwidth Management Interrupt Enable: Link Bandwidth Management Interrupt Enable - When Set, this bit enables the generation of an interrupt to indicate that the Link Bandwidth Management Status bit has been Set. This bit is not applicable and is reserved for Endpoint devices, PCI Express to PCI/PCI-X bridges, and Upstream Ports of Switches.	0h	RW
9	HAWD	Hardware Autonomous Width Disable: Hardware Autonomous Width Disable - When Set, this bit disables hardware from changing the Link width for reasons other than attempting to correct unreliable Link operation by reducing Link width. Devices that do not implement the ability autonomously to change Link width are permitted to hardwire this bit to 0b.	0h	RW

continued...

B/D/F/Type: 0/1/2/CFG			Access: RW; RO; RW_V	
Size: 16	Default Value: 0000h		Address Offset: B0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
8	ECPM	Enable Clock Power Management: Applicable only for form factors that support a "Clock Request" (CLKREQ#) mechanism, this enable functions as follows 0b - Clock power management is disabled and device must hold CLKREQ# signal low 1b - When this bit is set to 1 the device is permitted to use CLKREQ# signal to power manage link clock according to protocol defined in appropriate form factor specification. Default value of this field is 0b. Components that do not support Clock Power Management (as indicated by a 0b value in the Clock Power Management bit of the Link Capabilities Register) must hardwire this bit to 0b.	0h	RO
7	ES	Extended Synch: Extended synch 0: Standard Fast Training Sequence (FTS). 1: Forces the transmission of additional ordered sets when exiting the L0s state and when in the Recovery state. This mode provides external devices (e.g., logic analyzers) monitoring the Link time to achieve bit and symbol lock before the link enters L0 and resumes communication. This is a test mode only and may cause other undesired side effects such as buffer overflows or underruns.	0h	RW
6	CCC	Common Clock Configuration: 0: Indicates that this component and the component at the opposite end of this Link are operating with asynchronous reference clock. 1: Indicates that this component and the component at the opposite end of this Link are operating with a distributed common reference clock. The state of this bit affects the L0s Exit Latency reported in LCAP[14:12] and the N_FTS value advertised during link training. See PEGL0SLAT at offset 22Ch.	0h	RW
5	RL	Retrain Link: 0b Normal operation. 1b Full Link retraining is initiated by directing the Physical Layer LTSSM from L0, L0s, or L1 states to the Recovery state. This bit always returns 0 when read. This bit is cleared automatically (no need to write a 0). Intel Reserved Text: PEG10: This bit is set implicitly if CAPID0_B[SPEGFX1] or CAPID0_B[DPEGFX1] are set to 1b, and in addition BCTRL[VGAEN] is set to 1b. PEG11: This bit is set implicitly if CAPID0_B[DPEGFX1] is set to 1b and in addition BCTRL[VGAEN] is set to 1b. PEG12: This bit is set implicitly if CAPID0_B[DPEGFX1] is set to 1b and in addition BCTRL[VGAEN] is set to 1b.	0h	RW_V
4	LD	Link Disable: 0: Normal operation 1: Link is disabled. Forces the LTSSM to transition to the Disabled state (via Recovery) from L0, L0s, or L1 states. Link retraining happens automatically on 0 to 1 transition, just like when coming out of reset. Writes to this bit are immediately reflected in the value read from the bit, regardless of actual Link state.	0h	RW

continued...

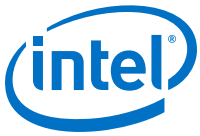

B/D/F/Type: 0/1/2/CFG			Access: RW; RO; RW_V	
Size: 16	**Default Value: 0000h**		**Address Offset: B0h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
3	RCB	Read Completion Boundary: Hardwired to 0 to indicate 64 byte.	0h	RO
2	RSVD	Reserved.	0h	RO
1:0	ASPM	Active State PM: Controls the level of active state power management supported on the given link. 00: Disabled 01: L0s Entry Supported 10: L1 Entry Supported 11: L0s and L1 Entry Supported	0h	RW

4.3.39 LSTS—Link Status

Indicates PCI Express link status.

B/D/F/Type: 0/1/2/CFG			Access: RO_V; RO; RW1C	
Size: 16	**Default Value: 1001h**		**Address Offset: B2h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15	LABWS	Link Autonomous Bandwidth Status: This bit is set to 1b by hardware to indicate that hardware has autonomously changed link speed or width, without the port transitioning through DL_Down status, for reasons other than to attempt to correct unreliable link operation. This bit must be set if the Physical Layer reports a speed or width change was initiated by the downstream component that was indicated as an autonomous change.	0h	RW1C
14	LBWMS	Link Bandwidth Management Status: This bit is set to 1b by hardware to indicate that either of the following has occurred without the port transitioning through DL_Down status: A link retraining initiated by a write of 1b to the Retrain Link bit has completed. Note: This bit is Set following any write of 1b to the Retrain Link bit, including when the Link is in the process of retraining for some other reason. Hardware has autonomously changed link speed or width to attempt to correct unreliable link operation, either through an LTSSM timeout or a higher level process This bit must be set if the Physical Layer reports a speed or width change was initiated by the downstream component that was not indicated as an autonomous change.	0h	RW1C
13	DLLLA	Data Link Layer Link Active (Optional): This bit indicates the status of the Data Link Control and Management State Machine. It returns a 1b to indicate the DL_Active state, 0b otherwise. This bit must be implemented if the corresponding Data Link Layer Active Capability bit is implemented. Otherwise, this bit must be hardwired to 0b.	0h	RO_V

continued...

B/D/F/Type: 0/1/2/CFG			Access: RO_V; RO; RW1C	
Size: 16	Default Value: 1001h		Address Offset: B2h	
Bit Range	Acronym	Description	Default	Access
12	SCC	Slot Clock Configuration: 0: The device uses an independent clock irrespective of the presence of a reference on the connector. 1: The device uses the same physical reference clock that the platform provides on the connector.	1h	RO
11	LTRN	Link Training: Indicates that the Physical Layer LTSSM is in the Configuration or Recovery state, or that 1b was written to the Retrain Link bit but Link training has not yet begun. Hardware clears this bit when the LTSSM exits the Configuration/Recovery state once Link training is complete.	0h	RO_V
10	RSVD	Reserved.	0h	RO
9:4	NLW	Negotiated Link Width: Indicates negotiated link width. This field is valid only when the link is in the L0, L0s, or L1 states (after link width negotiation is successfully completed). 00h: Reserved 01h: X1 02h: X2 04h: X4 08h: X8 10h: X16 All other encodings are reserved.	00h	RO_V
3:0	CLS	Current Link Speed: This field indicates the negotiated Link speed of the given PCI Express Link. The encoding is the binary value of the bit location in the Supported Link Speeds Vector (in the Link Capabilities 2 register) that corresponds to the current Link speed. For example, a value of 0010b in this field indicates that the current Link speed is that corresponding to bit 2 in the Supported Link Speeds Vector, which is 5.0 GT/s. The value in this field is undefined when the Link is not up.	1h	RO_V

4.3.40 SLOTCAP—Slot Capabilities

PCI Express Slot related registers allow for the support of Hot Plug.

B/D/F/Type: 0/1/2/CFG			Access: RO; RW_O	
Size: 32	Default Value: 00040000h		Address Offset: B4h	
Bit Range	Acronym	Description	Default	Access
31:19	PSN	Physical Slot Number: Indicates the physical slot number attached to this Port. BIOS Requirement: This field must be initialized by BIOS to a value that assigns a slot number that is globally unique within the chassis.	0000h	RW_O
18	NCCS	No Command Completed Support: When set to 1b, this bit indicates that this slot does not generate software notification when an issued command is completed by the Hot-Plug Controller. This bit is only permitted to be set to 1b if the hotplug capable port is able to accept writes to all fields of the Slot Control register without delay between successive writes.	1h	RO

continued...

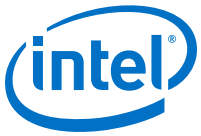

B/D/F/Type: 0/1/2/CFG			Access: RO; RW_O	
Size: 32	**Default Value: 00040000h**		**Address Offset: B4h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
17	EIP	Reserved for Electromechanical Interlock Present: When set to 1b, this bit indicates that an Electromechanical Interlock is implemented on the chassis for this slot.	0h	RO
16:15	SPLS	Slot Power Limit Scale: Specifies the scale used for the Slot Power Limit Value. 00: 1.0x 01: 0.1x 10: 0.01x 11: 0.001x If this field is written, the link sends a Set_Slot_Power_Limit message.	0h	RW_O
14:7	SPLV	Slot Power Limit Value: In combination with the Slot Power Limit Scale value, specifies the upper limit on power supplied by slot. Power limit (in Watts) is calculated by multiplying the value in this field by the value in the Slot Power Limit Scale field. If this field is written, the link sends a Set_Slot_Power_Limit message.	00h	RW_O
6	HPC	Reserved for Hot-plug Capable: When set to 1b, this bit indicates thta this slot is capable of supporting hot-plug operations.	0h	RO
5	HPS	Reserved for Hot-plug Surprise: When set to 1b, this bit indicates that an adapter present in this slot might be removed from the system without any prior notification. This is a form factor specific capability. this bit is an indication to the operating system to allow for such removal without impacting continued software operation.	0h	RO
4	PIP	Reserved for Power Indicator Present: When set to 1b, this bit indicates that a Power Indicator is electrically controlled by the chassis for this slot.	0h	RO
3	AIP	Reserved for Attention Indicator Present: When set to 1b, this bit indicates that an Attention Indicator is electrically controlled by the chassis.	0h	RO
2	MSP	Reserved for MRL Sensor Present: When set to 1b, this bit indicates that an MRL Sensor is implemented on the chassis for this slot.	0h	RO
1	PCP	Reserved for Power Controller Present: When set to 1b, this bit indicates that a software programmable Power Controller is implemented for this slot/adapter (depending on form factor).	0h	RO
0	ABP	Reserved for Attention Button Present: When set to 1b, this bit indicates that an Attention Button for this slot is electrically controlled by the chassis.	0h	RO

4.3.41 SLOTCTL—Slot Control

PCI Express Slot related registers allow for the support of Hot Plug.

B/D/F/Type: 0/1/2/CFG			Access: RO	
Size: 16	Default Value: 0000h		Address Offset: B8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:13	RSVD	Reserved.	0h	RO
12	DLLSCE	Reserved for Data Link Layer State Changed Enable: Reserved for Data Link Layer State Changed Enable (DLLSCE): If the Data Link Layer Link Active capability is implemented, when set to 1b, this field enables software notification when Data Link Layer Link Active field is changed. If the Data Link Layer Link Active capability is not implemented, this bit is permitted to be read-only with a value of 0b.	0h	RO
11	EIC	Reserved for Electromechanical Interlock Control: If an Electromechanical Interlock is implemented, a write of 1b to this field causes the state of the interlock to toggle. A write of 0b to this field has no effect. A read to this register always returns a 0.	0h	RO
10	PCC	Reserved for Power Controller Control: If a Power Controller is implemented, this field when written sets the power state of the slot per the defined encodings. Reads of this field must reflect the value from the latest write, even if the corresponding hotplug command is not complete, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined. Depending on the form factor, the power is turned on/off either to the slot or within the adapter. Note that in some cases the power controller may autonomously remove slot power or not respond to a power-up request based on a detected fault condition, independent of the Power Controller Control setting. The defined encodings are: 0b Power On 1b Power Off If the Power Controller Implemented field in the Slot Capabilities register is set to 0b, then writes to this field have no effect and the read value of this field is undefined.	0h	RO
9:8	PIC	Reserved Power Indicator Control: Reserved Power Indicator Control (PIC): If a Power Indicator is implemented, writes to this field set the Power Indicator to the written state. Reads of this field must reflect the value from the latest write, even if the corresponding hot-plug command is not complete, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined. 00: Reserved 01: On 10: Blink 11: Off If the Power Indicator Present bit in the Slot Capabilities register is 0b, this field is permitted to be read-only with a value of 00b.	0h	RO
7:6	AIC	Reserved for Attention Indicator Control: Reserved for Attention Indicator Control (AIC): If an Attention Indicator is implemented, writes to this	0h	RO

continued...

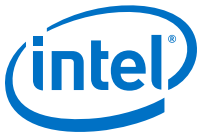

B/D/F/Type: 0/1/2/CFG			Access: RO	
Size: 16	**Default Value: 0000h**		**Address Offset: B8h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		field set the Attention Indicator to the written state. Reads of this field must reflect the value from the latest write, even if the corresponding hot-plug command is not complete, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined. If the indicator is electrically controlled by chassis, the indicator is controlled directly by the downstream port through implementation specific mechanisms. 00: Reserved 01: On 10: Blink 11: Off If the Attention Indicator Present bit in the Slot Capabilities register is 0b, this field is permitted to be readonly with a value of 00b.		
5	HPIE	Reserved for Hot-plug Interrupt Enable: When set to 1b, this bit enables generation of an interrupt on enabled hot-plug events Default value of this field is 0b. If the Hot Plug Capable field in the Slot Capabilities register is set to 0b, this bit is permitted to be read-only with a value of 0b.	0h	RO
4	CCI	Reserved for Command Completed Interrupt Enable: If Command Completed notification is supported (as indicated by No Command Completed Support field of Slot Capabilities Register), when set to 1b, this bit enables software notification when a hot-plug command is completed by the Hot-Plug Controller. Default value of this field is 0b. If Command Completed notification is not supported, this bit must be hardwired to 0b.	0h	RO
3	PDCE	Presence Detect Changed Enable: When set to 1b, this bit enables software notification on a presence detect changed event.	0h	RO
2	MSCE	Reserved for MRL Sensor Changed Enable: When set to 1b, this bit enables software notification on a MRL sensor changed event. Default value of this field is 0b. If the MRL Sensor Present field in the Slot Capabilities register is set to 0b, this bit is permitted to be read-only with a value of 0b.	0h	RO
1	PFDE	Reserved for Power Fault Detected Enable: When set to 1b, this bit enables software notification on a power fault event. Default value of this field is 0b. If Power Fault detection is not supported, this bit is permitted to be read-only with a value of 0b	0h	RO
0	ABPE	Reserved for Attention Button Pressed Enable: When set to 1b, this bit enables software notification on an attention button pressed event.	0h	RO

4.3.42 SLOTSTS—Slot Status

PCI Express Slot related registers.

B/D/F/Type: 0/1/2/CFG			Access:	RO; RW1C; RO_V
Size: 16	Default Value: 0000h		Address Offset:	BAh
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:9	RSVD	Reserved.	00h	RO
8	DLLSC	Reserved for Data Link Layer State Changed: This bit is set when the value reported in the Data Link Layer Link Active field of the Link Status register is changed. In response to a Data Link Layer State Changed event, software must read the Data Link Layer Link Active field of the Link Status register to determine if the link is active before initiating configuration cycles to the hot plugged device.	0h	RO
7	EIS	Reserved for Electromechanical Interlock Status: If an Electromechanical Interlock is implemented, this bit indicates the current status of the Electromechanical Interlock. Defined encodings are: 0b Electromechanical Interlock Disengaged 1b Electromechanical Interlock Engaged	0h	RO
6	PDS	Presence Detect State: --In band presence detect state: 0b: Slot Empty 1b: Card present in slot This bit indicates the presence of an adapter in the slot, reflected by the logical "OR" of the Physical Layer in-band presence detect mechanism and, if present, any out-of-band presence detect mechanism defined for the slot's corresponding form factor. Note that the in-band presence detect mechanism requires that power be applied to an adapter for its presence to be detected. Consequently, form factors that require a power controller for hot-plug must implement a physical pin presence detect mechanism. Defined encodings are: 0b Slot Empty 1b Card Present in slot This register must be implemented on all Downstream Ports that implement slots. For Downstream Ports not connected to slots (where the Slot Implemented bit of the PCI Express Capabilities Register is 0b), this bit must return 1b.	0h	RO_V
5	MSS	Reserved for MRL Sensor State: This register reports the status of the MRL sensor if it is implemented. Defined encodings are: 0b MRL Closed 1b MRL Open	0h	RO
4	CC	Reserved for Command Completed: If Command Completed notification is supported (as indicated by No Command Completed Support field of Slot Capabilities Register), this bit is set when a hot-plug command has completed and the Hot-Plug Controller is ready to accept a subsequent command. The Command Completed status bit is set as an indication to host software that the Hot-Plug Controller has processed the previous command and is ready to receive the next command; it provides no guarantee that the action corresponding to the command	0h	RO

continued...

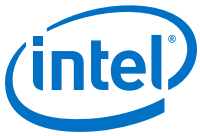

B/D/F/Type: 0/1/2/CFG			Access: RO; RW1C; RO_V	
Size: 16	Default Value: 0000h		Address Offset: BAh	
Bit Range	Acronym	Description	Default	Access
		is complete. If Command Completed notification is not supported, this bit must be hardwired to 0b.		
3	PDC	Presence Detect Changed: --A pulse indication that the inband presence detect state has changed This bit is set when the value reported in Presence Detect State is changed.	0h	RW1C
2	MSC	Reserved for MRL Sensor Changed: If an MRL sensor is implemented, this bit is set when a MRL Sensor state change is detected. If an MRL sensor is not implemented, this bit must not be set.	0h	RO
1	PFD	Reserved for Power Fault Detected: If a Power Controller that supports power fault detection is implemented, this bit is set when the Power Controller detects a power fault at this slot. Note that, depending on hardware capability, it is possible that a power fault can be detected at any time, independent of the Power Controller Control setting or the occupancy of the slot. If power fault detection is not supported, this bit must not be set.	0h	RO
0	ABP	Reserved for Attention Button Pressed: If an Attention Button is implemented, this bit is set when the attention button is pressed. If an Attention Button is not supported, this bit must not be set.	0h	RO

4.3.43 RCTL—Root Control

Allows control of PCI Express Root Complex specific parameters. The system error control bits in this register determine if corresponding SERRs are generated when our device detects an error (reported in this device's Device Status register) or when an error message is received across the link. Reporting of SERR as controlled by these bits takes precedence over the SERR Enable in the PCI Command Register.

B/D/F/Type: 0/1/2/CFG			Access: RW; RO; RWS_V; RO_V	
Size: 32	Default Value: 00000000h		Address Offset: BCh	
Bit Range	Acronym	Description	Default	Access
31:19	RSVD	Reserved.	0000h	RO
18	INSWL2	this bit indicates whether the current state of the LTSSM is L2, and it got there through the SWL2 flow. This bit is part of the RTD3 flow ,and may be used as indication for initiate power savings measures.	0h	RO_V
17:9	RSVD	Reserved.	000h	RO
8	SWL2DIS	SWL2 disable bit, when set writing to bits RCTL.L23ER or RCTL.L22DT wont have any affect over H/W	0h	RWS_V
7	RSVD	Reserved.	0h	RO

continued...

B/D/F/Type: 0/1/2/CFG			Access: RW; RO; RWS_V; RO_V	
Size: 32	**Default Value: 00000000h**		**Address Offset: BCh**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
6	L22DT	L23 to Detect Transition: When set by software, the LTSSM moves from L2 to Detect state. This bit is cleared by hardware once the LTSSM is in detect (and in general, while not in L2 state)	0h	RWS_V
5	L23ER	L23_Rdy Entry Request: When set by software, the corresponding PCIe root port will initiate the sequence to put the link into L2/L3 Ready state. PME_Turn_Off message will be sent and the corresponding PME_TO_Ack response will be returned by the device. Once the link enters L2, this bit will be cleared by hardware.	0h	RWS_V
4	CSVE	Reserved for CRS Software Visibility Enable: This bit, when set, enables the Root Port to return Configuration Request Retry Status (CRS) Completion Status to software. Root Ports that do not implement this capability must hardwire this bit to 0b.	0h	RO
3	PMEIE	PME Interrupt Enable: 0: No interrupts are generated as a result of receiving PME messages. 1: Enables interrupt generation upon receipt of a PME message as reflected in the PME Status bit of the Root Status Register. A PME interrupt is also generated if the PME Status bit of the Root Status Register is set when this bit is set from a cleared state. If the bit change from 1 to 0 and interrupt is pending than interrupt is deasserted	0h	RW
2	SEFEE	System Error on Fatal Error Enable: Controls the Root Complex's response to fatal errors. 0: No SERR generated on receipt of fatal error. 1: Indicates that an SERR should be generated if a fatal error is reported by any of the devices in the hierarchy associated with this Root Port, or by the Root Port itself.	0h	RW
1	SENFUEE	System Error on Non-Fatal Uncorrectable Error Enable: Controls the Root Complex's response to non-fatal errors. 0: No SERR generated on receipt of non-fatal error. 1: Indicates that an SERR should be generated if a non-fatal error is reported by any of the devices in the hierarchy associated with this Root Port, or by the Root Port itself.	0h	RW
0	SECEE	System Error on Correctable Error Enable: Controls the Root Complex's response to correctable errors. 0: No SERR generated on receipt of correctable error. 1: Indicates that an SERR should be generated if a correctable error is reported by any of the devices in the hierarchy associated with this Root Port, or by the Root Port itself.	0h	RW

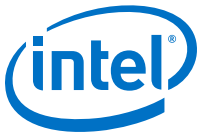

4.3.44 RSTS—Root Status

Provides information about PCI Express Root Complex specific parameters.

B/D/F/Type: 0/1/2/CFG			Access: RO_V; RW1C; RO	
Size: 32	**Default Value: 00000000h**		**Address Offset: C0h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:18	RSVD	Reserved.	0000h	RO
17	PMEP	PME Pending: Indicates that another PME is pending when the PME Status bit is set. When the PME Status bit is cleared by software; the PME is delivered by hardware by setting the PME Status bit again and updating the Requestor ID appropriately. The PME pending bit is cleared by hardware if no more PMEs are pending. Intel Reserved text: In order to simplify design and validation, if a 2nd PME arrives, before the 1st is processed, we will drop it. The initiating device must re-initiate the PME, as described in the spec so there should be no functional impact by taking this approach. This means this bit is hardwired to '0' and will never be set to '1'.	0h	RO
16	PMES	PME Status: Indicates that PME was asserted by the requestor ID indicated in the PME Requestor ID field. Subsequent PMEs are kept pending until the status register is cleared by writing a 1 to this field. An interrupt is asserted If PMEIE is asserted and PMES is changing from 0 to 1 An interrupt is deasserted If PMEIE is asserted and PMES is changing from 1 to 0 An Assert_PMEGPE is sent upstream If PMEGPEE in PEG Legacy cControl register (PEGLC) is asserted and PMES is changing from 0 to 1 An Deassert_PMEGPE is sent upstream If PMEGPEE in PEG Legacy cControl register (PEGLC) is asserted and PMES is changing from 1 to 0 An interrupt is deasserted If PMEIE is asserted and PMES is changing from 1 to 0 Intel Reserved text: In order to simplify design and validation, if a 2nd PME arrives, before the 1st is processed, we will drop it. The initiating device must re-initiate the PME, as described in the spec so there should be no functional impact by taking this approach.	0h	RW1C
15:0	PMERID	PME Requestor ID: Indicates the PCI requestor ID of the last PME requestor.	0000h	RO_V

4.3.45 DCAP2—Device Capabilites 2

B/D/F/Type: 0/1/2/CFG			Access: RO; RW_O	
Size: 32	**Default Value: 00000B80h**		**Address Offset: C4h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:20	RSVD	Reserved.	000h	RO
19:18	OBFF_SUPPORTED	OBFF Supported 00b OBFF Not Supported 01b OBFF supported using Message signaling only	0h	RW_O

continued...

B/D/F/Type: 0/1/2/CFG			Access: RO; RW_O	
Size: 32	Default Value: 00000B80h		Address Offset: C4h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		10b OBFF supported using WAKE# signaling only 11b OBFF supported using WAKE# and Message signaling The value reported in this field must indicate support for WAKE# signaling only if: - for a Downstream Port, driving the WAKE# signal for OBFF is supported and the connector or component connected Downstream is known to receive that same WAKE# signal - for an Upstream Port, receiving the WAKE# signal for OBFF is supported and, if the component is on an add-in-card, that the component is connected to the WAKE# signal on the connector. Root Ports, Switch Ports, and Endpoints are permitted to implement this capability. For a multi-Function device associated with an Upstream Port, each Function must report the same value for this field. For Bridges and Ports that do not implement this capability, this field must be hardwired to 00b.		
17:12	RSVD	Reserved.	00h	RO
11	LTRS	Latency Tolerance and BW reporting Mechanism Supported: A value of 1b indicates support for the optional Latency Tolerance & Bandwidth Requirement Reporting (LTBWR) mechanism capability. Root Ports, Switches and Endpoints are permitted to implement this capability. For Switches that implement LTBWR, this bit must be set only at the upstream port. For a multi-Function device, each Function must report the same value for this bit. For Bridges, Downstream Ports, and components that do not implement this capability, this bit must be hardwired to 0b.	1h	RO
10	RSVD	Reserved.	0h	RO
9	ATOMIC128SUP	128-bit CAS atomic operation completion support. This bit must be set to 1b if the Function supports this optional capability.	1h	RO
8	ATOMIC64SUP	64-bit atomic operation completion support. Includes FetchAdd, Swap, and CAS AtomicOps. This bit must be set to 1b if the Function supports this optional capability.	1h	RO
7	ATOMIC32SUP	32-bit atomic operation completion support. Includes FetchAdd, Swap, and CAS AtomicOps. This bit must be set to 1b if the Function supports this optional capability.	1h	RO
6	ATOMIC_OP_ROUTING_SUPPORT	Atomic Operation Routing Supported. If set then then atomic operations are supported.	0h	RO

continued...

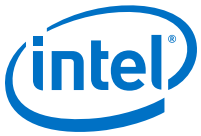

B/D/F/Type: 0/1/2/CFG			Access: RO; RW_O	
Size: 32	Default Value: 00000B80h		Address Offset: C4h	
Bit Range	Acronym	Description	Default	Access
5	ARIFS	ARI Forwarding Supported: Applicable only to Switch Downstream Ports and Root Ports; must be 0b for other Function types. This bit must be set to 1b if a Switch Downstream Port or Root Port supports this optional capability.	0h	RW_O
4	CTODS	Completion Timeout Disabled Supported: A value of 1b indicates support for the Completion Timeout Disable mechanism. The Completion Timeout Disable mechanism is required for Endpoints that issue Requests on their own behalf and PCI Express to PCI/PCI-X Bridges that take ownership of Requests issued on PCI Express. This mechanism is optional for Root Ports. The Root port does nopt support completion timeout disable	0h	RO
3:0	CTOR	Completion Timer Ranges Supported: device Function support for the optional Completion Timeout programmability mechanism. This mechanism allows system software to modify the Completion Timeout value. This field is applicable only to Root Ports, Endpoints that issue Requests on their own behalf, and PCI Express to PCI/PCI-X Bridges that take ownership of Requests issued on PCI Express. For all other Functions this field is reserved and must be hardwired to 0000b. 0000b Completion Timeout programming not supported - the Function must implement a timeout value in the range 50 us to 50 ms.	0h	RO

4.3.46 DCTL2—Device Control 2

B/D/F/Type: 0/1/2/CFG			Access: RW; RO; RW_V	
Size: 16	Default Value: 0000h		Address Offset: C8h	
Bit Range	Acronym	Description	Default	Access
15	RSVD	Reserved.	0h	RO
14:13	OBFFEN		0h	RW
12:11	RSVD	Reserved.	0h	RO
10	LTREN	Latency Tolerance Reporting Mechansim Enable: When Set to 1b, this bit enables the Latency Tolerance & Reporting (LTR) mechanism. This bit is required for all Functions that support the LTR Capability. For a Multi-Function device associated with an upstream port of a device that implements LTBWR, the bit in Function 0 is of type RW, and only Function 0 controls the components Link behavior. In all other Functions of that device, this bit is of type RsvdP. Components that do not implement LTR are permitted to hardwire this bit to 0b. Default value of this bit is 0b.	0h	RW_V

continued...

B/D/F/Type: 0/1/2/CFG			Access: RW; RO; RW_V	
Size: 16	Default Value: 0000h		Address Offset: C8h	
Bit Range	Acronym	Description	Default	Access
		This bit is cleared when the port goes to DL_down state. HW ignores the value of this bit.		
9:7	RSVD	Reserved.	0h	RO
6	ATOMIC_OP_REQUESTER_EN	AtomicOp Requester Enable Applicable only to Endpoints and Root Ports; must be hardwired to 0b for other Function types. The Function is allowed to initiate AtomicOp Requests only if this bit and the Bus Master Enable bit in the Command register are both Set. This bit is required to be RW if the Endpoint or Root Port is capable of initiating AtomicOp Requests, but otherwise is permitted to be hardwired to 0b. This bit does not serve as a capability bit. This bit is permitted to be RW even if no AtomicOp Requester capabilities are supported by the Endpoint or Root Port.	0h	RO
5	ARIFEN	ARI Forward Enable: When set, the Downstream Port disables its traditional Device Number field being 0 enforcement when turning a Type 1 Configuration Request into a Type 0 Configuration Request, permitting access to Extended Functions in an ARI Device immediately below the Port. Default value of this bit is 0b. Must be hardwired to 0b if the ARI Forwarding Supported bit is 0b.	0h	RW
4:0	RSVD	Reserved.	00h	RO

4.3.47 LCTL2—Link Control 2

B/D/F/Type: 0/1/2/CFG			Access: RWS; RWS_V	
Size: 16	Default Value: 0003h		Address Offset: D0h	
Bit Range	Acronym	Description	Default	Access
15:12	ComplianceDeemphasis	Compliance De-emphasis: For 8 GT/s Data Rate: This field sets the Transmitter Preset level in Polling.Compliance state if the entry occurred due to the Enter Compliance bit being 1b. This bit sets the de-emphasis level in Polling.Compliance state if the entry occurred due to the Enter Compliance bit being 1b. Defined encodings are: 0001b -3.5 dB 0000b -6 dB When the Link is operating at 2.5 GT/s, the setting of this bit has	0h	RWS

continued...

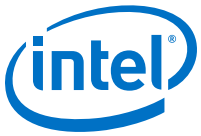

B/D/F/Type: 0/1/2/CFG			Access: RWS; RWS_V	
Size: 16	**Default Value: 0003h**		**Address Offset: D0h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		no effect. Components that support only 2.5 GT/s speed are permitted to hardwire this bit to 0b. For a Multi-Function device associated with an Upstream Port, the bit in Function 0 is of type RWS, and only Function 0 controls the component's Link behavior. In all other Functions of that device, this bit is of type RsvdP. The default value of this bit is 0000b. This bit is intended for debug, compliance testing purposes. System firmware and software is allowed to modify this bit only during debug or compliance testing.		
11	compsos	Compliance SOS: When set to 1b, the LTSSM is required to send SKP Ordered Sets periodically in between the (modified) compliance patterns. For a Multi-Function device associated with an Upstream Port, the bit in Function 0 is of type RWS, and only Function 0 controls the component's Link behavior. In all other Functions of that device, this bit is of type RsvdP. The default value of this bit is 0b. This bit is applicable when the Link is operating at 2.5 GT/s or 5 GT/s data rates only. Components that support only the 2.5 GT/s speed are permitted to hardwire this field to 0b.	0h	RWS
10	entermodcompliance	Enter Modified Compliance: When this bit is set to 1b, the device transmits modified compliance pattern if the LTSSM enters Polling.Compliance state. Components that support only the 2.5GT/s speed are permitted to hardwire this bit to 0b. Default value of this field is 0b.	0h	RWS
9:7	txmargin	Transmit Margin: This field controls the value of the non-deemphasized voltage level at the Transmitter pins. This field is reset to 000b on entry to the LTSSM Polling.Configuration substate (see Chapter 4 for details of how the transmitter voltage level is determined in various states). Encodings: 000: Normal operating range 001: 800-1200 mV for full swing and 400-700 mV for half-swing 010 - (n-1): Values must be monotonic with a non-zero slope. The value of n must be greater than 3 and less than 7. At least two of these must be below the normal operating range n : 200-400 mV for full-swing and 100-200 mV for half-swing n -111: reserved Default value is 000b. Components that support only the 2.5GT/s speed are	0h	RWS_V

continued...

B/D/F/Type: 0/1/2/CFG			Access: RWS; RWS_V	
Size: 16	Default Value: 0003h		Address Offset: D0h	
Bit Range	Acronym	Description	Default	Access
		permitted to hardwire this bit to 0b. When operating in 5GT/s mode with full swing, the deemphasis ratio must be maintained within +/- 1dB from the spec defined operational value (either -3.5 or -6 dB).		
6	selectabledeemphasis	Selectable De-emphasis: When the Link is operating at 5GT/s speed, selects the level of de-emphasis. Encodings: 1b -3.5 dB 0b -6 dB Default value is implementation specific, unless a specific value is required for a selected form factor or platform. When the Link is operating at 2.5GT/s speed, the setting of this bit has no effect. Components that support only the 2.5GT/s speed are permitted to hardwire this bit to 0b.	0h	RWS
5	HASD	Hardware Autonomous Speed Disable: When set to 1b this bit disables hardware from changing the link speed for reasons other than attempting to correct unreliable link operation by reducing link speed.	0h	RWS
4	EC	Enter Compliance: Software is permitted to force a link to enter Compliance mode at the speed indicated in the Target Link Speed field by setting this bit to 1b in both components on a link and then initiating a hot reset on the link.	0h	RWS
3:0	TLS	Target Link Speed: For Downstream Ports, this field sets an upper limit on Link operational speed by restricting the values advertised by the Upstream component in its training sequences. The encoding is the binary value of the bit in the Supported Link Speeds Vector (in the Link Capabilities 2 register) that corresponds to the desired target Link speed. All other encodings are reserved. For example, 5.0 GT/s corresponds to bit 2 in the Supported Link Speeds Vector, so the encoding for a 5.0 GT/s target Link speed in this field is 0010b. If a value is written to this field that does not correspond to a supported speed (as indicated by the Max Link Speed Vector), the result is undefined. The default value of this field is the highest Link speed supported by the component (as reported in the Max Link Speed field of the Link Capabilities register) unless the corresponding platform/form factor requires a different default value. For both Upstream and Downstream Ports, this field is used to set the target compliance mode speed when software is using the Enter Compliance bit to force a Link into compliance mode. For a Multi-Function device associated with an Upstream Port, the field in Function 0 is of type RWS, and only Function 0	3h	RWS

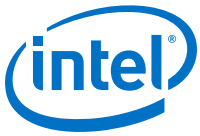

B/D/F/Type: 0/1/2/CFG			Access: RWS; RWS_V	
Size: 16	**Default Value: 0003h**		**Address Offset: D0h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		controls the components Link behavior. In all other Functions of that device, this field is of type RsvdP.		

4.3.48 LSTS2—Link Status 2

B/D/F/Type: 0/1/2/CFG			Access: RO_V; RW1C	
Size: 16	**Default Value: 0000h**		**Address Offset: D2h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:6	RSVD	Reserved.	000h	RO
5	LNKEQREQ	This bit is Set by hardware to request the Link equalization process to be performed on the Link.	0h	RW1C
4	EQPH3SUCC	Equalization Phase 3 Successful When set to 1b, this bit indicates that Phase 3 of the Transmitter Equalization procedure has successfully completed.	0h	RO_V
3	EQPH2SUCC	Equalization Phase 2 Successful When set to 1b, this bit indicates that Phase 2 of the Transmitter Equalization procedure has successfully completed.	0h	RO_V
2	EQPH1SUCC	Equalization Phase 1 Successful When set to 1b, this bit indicates that Phase 1 of the Transmitter Equalization procedure has successfully completed.	0h	RO_V
1	EQCOMPLETE	Equalization Complete When set to 1b, this bit indicates that the Transmitter Equalization procedure has completed.	0h	RO_V
0	CURDELVL	Current De-emphasis Level: Current De-emphasis Level - When the Link is operating at 5 GT/s speed, this reflects the level of de-emphasis. Encodings: 1b -3.5 dB 0b -6 dB When the Link is operating at 2.5 GT/s speed, this bit is 0b.	0h	RO_V

4.3.49 PVCCAP1—Port VC Capability Register 1

Describes the configuration of PCI Express Virtual Channels associated with this port.

B/D/F/Type: 0/1/2/CFG			Access: RO	
Size: 32	Default Value: 00000000h		Address Offset: 104h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:7	RSVD	Reserved.	0000000h	RO
6:4	LPEVCC	Low Priority Extended VC Count: Indicates the number of (extended) Virtual Channels in addition to the default VC belonging to the low-priority VC (LPVC) group that has the lowest priority with respect to other VC resources in a strict-priority VC Arbitration. The value of 0 in this field implies strict VC arbitration.	0h	RO
3	RSVD	Reserved.	0h	RO
2:0	EVCC	Extended VC Count: Indicates the number of (extended) Virtual Channels in addition to the default VC supported by the device.	0h	RO

4.3.50 PVCCAP2—Port VC Capability Register 2

Describes the configuration of PCI Express Virtual Channels associated with this port.

B/D/F/Type: 0/1/2/CFG			Access: RO	
Size: 32	Default Value: 00000000h		Address Offset: 108h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:24	VCATO	VC Arbitration Table Offset: Indicates the location of the VC Arbitration Table. This field contains the zero-based offset of the table in DQWORDS (16 bytes) from the base address of the Virtual Channel Capability Structure. A value of 0 indicates that the table is not present (due to fixed VC priority).	00h	RO
23:8	RSVD	Reserved.	0000h	RO
7:0	VCAC	Reserved for VC Arbitration Capability:	00h	RO

4.3.51 PVCCTL—Port VC Control

B/D/F/Type: 0/1/2/CFG			Access: RO; RW	
Size: 16	Default Value: 0000h		Address Offset: 10Ch	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:4	RSVD	Reserved.	000h	RO
3:1	VCAS	VC Arbitration Select: This field will be programmed by software to the only possible value as indicated in the VC Arbitration Capability field. Since there is no other VC supported than the default, this field is reserved.	0h	RW
0	VCARB	Reserved for Load VC Arbitration Table: Used for software to update the VC Arbitration Table when VC arbitration uses the VC Arbitration Table. As a VC Arbitration Table is never used by this component this field will never be used.	0h	RO

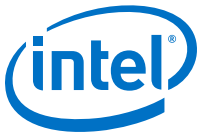

4.3.52 VC0RCAP—VC0 Resource Capability

B/D/F/Type: 0/1/2/CFG			Access: RO	
Size: 32	**Default Value: 00000001h**		**Address Offset: 110h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:24	PATO	Reserved for Port Arbitration Table Offset:	00h	RO
23	RSVD	Reserved.	0h	RO
22:16	MTS	Reserved for Maximum Time Slots:	00h	RO
15	RSNPT	Reject Snoop Transactions: Reject Snoop Transactions (RSNPT): 0: Transactions with or without the No Snoop bit set within the TLP header are allowed on this VC. 1: When Set, any transaction for which the No Snoop attribute is applicable but is not Set within the TLP Header will be rejected as an Unsupported Request	0h	RO
14:8	RSVD	Reserved.	00h	RO
7:0	PAC	Port Arbitration Capability: Port Arbitration Capability - Indicates types of Port Arbitration supported by the VC resource. This field is valid for all Switch Ports, Root Ports that support peer-to to-peer traffic, and RCRBs, but not for PCI Express Endpoint devices or Root Ports that do not support peer to peer traffic. Each bit location within this field corresponds to a Port Arbitration Capability defined below. When more than one bit in this field is Set, it indicates that the VC resource can be configured to provide different arbitration services. Software selects among these capabilities by writing to the Port Arbitration Select field (see below). Defined bit positions are: Bit 0 Non-configurable hardware-fixed arbitration scheme, e.g., Round Robin (RR) Bit 1 Weighted Round Robin (WRR) arbitration with 32 phases Bit 2 WRR arbitration with 64 phases Bit 3 WRR arbitration with 128 phases Bit 4 Time-based WRR with 128 phases Bit 5 WRR arbitration with 256 phases Bits 6-7 Reserved Processor only supported arbitration indicates "Non-configurable hardware-fixed arbitration scheme".	01h	RO

4.3.53 VC0RCTL—VC0 Resource Control

Controls the resources associated with PCI Express Virtual Channel 0.

B/D/F/Type: 0/1/2/CFG			Access: RO; RW	
Size: 32	**Default Value: 800000FFh**		**Address Offset: 114h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31	VC0E	VC0 Enable: For VC0 this is hardwired to 1 and read only as VC0 can never be disabled.	1h	RO
30:27	RSVD	Reserved.	0h	RO
26:24	VC0ID	VC0 ID: Assigns a VC ID to the VC resource. For VC0 this is hardwired to 0 and read only.	0h	RO

continued...

B/D/F/Type: 0/1/2/CFG			Access: RO; RW	
Size: 32	Default Value: 800000FFh		Address Offset: 114h	
Bit Range	Acronym	Description	Default	Access
23:20	RSVD	Reserved.	0h	RO
19:17	PAS	Port Arbitration Select: Port Arbitration Select - This field configures the VC resource to provide a particular Port Arbitration service. This field is valid for RCRBs, Root Ports that support peer to peer traffic, and Switch Ports, but not for PCI Express Endpoint devices or Root Ports that do not support peer to peer traffic. The permissible value of this field is a number corresponding to one of the asserted bits in the Port Arbitration Capability field of the VC resource. This field does not affect the root port behavior.	0h	RW
16	RSVD	Reserved.	0h	RO
15:8	TCHVC0M	TC High VC0 Map: Allow usage of high order TCs. BIOS should keep this field zeroed to allow usage of the reserved TC[3] for other purposes	00h	RW
7:1	TCVC0M	TC/VC0 Map: Indicates the TCs (Traffic Classes) that are mapped to the VC resource. Bit locations within this field correspond to TC values. For example, when bit 7 is set in this field, TC7 is mapped to this VC resource. When more than one bit in this field is set, it indicates that multiple TCs are mapped to the VC resource. In order to remove one or more TCs from the TC/VC Map of an enabled VC, software must ensure that no new or outstanding transactions with the TC labels are targeted at the given Link.	7Fh	RW
0	TC0VC0M	TC0/VC0 Map: Traffic Class 0 is always routed to VC0.	1h	RO

4.3.54 VC0RSTS—VC0 Resource Status

Reports the Virtual Channel specific status.

B/D/F/Type: 0/1/2/CFG			Access: RO_V	
Size: 16	Default Value: 0002h		Address Offset: 11Ah	
Bit Range	Acronym	Description	Default	Access
15:2	RSVD	Reserved.	0000h	RO
1	VC0NP	VC0 Negotiation Pending: 0: The VC negotiation is complete. 1: The VC resource is still in the process of negotiation (initialization or disabling). This bit indicates the status of the process of Flow Control initialization. It is set by default on Reset, as well as whenever the corresponding Virtual Channel is Disabled or the Link is in the DL_Down state. It is cleared when the link successfully exits the FC_INIT2 state. Before using a Virtual Channel, software must check whether the VC Negotiation Pending fields for that Virtual Channel are cleared in both Components on a Link.	1h	RO_V
0	RSVD	Reserved.	0h	RO

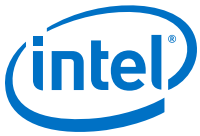

5.0 Memory Configuration Registers

5.1 DMIBAR Registers Summary

Offset	Register ID—Description	Default Value	Access
0	DMIVCECH—DMI Virtual Channel Enhanced Capability on page 240	04010002h	RO
4	DMIPVCCAP1—DMI Port VC Capability Register 1 on page 240	00000000h	RW_O; RO
8	DMIPVCCAP2—DMI Port VC Capability Register 2 on page 241	00000000h	RO
C	DMIPVCCTL—DMI Port VC Control on page 241	0000h	RO; RW
10	DMIVC0RCAP—DMI VC0 Resource Capability on page 241	00000001h	RO
14	DMIVC0RCTL—DMI VC0 Resource Control on page 242	8000007Fh	RO; RW
1A	DMIVC0RSTS—DMI VC0 Resource Status on page 242	0002h	RO_V
1C	DMIVC1RCAP—DMI VC1 Resource Capability on page 243	00008001h	RO
20	DMIVC1RCTL—DMI VC1 Resource Control on page 243	01000000h	RO; RW
26	DMIVC1RSTS—DMI VC1 Resource Status on page 244	0002h	RO_V
28	DMIVCPRCAP—DMI VCp Resource Capability on page 245	00000001h	RO
2C	DMIVCPRCTL—DMI VCp Resource Control on page 245	02000000h	RO; RW
32	DMIVCPRSTS—DMI VCp Resource Status on page 246	0002h	RO_V
34	DMIVCMRCAP—DMI VCm Resource Capability on page 247	00008000h	RO
38	DMIVCMRCTL—DMI VCm Resource Control on page 247	07000080h	RO; RW
3E	DMIVCMRSTS—DMI VCm Resource Status on page 248	0002h	RO_V
40	DMIRCLDECH—DMI Root Complex Link Declaration on page 248	08010005h	RO
44	DMIESD—DMI Element Self Description on page 249	01000202h	RO; RW_O
50	DMILE1D—DMI Link Entry 1 Description on page 249	00000000h	RW_O; RO
58	DMILE1A—DMI Link Entry 1 Address on page 250	00000000h	RW_O
5C	DMILUE1A—DMI Link Upper Entry 1 Address on page 250	00000000h	RW_O
60	DMILE2D—DMI Link Entry 2 Description on page 251	00000000h	RW_O; RO
68	DMILE2A—DMI Link Entry 2 Address on page 251	00000000h	RW_O
84	LCAP—Link Capabilities on page 251	0041AC42h	RW_OV; RO; RW_O
88	LCTL—Link Control on page 253	0000h	RW; RW_V
8A	LSTS—DMI Link Status on page 253	0001h	RO_V
98	LCTL2—Link Control 2 on page 254	0002h	RWS; RWS_V
9A	LSTS2—Link Status 2 on page 257	0000h	RO_V; RW1C

continued...

Offset	Register ID—Description	Default Value	Access
1C4	DMIUESTS—DMI Uncorrectable Error Status on page 257	00000000h	RW1CS
1C8	DMIUEMSK—DMI Uncorrectable Error Mask on page 258	00000000h	RWS
1CC	DMIUESEV—DMI Uncorrectable Error Severity on page 259	00060010h	RWS; RO
1D0	DMICESTS—DMI Correctable Error Status on page 259	00000000h	RW1CS
1D4	DMICEMSK—DMI Correctable Error Mask on page 260	00002000h	RWS

5.1.1 DMIVCECH—DMI Virtual Channel Enhanced Capability

Indicates DMI Virtual Channel capabilities.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO	
Size: 32	Default Value: 04010002h		Address Offset: 0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:20	PNC	Pointer to Next Capability: This field contains the offset to the next PCI Express capability structure in the linked list of capabilities (Link Declaration Capability).	040h	RO
19:16	PCIEVCCV	PCI Express Virtual Channel Capability Version: Hardwired to 1 to indicate compliances with the 1.1 version of the PCI Express specification. Note: This version does not change for 2.0 compliance.	1h	RO
15:0	ECID	Extended Capability ID: Value of 0002h identifies this linked list item (capability structure) as being for PCI Express Virtual Channel registers.	0002h	RO

5.1.2 DMIPVCCAP1—DMI Port VC Capability Register 1

Describes the configuration of PCI Express Virtual Channels associated with this port.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RW_O; RO	
Size: 32	Default Value: 00000000h		Address Offset: 4h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:7	RSVD	Reserved.	0000000h	RO
6:4	LPEVCC	Low Priority Extended VC Count: Indicates the number of (extended) Virtual Channels in addition to the default VC belonging to the low-priority VC (LPVC) group that has the lowest priority with respect to other VC resources in a strict-priority VC Arbitration. The value of 0 in this field implies strict VC arbitration.	0h	RO
3	RSVD	Reserved.	0h	RO
2:0	EVCC	Extended VC Count: Indicates the number of (extended) Virtual Channels in addition to the default VC supported by the device. The Private Virtual Channel, VC1 and the Manageability Virtual Channel are not included in this count.	0h	RW_O

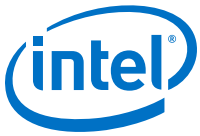

5.1.3 DMIPVCCAP2—DMI Port VC Capability Register 2

Describes the configuration of PCI Express Virtual Channels associated with this port.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO	
Size: 32	Default Value: 00000000h		Address Offset: 8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:24	VCATO	Reserved for VC Arbitration Table Offset:	00h	RO
23:8	RSVD	Reserved.	0000h	RO
7:0	VCAC	Reserved for VC Arbitration Capability:	00h	RO

5.1.4 DMIPVCCTL—DMI Port VC Control

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO; RW	
Size: 16	Default Value: 0000h		Address Offset: Ch	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:4	RSVD	Reserved.	000h	RO
3:1	VCAS	VC Arbitration Select: This field will be programmed by software to the only possible value as indicated in the VC Arbitration Capability field. The value 000b when written to this field will indicate the VC arbitration scheme is hardware fixed (in the root complex). This field cannot be modified when more than one VC in the LPVC group is enabled. 000: Hardware fixed arbitration scheme. E.G. Round Robin Others: Reserved See the PCI express specification for more details.	0h	RW
0	LVCAT	Reserved for Load VC Arbitration Table:	0h	RO

5.1.5 DMIVC0RCAP—DMI VC0 Resource Capability

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO	
Size: 32	Default Value: 00000001h		Address Offset: 10h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:24	PATO	Reserved for Port Arbitration Table Offset:	00h	RO
23	RSVD	Reserved.	0h	RO
22:16	MTS	Reserved for Maximum Time Slots:	00h	RO
15	REJSNPT	Reject Snoop Transactions: 0: Transactions with or without the No Snoop bit set within the TLP header are allowed on this VC. 1: Any transaction for which the No Snoop attribute is applicable but is not set within the TLP Header will be rejected as an Unsupported Request.	0h	RO
14:8	RSVD	Reserved.	00h	RO
7:0	PAC	Port Arbitration Capability: Having only bit 0 set indicates that the only supported arbitration scheme for this VC is non-configurable hardware-fixed.	01h	RO

5.1.6 DMIVC0RCTL—DMI VC0 Resource Control

Controls the resources associated with PCI Express Virtual Channel 0.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO; RW	
Size: 32	Default Value: 8000007Fh		Address Offset: 14h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31	VC0E	Virtual Channel 0 Enable: For VC0 this is hardwired to 1 and read only as VC0 can never be disabled.	1h	RO
30:27	RSVD	Reserved.	0h	RO
26:24	VC0ID	Virtual Channel 0 ID: Assigns a VC ID to the VC resource. For VC0 this is hardwired to 0 and read only.	0h	RO
23:20	RSVD	Reserved.	0h	RO
19:17	PAS	Port Arbitration Select: Configures the VC resource to provide a particular Port Arbitration service. Valid value for this field is a number corresponding to one of the asserted bits in the Port Arbitration Capability field of the VC resource. Because only bit 0 of that field is asserted. This field will always be programmed to '1'.	0h	RW
16:8	RSVD	Reserved.	000h	RO
7	TCMVC0M	Traffic Class m / Vritual Channel 0 Map:	0h	RO
6:1	TCVC0M	Traffic Class / Virtual Channel 0 Map: Indicates the TCs (Traffic Classes) that are mapped to the VC resource. Bit locations within this field correspond to TC values. For example, when bit 7 is set in this field, TC7 is mapped to this VC resource. When more than one bit in this field is set, it indicates that multiple TCs are mapped to the VC resource. In order to remove one or more TCs from the TC/VC Map of an enabled VC, software must ensure that no new or outstanding transactions with the TC labels are targeted at the given Link.	3Fh	RW
0	TC0VC0M	Traffic Class 0 / Virtual Channel 0 Map: Traffic Class 0 is always routed to VC0.	1h	RO

5.1.7 DMIVC0RSTS—DMI VC0 Resource Status

Reports the Virtual Channel specific status.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO_V	
Size: 16	Default Value: 0002h		Address Offset: 1Ah	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:2	RSVD	Reserved.	0000h	RO
1	VC0NP	Virtual Channel 0 Negotiation Pending: 0: The VC negotiation is complete. 1: The VC resource is still in the process of negotiation (initialization or disabling). This bit indicates the status of the process of Flow Control initialization. It is set by default on Reset, as well as whenever the corresponding Virtual Channel is Disabled or the Link is in the DL_Down state. It is cleared when the link successfully exits the FC_INIT2 state. BIOS Requirement: Before using a Virtual Channel,	1h	RO_V

continued...

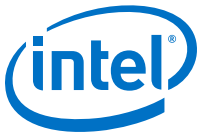

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO_V	
Size: 16	Default Value: 0002h		Address Offset: 1Ah	
Bit Range	Acronym	Description	Default	Access
		software must check whether the VC Negotiation Pending fields for that Virtual Channel are cleared in both Components on a Link.		
0	RSVD	Reserved.	0h	RO

5.1.8 DMIVC1RCAP—DMI VC1 Resource Capability

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO	
Size: 32	Default Value: 00008001h		Address Offset: 1Ch	
Bit Range	Acronym	Description	Default	Access
31:24	PATO	Reserved for Port Arbitration Table Offset:	00h	RO
23	RSVD	Reserved.	0h	RO
22:16	MTS	Reserved for Maximum Time Slots:	00h	RO
15	REJSNPT	Reject Snoop Transactions: 0: Transactions with or without the No Snoop bit set within the TLP header are allowed on this VC. 1: When Set, any transaction for which the No Snoop attribute is applicable but is not Set within the TLP Header will be rejected as an Unsupported Request.	1h	RO
14:8	RSVD	Reserved.	00h	RO
7:0	PAC	Port Arbitration Capability: Having only bit 0 set indicates that the only supported arbitration scheme for this VC is non-configurable hardware-fixed.	01h	RO

5.1.9 DMIVC1RCTL—DMI VC1 Resource Control

Controls the resources associated with PCI Express Virtual Channel 1.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO; RW	
Size: 32	Default Value: 01000000h		Address Offset: 20h	
Bit Range	Acronym	Description	Default	Access
31	VC1E	Virtual Channel 1 Enable: 0: Virtual Channel is disabled. 1: Virtual Channel is enabled. See exceptions below. Software must use the VC Negotiation Pending bit to check whether the VC negotiation is complete. When VC Negotiation Pending bit is cleared, a 1 read from this VC Enable bit indicates that the VC is enabled (Flow Control Initialization is completed for the PCI Express port). A 0 read from this bit indicates that the Virtual Channel is currently disabled. BIOS Requirement: 1. To enable a Virtual Channel, the VC Enable bits for that Virtual Channel must be set in both Components on a Link. 2. To disable a Virtual Channel, the VC Enable bits for that Virtual Channel must be cleared in both Components on a Link. 3. Software must ensure that no traffic is using a Virtual	0h	RW

continued...

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO; RW	
Size: 32	Default Value: 01000000h		Address Offset: 20h	
Bit Range	Acronym	Description	Default	Access
		Channel at the time it is disabled. 4. Software must fully disable a Virtual Channel in both Components on a Link before re-enabling the Virtual Channel.		
30:27	RSVD	Reserved.	0h	RO
26:24	VC1ID	Virtual Channel 1 ID: Assigns a VC ID to the VC resource. Assigned value must be non-zero. This field can not be modified when the VC is already enabled.	1h	RW
23:20	RSVD	Reserved.	0h	RO
19:17	PAS	Port Arbitration Select: Configures the VC resource to provide a particular Port Arbitration service. Valid value for this field is a number corresponding to one of the asserted bits in the Port Arbitration Capability field of the VC resource.	0h	RW
16:8	RSVD	Reserved.	000h	RO
7	TCMVC1M	Traffic Class m / Virtual Channel 1:	0h	RO
6:1	TCVC1M	Traffic Class / Virtual Channel 1 Map: Indicates the TCs (Traffic Classes) that are mapped to the VC resource. Bit locations within this field correspond to TC values. For example, when bit 6 is set in this field, TC6 is mapped to this VC resource. When more than one bit in this field is set, it indicates that multiple TCs are mapped to the VC resource. In order to remove one or more TCs from the TC/VC Map of an enabled VC, software must ensure that no new or outstanding transactions with the TC labels are targeted at the given Link. BIOS Requirement: Program this field with the value 010001b, which maps TC1 and TC5 to VC1.	00h	RW
0	TC0VC1M	Traffic Class 0 / Virtual Channel 1 Map: Traffic Class 0 is always routed to VC0.	0h	RO

5.1.10 DMIVC1RSTS—DMI VC1 Resource Status

Reports the Virtual Channel specific status.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO_V	
Size: 16	Default Value: 0002h		Address Offset: 26h	
Bit Range	Acronym	Description	Default	Access
15:2	RSVD	Reserved.	0000h	RO
1	VC1NP	Virtual Channel 1 Negotiation Pending: 0: The VC negotiation is complete. 1: The VC resource is still in the process of negotiation (initialization or disabling). Software may use this bit when enabling or disabling the VC. This bit indicates the status of the process of Flow Control initialization. It is set by default on Reset, as well as whenever the corresponding Virtual Channel is Disabled or the Link is in the DL_Down state. It is cleared when the link successfully exits the FC_INIT2 state.	1h	RO_V

continued...

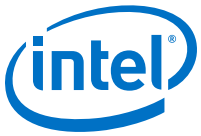

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO_V	
Size: 16	Default Value: 0002h		Address Offset: 26h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		Before using a Virtual Channel, software must check whether the VC Negotiation Pending fields for that Virtual Channel are cleared in both Components on a Link.		
0	RSVD	Reserved.	0h	RO

5.1.11 DMIVCPRCAP—DMI VCp Resource Capability

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO	
Size: 32	Default Value: 00000001h		Address Offset: 28h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:24	PATO	Reserved for Port Arbitration Table Offset:	00h	RO
23	RSVD	Reserved.	0h	RO
22:16	MTS	Reserved for Maximum Time Slots:	00h	RO
15	REJSNPT	Reject Snoop Transactions: 0: Transactions with or without the No Snoop bit set within the TLP header are allowed on this VC. 1: When Set, any transaction for which the No Snoop attribute is applicable but is not Set within the TLP Header will be rejected as an Unsupported Request.	0h	RO
14:8	RSVD	Reserved.	00h	RO
7:0	PAC	Reserved for Port Arbitration Capability:	01h	RO

5.1.12 DMIVCPRCTL—DMI VCp Resource Control

Controls the resources associated with the DMI Private Channel (VCp).

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO; RW	
Size: 32	Default Value: 02000000h		Address Offset: 2Ch	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31	VCPE	Virtual Channel private Enable: 0: Virtual Channel is disabled. 1: Virtual Channel is enabled. See exceptions below. Software must use the VC Negotiation Pending bit to check whether the VC negotiation is complete. When VC Negotiation Pending bit is cleared, a 1 read from this VC Enable bit indicates that the VC is enabled (Flow Control Initialization is completed for the PCI Express port). A 0 read from this bit indicates that the Virtual Channel is currently disabled. BIOS Requirement: 1. To enable a Virtual Channel, the VC Enable bits for that Virtual Channel must be set in both Components on a Link. 2. To disable a Virtual Channel, the VC Enable bits for that Virtual Channel must be cleared in both Components on a Link. 3. Software must ensure that no traffic is using a Virtual Channel at the time it is disabled.	0h	RW

continued...

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO; RW	
Size: 32	Default Value: 02000000h		Address Offset: 2Ch	
Bit Range	Acronym	Description	Default	Access
		4. Software must fully disable a Virtual Channel in both Components on a Link before re-enabling the Virtual Channel.		
30:27	RSVD	Reserved.	0h	RO
26:24	VCPID	Virtual Channel private ID: Assigns a VC ID to the VC resource. This field can not be modified when the VC is already enabled.	2h	RW
23:8	RSVD	Reserved.	0000h	RO
7	TCMVCPM	Traffic Class m / Vritual Channel private Map:	0h	RO
6:1	TCVCPM	Traffic Class / Virtual Channel private Map: It is recommended that private TC6 (01000000b) is the only value that should be programmed into this field for VCp traffic which will be translated by a virtualization engine, and TC2 (00000010b) is the only value that should be programmed into this field for VCp traffic which will not be translated by a virtualization engine. This strategy can simplify debug and limit validation permutations. BIOS Requirement: Program this field with the value 100010b, which maps TC2 and TC6 to VCp.	00h	RW
0	TC0VCPM	Tc0 VCp Map:	0h	RO

5.1.13 DMIVCPRSTS—DMI VCp Resource Status

Reports the Virtual Channel specific status.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO_V	
Size: 16	Default Value: 0002h		Address Offset: 32h	
Bit Range	Acronym	Description	Default	Access
15:2	RSVD	Reserved.	0000h	RO
1	VCPNP	Virtual Channel private Negotiation Pending: 0: The VC negotiation is complete. 1: The VC resource is still in the process of negotiation (initialization or disabling). Software may use this bit when enabling or disabling the VC. This bit indicates the status of the process of Flow Control initialization. It is set by default on Reset, as well as whenever the corresponding Virtual Channel is Disabled or the Link is in the DL_Down state. It is cleared when the link successfully exits the FC_INIT2 state. Before using a Virtual Channel, software must check whether the VC Negotiation Pending fields for that Virtual Channel are cleared in both Components on a Link.	1h	RO_V
0	RSVD	Reserved.	0h	RO

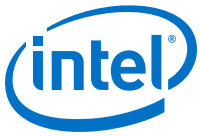

5.1.14 DMIVCMRCAP—DMI VCm Resource Capability

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO	
Size: 32	Default Value: 00008000h		Address Offset: 34h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15	REJSNPT	Reject Snoop Transactions: 0: Transactions with or without the No Snoop bit set within the TLP header are allowed on the VC. 1: When Set, any transaction for which the No Snoop attribute is applicable but is not Set within the TLP Header will be rejected as an Unsupported Request	1h	RO
14:0	RSVD	Reserved.	0000h	RO

5.1.15 DMIVCMRCTL—DMI VCm Resource Control

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO; RW	
Size: 32	Default Value: 07000080h		Address Offset: 38h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31	VCMEN	Virtual Channel enable: 0: Virtual Channel is disabled. 1: Virtual Channel is enabled. See exceptions below. Software must use the VC Negotiation Pending bit to check whether the VC negotiation is complete. When VC Negotiation Pending bit is cleared, a 1 read from this VC Enable bit indicates that the VC is enabled (Flow Control Initialization is completed for the PCI Express port). A 0 read from this bit indicates that the Virtual Channel is currently disabled. BIOS Requirement: 1. To enable a Virtual Channel, the VC Enable bits for that Virtual Channel must be set in both Components on a Link. 2. To disable a Virtual Channel, the VC Enable bits for that Virtual Channel must be cleared in both Components on a Link. 3. Software must ensure that no traffic is using a Virtual Channel at the time it is disabled. 4. Software must fully disable a Virtual Channel in both Components on a Link before re-enabling the Virtual Channel.	0h	RW
30:27	RSVD	Reserved.	0h	RO
26:24	VCID	Virtual Channel ID: Assigns a VC ID to the VC resource. Assigned value must be non-zero. This field can not be modified when the VC is already enabled.	7h	RW
23:8	RSVD	Reserved.	0000h	RO
7:0	TCVCMMAP	Traffic Class/Virtual Channel Map: Indicates the TCs (Traffic Classes) that are mapped to the VC resource. Bit locations within this field correspond to TC values. For example, when bit 7 is set in this field, TC7 is mapped to this VC resource. When more than one bit in this field is set, it indicates that multiple TCs are mapped to the VC resource. In order to remove one or more TCs from the TC/VC Map of an enabled VC, software must ensure that no new or outstanding transactions with the TC labels are targeted at the given Link.	80h	RO

5.1.16 DMIVCMRSTS—DMI VCm Resource Status

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO_V	
Size: 16	Default Value: 0002h		Address Offset: 3Eh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:2	RSVD	Reserved.	0000h	RO
1	VCNEGPND	Virtual Channel Negotiation Pending: 0: The VC negotiation is complete. 1: The VC resource is still in the process of negotiation (initialization or disabling). Software may use this bit when enabling or disabling the VC. This bit indicates the status of the process of Flow Control initialization. It is set by default on Reset, as well as whenever the corresponding Virtual Channel is Disabled or the Link is in the DL_Down state. It is cleared when the link successfully exits the FC_INIT2 state. Before using a Virtual Channel, software must check whether the VC Negotiation Pending fields for that Virtual Channel are cleared in both Components on a Link.	1h	RO_V
0	RSVD	Reserved.	0h	RO

5.1.17 DMIRCLDECH—DMI Root Complex Link Declaration

This capability declares links from the respective element to other elements of the root complex component to which it belongs and to an element in another root complex component. See PCI Express specification for link/topology declaration requirements.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO	
Size: 32	Default Value: 08010005h		Address Offset: 40h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:20	PNC	Pointer to Next Capability: This field contains the offset to the next PCI Express capability structure in the linked list of capabilities (Internal Link Control Capability).	080h	RO
19:16	LDCV	Link Declaration Capability Version: Hardwired to 1 to indicate compliances with the 1.1 version of the PCI Express specification. Note: This version does not change for 2.0 compliance.	1h	RO
15:0	ECID	Extended Capability ID: Value of 0005h identifies this linked list item (capability structure) as being for PCI Express Link Declaration Capability.	0005h	RO

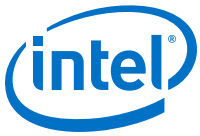

5.1.18 DMIESD—DMI Element Self Description

Provides information about the root complex element containing this Link Declaration Capability.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO; RW_O	
Size: 32	**Default Value: 01000202h**		**Address Offset: 44h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:24	PORTNUM	Port Number: Specifies the port number associated with this element with respect to the component that contains this element. This port number value is utilized by the egress port of the component to provide arbitration to this Root Complex Element.	01h	RO
23:16	CID	Component ID: Identifies the physical component that contains this Root Complex Element. BIOS Requirement: Must be initialized according to guidelines in the PCI Express* Isochronous/Virtual Channel Support Hardware Programming Specification (HPS).	00h	RW_O
15:8	NLE	Number of Link Entries: Indicates the number of link entries following the Element Self Description. This field reports 2 (one for MCH egress port to main memory and one to egress port belonging to ICH on other side of internal link).	02h	RO
7:4	RSVD	Reserved.	0h	RO
3:0	ETYP	Element Type: Indicates the type of the Root Complex Element. Value of 2h represents an Internal Root Complex Link (DMI).	2h	RO

5.1.19 DMILE1D—DMI Link Entry 1 Description

First part of a Link Entry which declares an internal link to another Root Complex Element.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RW_O; RO	
Size: 32	**Default Value: 00000000h**		**Address Offset: 50h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:24	TPN	Target Port Number: Specifies the port number associated with the element targeted by this link entry (egress port of PCH). The target port number is with respect to the component that contains this element as specified by the target component ID. This can be programmed by BIOS, but the default value will likely be correct because the DMI RCRB in the PCH will likely be associated with the default egress port for the PCH meaning it will be assigned port number 0.	00h	RW_O
23:16	TCID	Target Component ID: Identifies the physical component that is targeted by this link entry. BIOS Requirement: Must be initialized according to guidelines in the PCI Express* Isochronous/Virtual Channel Support Hardware Programming Specification (HPS).	00h	RW_O

continued...

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RW_O; RO	
Size: 32	Default Value: 00000000h		Address Offset: 50h	
Bit Range	Acronym	Description	Default	Access
15:2	RSVD	Reserved.	0000h	RO
1	LTYP	Link Type: Indicates that the link points to memory-mapped space (for RCRB). The link address specifies the 64-bit base address of the target RCRB.	0h	RO
0	LV	Link Valid: 0: Link Entry is not valid and will be ignored. 1: Link Entry specifies a valid link.	0h	RW_O

5.1.20 DMILE1A—DMI Link Entry 1 Address

Second part of a Link Entry which declares an internal link to another Root Complex Element.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RW_O	
Size: 32	Default Value: 00000000h		Address Offset: 58h	
Bit Range	Acronym	Description	Default	Access
31:12	LA	Link Address: Memory mapped base address of the RCRB that is the target element (egress port of PCH) for this link entry.	00000h	RW_O
11:0	RSVD	Reserved.	000h	RO

5.1.21 DMILUE1A—DMI Link Upper Entry 1 Address

Second part of a Link Entry which declares an internal link to another Root Complex Element.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RW_O	
Size: 32	Default Value: 00000000h		Address Offset: 5Ch	
Bit Range	Acronym	Description	Default	Access
31:8	RSVD	Reserved.	000000h	RO
7:0	ULA	Upper Link Address: Memory mapped base address of the RCRB that is the target element (egress port of PCH) for this link entry.	00h	RW_O

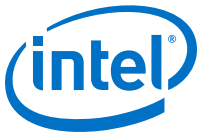

5.1.22 DMILE2D—DMI Link Entry 2 Description

First part of a Link Entry which declares an internal link to another Root Complex Element.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access:	RW_O; RO
Size: 32	Default Value: 00000000h		Address Offset:	60h
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:24	TPN	Target Port Number: Specifies the port number associated with the element targeted by this link entry (Egress Port). The target port number is with respect to the component that contains this element as specified by the target component ID.	00h	RO
23:16	TCID	Target Component ID: Identifies the physical or logical component that is targeted by this link entry. BIOS Requirement: Must be initialized according to guidelines in the PCI Express* Isochronous/Virtual Channel Support Hardware Programming Specification (HPS).	00h	RW_O
15:2	RSVD	Reserved.	0000h	RO
1	LTYP	Link Type: Indicates that the link points to memory-mapped space (for RCRB). The link address specifies the 64-bit base address of the target RCRB.	0h	RO
0	LV	Link Valid: 0: Link Entry is not valid and will be ignored. 1: Link Entry specifies a valid link.	0h	RW_O

5.1.23 DMILE2A—DMI Link Entry 2 Address

Second part of a Link Entry which declares an internal link to another Root Complex Element.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access:	RW_O
Size: 32	Default Value: 00000000h		Address Offset:	68h
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:12	LA	Link Address: Memory mapped base address of the RCRB that is the target element (Egress Port) for this link entry.	00000h	RW_O
11:0	RSVD	Reserved.	000h	RO

5.1.24 LCAP—Link Capabilities

Indicates DMI specific capabilities.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access:	RW_OV; RO; RW_O
Size: 32	Default Value: 0041AC42h		Address Offset:	84h
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:23	RSVD	Reserved.	000h	RO
22	ASPM_OPT_COMPLIANCE	ASPM Optionality Compliance. This bit must be set to 1b in all Functions. Components implemented against certain earlier	1h	RO

continued...

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RW_OV; RO; RW_O	
Size: 32	**Default Value: 0041AC42h**		**Address Offset: 84h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		versions of this specification will have this bit set to 0b. Software is permitted to use the value of this bit to help determine whether to enable ASPM or whether to run ASPM compliance tests.		
21:18	RSVD	Reserved.	0h	RO
17:15	L1SELAT	L1 Exit Latency: Indicates the length of time this Port requires to complete the transition from L1 to L0. The value 010b indicates the range of 2 us to less than 4 us. 000: Less than 1 us 001: 1 us to less than 2 us 010: 2 us to less than 4 us 011: 4 us to less than 8 us 100: 8 us to less than 16 us 101: 16 us to less than 32 us 110: 32 us-64 us 111: More than 64 us Both bytes of this register that contain a portion of this field must be written simultaneously in order to prevent an intermediate (and undesired) value from ever existing.	3h	RW_O
14:12	L0SELAT	L0s Exit Latency: Indicates the length of time this Port requires to complete the transition from L0s to L0. 000: Less than 64 ns 001: 64 ns to less than 128 ns 010: 128 ns to less than 256 ns 011: 256 ns to less than 512 ns 100: 512 ns to less than 1 us 101: 1 us to less than 2 us 110: 2 us-4 us 111: More than 4 us	2h	RW_O
11:10	ASLPMS	Active State Link PM Support: L0s & L1 entry supported.	3h	RO
9:4	MLW	Indicates the maximum number of lanes supported for this link.	04h	RO
3:0	MLS	This default value reflects gen1. Later the field may be changed by BIOS to allow gen2 subject to Fuse enabled. Defined encodings are: 0001b 2.5 GT/s Link speed supported 0010b 5.0 GT/s and 2.5 GT/s Link speeds supported	2h	RW_OV

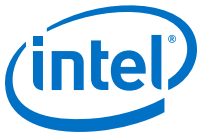

5.1.25 LCTL—Link Control

Allows control of PCI Express link.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RW; RW_V	
Size: 16	**Default Value: 0000h**		**Address Offset: 88h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:10	RSVD	Reserved.	00h	RO
9	HAWD	OPI - N/A Hardware Autonomous Width Disable: Hardware Autonomous Width Disable - When Set, this bit disables hardware from changing the Link width for reasons other than attempting to correct unreliable Link operation by reducing Link width. Devices that do not implement the ability autonomously to change Link width are permitted to hardwire this bit to 0b.	0h	RW
8	RSVD	Reserved.	0h	RO
7	ES	OPI - N/A Extended Synch: Extended synch 0: Standard Fast Training Sequence (FTS). 1: Forces the transmission of additional ordered sets when exiting the L0s state and when in the Recovery state. This mode provides external devices (e.g., logic analyzers) monitoring the Link time to achieve bit and symbol lock before the link enters L0 and resumes communication. This is a test mode only and may cause other undesired side effects such as buffer overflows or underruns.	0h	RW
6	RSVD	Reserved.	0h	RO
5	RL	Retrain Link: 0: Normal operation. 1: Full Link retraining is initiated by directing the Physical Layer LTSSM from L0, L0s, or L1 states to the Recovery state. This bit always returns 0 when read. This bit is cleared automatically (no need to write a 0).	0h	RW_V
4:2	RSVD	Reserved.	0h	RO
1:0	ASPM	Active State PM: Controls the level of active state power management supported on the given link. 00: Disabled 01: L0s Entry Supported 10: L1 Entry Supported 11: L0s and L1 Entry Supported	0h	RW

5.1.26 LSTS—DMI Link Status

Indicates DMI status.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO_V	
Size: 16	**Default Value: 0001h**		**Address Offset: 8Ah**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:12	RSVD	Reserved.	0h	RO
11	LTRN	Link Training: Indicates that the Physical Layer LTSSM is in the Configuration or Recovery state, or that 1b was written to the Retrain Link bit but Link training has not yet begun.	0h	RO_V

continued...

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO_V	
Size: 16	Default Value: 0001h		Address Offset: 8Ah	
Bit Range	Acronym	Description	Default	Access
		Hardware clears this bit when the LTSSM exits the Configuration/Recovery state once Link training is complete.		
10	RSVD	Reserved.	0h	RO
9:4	NWID	Negotiated Width: Indicates negotiated link width. This field is valid only when the link is in the L0, L0s, or L1 states (after link width negotiation is successfully completed). 00h: Reserved 01h: X1 02h: X2 04h: X4 All other encodings are reserved.	00h	RO_V
3:0	NSPD	Negotiated Speed: Indicates negotiated link speed. 1h: 2.5 Gb/s 2h: 5.0 Gb/s All other encodings are reserved. The value in this field is undefined when the Link is not up.	1h	RO_V

5.1.27 LCTL2—Link Control 2

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RWS; RWS_V	
Size: 16	Default Value: 0002h		Address Offset: 98h	
Bit Range	Acronym	Description	Default	Access
15:12	ComplianceDeemphasis	Compliance De-emphasis: For 8 GT/s Data Rate: This field sets the Transmitter Preset level in Polling.Compliance state if the entry occurred due to the Enter Compliance bit being 1b. This bit sets the de-emphasis level in Polling.Compliance state if the entry occurred due to the Enter Compliance bit being 1b. Defined encodings are: 0001b -3.5 dB 0000b -6 dB When the Link is operating at 2.5 GT/s, the setting of this bit has no effect. Components that support only 2.5 GT/s speed are permitted to hardwire this bit to 0b. For a Multi-Function device associated with an Upstream Port, the bit in Function 0 is of type RWS, and only Function 0 controls the component's Link behavior. In all other Functions of that device, this bit is of type RsvdP. The default value of this bit is 0000b. This bit is intended for debug, compliance testing purposes.	0h	RWS

continued...

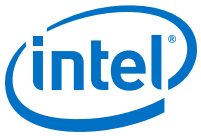

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RWS; RWS_V	
Size: 16	**Default Value: 0002h**		**Address Offset: 98h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		System firmware and software is allowed to modify this bit only during debug or compliance testing.		
11	compsos	Compliance SOS: When set to 1b, the LTSSM is required to send SKP Ordered Sets periodically in between the (modified) compliance patterns. For a Multi-Function device associated with an Upstream Port, the bit in Function 0 is of type RWS, and only Function 0 controls the component's Link behavior. In all other Functions of that device, this bit is of type RsvdP. The default value of this bit is 0b. This bit is applicable when the Link is operating at 2.5 GT/s or 5 GT/s data rates only. Components that support only the 2.5 GT/s speed are permitted to hardwire this field to 0b.	0h	RWS
10	entermodcompliance	Enter Modified Compliance: When this bit is set to 1b, the device transmits modified compliance pattern if the LTSSM enters Polling.Compliance state. Components that support only the 2.5GT/s speed are permitted to hardwire this bit to 0b. Default value of this field is 0b.	0h	RWS
9:7	txmargin	Transmit Margin: This field controls the value of the non-deemphasized voltage level at the Transmitter pins. This field is reset to 000b on entry to the LTSSM Polling.Configuration substate (see Chapter 4 for details of how the transmitter voltage level is determined in various states). Encodings: 000: Normal operating range 001: 800-1200 mV for full swing and 400-700 mV for half-swing 010 - (n-1): Values must be monotonic with a non-zero slope. The value of n must be greater than 3 and less than 7. At least two of these must be below the normal operating range n : 200-400 mV for full-swing and 100-200 mV for half-swing n -111: reserved Default value is 000b. Components that support only the 2.5GT/s speed are permitted to hardwire this bit to 0b. When operating in 5GT/s mode with full swing, the deemphasis ratio must be maintained within +/- 1dB from the spec defined operational value (either -3.5 or -6 dB).	0h	RWS_V
6	selectabledeemphasis	Selectable De-emphasis: When the Link is operating at 5GT/s speed, selects the level of de-emphasis. Encodings: 1b -3.5 dB 0b -6 dB Default value is implementation specific, unless a specific value is required for a selected form factor or platform.	0h	RWS

continued...

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RWS; RWS_V	
Size: 16	Default Value: 0002h		Address Offset: 98h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		When the Link is operating at 2.5GT/s speed, the setting of this bit has no effect. Components that support only the 2.5GT/s speed are permitted to hardwire this bit to 0b.		
5	HASD	Hardware Autonomous Speed Disable: When set to 1b this bit disables hardware from changing the link speed for reasons other than attempting to correct unreliable link operation by reducing link speed.	0h	RWS
4	EC	Enter Compliance: Software is permitted to force a link to enter Compliance mode at the speed indicated in the Target Link Speed field by setting this bit to 1b in both components on a link and then initiating a hot reset on the link.	0h	RWS
3:0	TLS	Target Link Speed: For Downstream Ports, this field sets an upper limit on Link operational speed by restricting the values advertised by the Upstream component in its training sequences. The encoding is the binary value of the bit in the Supported Link Speeds Vector (in the Link Capabilities 2 register) that corresponds to the desired target Link speed. All other encodings are reserved. For example, 5.0 GT/s corresponds to bit 2 in the Supported Link Speeds Vector, so the encoding for a 5.0 GT/s target Link speed in this field is 0010b. If a value is written to this field that does not correspond to a supported speed (as indicated by the Max Link Speed Vector), the result is undefined. The default value of this field is the highest Link speed supported by the component (as reported in the Max Link Speed field of the Link Capabilities register) unless the corresponding platform/form factor requires a different default value. For both Upstream and Downstream Ports, this field is used to set the target compliance mode speed when software is using the Enter Compliance bit to force a Link into compliance mode. For a Multi-Function device associated with an Upstream Port, the field in Function 0 is of type RWS, and only Function 0 controls the components Link behavior. In all other Functions of that device, this field is of type RsvdP.	2h	RWS

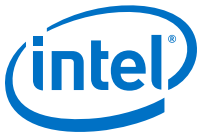

5.1.28 LSTS2—Link Status 2

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RO_V; RW1C	
Size: 16	**Default Value: 0000h**		**Address Offset: 9Ah**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:6	RSVD	Reserved.	000h	RO
5	LNKEQREQ	This bit is Set by hardware to request the Link equalization process to be performed on the Link.	0h	RW1C
4	EQPH3SUCC	Equalization Phase 3 Successful When set to 1b, this bit indicates that Phase 3 of the Transmitter Equalization procedure has successfully completed.	0h	RO_V
3	EQPH2SUCC	Equalization Phase 2 Successful When set to 1b, this bit indicates that Phase 2 of the Transmitter Equalization procedure has successfully completed.	0h	RO_V
2	EQPH1SUCC	Equalization Phase 1 Successful When set to 1b, this bit indicates that Phase 1 of the Transmitter Equalization procedure has successfully completed.	0h	RO_V
1	EQCOMPLETE	Equalization Complete When set to 1b, this bit indicates that the Transmitter Equalization procedure has completed.	0h	RO_V
0	CURDELVL	Current De-emphasis Level: Current De-emphasis Level - When the Link is operating at 5 GT/s speed, this reflects the level of de-emphasis. Encodings: 1b -3.5 dB 0b -6 dB When the Link is operating at 2.5 GT/s speed, this bit is 0b.	0h	RO_V

5.1.29 DMIUESTS—DMI Uncorrectable Error Status

DMI Uncorrectable Error Status register. This register is for test and debug purposes only.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RW1CS	
Size: 32	**Default Value: 00000000h**		**Address Offset: 1C4h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:21	RSVD	Reserved.	000h	RO
20	URES	Unsupported Request Error Status:	0h	RW1CS
19	RSVD	Reserved.	0h	RO
18	MTLPS	Malformed TLP Status:	0h	RW1CS
17	ROS	Receiver Overflow Status:	0h	RW1CS
16	UCS	Unexpected Completion Status:	0h	RW1CS

continued...

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RW1CS	
Size: 32	Default Value: 00000000h		Address Offset: 1C4h	
Bit Range	Acronym	Description	Default	Access
15	RSVD	Reserved.	0h	RO
14	CTS	Completion Timeout Status:	0h	RW1CS
13	RSVD	Reserved.	0h	RO
12	PTLPS	Poisoned TLP Status:	0h	RW1CS
11:5	RSVD	Reserved.	00h	RO
4	DLPES	Data Link Protocol Error Status:	0h	RW1CS
3:0	RSVD	Reserved.	0h	RO

5.1.30 DMIUEMSK—DMI Uncorrectable Error Mask

DMI Uncorrectable Error Mask register. This register is for test and debug purposes only.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RWS	
Size: 32	Default Value: 00000000h		Address Offset: 1C8h	
Bit Range	Acronym	Description	Default	Access
31:23	RSVD	Reserved.	000h	RO
22	ECCERRM	2 Bit OPI Error Mask: This register bit only exists, i.e. CR readable, when OPI has been enabled	0h	RWS
21	RSVD	Reserved.	0h	RO
20	UREM	Unsupported Request Error Mask:	0h	RWS
19	RSVD	Reserved.	0h	RO
18	MTLPM	Malformed TLP Mask:	0h	RWS
17	ROM	Receiver Overflow Mask:	0h	RWS
16	UCM	Unexpected Completion Mask:	0h	RWS
15	RSVD	Reserved.	0h	RO
14	CPLTM	Completion Timeout Mask:	0h	RWS
13	RSVD	Reserved.	0h	RO
12	PTLPM	Poisoned TLP Mask:	0h	RWS
11:5	RSVD	Reserved.	00h	RO
4	DLPEM	Data Link Protocol Error Mask:	0h	RWS
3:0	RSVD	Reserved.	0h	RO

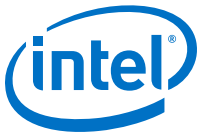

5.1.31 DMIUESEV—DMI Uncorrectable Error Severity

DMI Uncorrectable Error Severity register. This register controls whether an individual error is reported as a non-fatal or fatal error. An error is reported as fatal when the corresponding error bit in the severity register is set. If the bit is cleared, the corresponding error is considered nonfatal. It is for test and debug purposes only.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RWS; RO	
Size: 32	Default Value: 00060010h		Address Offset: 1CCh	
Bit Range	Acronym	Description	Default	Access
31:23	RSVD	Reserved.	000h	RO
22	ECCERRS	2 Bit Error Mask: This register bit only exists, i.e. CR readable, when OPI has been enabled	0h	RWS
21	RSVD	Reserved.	0h	RO
20	URES	Unsupported Request Error Severity:	0h	RWS
19	ECRCES	Reserved for ECRC Error Severity:	0h	RO
18	MTLPES	Malformed TLP Error Severity:	1h	RWS
17	ROEV	Receiver Overflow Error Severity:	1h	RWS
16	UCES	Unexpected Completion Error Severity:	0h	RWS
15	CAES	Reserved for Completer Abort Error Severity:	0h	RO
14	CTES	Completion Timeout Error Severity:	0h	RWS
13	FCPES	Reserved for Flow Control Protocol Error Severity:	0h	RO
12	PTLPES	Poisoned TLP Error Severity:	0h	RWS
11:5	RSVD	Reserved.	00h	RO
4	DLPES	Data Link Protocol Error Severity:	1h	RWS
3:0	RSVD	Reserved.	0h	RO

5.1.32 DMICESTS—DMI Correctable Error Status

DMI Correctable Error Status Register. This register is for test and debug purposes only.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RW1CS	
Size: 32	Default Value: 00000000h		Address Offset: 1D0h	
Bit Range	Acronym	Description	Default	Access
31:14	RSVD	Reserved.	00000h	RO
13	ANFES	Advisory Non-Fatal Error Status: When set, indicates that an Advisory Non-Fatal Error occurred.	0h	RW1CS
12	RTTS	Replay Timer Timeout Status:	0h	RW1CS
11:9	RSVD	Reserved.	0h	RO
8	RNRS	REPLAY_NUM Rollover Status:	0h	RW1CS
7	BDLLPS	Bad DLLP Status:	0h	RW1CS

continued...

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RW1CS	
Size: 32	Default Value: 00000000h		Address Offset: 1D0h	
Bit Range	Acronym	Description	Default	Access
6	BTLPS	Bad TLP Status:	0h	RW1CS
5:1	RSVD	Reserved.	00h	RO
0	RES	Receiver Error Status: Physical layer receiver Error occurred. These errors include: elastic Buffer Collision, 8b/10b error, De-skew Timeout Error.	0h	RW1CS

5.1.33 DMICEMSK—DMI Correctable Error Mask

DMI Correctable Error Mask register. This register is for test and debug purposes only.

B/D/F/Type: 0/0/0/MEM/DMIBAR			Access: RWS	
Size: 32	Default Value: 00002000h		Address Offset: 1D4h	
Bit Range	Acronym	Description	Default	Access
31:14	RSVD	Reserved.	00000h	RO
13	ANFEM	Advisory Non-Fatal Error Mask: When set, masks Advisory Non-Fatal errors from (a) signaling ERR_COR to the device control register, and (b) updating the Uncorrectable Error Status register. This register is set by default to enable compatibility with software that does not comprehend Role-Based Error Reporting.	1h	RWS
12:0	RSVD	Reserved.	0000h	RO

5.2 MCHBAR Registers Summary

Offset	Register ID—Description	Default Value	Access
40C8	ECCERRLOG0—ECC Error Log 0 on page 262	00000000h	RWS_LV
40CC	ECCERRLOG1—ECC Error Log 1 on page 263	00000000h	RWS_LV
44C8	ECCERRLOG0—ECC Error Log 0 on page 263	00000000h	RWS_LV
44CC	ECCERRLOG1—ECC Error Log 1 on page 264	00000000h	RWS_LV
4C04	TC—DDR Bank or Rank Timing parameters on page 264	00224000h	RW_L
4C08	TC—DDR Bank or Rank Timing Parametrs on page 265	00AB5545h	RW_L
4C0C	TC—Bank or Rank timing parameters on page 265	05294000h	RW_L
4CB0	PM—Power-down configuration register on page 266	00000000h	RW_L
4CC8	ECCERRLOG0—ECC Error Log 0 on page 266	00000000h	RWS_LV
4CCC	ECCERRLOG1—ECC Error Log 1 on page 267	00000000h	RWS_LV
4E94	TC—Refresh parameters on page 268	0000980Fh	RW_L
4E98	TC—Refresh timing parameters on page 268	46B41004h	RW_L
4EEC	PM—Power Management DIMM Idle Energy on page 268	00000000h	RW_L
4EF0	PM—Power Management DIMM Power Down Energy on page 269	00000000h	RW_L

continued...

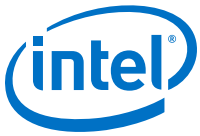

Offset	Register ID—Description	Default Value	Access
4EF4	PM—Power Management DIMM Activate Energy on page 269	00000000h	RW_L
4EF8	PM—Power Management DIMM RdCas Energy on page 269	00000000h	RW_L
4EFC	PM—Power Management DIMM WrCas Energy on page 270	00000000h	RW_L
5000	MAD—Address decoder Channel configuration register on page 270	00000024h	RW_L
5004	MAD—Address decode channel 0 on page 271	00600000h	RW_L
5008	MAD—Address decode channel 1 on page 272	00600000h	RW_L
5060	PM—Self refresh config. register on page 273	00010200h	RW_L
5090	ECC—Address compare for ECC error injection on page 273	00000000h	RW_L
5094	ECC—Address mask for ECC error injection on page 273	FFFFFFFFh	RW_L
5880	DDR—DDR_PTM_CTL_0_0_0_MCHBAR_PCU on page 274	00000000h	RW; RW_KL
5884	DRAM—DRAM_ENERGY_SCALEFACTOR_0_0_0_MCHBAR on page 275	00000003h	RW
5888	DRAM—DRAM_RAPL_CHANNEL_POWER_FLOOR_0_0_0_MCHBAR on page 276	00000000h	RW
588C	DDR—DDR_THERM_PERDIMM_STATUS_0_0_0_MCHBAR_PCU on page 276	00000000h	RO
5890	DDR—DDR_WARM_THRESHOLD_CH0_0_0_0_MCHBAR_PCU on page 276	0000FFFFh	RWS_L
5894	DDR—DDR_WARM_THRESHOLD_CH1_0_0_0_MCHBAR_PCU on page 277	0000FFFFh	RWS_L
5898	DDR—DDR_HOT_THRESHOLD_CH0_0_0_0_MCHBAR_PCU on page 277	0000FFFFh	RWS_L
589C	DDR—DDR_HOT_THRESHOLD_CH1_0_0_0_MCHBAR_PCU on page 277	0000FFFFh	RWS_L
58A0	DDR—DDR_THERM_INTERRUPT_CONFIG on page 278	00000000h	RW
58A8	PACKAGE—PACKAGE_THERM_MARGIN_0_0_0_MCHBAR_PCU on page 279	00007F00h	RO_V
58B0	DDR—DDR_DIMM_TEMPERATURE_CH0_0_0_0_MCHBAR_PCU on page 279	00000000h	RO
58B4	DDR—DDR_DIMM_TEMPERATURE_CH1_0_0_0_MCHBAR_PCU on page 280	00000000h	RO
58C0	DDR—DDR_THROTTLE_DURATION_CH0_0_0_0_MCHBAR_PCU on page 280	0000000000000000h	RO
58C8	DDR—DDR_THROTTLE_DURATION_CH1_0_0_0_MCHBAR_PCU on page 280	0000000000000000h	RO
58D0	DDR—DDR_WARM_BUDGET_CH0_0_0_0_MCHBAR_PCU on page 281	0000FFFFh	RWS_L
58D4	DDR—DDR_WARM_BUDGET_CH1_0_0_0_MCHBAR_PCU on page 281	0000FFFFh	RWS_L
58D8	DDR—DDR_HOT_BUDGET_CH0_0_0_0_MCHBAR_PCU on page 281	0000FFFFh	RWS_L
58DC	DDR—DDR_HOT_BUDGET_CH1_0_0_0_MCHBAR_PCU on page 281	0000FFFFh	RWS_L
58E0	DRAM—DRAM_POWER_LIMIT on page 282	0000000000000000h	RWS_L; RWS_KL
58E8	DRAM—DRAM_ENERGY_STATUS on page 282	00000000h	ROS_V
58EC	DRAM—DRAM_RAPL_PERF_STATUS on page 283	00000000h	ROS_V
58F0	PACKAGE—PACKAGE_RAPL_PERF_STATUS_0_0_0_MCHBAR_PCU on page 283	00000000h	ROS_V
58FC	CORE—CORE_PERF_LIMIT_REASONS on page 283	00000000h	ROV; RW0C
5900	GRAPHICS—GRAPHICS_PERF_LIMIT_REASONS on page 285	00000000h	ROV; RW0C

continued...

Offset	Register ID—Description	Default Value	Access
5904	RING—RING_PERF_LIMIT_REASONS on page 287	00000000h	ROV; RW0C
5928	PRIMARY—PRIMARY_PLANE_ENERGY_STATUS on page 288	00000000h	RO_V
592C	SECONDARY—SECONDARY_PLANE_ENERGY_STATUS on page 288	00000000h	RO_V
5930	PACKAGE—PACKAGE_POWER_SKU on page 288	0012024000600000h	ROS_V
5938	PACKAGE—PACKAGE_POWER_SKU_UNIT on page 289	000A0E03h	RO_V
593C	PACKAGE—PACKAGE_ENERGY_STATUS on page 289	00000000h	RO_V
5948	GT—GT_PERF_STATUS_0_0_0_MCHBAR_PCU on page 290	00000000h	RO_V
5950	IA32—IA32_PLATFORM_ID on page 290	0000000000000000h	ROS_V
5994	RP—RP_STATE_LIMITS_0_0_0_MCHBAR_PCU on page 290	000000FFh	RW
5998	RP—RP_STATE_CAP_0_0_0_MCHBAR_PCU on page 291	00000000h	RO
599C	TEMPERATURE—TEMPERATURE_TARGET on page 291	00000000h	RO_V
59B0	VR—VR_CURRENT_CONFIG on page 292	0000000000000190h	RO_V; ROS_V
59C0	IA32—IA32_THERM_STATUS on page 292	08000000h	ROV; RW0C; RO
59C4	IA32—IA32_THERM_INTERRUPT on page 293	00000000h	RW
5D10	SSKPD—SSKPD_0_0_0_MCHBAR_PCU on page 294	0000000000000000h	RWS
5F3C	CONFIG—CONFIG_TDP_NOMINAL_0_0_0_MCHBAR_PCU on page 294	00000000h	RO_V
5F40	CONFIG—CONFIG_TDP_LEVEL1_0_0_0_MCHBAR_PCU on page 295	0000000000000000h	RO_V
5F48	CONFIG—CONFIG_TDP_LEVEL2_0_0_0_MCHBAR_PCU on page 295	0000000000000000h	RO_V
5F50	CONFIG—CONFIG_TDP_CONTROL_0_0_0_MCHBAR_PCU on page 296	00000000h	RWS_L; RW_KL
5F54	TURBO—TURBO_ACTIVATION_RATIO_0_0_0_MCHBAR_PCU on page 296	00000000h	RWS_L; RWS_KL
6204	DDR—Memory Thermal Camarillo Status on page 297	00000000h	ROV; RW0C
7410	CRDTCTL4—IOTrk and RRTrk shared credits on page 298	00000856h	RW_L

5.2.1 ECCERRLOG0—ECC Error Log 0

This register logs ECC error information.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_LV	
Size: 32	Default Value: 00000000h		Address Offset: 40C8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:27	ERRBANKID	This field holds the Bank ID Address of the read transaction that had the ECC error. In DDR3/LPDDR the bit mapping is as follows: [31:30] = Rank, [29:27] = Bank	00h	RWS_LV
26:24	ERRCHUNK	Holds the chunk number of the error stored in the register.	0h	RWS_LV
23:16	ERRSYND	This field contains the error syndrome. A value of 0xFF indicates that the error is due to poisoning.	00h	RWS_LV

continued...

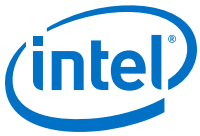

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_LV	
Size: 32	Default Value: 00000000h		Address Offset: 40C8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
15:2	RSVD	Reserved.	0000h	RO
1	MERRSTS	This bit is set when an uncorrectable multiple-bit error occurs on a memory read data transfer. When this bit is set, the address that caused the error and the error syndrome are also logged and they are locked until this bit is cleared. This bit is cleared when the corresponding bit in 0.0.0.PCI.ERRSTS is cleared.	0h	RWS_LV
0	CERRSTS	This bit is set when a correctable single-bit error occurs on a memory read data transfer. When this bit is set, the address that caused the error and the error syndrome are also logged and they are locked to further single bit errors, until this bit is cleared. A multiple bit error that occurs after this bit is set will override the address/error syndrome information. This bit is cleared when the corresponding bit in 0.0.0.PCI.ERRSTS is cleared.	0h	RWS_LV

5.2.2 ECCERRLOG1—ECC Error Log 1

This register logs ECC error information.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_LV	
Size: 32	Default Value: 00000000h		Address Offset: 40CCh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	ERRCOL	This field holds the DRAM column address of the read transaction that had the ECC error.	0000h	RWS_LV
15:0	ERRROW	This field holds the DRAM row (page) address of the read transaction that had the ECC error.	0000h	RWS_LV

5.2.3 ECCERRLOG0—ECC Error Log 0

This register logs ECC error information.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_LV	
Size: 32	Default Value: 00000000h		Address Offset: 44C8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:27	ERRBANKID	This field holds the Bank ID Address of the read transaction that had the ECC error. In DDR3/LPDDR the bit mapping is as follows: [31:30] = Rank, [29:27] = Bank	00h	RWS_LV
26:24	ERRCHUNK	Holds the chunk number of the error stored in the register.	0h	RWS_LV
23:16	ERRSYND	This field contains the error syndrome. A value of 0xFF indicates that the error is due to poisoning.	00h	RWS_LV

continued...

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_LV	
Size: 32	Default Value: 00000000h		Address Offset: 44C8h	
Bit Range	Acronym	Description	Default	Access
15:2	RSVD	Reserved.	0000h	RO
1	MERRSTS	This bit is set when an uncorrectable multiple-bit error occurs on a memory read data transfer. When this bit is set, the address that caused the error and the error syndrome are also logged and they are locked until this bit is cleared. This bit is cleared when the corresponding bit in 0.0.0.PCI.ERRSTS is cleared.	0h	RWS_LV
0	CERRSTS	This bit is set when a correctable single-bit error occurs on a memory read data transfer. When this bit is set, the address that caused the error and the error syndrome are also logged and they are locked to further single bit errors, until this bit is cleared. A multiple bit error that occurs after this bit is set will override the address/error syndrome information. This bit is cleared when the corresponding bit in 0.0.0.PCI.ERRSTS is cleared.	0h	RWS_LV

5.2.4 ECCERRLOG1—ECC Error Log 1

This register logs ECC error information.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_LV	
Size: 32	Default Value: 00000000h		Address Offset: 44CCh	
Bit Range	Acronym	Description	Default	Access
31:16	ERRCOL	This field holds the DRAM column address of the read transaction that had the ECC error.	0000h	RWS_LV
15:0	ERRROW	This field holds the DRAM row (page) address of the read transaction that had the ECC error.	0000h	RWS_LV

5.2.5 TC—DDR Bank or Rank Timing parameters

DDR Bank or Rank timing parameters in DCLK cycles

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 00224000h		Address Offset: 4C04h	
Bit Range	Acronym	Description	Default	Access
31:23	RSVD	Reserved.	000h	RO
22:19	tRDRD_dd	The minimum time between RD-CAS to RD-CAS to different dimms in DCLK cycles. The minimum is 4 DCLKS.	4h	RW_L

continued...

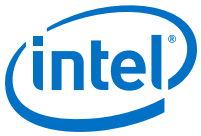

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 00224000h		Address Offset: 4C04h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
18:15	tRDRD_dr	The minimum time between RD-CAS to RD-CAS to different ranks on the same dimm in DCLK cycles. The minimum is 4 DCLKS.	4h	RW_L
14:12	tRDRD	The minimum time between RD-CAS to RD-CAS to different banks on the same rank in DCLKS. The minimum is 4 DCLKS.	4h	RW_L
11:0	RSVD	Reserved.	000h	RO

5.2.6 TC—DDR Bank or Rank Timing Parametrs

DDR Bank or Rank timing parameters in DCLK cycles

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 00AB5545h		Address Offset: 4C08h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:25	RSVD	Reserved.	00h	RO
24:21	tWRWR_dd	Minimum time between WR-CAS to WR-CAS to different dimms in DCLKS.	5h	RW_L
20:17	tWRWR_dr	Minimum time between WR-CAS to WR-CAS to different ranks on the same dimm in DCLKS.	5h	RW_L
16:14	tWRWR	Minimum time between WR-CAS to WR-CAS to different banks on the same rank in DCLKS.	5h	RW_L
13:10	tWRRD_dd	Minimum time between WR-CAS to RD-CAS to different dimms in DCLKS.	5h	RW_L
9:6	tWRRD_dr	Minimum time between WR-CAS to RD-CAS to different ranks on the same dimm in DCLKS.	5h	RW_L
5:0	tWRRD	Minimum time between WR-CAS to RD-CAS to different banks on the same rank in DCLKS.	05h	RW_L

5.2.7 TC—Bank or Rank timing parameters

Bank or Rank timing parameters in DCLK cycles

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 05294000h		Address Offset: 4C0Ch	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:29	RSVD	Reserved.	0h	RO
28:24	tRDWR_dd	Minimum time between RD-CAS to WR-CAS to different dimms in DCLK cycles. NOTE: The maximum value that can be programmed into this field is 30 to prevent overflow as hardware can conditionally add 1 to this field in some cases.	05h	RW_L

continued...

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 05294000h		Address Offset: 4C0Ch	
Bit Range	Acronym	Description	Default	Access
23:19	tRDWR_dr	Minimum time between RD-CAS to WR-CAS to different ranks on the same dimm in DCLK cycles. NOTE: The maximum value that can be programmed into this field is 30 to prevent overflow as hardware can conditionally add 1 to this field in some cases.	05h	RW_L
18:14	tRDWR	Minimum time between RD-CAS to WR-CAS to different banks on the same rank in DCLK cycles. NOTE: The maximum value that can be programmed into this field is 30 to prevent overflow as hardware can conditionally add 1 to this field in some cases.	05h	RW_L
13:0	RSVD	Reserved.	0000h	RO

5.2.8 PM—Power-down configuration register

This register defines the power-down (CKE-off) operation - power-down mode, idle timer and global / per rank decision.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 00000000h		Address Offset: 4CB0h	
Bit Range	Acronym	Description	Default	Access
31:16	RSVD	Reserved.	0000h	RO
15:12	PDWN_mode	selects the mode of power-down: 0x0: no power-down 0x1: APD 0x2: PPD 0x3 - 0x5: Reserved 0x6: PPD-DLLoff 0x7 - 0xf: Reserved	0h	RW_L
11:0	PDWN_idle_counter	This defines the rank idle period in DCLK cycles that causes power-down entrance.The minimum value for this filed should be greater then or equal to the worst case Roundtrip delay defined in SC_Roundt_lat_0_0_0_MCHBAR plus Burst_Length	000h	RW_L

5.2.9 ECCERRLOG0—ECC Error Log 0

This register logs ECC error information.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_LV	
Size: 32	Default Value: 00000000h		Address Offset: 4CC8h	
Bit Range	Acronym	Description	Default	Access
31:27	ERRBANKID	This field holds the Bank ID Address of the read transaction that had the ECC error. In DDR3/LPDDR the bit mapping is as follows: [31:30] = Rank, [29:27] = Bank	00h	RWS_LV
26:24	ERRCHUNK	Holds the chunk number of the error stored in the register.	0h	RWS_LV

continued...

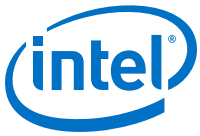

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_LV	
Size: 32	**Default Value: 00000000h**		**Address Offset: 4CC8h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
23:16	ERRSYND	This field contains the error syndrome. A value of 0xFF indicates that the error is due to poisoning.	00h	RWS_LV
15:2	RSVD	Reserved.	0000h	RO
1	MERRSTS	This bit is set when an uncorrectable multiple-bit error occurs on a memory read data transfer. When this bit is set, the address that caused the error and the error syndrome are also logged and they are locked until this bit is cleared. This bit is cleared when the corresponding bit in 0.0.0.PCI.ERRSTS is cleared.	0h	RWS_LV
0	CERRSTS	This bit is set when a correctable single-bit error occurs on a memory read data transfer. When this bit is set, the address that caused the error and the error syndrome are also logged and they are locked to further single bit errors, until this bit is cleared. A multiple bit error that occurs after this bit is set will override the address/error syndrome information. This bit is cleared when the corresponding bit in 0.0.0.PCI.ERRSTS is cleared.	0h	RWS_LV

5.2.10 ECCERRLOG1—ECC Error Log 1

This register logs ECC error information.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_LV	
Size: 32	**Default Value: 00000000h**		**Address Offset: 4CCCh**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	ERRCOL	This field holds the DRAM column address of the read transaction that had the ECC error.	0000h	RWS_LV
15:0	ERRROW	This field holds the DRAM row (page) address of the read transaction that had the ECC error.	0000h	RWS_LV

5.2.11 TC—Refresh parameters

Refresh parameters

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 0000980Fh		Address Offset: 4E94h	
Bit Range	Acronym	Description	Default	Access
31:16	RSVD	Reserved.	0000h	RO
15:12	Refresh_panic_wm	tREFI count level in which the refresh priority is panic (default is 9). The Maximum value for this field is 9.	9h	RW_L
11:8	Refresh_HP_WM	tREFI count level that turns the refresh priority to high (default is 8)	8h	RW_L
7:0	OREF_RI	Rank idle period that defines an opportunity for refresh, in DCLK cycles	0Fh	RW_L

5.2.12 TC—Refresh timing parameters

Refresh timing parameters

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 46B41004h		Address Offset: 4E98h	
Bit Range	Acronym	Description	Default	Access
31:25	tREFIx9	Maximum time allowed between refreshes to a rank (in intervals of 1024 DCLK cycles). Should be programmed to 8.9*tREFI/1024 (to allow for possible delays from ZQ or isoc).	23h	RW_L
24:16	tRFC	Time of refresh - from beginning of refresh until next ACT or refresh is allowed (in DCLK cycles, default is 180)	0B4h	RW_L
15:0	tREFI	defines the average period between refreshes, and the rate that tREFI counter is incremented (in DCLK cycles, default is 4100)	1004h	RW_L

5.2.13 PM—Power Management DIMM Idle Energy

This register defines the energy of an idle DIMM with CKE on. Each 6-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 6-bit fields, one per DIMM.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 00000000h		Address Offset: 4EECh	
Bit Range	Acronym	Description	Default	Access
31:14	RSVD	Reserved.	00000h	RO
13:8	DIMM1_IDLE_ENERGY	This register defines the energy consumed by DIMM1 for one clock cycle when the DIMM is idle with cke on	00h	RW_L
7:6	RSVD	Reserved.	0h	RO
5:0	DIMM0_IDLE_ENERGY	This register defines the energy consumed by DIMM0 for one clock cycle when the DIMM is idle with cke on.	00h	RW_L

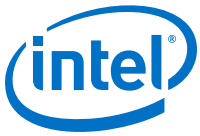

5.2.14 PM—Power Management DIMM Power Down Energy

This register defines the energy of an idle DIMM with CKE off. Each 6-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 6-bit fields, one per DIMM.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 00000000h		Address Offset: 4EF0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:14	RSVD	Reserved.	00000h	RO
13:8	DIMM1_PD_ENERGY	This register defines the energy consumed by DIMM1 for one clock cycle when the DIMM is idle with cke off	00h	RW_L
7:6	RSVD	Reserved.	0h	RO
5:0	DIMM0_PD_ENERGY	This register defines the energy consumed by DIMM0 for one clock cycle when the DIMM is idle with cke off	00h	RW_L

5.2.15 PM—Power Management DIMM Activate Energy

This register defines the combined energy contribution of activate and precharge commands. Each 8-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 8-bit fields, one per DIMM.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 00000000h		Address Offset: 4EF4h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:8	DIMM1_ACT_ENERGY	This register defines the combined energy contribution of activate and precharge commands.	00h	RW_L
7:0	DIMM0_ACT_ENERGY	This register defines the combined energy contribution of activate and precharge commands.	00h	RW_L

5.2.16 PM—Power Management DIMM RdCas Energy

This register defines the energy contribution of a read CAS command. Each 8-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 8-bit fields, one per DIMM.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 00000000h		Address Offset: 4EF8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:8	DIMM1_RD_ENERGY	This register defines the energy contribution of a read CAS command.	00h	RW_L
7:0	DIMM0_RD_ENERGY	This register defines the energy contribution of a read CAS command.	00h	RW_L

5.2.17 PM—Power Management DIMM WrCas Energy

This register defines the energy contribution of a write CAS command. Each 8-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 8-bit fields, one per DIMM.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 00000000h		Address Offset: 4EFCh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:8	DIMM1_WR_ENERGY	This register defines the energy contribution of a write CAS command.	00h	RW_L
7:0	DIMM0_WR_ENERGY	This register defines the energy contribution of a write CAS command.	00h	RW_L

5.2.18 MAD—Address decoder Channel configuration register

This register defines which channel is assigned to be channel A, channel B and channel C according to the rule:
size(A) >= size (B) >= size(C)
Since the Processor implements only two channels, channel C is always channel 2, and its size is always 0

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 00000024h		Address Offset: 5000h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:12	RSVD	Reserved.	00000h	RO
11	RSVD	Reserved.	0h	RO
10	LPDDR	LPDDR - Indicate that LPDDR devices are connected used on the system rather than DDR3 devices.	0h	RW_L
9:6	RSVD	Reserved.	0h	RO
5:4	CH_C	CH_C - defines the smallest channel: 00: Channel 0 01: Channel 1 10: Channel 2	2h	RW_L
3:2	CH_B	CH_B - defines the mid-size channel: 00: Channel 0 01: Channel 1 10: Channel 2	1h	RW_L
1:0	CH_A	CH_A - defines the largest channel: 00: Channel 0 01: Channel 1 10: Channel 2	0h	RW_L

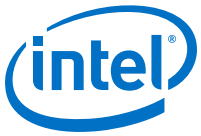

5.2.19 MAD—Address decode channel 0

This register defines channel characteristics - number of DIMMs, number of ranks, size, ECC, interleave options and ECC options

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 00600000h		Address Offset: 5004h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:30	RSVD	Reserved.	0h	RO
29:27	HORIAddr	High Order Rank Interleave Address. Specifies which DIMM address bit 20-27 to use as the rank interleave bit 000 - bit 20 001 - bit 21 ... 111 - bit 27	0h	RW_L
26	HORI	High Order Rank Interleave 0 - off 1 - on High Order Rank Interleave (HORI) is mutually exclusive with Rank Interleave (RI)	0h	RW_L
25:24	ECC	ECC configuration in the channel: 00: No ECC active in the channel. 11: ECC active in both IO and ECC logic. Note: This field must be programmed identically for all populated channels.	0h	RW_L
23	RSVD	Reserved.	0h	RO
22	Enh_Interleave	Enhanced interleave mode 0 - off 1 - on	1h	RW_L
21	RI	Rank Interleave 0 - off 1 - on	1h	RW_L
20	DBW	DBW: DIMM B width of DDR chips 0 - Other than X16 chips 1 - X16 chips	0h	RW_L
19	DAW	DAW: DIMM A width of DDR chips 0 - Other than X16 chips 1 - X16 chips	0h	RW_L
18	DBNOR	DIMM B number of ranks: 0 - single rank 1 - dual rank	0h	RW_L
17	DANOR	DIMM A number of ranks: 0 - single rank 1 - dual rank	0h	RW_L
16	DAS	Selects which of the DIMMs is DIMM A - should be the larger DIMM: 0 - DIMM 0 1 - DIMM 1	0h	RW_L
15:8	DIMM_B_Size	Size of DIMM B in 256 MB multiples	00h	RW_L
7:0	DIMM_A_Size	Size of DIMM A in 256 MB multiples	00h	RW_L

5.2.20 MAD—Address decode channel 1

This register defines channel charesteristics - number of DIMMs, number of ranks, size, ECC, interleave options and ECC options

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 00600000h		Address Offset: 5008h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:30	RSVD	Reserved.	0h	RO
29:27	HORIAddr	High Order Rank Interleave Address. Specifies which address bit 20-27 to use as the rank interleave bit 000 - bit 20 001 - bit 21 ... 111 - bit 27	0h	RW_L
26	HORI	High Order Rank Interleave 0 - off 1 - on	0h	RW_L
25:24	ECC	ECC configuration in the channel: 00: No ECC active in the channel. 11: ECC active in both IO and ECC logic. Note: This field must be programmed identically for all populated channels.	0h	RW_L
23	RSVD	Reserved.	0h	RO
22	Enh_Interleave	Enhanced interleave mode 0 - off 1 - on	1h	RW_L
21	RI	Rank Interleave 0 - off 1 - on	1h	RW_L
20	DBW	DBW: DIMM B width of DDR chips 0 - Other than X16 chips 1 - X16 chips	0h	RW_L
19	DAW	DAW: DIMM A width of DDR chips 0 - Other than X16 chips 1 - X16 chips	0h	RW_L
18	DBNOR	DIMM B number of ranks: 0 - single rank 1 - dual rank	0h	RW_L
17	DANOR	DIMM A number of ranks: 0 - single rank 1 - dual rank	0h	RW_L
16	DAS	Selects which of the DIMMs is DIMM A - should be the larger DIMM: 0 - DIMM 0 1 - DIMM 1	0h	RW_L
15:8	DIMM_B_Size	Size of DIMM B in 256 MB multiples	00h	RW_L
7:0	DIMM_A_Size	Size of DIMM A in 256 MB multiples	00h	RW_L

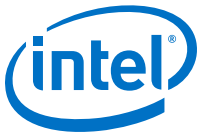

5.2.21 PM—Self refresh config. register

Self refresh mode control register - defines if and when DDR can go into SR

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 00010200h		Address Offset: 5060h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:17	RSVD	Reserved.	0000h	RO
16	SR_Enable	enables or disables self-refresh mechanism. In order to allow SR, both SREF_en bit should be set and SREF_exit signal should be cleared. PM_SREF_config may be updated in run-time	1h	RW_L
15:0	Idle_timer	This value is used when the SREF_enable field is set. It defines the # of cycles that there should not be any transaction in order to enter self-refresh. It is programmable from 512 to 64K-1. In DCLK=800 it determines time of up to 82 us. This parameter has been adjusted to protect ODTLoff + 1 to MRS command timing. As part of the bug fix for bug 3138064/3538082 the minimum time has been increased to 512. See the bug for details.	0200h	RW_L

5.2.22 ECC—Address compare for ECC error injection

Address compare for ECC error inject. Error injection is issued when ECC_Inj_Addr_Compare[31:0] = ADDR[37:6] & ECC_Inj_Addr_Mask[31:0]

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 00000000h		Address Offset: 5090h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:0	Address	Inject error when ECC_Inj_Addr_Compare[31:0] = ADDR[37:6]	00000000h	RW_L

5.2.23 ECC—Address mask for ECC error injection

Address compare for ECC error inject. Error injection is issued when ECC_Inj_Addr_Compare[31:0] = ADDR[37:6] & ECC_Inj_Addr_Mask[31:0]

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: FFFFFFFFh		Address Offset: 5094h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:0	Mask	Inject error when ECC_Inj_Addr_Compare[31:0] = ADDR[37:6] & & ECC_Inj_Addr_Mask[31:0]	FFFFFFFFh	RW_L

5.2.24 DDR—DDR_PTM_CTL_0_0_0_MCHBAR_PCU

Mode control bits for DDR power and thermal management features.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access:	RW; RW_KL
Size: 32	Default Value: 00000000h		Address Offset:	5880h
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:9	RSVD	Reserved.	000000h	RO
8	RSVDESH_EN	Reserved.	0h	RO
7	DISABLE_DRAM_TS	When this bit is zero and MAD_CHNL.LPDDR=1, the processor will use DDR MR4 for DIMM thermal status purposes. Otherwise, the processor will ignore MR4 data and use the legacy CLTM/OLTM/EXTTS algorithms for computing DIMM thermal status.	0h	RW
6	PDWN_CONFIG_CTL	This bit determined whether BIOS or the processor will control DDR powerdown modes and idle counter (via programming the PM_PDWN_config regs in iMC). When clear, the processor will manage the modes based on either core P-states or IA32_ENERGY_PERFORMANCE_BIAS MSR value (when enabled). When set, BIOS is in control of DDR CKE mode and idle timer value, and the processor algorithm does not run.	0h	RW
5	LOCK_PTM_REGS_PCU	When set, several processor registers related to DDR power/thermal management all become unwritable (writes will be silently ignored). List of registered locked by this bit is: DDR_WARM_THRESHOLD_CH*, DDR_HOT_THRESHOLD_CH*, DDR_WARM_BUDGET_CH*, DDR_HOT_BUDGET_CH*, (note that RAPL regs, such as RAPL_LIMIT, are NOT included as those have separate lock bit). Note that BIOS should complete its writes to all of the locked registers prior to setting this bit, since it can only be cleared via reset.	0h	RW_KL
4	EXTTS_ENABLE	When clear (default), the processor ignores the EXTTS (external thermal status) indication which is obtained from the PCH (via PM_SYNC). When set, the value from EXTTS is used only when it is hotter than the thermal status reported by OLTM/CLTM algorithm (or used all of the time if neither of those modes is enabled).	0h	RW
3:2	REFRESH_2X_MODE	These bits are read at reset and later broadcast (together with the thermal status) into the iMC cregs that control 2x refresh modes. A value of 00 means the iMC 2x refresh is disabled. A value of 01 means that the iMC will enable 2x refresh whenever thermal status is WARM or HOT. A value of 10 means the iMC will enable 2x refresh only when HOT. The value 11 is illegal. This field is ignored for LPDDR when DISABLE_DRAM_TS is zero, in which case refresh rates in the MC are controlled by MR4 coming directly from DIMMs.	0h	RW
1	CLTM_ENABLE	A value of 1 means CLTM (Closed Loop Thermal Management) algorithm will be used to compute the memory thermal status (which will be written to the iMC). Note that OLTM and CLTM modes are mutex, so if both OLTM_ENABLE and CLTM_ENABLE are set, the OLTM_ENABLE will be ignored and CTLM mode will be	0h	RW

continued...

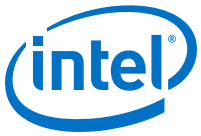

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW; RW_KL	
Size: 32	**Default Value: 00000000h**		**Address Offset: 5880h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		active. BIOS should enable CLTM whenever DIMM thermal sensor data is available and memory thermal management is desired.		
0	OLTM_ENABLE	A value of 1 means OLTM (Open Loop Thermal Management) algorithm will be used to compute the memory thermal status (which will be written to the iMC). Note that OLTM and CLTM modes are mutex, so if both OLTM_ENABLE and CLTM_ENABLE are set, the OLTM_ENABLE will be ignored and CTLM mode will be active. BIOS should enable CLTM whenever DIMM thermal sensor data is not available, but memory thermal management is desired. Obviously lack of real temperature data means this mode will be somewhat conservative, and may result in the iMC throttling more often than necessary. Thus for perf reasons CLTM is preferred on systems with available DIMM thermal sensor data.	0h	RW

5.2.25 DRAM—DRAM_ENERGY_SCALEFACTOR_0_0_0_MCHBAR

Defines the base energy unit for DDR energy values in iMC command energy config regs, iMC rank energy counters (used for OLTM and Memory RAPL), OLTM thresholds, etc.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW	
Size: 32	**Default Value: 00000003h**		**Address Offset: 5884h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:3	RSVD	Reserved.	00000000h	RO
2:0	SCALEFACTOR	Defines the base DDR energy unit of 2^(-30-scalefactor) Joules. The values are defined as follows: 0d0 = 3'b000 = 931.3pJ, 0d1 = 3'b001 = 465.7pJ, 0d2 = 3'b010 = 232.8pJ, 0d3 = 3'b011 = 116.4pJ, 0d4 = 3'b100 = 58.2pJ, 0d5 = 3'b101 = 29.1pJ, 0d6 = 3'b110 = 14.6pJ, 0d7 = 3'b111 = 7.3pJ. The default reset value is 0d3 = 3'b011 = 116.4pJ.	3h	RW

5.2.26 DRAM—DRAM_RAPL_CHANNEL_POWER_FLOOR_0_0_0_MCHBAR

Defines the minimum required power consumption of each DDR channel, in order to satisfy minimum memory bandwidth requirements for the platform. DDR RAPL should never throttle below the levels defined here. It is the responsibility of BIOS to comprehend the power consumption on each channel in order to write meaningful values into this register.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 5888h	
Bit Range	Acronym	Description	Default	Access
31:16	RSVD	Reserved.	0000h	RO
15:8	CH1	Minimum power level (in format of 5.3 W) used to clip DDR RAPL power budget for channel 1.	00h	RW
7:0	CH0	Minimum power level (in format of 5.3 W) used to clip DDR RAPL power budget for channel 0.	00h	RW

5.2.27 DDR—DDR_THERM_PERDIMM_STATUS_0_0_0_MCHBAR_PCU

Per-DIMM thermal status values. The encoding of each DIMM thermal status is the same: 2'b00 = COLD, 2'b01 = WARM, 2'b11 = HOT, 2'b10 == Reserved.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO	
Size: 32	Default Value: 00000000h		Address Offset: 588Ch	
Bit Range	Acronym	Description	Default	Access
31:12	RSVD	Reserved.	00000h	RO
11:10	CH1_DIMM1	Thermal Status for Channel 1, DIMM1	0h	RO
9:8	CH1_DIMM0	Thermal Status for Channel 1, DIMM0	0h	RO
7:4	RSVD	Reserved.	0h	RO
3:2	CH0_DIMM1	Thermal Status for Channel 0, DIMM1	0h	RO
1:0	CH0_DIMM0	Thermal Status for Channel 0, DIMM0	0h	RO

5.2.28 DDR—DDR_WARM_THRESHOLD_CH0_0_0_0_MCHBAR_PCU

Per-DIMM temp/power thresholds used for CLTM/OLTM thermal status computation. These values can impact iMC throttling and memory thermal interrrupts.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_L	
Size: 32	Default Value: 0000FFFFh		Address Offset: 5890h	
Bit Range	Acronym	Description	Default	Access
31:16	RSVD	Reserved.	0000h	RO
15:8	DIMM1	WARM_THRESHOLD for DIMM1 on this channel.	FFh	RWS_L
7:0	DIMM0	WARM_THRESHOLD for DIMM0 on this channel.	FFh	RWS_L

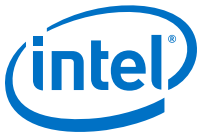

5.2.29 DDR—DDR_WARM_THRESHOLD_CH1_0_0_0_MCHBAR_PCU

Per-DIMM temp/power thresholds used for CLTM/OLTM thermal status computation. These values can impact iMC throttling and memory thermal interrrupts.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_L	
Size: 32	Default Value: 0000FFFFh		Address Offset: 5894h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:8	DIMM1	WARM_THRESHOLD for DIMM1 on this channel.	FFh	RWS_L
7:0	DIMM0	WARM_THRESHOLD for DIMM0 on this channel.	FFh	RWS_L

5.2.30 DDR—DDR_HOT_THRESHOLD_CH0_0_0_0_MCHBAR_PCU

Per-DIMM temp/power thresholds used for CLTM/OLTM thermal status computation. These values can impact iMC throttling and memory thermal interrrupts.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_L	
Size: 32	Default Value: 0000FFFFh		Address Offset: 5898h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:8	DIMM1	HOT_THRESHOLD for DIMM1 on this channel.	FFh	RWS_L
7:0	DIMM0	HOT_THRESHOLD for DIMM0 on this channel.	FFh	RWS_L

5.2.31 DDR—DDR_HOT_THRESHOLD_CH1_0_0_0_MCHBAR_PCU

Per-DIMM temp/power thresholds used for CLTM/OLTM thermal status computation. These values can impact iMC throttling and memory thermal interrrupts.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_L	
Size: 32	Default Value: 0000FFFFh		Address Offset: 589Ch	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:8	DIMM1	HOT_THRESHOLD for DIMM1 on this channel.	FFh	RWS_L
7:0	DIMM0	HOT_THRESHOLD for DIMM0 on this channel.	FFh	RWS_L

5.2.32 DDR—DDR_THERM_INTERRUPT_CONFIG

Enable bits and policy-free thresholds used for controlling memory thermal interrupt generation.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 58A0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:24	POLICY_FREE_THRESHOLD2	A threshold temperature value used only for interrupt generation. No iMC throttling or other actions should be directly affected by this value. This only works when CLTM is enabled. This is an 8-bit unsigned value in 7.1 format, 0.5C increment. THRESHOLD1 and THRESHOLD2 values and enables are fully independent from each other.	00h	RW
23:16	POLICY_FREE_THRESHOLD1	A threshold temperature value used only for interrupt generation. No iMC throttling or other actions should be directly affected by this value. This only works when CLTM is enabled. This is an 8-bit unsigned value in 7.1 format, 0.5C increment. THRESHOLD1 and THRESHOLD2 values and enables are fully independent from each other.	00h	RW
15:11	RSVD	Reserved.	00h	RO
10	ENABLE_THRESHOLD2_INTERRUPT	When set, interrupts will be generated on both rising and falling transition of the hottest absolute DIMM temperature across the POLICY_FREE_THRESHOLD2 value. This interrupt will never get triggered by the processor in cases where CLTM is not enabled (i.e. does not work with OLTM). THRESHOLD1 and THRESHOLD2 values and enables are fully independent from each other.	0h	RW
9	RSVD	Reserved.	0h	RO
8	ENABLE_THRESHOLD1_INTERRUPT	When set, interrupts will be generated on both rising and falling transition of the hottest absolute DIMM temperature across the POLICY_FREE_THRESHOLD1 value. This interrupt will never get triggered by the processor in cases where CLTM is not enabled (i.e. does not work with OLTM). THRESHOLD1 and THRESHOLD2 values and enables are fully independent from each other.	0h	RW
7	RSVD	Reserved.	0h	RO
6	ENABLE_OOS_TEMP_INTERRUPT	When set, interrupts will be generated on a rising transition of hottest MR4 to 3'b111. This interrupt will never get triggered by the processor in cases where MAD_CHNL.LPDDR is zero or DISABLE_DRAM_TS is set.	0h	RW
5	RSVD	Reserved.	0h	RO
4	ENABLE_2X_REFRESH_INTERRUPT	When set, interrupts will be generated on a rising transition of the hottest DIMM thermal status across whichever threshold 2x refresh is configured for (WARM_THRESHOLD, HOT_THRESHOLD, or never, depending on DDR_PTM_CTL.REFRESH_2X_MODE). This interrupt will never be triggered by the processor in cases where 2X refresh is disabled OR when no thermal status updates are being performed because CLTM, OLTM, and EXTTS are all disabled. In the case of LPDDR when DISABLE_DRAM_TS is zero, MR4 is used for refresh rate control and this interrupt still exists but its name becomes slightly misleading... in that case it is triggered whenever MR4 changes such that DIMM refresh rate has crossed the boundary (in either	0h	RW

continued...

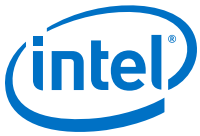

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 58A0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		direction) between 1x or lower refresh rate, and higher than 1x refresh rate. LPDDR might go above the 2x refresh rate and still generate this interrupt, for example.		
3	RSVD	Reserved.	0h	RO
2	ENABLE_HOT_INTERRUPT	When set, interrupts will be generated on a rising transition of the hottest DIMM thermal status from WARM to HOT (i.e. rise to or above HOT_THRESHOLD). This interrupt will never get triggered by the processor in cases where CLTM, OLTM, and EXTTS are all disabled.	0h	RW
1	RSVD	Reserved.	0h	RO
0	ENABLE_WARM_INTERRUPT	When set, interrupts will be generated on a rising transition of the hottest DIMM thermal status from COLD to WARM (i.e. rise to or above WARM_THRESHOLD). This interrupt will never get triggered by the processor in cases where CLTM, OLTM, and EXTTS are all disabled.	0h	RW

5.2.33 PACKAGE—PACKAGE_THERM_MARGIN_0_0_0_MCHBAR_PCU

Temperature margin in PECI temperature counts from the thermal profile specification. Platform fan control SW is expected to read therm_margin value to control fan or blower speed.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO_V	
Size: 32	Default Value: 00007F00h		Address Offset: 58A8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:0	THERM_MARGIN	Temperature margin in PECI temperature counts from the thermal profile specification. THERM_MARGIN is in 2's complement format (8.8 format where MSB equals 1 Sign bit + 7 bits of integer temperature value and the LSB equals 8 precison bits of temperature value). A value of zero indicates the hottest processor die temperature is on the thermal profile line. A negative value indicates gap to the thermal profile that platform SW should increase cooling capacity. A sustained negative value should be avoided as it may impact part reliability.	7F00h	RO_V

5.2.34 DDR—DDR_DIMM_TEMPERATURE_CH0_0_0_0_MCHBAR_PCU

Per-DIMM temperature values.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO	
Size: 32	Default Value: 00000000h		Address Offset: 58B0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:8	DIMM1	Temperature of DIMM1 on this channel.	00h	RO
7:0	DIMM0	Temperature of DIMM0 on this channel.	00h	RO

5.2.35 DDR—DDR_DIMM_TEMPERATURE_CH1_0_0_0_MCHBAR_PCU

Per-DIMM temperature values.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO	
Size: 32	Default Value: 00000000h		Address Offset: 58B4h	
Bit Range	Acronym	Description	Default	Access
31:16	RSVD	Reserved.	0000h	RO
15:8	DIMM1	Temperature of DIMM1 on this channel.	00h	RO
7:0	DIMM0	Temperature of DIMM0 on this channel.	00h	RO

5.2.36 DDR—DDR_THROTTLE_DURATION_CH0_0_0_0_MCHBAR_PCU

Per-DIMM throttle duration counters. These accumulate the duration (in absolute wall clock time) that the iMC rank throttlers have been blocking memory traffic due to OLTM/CLTM/EXTTS thermal status. Note that RAPL throttling is done at the channel level, and thus is NOT included in these values.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO	
Size: 64	Default Value: 0000000000000000h		Address Offset: 58C0h	
Bit Range	Acronym	Description	Default	Access
63:32	RSVD	Reserved.	00000000h	RO
31:16	DIMM1	Throttle duration of DIMM 1 on this channel, in units of 1/1024 seconds.	0000h	RO
15:0	DIMM0	Throttle duration of DIMM 0 on this channel, in units of 1/1024 seconds.	0000h	RO

5.2.37 DDR—DDR_THROTTLE_DURATION_CH1_0_0_0_MCHBAR_PCU

Per-DIMM throttle duration counters. These accumulate the duration (in absolute wall clock time) that the iMC rank throttlers have been blocking memory traffic due to OLTM/CLTM/EXTTS thermal status. Note that RAPL throttling is done at the channel level, and thus is NOT included in these values.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO	
Size: 64	Default Value: 0000000000000000h		Address Offset: 58C8h	
Bit Range	Acronym	Description	Default	Access
63:32	RSVD	Reserved.	00000000h	RO
31:16	DIMM1	Throttle duration of DIMM 1 on this channel, in units of 1/1024 seconds.	0000h	RO
15:0	DIMM0	Throttle duration of DIMM 0 on this channel, in units of 1/1024 seconds.	0000h	RO

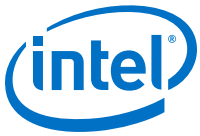

5.2.38 DDR—DDR_WARM_BUDGET_CH0_0_0_0_MCHBAR_PCU

Per-DIMM power budget for MC thermal throttling when thermal status is WARM.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_L	
Size: 32	Default Value: 0000FFFFh		Address Offset: 58D0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:8	DIMM1	WARM_BUDGET for DIMM1 on this channel.	FFh	RWS_L
7:0	DIMM0	WARM_BUDGET for DIMM0 on this channel.	FFh	RWS_L

5.2.39 DDR—DDR_WARM_BUDGET_CH1_0_0_0_MCHBAR_PCU

Per-DIMM power budget for MC thermal throttling when thermal status is WARM.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_L	
Size: 32	Default Value: 0000FFFFh		Address Offset: 58D4h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:8	DIMM1	WARM_BUDGET for DIMM1 on this channel.	FFh	RWS_L
7:0	DIMM0	WARM_BUDGET for DIMM0 on this channel.	FFh	RWS_L

5.2.40 DDR—DDR_HOT_BUDGET_CH0_0_0_0_MCHBAR_PCU

Per-DIMM power budget for MC thermal throttling when thermal status is HOT.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_L	
Size: 32	Default Value: 0000FFFFh		Address Offset: 58D8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:8	DIMM1	HOT_BUDGET for DIMM1 on this channel.	FFh	RWS_L
7:0	DIMM0	HOT_BUDGET for DIMM0 on this channel.	FFh	RWS_L

5.2.41 DDR—DDR_HOT_BUDGET_CH1_0_0_0_MCHBAR_PCU

Per-DIMM power budget for MC thermal throttling when thermal status is HOT.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_L	
Size: 32	Default Value: 0000FFFFh		Address Offset: 58DCh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:8	DIMM1	HOT_BUDGET for DIMM1 on this channel.	FFh	RWS_L
7:0	DIMM0	HOT_BUDGET for DIMM0 on this channel.	FFh	RWS_L

5.2.42 DRAM—DRAM_POWER_LIMIT

Allows software to set power limits for
the DRAM domain and measurement attributes associated with each limit.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access:	RWS_L; RWS_KL
Size: 64	Default Value: 0000000000000000h		Address Offset:	58E0h
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63	LOCKED	When set, this entire register becomes read-only. This bit will typically be set by BIOS during boot.	0h	RWS_KL
62:56	RSVD	Reserved.	00h	RO
55:54	LIMIT2_TIME_WINDOW_X	Power Limit[1] time window X value, for DDR domain. Actual time_window for RAPL is: (1/1024 seconds) * (1+(x/4)) * (2^y)	0h	RWS_L
53:49	LIMIT2_TIME_WINDOW_Y	Power Limit[1] time window Y value, for DDR domain. Actual time_window for RAPL is: (1/1024 seconds) * (1+(x/4)) * (2^y)	00h	RWS_L
48	RSVD	Reserved.	0h	RO
47	LIMIT2_ENABLE	Power Limit[1] enable bit for DDR domain.	0h	RWS_L
46:32	LIMIT2_POWER	Power Limit[1] for DDR domain. Units=Watts, Format=11.3, Resolution=0.125W, Range=0-2047.875W.	0000h	RWS_L
31:24	RSVD	Reserved.	00h	RO
23:22	LIMIT1_TIME_WINDOW_X	Power Limit[0] time window X value, for DDR domain. Actual time_window for RAPL is: (1/1024 seconds) * (1+(x/4)) * (2^y)	0h	RWS_L
21:17	LIMIT1_TIME_WINDOW_Y	Power Limit[0] time window Y value, for DDR domain. Actual time_window for RAPL is: (1/1024 seconds) * (1+(x/4)) * (2^y)	00h	RWS_L
16	RSVD	Reserved.	0h	RO
15	LIMIT1_ENABLE	Power Limit[0] enable bit for DDR domain.	0h	RWS_L
14:0	LIMIT1_POWER	Power Limit[0] for DDR domain. Units=Watts, Format=11.3, Resolution=0.125W, Range=0-2047.875W.	0000h	RWS_L

5.2.43 DRAM—DRAM_ENERGY_STATUS

Accumulates the energy consumed by the DIMMs (summed across all channels).

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access:	ROS_V
Size: 32	Default Value: 00000000h		Address Offset:	58E8h
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:0	JOULES_CONSUMED	Total Joules of energy consumed by all DIMMs: Format = 18.14, Resolution = ~61uJ, Range = 0 to 2.62e5 J.	00000000h	ROS_V

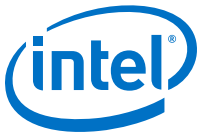

5.2.44 DRAM—DRAM_RAPL_PERF_STATUS

Memory RAPL performance excursion counter. This register can report the performance impact of power limiting.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: ROS_V	
Size: 32	Default Value: 00000000h		Address Offset: 58ECh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:0	DURATION	Throttle duration due to RAPL (sum across all channels), in units of 1/1024 seconds. This data can serve as a proxy for the potential performance impacts of RAPL on memory accesses. This is a real time accumulator that is based on iMC counters at QCLK granularity, thus this register is more accurate than PACKAGE_RAPL_PERF_STATUS.	00000000h	ROS_V

5.2.45 PACKAGE—PACKAGE_RAPL_PERF_STATUS_0_0_0_MCHBAR_PCU

Package RAPL Performance Status Register. This register provides information on the performance impact of the RAPL power limit and indicates the duration for processor went below the requested P-state due to package power constraint.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: ROS_V	
Size: 32	Default Value: 00000000h		Address Offset: 58F0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:0	COUNTS	Counter of the time units within which RAPL was limiting P-states. If limitation occurred anywhere within the time window of 1/1024 seconds, the count will be incremented (limitation on accuracy). This data can serve as a proxy for the potential performance impacts of RAPL on cores performance.	00000000h	ROS_V

5.2.46 CORE—CORE_PERF_LIMIT_REASONS

Indicator of Frequency Clipping in Processor Cores (frequency refers to processor core frequency)

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: ROV; RW0C	
Size: 32	Default Value: 00000000h		Address Offset: 58FCh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:30	RSVD	Reserved.	0h	RO
29	TURBO_TRANSITION_ATTENUATION_LOG	When set, indicates that the corresponding Turbo Transition Attenuation Status bit was set since it was last cleared by software. Software can write 0 to this bit to clear Turbo Transition Attenuation Status.	0h	RW0C
28	MAX_TURBO_LIMIT_LOG	When set, indicates that the corresponding Max Turbo Limit Status bit was set since it was last cleared by software. Software can write 0 to this bit to clear Max Turbo Limit Status.	0h	RW0C

continued...

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: ROV; RW0C	
Size: 32	Default Value: 00000000h		Address Offset: 58FCh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
27	PACKAGE_LEVEL_POWER_LIMITING_PL2_LOG	When set, indicates that the corresponding Package-level Power Limiting PL2 Status bit was set since it was last cleared by software. Software can write 0 to this bit to clear Package-level Power Limiting PL2 Status.	0h	RW0C
26	PACKAGE_LEVEL_POWER_LIMITING_PL1_LOG	When set, indicates that the corresponding Package-level Power Limiting PL1 Status bit was set since it was last cleared by software. Software can write 0 to this bit to clear Package-level Power Limiting PL1 Status.	0h	RW0C
25	CORE_POWER_LIMITING_LOG	When set, indicates that the corresponding Core Power Limiting Status bit was set since it was last cleared by software. Software can write 0 to this bit to clear Core Power Limiting Status.	0h	RW0C
24	EDP_LOG	When set, indicates that the corresponding EDP Status bit was set since it was last cleared by software. Software can write 0 to this bit to clear EDP Status.	0h	RW0C
23	RSVD	Reserved.	0h	RO
22	VR_THERMALERT_LOG	When set, indicates that the corresponding VR Therm Alert Status bit was set since it was last cleared by software. Software can write 0 to this bit to clear VR Therm Alert Status.	0h	RW0C
21	RSVD	Reserved.	0h	RO
20	GRAPHICS_DRIVER_LOG	When set, indicates that the corresponding Graphics Driver status bit was set since it was last cleared by software. Software can write 0 to this bit to clear Graphics Driver Status.	0h	RW0C
19:18	RSVD	Reserved.	0h	RO
17	THERMAL_LOG	When set, indicates that the corresponding Thermal status bit was set since it was last cleared by software. Software can write 0 to this bit to clear Thermal Status.	0h	RW0C
16	PROCHOT_LOG	When set, indicates that the corresponding PROCHOT Status bit is set. Software can write 0 to this bit to clear PROCHOT Status.	0h	RW0C
15:14	RSVD	Reserved.	0h	RO
13	TURBO_TRANSITION_ATTENUATION_STATUS	When set, frequency is reduced below the operating system request due to Turbo transition attenuation. This prevents performance degradation due to frequent operating ratio changes.	0h	ROV
12	MAX_TURBO_LIMIT_STATUS	When set, frequency is reduced below the operating system request due to multi-core turbo limits.	0h	ROV
11	PACKAGE_LEVEL_POWER_LIMITING_PL2_STATUS	When set, frequency is reduced below the operating system request due to package-level power limiting PL2.	0h	ROV
10	PACKAGE_LEVEL_POWER_LIMITING_PL1_STATUS	When set, frequency is reduced below the operating system request due to package-level power limiting PL1.	0h	ROV

continued...

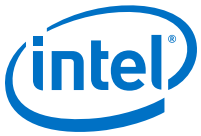

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: ROV; RW0C	
Size: 32	**Default Value: 00000000h**		**Address Offset: 58FCh**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
9	CORE_POWER_LIMITING_STATUS	When set, frequency is reduced below the operating system request due to domain-level power limiting.	0h	ROV
8	EDP_STATUS	When set, frequency is reduced below the operating system request due to electrical design point constraints (e.g. maximum electrical current consumption).	0h	ROV
7	RSVD	Reserved.	0h	RO
6	VR_THERMALERT_STATUS	When set, frequency is reduced below the operating system request due to a thermal alert from the Voltage Regulator.	0h	ROV
5	UTILIZATION_STATUS	When set, frequency is reduced below the operating system request because the processor has detected that utilization is low.	0h	ROV
4	GRAPHICS_DRIVER_STATUS	When set, frequency is reduced below the operating system request due to Processor Graphics driver override.	0h	ROV
3:2	RSVD	Reserved.	0h	RO
1	THERMAL_STATUS	When set, frequency is reduced below the operating system request due to a thermal event.	0h	ROV
0	PROCHOT_STATUS	When set, processor core frequency is reduced below the operating system request due to assertion of external PROCHOT.	0h	ROV

5.2.47 GRAPHICS—GRAPHICS_PERF_LIMIT_REASONS

Indicator of Frequency Clipping in the Processor Graphics (frequency refers to processor graphics frequency)

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: ROV; RW0C	
Size: 32	**Default Value: 00000000h**		**Address Offset: 5900h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:28	RSVD	Reserved.	0h	RO
27	PACKAGE_LEVEL_POWER_LIMITING_PL2_LOG	When set, indicates that the corresponding Package-level Power Limiting PL2 Status bit was set since it was last cleared by software. Software can write 0 to this bit to clear Package-level Power Limiting PL2 Status.	0h	RW0C
26	PACKAGE_LEVEL_POWER_LIMITING_PL1_LOG	When set, indicates that the corresponding Package-level Power Limiting PL1 Status bit was set since it was last cleared by software. Software can write 0 to this bit to clear Package-level Power Limiting PL1 Status.	0h	RW0C
25	GRAPHICS_POWER_LIMITING_LOG	When set, indicates that the corresponding Graphics Power Limiting Status bit was set since it was last cleared by software. Software can write 0 to this bit to clear Graphics Power Limiting Status.	0h	RW0C
24	EDP_LOG	When set, indicates that the corresponding EDP Status bit was set since it was last cleared by software. Software can write 0 to this bit to clear EDP Status.	0h	RW0C

continued...

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: ROV; RW0C	
Size: 32	**Default Value: 00000000h**		**Address Offset: 5900h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
23	RSVD	Reserved.	0h	RO
22	VR_THERMALERT_LOG	When set, indicates that the corresponding VR Therm Alert Status bit was set since it was last cleared by software. Software can write 0 to this bit to clear VR Therm Alert Status.	0h	RW0C
21	RSVD	Reserved.	0h	RO
20	GRAPHICS_DRIVER_LOG	When set, indicates that the corresponding Graphics Driver status bit was set since it was last cleared by software. Software can write 0 to this bit to clear Graphics Driver Status.	0h	RW0C
19:18	RSVD	Reserved.	0h	RO
17	THERMAL_LOG	When set, indicates that the corresponding Thermal status bit was set since it was last cleared by software. Software can write 0 to this bit to clear Thermal Status.	0h	RW0C
16	PROCHOT_LOG	When set, indicates that the corresponding PROCHOT Status bit is set. Software can write 0 to this bit to clear PROCHOT Status.	0h	RW0C
15:12	RSVD	Reserved.	0h	RO
11	PACKAGE_LEVEL_POWER_LIMITING_PL2_STATUS	When set, frequency is reduced below the operating system request due to package-level power limiting PL2.	0h	ROV
10	PACKAGE_LEVEL_POWER_LIMITING_PL1_STATUS	When set, frequency is reduced below the operating system request due to package-level power limiting PL1.	0h	ROV
9	GRAPHICS_POWER_LIMITING_STATUS	When set, frequency is reduced below the operating system request due to domain-level power limiting.	0h	ROV
8	EDP_STATUS	When set, frequency is reduced below the operating system request due to electrical design point constraints (e.g. maximum electrical current consumption).	0h	ROV
7	RSVD	Reserved.	0h	RO
6	VR_THERMALERT_STATUS	When set, frequency is reduced below the operating system request due to a thermal alert from the Voltage Regulator.	0h	ROV
5	RSVD	Reserved.	0h	RO
4	GRAPHICS_DRIVER_STATUS	When set, frequency is reduced below the operating system request due to Processor Graphics driver override.	0h	ROV
3:2	RSVD	Reserved.	0h	RO
1	THERMAL_STATUS	When set, frequency is reduced below the operating system request due to a thermal event.	0h	ROV
0	PROCHOT_STATUS	When set, frequency is reduced below the operating system request due to assertion of external PROCHOT.	0h	ROV

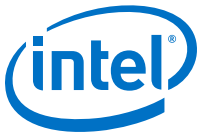

5.2.48 RING—RING_PERF_LIMIT_REASONS

Indicator of Frequency Clipping in the Ring Interconnect (frequency refers to ring interconnect in the uncore)

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: ROV; RW0C	
Size: 32	**Default Value: 00000000h**		**Address Offset: 5904h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:28	RSVD	Reserved.	0h	RO
27	PACKAGE_LEVEL_POWER_LIMITING_PL2_LOG	When set, indicates that the corresponding Package-level Power Limiting PL2 Status bit was set since it was last cleared by software. Software can write 0 to this bit to clear Package-level Power Limiting PL2 Status.	0h	RW0C
26	PACKAGE_LEVEL_POWER_LIMITING_PL1_LOG	When set, indicates that the corresponding Package-level Power Limiting PL1 Status bit was set since it was last cleared by software. Software can write 0 to this bit to clear Package-level Power Limiting PL1 Status.	0h	RW0C
25	RSVD	Reserved.	0h	RO
24	EDP_LOG		0h	RW0C
23	RSVD	Reserved.	0h	RO
22	VR_THERMALERT_LOG	When set, indicates that the corresponding VR Therm Alert Status bit was set since it was last cleared by software. Software can write 0 to this bit to clear VR Therm Alert Status.	0h	RW0C
21:18	RSVD	Reserved.	0h	RO
17	THERMAL_LOG	When set, indicates that the corresponding Thermal status bit was set since it was last cleared by software. Software can write 0 to this bit to clear Thermal Status.	0h	RW0C
16	PROCHOT_LOG	When set, indicates that the corresponding PROCHOT Status bit is set. Software can write 0 to this bit to clear PROCHOT Status.	0h	RW0C
15:12	RSVD	Reserved.	0h	RO
11	PACKAGE_LEVEL_POWER_LIMITING_PL2_STATUS	When set, frequency is reduced below the operating system request due to package-level power limiting PL2.	0h	ROV
10	PACKAGE_LEVEL_POWER_LIMITING_PL1_STATUS	When set, frequency is reduced below the operating system request due to package-level power limiting PL1.	0h	ROV
9	RSVD	Reserved.	0h	RO
8	EDP_STATUS	When set, frequency is reduced below the operating system request due to electrical design point constraints (e.g. maximum electrical current consumption).	0h	ROV
7	RSVD	Reserved.	0h	RO
6	VR_THERMALERT_STATUS	When set, frequency is reduced below the operating system request due to a thermal alert from the Voltage Regulator.	0h	ROV

continued...

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: ROV; RW0C	
Size: 32	Default Value: 00000000h		Address Offset: 5904h	
Bit Range	Acronym	Description	Default	Access
5:2	RSVD	Reserved.	0h	RO
1	THERMAL_STATUS	When set, frequency is reduced below the operating system request due to a thermal event.	0h	ROV
0	PROCHOT_STATUS	When set, frequency is reduced below the operating system request due to assertion of external PROCHOT.	0h	ROV

5.2.49 PRIMARY—PRIMARY_PLANE_ENERGY_STATUS

Reports total energy consumed. The counter will wrap around and continue counting when it reaches its limit.
The energy status is reported in units which are defined in PACKAGE_POWER_SKU_UNIT_MSR[ENERGY_UNIT].

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO_V	
Size: 32	Default Value: 00000000h		Address Offset: 5928h	
Bit Range	Acronym	Description	Default	Access
31:0	DATA	Energy Value	00000000h	RO_V

5.2.50 SECONDARY—SECONDARY_PLANE_ENERGY_STATUS

Reports total energy consumed. The counter will wrap around and continue counting when it reaches its limit.
The energy status is reported in units which are defined in PACKAGE_POWER_SKU_UNIT_MSR[ENERGY_UNIT].

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO_V	
Size: 32	Default Value: 00000000h		Address Offset: 592Ch	
Bit Range	Acronym	Description	Default	Access
31:0	DATA	Energy Value	00000000h	RO_V

5.2.51 PACKAGE—PACKAGE_POWER_SKU

Defines allowed SKU power and timing parameters.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: ROS_V	
Size: 64	Default Value: 0012024000600000h		Address Offset: 5930h	
Bit Range	Acronym	Description	Default	Access
63:55	RSVD	Reserved.	000h	RO
54:48	PKG_MAX_WIN	The maximal time window allowed for the SKU. Higher values will be clamped to this value. x = PKG_MAX_WIN[54:53] y = PKG_MAX_WIN[52:48] The timing interval window is Floating Point number given	12h	ROS_V

continued...

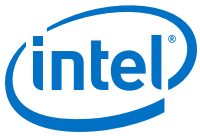

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: ROS_V	
Size: 64	Default Value: 0012024000600000h		Address Offset: 5930h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		by 1.x * power(2,y). The unit of measurement is defined in PACKAGE_POWER_SKU_UNIT_MSR[TIME_UNIT].		
47	RSVD	Reserved.	0h	RO
46:32	PKG_MAX_PWR	The maximal package power setting allowed for the SKU. Higher values will be clamped to this value. The maximum setting is typical (not guaranteed). The units for this value are defined in PACKAGE_POWER_SKU_MSR[PWR_UNIT].	0240h	ROS_V
31	RSVD	Reserved.	0h	RO
30:16	PKG_MIN_PWR	The minimal package power setting allowed for this part. Lower values will be clamped to this value. The minimum setting is typical (not guaranteed). The units for this value are defined in PACKAGE_POWER_SKU_MSR[PWR_UNIT].	0060h	ROS_V
15:0	RSVD	Reserved.	0000h	RO

5.2.52 PACKAGE—PACKAGE_POWER_SKU_UNIT

Defines units for calculating SKU power and timing parameters.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO_V	
Size: 32	Default Value: 000A0E03h		Address Offset: 5938h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:20	RSVD	Reserved.	000h	RO
19:16	TIME_UNIT	Time Units used for power control registers. The actual unit value is calculated by 1 s / Power(2,TIME_UNIT). The default value of Ah corresponds to 976 usec.	Ah	RO_V
15:13	RSVD	Reserved.	0h	RO
12:8	ENERGY_UNIT	Energy Units used for power control registers. The actual unit value is calculated by 1 J / Power(2,ENERGY_UNIT). The default value of 14 corresponds to Ux.14 number.	0eh	RO_V
7:4	RSVD	Reserved.	0h	RO
3:0	PWR_UNIT	Power Units used for power control registers. The actual unit value is calculated by 1 W / Power(2,PWR_UNIT). The default value of 0011b corresponds to 1/8 W.	3h	RO_V

5.2.53 PACKAGE—PACKAGE_ENERGY_STATUS

Package energy consumed by the entire processor (including IA, Integrated Graphics and Uncore). The counter will wrap around and continue counting when it reaches its limit.
The energy status is reported in units which are defined in PACKAGE_POWER_SKU_UNIT_MSR[ENERGY_UNIT].

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO_V	
Size: 32	Default Value: 00000000h		Address Offset: 593Ch	
Bit Range	Acronym	Description	Default	Access
31:0	DATA	Energy Value	00000000h	RO_V

5.2.54 GT—GT_PERF_STATUS_0_0_0_MCHBAR_PCU

P-state encoding for the Secondary Power Plane's current PLL frequency and the current VID.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO_V	
Size: 32	Default Value: 00000000h		Address Offset: 5948h	
Bit Range	Acronym	Description	Default	Access
31:16	RSVD	Reserved.	0000h	RO
15:8	RP_STATE_RATIO	Ratio of the current RP-state. When the graphics engine is in RC6, this field reflects the last ratio in use.	00h	RO_V
7:0	RP_STATE_VOLTAGE	Voltage of the current RP-state.	00h	RO_V

5.2.55 IA32—IA32_PLATFORM_ID

Indicates the platform that the processor is intended for.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: ROS_V	
Size: 64	Default Value: 0000000000000000h		Address Offset: 5950h	
Bit Range	Acronym	Description	Default	Access
63:53	RSVD	Reserved.	000h	RO
52:50	PLATFORMID	Platform ID Bits (RO) The field gives information concerning the intended platform for the processor. Bits 52, 51, 50 are concatenated to form the platform ID as shown below: 000b Processor Flag 0, 001b Processor Flag 1, 010b Processor Flag 2, 011b Processor Flag 3, 100b Processor Flag 4, 101b Processor Flag 5, 110b Processor Flag 6, 111b Processor Flag 7	0h	ROS_V
49:0	RSVD	Reserved.	0000000000000h	RO

5.2.56 RP—RP_STATE_LIMITS_0_0_0_MCHBAR_PCU

This register allows SW to limit the maximum base frequency for the Integrated GFX Engine (GT) allowed during run-time.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW	
Size: 32	Default Value: 000000FFh		Address Offset: 5994h	
Bit Range	Acronym	Description	Default	Access
31:8	RSVD	Reserved.	000000h	RO
7:0	RPSTT_LIM	This field indicates the maximum base frequency limit for the Integrated GFX Engine (GT) allowed during run-time.	FFh	RW

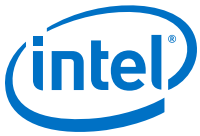

5.2.57 RP—RP_STATE_CAP_0_0_0_MCHBAR_PCU

This register contains the maximum base frequency capability for the Integrated GFX Engine (GT).

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO	
Size: 32	Default Value: 00000000h		Address Offset: 5998h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:24	RSVD	Reserved.	00h	RO
23:16	RPN_CAP	This field indicates the maximum RPN base frequency capability for the Integrated GFX Engine (GT). Values are in units of 100 MHz.	00h	RO
15:8	RP1_CAP	This field indicates the maximum RP1 base frequency capability for the Integrated GFX Engine (GT). Values are in units of 100 MHz.	00h	RO
7:0	RP0_CAP	This field indicates the maximum RP0 base frequency capability for the Integrated GFX Engine (GT). Values are in units of 100 MHz.	00h	RO

5.2.58 TEMPERATURE—TEMPERATURE_TARGET

Register to control the temperature at which the Thermal Control Circuit (TCC) must be activated and monitor temperature levels.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO_V	
Size: 32	Default Value: 00000000h		Address Offset: 599Ch	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:30	RSVD	Reserved.	0h	RO
29:24	TCC_ACTIVATION_OFFSET	This field indicates to the processor the offset from the factory set TCC activation temperature at which the Thermal Control Circuit (TCC) must be activated. TCC will be activated at a temp (TCC Activation Temperature - TCC Activation Offset). Default value is 0 causing TCC to activate at TCC Activation temperature. This field is valid only when PLATFORM_INFO[30] is set indicating feature availability. Note: This field needs to be programmed prior to setting the BIOS_RESET_CPL bit.	00h	RO_V
23:16	TCC_ACTIVATION_TEMPERATURE	This is the factory set temperature at which the Thermal Control Circuit (TCC) will assert the PROCHOT# signal and activate the Adaptive Thermal monitor.	00h	RO_V
15:8	TEMP_CONTROL_OFFSET	Temperature Control Offset (RO) When the processor temperature is read by PECI, this bit field specifies a temperature value relative to the PROCHOT# signal activation temperature. Using the appropriate temperature sensing feature, when the reported temperature is less than the value in this field, the Advanced Fan Speed Controller must operate the processor fan at the maximum RPM. This is an unsigned value which is measured in 1°C increments clipped at 0°.	00h	RO_V
7:0	RSVD	Reserved.	00h	RO

5.2.59 VR—VR_CURRENT_CONFIG

Limitation on the maximum current consumption of the primary power plane.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO_V; ROS_V	
Size: 64	Default Value: 0000000000000190h		Address Offset: 59B0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63	RSVD	Reserved.	0h	RO
62	PS4_ENABLE	Enable for PS4 (current is regulated to ~0A). 0 - PS4 Disabled. If Disabled, C10 will use VR_ENABLE=0 instead. 1 - PS4 Enabled. Default is disabled.	0h	RO_V
61:52	PSI3_THRESHOLD	Maximum current in 1A units supported at external voltage regulator PS3	000h	RO_V
51:42	PSI2_THRESHOLD	Maximum current in 1A units supported at external voltage regulator PS2	000h	RO_V
41:32	PSI1_THRESHOLD	Maximum current in 1A units supported at external voltage regulator PS1	000h	RO_V
31	LOCK	This bit will lock the CURRENT_LIMIT settings in this register and will also lock this setting. This means that once set to 1b, the CURRENT_LIMIT setting and this bit become Read Only until the next Warm Reset.	0h	ROS_V
30:13	RSVD	Reserved.	00000h	RO
12:0	CURRENT_LIMIT	Current limitation in 0.125 A increments. This field is locked by VR_CURRENT_CONFIG[LOCK]. When the LOCK bit is set to 1b, this field becomes Read Only.	0190h	RO_V

5.2.60 IA32—IA32_THERM_STATUS

Contains status information about the processor's thermal sensor and automatic thermal monitoring facilities.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: ROV; RW0C; RO	
Size: 32	Default Value: 08000000h		Address Offset: 59C0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31	VALID	This bit indicates that the TEMPERATURE field is valid.	0h	ROV
30:27	RESOLUTION	Supported resolution in degrees C.	1h	RO
26:23	RSVD	Reserved.	0h	RO
22:16	TEMPERATURE	This is a temperature offset in degrees C below theTJ Max temperature. This number is meaningful only if VALID bit in this register is set.	00h	ROV
15:12	RSVD	Reserved.	0h	RO
11	POWER_LIMITATION_LOG	Sticky log bit that asserts when PP P-State is below the (max P-State - offset). Set by HW cleared by SW. Not Supported in A-Step.	0h	RW0C

continued...

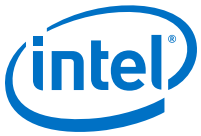

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: ROV; RW0C; RO	
Size: 32	**Default Value: 08000000h**		**Address Offset: 59C0h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
10	POWER_LIMITATION_STATUS	Status log bit that notifies if the PP P-state is below the (max P-state - offset) Not supported in A-Step.	0h	ROV
9	THRESHOLD2_LOG	Sticky log bit that asserts on a 0 to 1 or a 1 to 0 transition of the THRESHOLD2_STATUS bit. This bit is set by HW and cleared by SW.	0h	RW0C
8	THRESHOLD2_STATUS	Indicates that the current temperature is higher than or equal to Threshold 2 temperature.	0h	ROV
7	THRESHOLD1_LOG	Sticky log bit that asserts on a 0 to 1 or a 1 to 0 transition of the THRESHOLD1_STATUS bit. This bit is set by HW and cleared by SW.	0h	RW0C
6	THRESHOLD1_STATUS	Indicates that the current temperature is higher than or equal to Threshold 1 temperature.	0h	ROV
5	CRIT_TEMP_LOG	Sticky log bit indicating that the processor operating out of its thermal specification since the last time this bit was cleared. This bit is set by HW on a 0 to 1 transition of OUT_OF_SPEC_STATUS.	0h	RW0C
4	CRIT_TEMP_STATUS	Status bit indicating that the processor is operating out of its thermal specification. Once set, this bit should only clear on a reset.	0h	ROV
3	PROCHOT_LOG	Sticky log bit indicating that xxPROCHOT# has been asserted since the last time this bit was cleared by SW. This bit is set by HW on a 0 to 1 transition of PROCHOT_STATUS.	0h	RW0C
2	PROCHOT_STATUS	Status bit indicating that xxPROCHOT# is currently being asserted.	0h	ROV
1	THERMAL_MONITOR_LOG	Sticky log bit indicating that the core has seen a thermal monitor event since the last time SW cleared this bit. This bit is set by HW on a 0 to 1 transition of THERMAL_MONITOR_STATUS.	0h	RW0C
0	THERMAL_MONITOR_STATUS	Status bit indicating that the Thermal Monitor has tripped and is currently thermally throttling.	0h	ROV

5.2.61 IA32—IA32_THERM_INTERRUPT

Enables and disables the generation of an interrupt on temperature transitions detected with the processor's thermal sensors and thermal monitor.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW	
Size: 32	**Default Value: 00000000h**		**Address Offset: 59C4h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:25	RSVD	Reserved.	00h	RO
24	POWER_INT_ENABLE	When this bit is set, a thermal interrupt will be sent upon throttling due to power limitations.	0h	RW

continued...

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 59C4h	
Bit Range	Acronym	Description	Default	Access
23	THRESHOLD_2_INT_ENABLE	Controls the generation of a thermal interrupt whenever the Thermal Threshold 2 Temperature is crossed.	0h	RW
22:16	THRESHOLD_2_REL_TEMP	This value indicates the offset in degrees below TJ Max Temperature that should trigger a Thermal Threshold 2 trip.	00h	RW
15	THRESHOLD_1_INT_ENABLE	Controls the generation of a thermal interrupt whenever the Thermal Threshold 1 Temperature is crossed.	0h	RW
14:8	THRESHOLD_1_REL_TEMP	This value indicates the offset in degrees below TJ Max Temperature that should trigger a Thermal Threshold 1 trip.	00h	RW
7:5	RSVD	Reserved.	0h	RO
4	CRIT_TEMP_INT_ENABLE	Thermal interrupt enable for the critical temperature condition which is stored in the Critical Temperature Status bit in IA32_THERM_STATUS.	0h	RW
3	RSVD	Reserved.	0h	RO
2	PROCHOT_INT_ENABLE	Bidirectional PROCHOT# assertion interrupt enable. If set, a thermal interrupt is delivered on the rising edge of xxPROCHOT#.	0h	RW
1	LOW_TEMP_INT_ENABLE	Enables a thermal interrupt to be generated on the transition from a high-temperature to a low-temperature when set, where 'high temperature' is dictated by the thermal monitor trip temperature.	0h	RW
0	HIGH_TEMP_INT_ENABLE	Enables a thermal interrupt to be generated on the transition from a low-temperature to a high-temperature when set, where 'high temperature' is dictated by the thermal monitor trip temperature.	0h	RW

5.2.62 SSKPD—SSKPD_0_0_0_MCHBAR_PCU

This register holds 64 writable bits with no functionality behind them. It is for the convenience of BIOS and graphics drivers.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS	
Size: 64	Default Value: 0000000000000000h		Address Offset: 5D10h	
Bit Range	Acronym	Description	Default	Access
63:0	SKPD	4 WORDs of data storage.	0000000000000000h	RWS

5.2.63 CONFIG—CONFIG_TDP_NOMINAL_0_0_0_MCHBAR_PCU

This register is used to indicate the Nominal Configurable TDP ratio available for this specific sku. System BIOS must use this value while building the _PSS table if the feature is enabled.

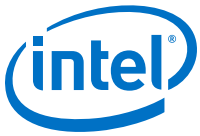

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO_V	
Size: 32	Default Value: 00000000h		Address Offset: 5F3Ch	
Bit Range	Acronym	Description	Default	Access
31:8	RSVD	Reserved.	000000h	RO
7:0	TDP_RATIO	Nominal TDP level ratio to be used for this specific processor (in units of 100 MHz). Note: A value of 0 in this field indicates invalid/undefined TDP point	00h	RO_V

5.2.64 CONFIG—CONFIG_TDP_LEVEL1_0_0_0_MCHBAR_PCU

Level 1 configurable TDP settings

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO_V	
Size: 64	Default Value: 0000000000000000h		Address Offset: 5F40h	
Bit Range	Acronym	Description	Default	Access
63	RSVD	Reserved.	0h	RO
62:48	PKG_MIN_PWR	Min pkg power setting allowed for this config TDP level. Lower values will be clamped up to this value. Units defined in PACKAGE_POWER_SKU_MSR[PWR_UNIT]. Similar to PACKAGE_POWER_SKU[PKG_MIN_PWR].	0000h	RO_V
47	RSVD	Reserved.	0h	RO
46:32	PKG_MAX_PWR	Max pkg power setting allowed for this config TDP level1. Higher values will be clamped down to this value. Units defined in PACKAGE_POWER_SKU_MSR[PWR_UNIT]. Similar to PACKAGE_POWER_SKU[PKG_MAX_PWR].	0000h	RO_V
31:24	RSVD	Reserved.	00h	RO
23:16	TDP_RATIO	TDP ratio for config tdp level 1.	00h	RO_V
15	RSVD	Reserved.	0h	RO
14:0	PKG_TDP	Power for this TDP level. Units defined in PACKAGE_POWER_SKU_MSR[PWR_UNIT] Similar to PACKAGE_POWER_SKU[PKG_TDP]	0000h	RO_V

5.2.65 CONFIG—CONFIG_TDP_LEVEL2_0_0_0_MCHBAR_PCU

Level 2 configurable TDP settings

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO_V	
Size: 64	Default Value: 0000000000000000h		Address Offset: 5F48h	
Bit Range	Acronym	Description	Default	Access
63	RSVD	Reserved.	0h	RO
62:48	PKG_MIN_PWR	Min pkg power setting allowed for this config TDP level 2. Lower values will be clamped up to this value. Units defined in PACKAGE_POWER_SKU_MSR[PWR_UNIT]. Similar to PACKAGE_POWER_SKU[PKG_MIN_PWR].	0000h	RO_V
47	RSVD	Reserved.	0h	RO

continued...

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RO_V	
Size: 64	Default Value: 0000000000000000h		Address Offset: 5F48h	
Bit Range	Acronym	Description	Default	Access
46:32	PKG_MAX_PWR	Max pkg power setting allowed for config TDP level 2. Higher values will be clamped down to this value. Units defined in PACKAGE_POWER_SKU_MSR[PWR_UNIT]. Similar to PACKAGE_POWER_SKU[PKG_MAX_PWR].	0000h	RO_V
31:24	RSVD	Reserved.	00h	RO
23:16	TDP_RATIO	TDP ratio for level 2.	00h	RO_V
15	RSVD	Reserved.	0h	RO
14:0	PKG_TDP	Power for this TDP level. Units defined in PACKAGE_POWER_SKU_MSR[PWR_UNIT] Similar to PACKAGE_POWER_SKU[PKG_TDP].	0000h	RO_V

5.2.66 CONFIG—CONFIG_TDP_CONTROL_0_0_0_MCHBAR_PCU

Rd/Wr register to allow platform SW to select TDP point and set lock

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_L; RW_KL	
Size: 32	Default Value: 00000000h		Address Offset: 5F50h	
Bit Range	Acronym	Description	Default	Access
31	CONFIG_TDP_LOCK	Config TDP level select lock 0 - unlocked. 1 - locked till next reset.	0h	RW_KL
30:2	RSVD	Reserved.	00000000h	RO
1:0	TDP_LEVEL	Config TDP level selected 0 = nominal TDP level (default) 1 = Level from CONFIG_TDP_LEVEL_1 2 = Level from CONFIG_TDP_LEVEL_2 3 = reserved	0h	RWS_L

5.2.67 TURBO—TURBO_ACTIVATION_RATIO_0_0_0_MCHBAR_PCU

Read/write register to allow MSR/MMIO access to ACPI P-state notify (PCS 33).

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RWS_L; RWS_KL	
Size: 32	Default Value: 00000000h		Address Offset: 5F54h	
Bit Range	Acronym	Description	Default	Access
31	TURBO_ACTIVATION_RATIO_LOCK	Lock this MSR until next reset 0 - unlocked 1 - locked	0h	RWS_KL
30:8	RSVD	Reserved.	000000h	RO
7:0	MAX_NON_TURBO_RATIO	Processor will treat any P-state request above this ratio as a request for max turbo 0 is special encoding which disables the feature.	00h	RWS_L

5.2.68 DDR—Memory Thermal Camarillo Status

Status and log bits of memory thermal interrupt enabled through configuration of DDR_THERM_THRESHOLDS_CONFIG.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: ROV; RW0C	
Size: 32	**Default Value: 00000000h**		**Address Offset: 6204h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:12	RSVD	Reserved.	00000h	RO
11	THRESHOLD2_LOG	Sticky log bit that asserts on a 0 to 1 transition of the THRESHOLD2_STATUS bit. HW controls this transition.	0h	RW0C
10	THRESHOLD2_STATUS	Status bit indicating that the hottest DIMM has crossed the THRESHOLD2 value programmed in bits 20:13 of DDR_THERM_CAMARILLO_INTERRUPT.	0h	ROV
9	THRESHOLD1_LOG	Sticky log bit that asserts on a 0 to 1 transition of the THRESHOLD1_STATUS bit. HW controls this transition.	0h	RW0C
8	THRESHOLD1_STATUS	Status bit indicating that the hottest DIMM has crossed the THRESHOLD1 value programmed in bits 11:4 of DDR_THERM_CAMARILLO_INTERRUPT.	0h	ROV
7	OOS_TEMP_LOG	Sticky log bit that asserts on a 0 to 1 transition of the OOS_TEMP_STATUS bit. HW controls this transition.	0h	RW0C
6	OOS_TEMP_STATUS	Status bit indicating that MR4 is currently indicating at least one DRAM with high temperature which is beyond the operating range. This can only occur currently when MAD_CHNL.LPDDR=1 and DDR_PTM_CTL.DISABLE_DRAM_TS=0.	0h	ROV
5	REFRESH2X_LOG	Sticky log bit that asserts on a 0 to 1 transition of the REFRESH2X_STATUS bit. HW controls this transition.	0h	RW0C
4	REFRESH2X_STATUS	Status bit indicating that the DIMM refresh rate has crossed the boundary (in either direction) between 1x or lower refresh rate, and higher than 1x refresh rate. The name is misleading for LPDDR where we may go above 2x refresh rate.	0h	ROV
3	HOT_THRESHOLD_LOG	Sticky log bit that asserts on a 0 to 1 transition of the HOT_THRESHOLD_STATUS bit. HW controls this transition.	0h	RW0C
2	HOT_THRESHOLD_STATUS	Status bit indicating that the DDR temperature is higher than or equal to the DDR Hot threshold defined in DDR_THERM_THRESHOLDS_CONFIG.	0h	ROV
1	WARM_THRESHOLD_LOG	Sticky log bit that asserts on a 0 to 1 transition of the WARM_THRESHOLD_STATUS bit. HW controls this transition.	0h	RW0C
0	WARM_THRESHOLD_STATUS	Status bit indicating that the DDR temperature is higher than or equal to the DDR Warm threshold defined in DDR_THERM_THRESHOLDS_CONFIG.	0h	ROV

5.2.69 CRDTCTL4—IOTrk and RRTrk shared credits

This register will have the minimum Read Return Tracker credits for each of the PEG/DMI/GSA streams.

B/D/F/Type: 0/0/0/MEM/MCHBAR			Access: RW_L	
Size: 32	Default Value: 00000856h		Address Offset: 7410h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:13	RSVD	Reserved.	00000h	RO
12:6	RRTRK_SHRD	Number of RRTrk entries available to be shared across all VC.	21h	RW_L
5:0	IOTRK_SHRD	Number of IOTrk entries available to be shared across all VCs.	16h	RW_L

5.3 GFXVTBAR Registers Summary

Offset	Register ID—Description	Default Value	Access
0	VER—Version Register on page 299	00000010h	RO
8	CAP—Capability Register on page 299	01C0000C40660462h	RO; ROV
10	ECAP—Extended Capability Register on page 302	0000007E3FF0505Eh	RO; ROV
18	GCMD—Global Command Register on page 304	00000000h	WO; RO
1C	GSTS—Global Status Register on page 307	00000000h	RO_V; ROV; RO
20	RTADDR—Root-Entry Table Address Register on page 308	0000000000000000h	RW_V; RW
28	CCMD—Context Command Register on page 309	0800000000000000h	RW; ROV; RW_V
34	FSTS—Fault Status Register on page 311	00000000h	RW1CS; ROSV; RO
38	FECTL—Fault Event Control Register on page 312	80000000h	ROV; RW
3C	FEDATA—Fault Event Data Register on page 313	00000000h	RW
40	FEADDR—Fault Event Address Register on page 313	00000000h	RW
44	FEUADDR—Fault Event Upper Address Register on page 314	00000000h	RW
58	AFLOG—Advanced Fault Log Register on page 314	0000000000000000h	RO
64	PMEN—Protected Memory Enable Register on page 314	00000000h	ROV; RW
68	PLMBASE—Protected Low-Memory Base Register on page 315	00000000h	RW
6C	PLMLIMIT—Protected Low-Memory Limit Register on page 316	00000000h	RW
70	PHMBASE—Protected High-Memory Base Register on page 316	0000000000000000h	RW
78	PHMLIMIT—Protected High-Memory Limit Register on page 317	0000000000000000h	RW
80	IQH—Invalidation Queue Head Register on page 318	0000000000000000h	ROV
88	IQT—Invalidation Queue Tail Register on page 318	0000000000000000h	RW_L
90	IQA—Invalidation Queue Address Register on page 319	0000000000000000h	RW_L
9C	ICS—Invalidation Completion Status Register on page 319	00000000h	RW1CS

continued...

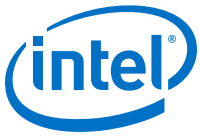

Offset	Register ID—Description	Default Value	Access
A0	IECTL—Invalidation Event Control Register on page 319	80000000h	ROV; RW_L
A4	IEDATA—Invalidation Event Data Register on page 320	00000000h	RW_L
A8	IEADDR—Invalidation Event Address Register on page 320	00000000h	RW_L
AC	IEUADDR—Invalidation Event Upper Address Register on page 321	00000000h	RW_L
B8	IRTA—Interrupt Remapping Table Address Register on page 321	0000000000000000h	RW_L; ROV
400	FRCDL—Fault Recording Low Register on page 322	0000000000000000h	ROSV
408	FRCDH—Fault Recording High Register on page 322	0000000000000000h	ROSV; ROV; RW1CS
500	IVA—Invalidate Address Register on page 324	0000000000000000h	RW
508	IOTLB—IOTLB Invalidate Register on page 325	0200000000000000h	RW; ROV; RW_V
FF0	ARCHDIS—DMA Remap Engine Policy Control on page 327	00000001h	RW_L; RO_KFW; RO; RW_KL
FF4	UARCHDIS—DMA Remap Engine Policy Control on page 329	00000000h	RO; RW_L

5.3.1 VER—Version Register

Register to report the architecture version supported. Backward compatibility for the architecture is maintained with new revision numbers, allowing software to load remapping hardware drivers written for prior architecture versions.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RO	
Size: 32	Default Value: 00000010h		Address Offset: 0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:8	RSVD	Reserved.	000000h	RO
7:4	MAJOR	Indicates supported architecture version.	1h	RO
3:0	MINOR	Indicates supported architecture minor version.	0h	RO

5.3.2 CAP—Capability Register

Register to report general remapping hardware capabilities

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RO; ROV	
Size: 64	Default Value: 01C0000C40660462h		Address Offset: 8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:59	RSVD	Reserved.	00h	RO
58	SL64KP	A value of 1 in this field indicates 64-KByte page size is supported for second-level translation.	0h	RO
57	FL64KP	A value of 1 in this field indicates 64-KByte page size is supported for first-level translation.	0h	ROV
56	FL1GP	A value of 1 in this field indicates 1-GByte page size is supported for first-level translation.	1h	ROV

continued...

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RO; ROV	
Size: 64	Default Value: 01C0000C40660462h		Address Offset: 8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
55	DRD	0: Hardware does not support draining of DMA read requests. 1: Hardware supports draining of DMA read requests.	1h	RO
54	DWD	0: Hardware does not support draining of DMA write requests. 1: Hardware supports draining of DMA write requests.	1h	RO
53:48	MAMV	The value in this field indicates the maximum supported value for the Address Mask (AM) field in the Invalidation Address register (IVA_REG) and IOTLB Invalidation Descriptor (iotlb_inv_dsc). This field is valid only when the PSI field in Capability register is reported as Set.	00h	RO
47:40	NFR	Number of fault recording registers is computed as N+1, where N is the value reported in this field. Implementations must support at least one fault recording register (NFR = 0) for each remapping hardware unit in the platform. The maximum number of fault recording registers per remapping hardware unit is 256.	00h	RO
39	PSI	0: Hardware supports only domain and global invalidates for IOTLB 1: Hardware supports page selective, domain and global invalidates for IOTLB. Hardware implementations reporting this field as set are recommended to support a Maximum Address Mask Value (MAMV) value of at least 9.	0h	RO
38	RSVD	Reserved.	0h	RO
37:34	SLLPS	This field indicates the super page sizes supported by hardware. A value of 1 in any of these bits indicates the corresponding super-page size is supported. The super-page sizes corresponding to various bit positions within this field are: 0: 21-bit offset to page frame (2MB) 1: 30-bit offset to page frame (1GB) 2: 39-bit offset to page frame (512GB) 3: 48-bit offset to page frame (1TB) Hardware implementations supporting a specific super-page size must support all smaller super-page sizes, i.e. only valid values for this field are 0000b, 0001b, 0011b, 0111b, 1111b.	3h	ROV
33:24	FRO	This field specifies the location to the first fault recording register relative to the register base address of this remapping hardware unit. If the register base address is X, and the value reported in this field is Y, the address for the first fault recording register is calculated as X+(16*Y).	040h	RO
23	RSVD	Reserved.	0h	RO
22	ZLR	0: Indicates the remapping hardware unit blocks (and treats as fault) zero length DMA read requests to write-only pages. 1: Indicates the remapping hardware unit supports zero	1h	RO

continued...

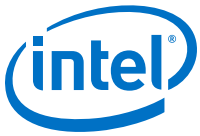

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RO; ROV	
Size: 64	**Default Value: 01C0000C40660462h**		**Address Offset: 8h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		length DMA read requests to write-only pages. DMA remapping hardware implementations are recommended to report ZLR field as Set.		
21:16	MGAW	This field indicates the maximum DMA virtual addressability supported by remapping hardware. The Maximum Guest Address Width (MGAW) is computed as (N +1), where N is the value reported in this field. For example, a hardware implementation supporting 48-bit MGAW reports a value of 47 (101111b) in this field. If the value in this field is X, untranslated and translated DMA requests to addresses above 2^(x+1)-1 are always blocked by hardware. Translations requests to address above 2^(x+1)-1 from allowed devices return a null Translation Completion Data Entry with R=W=0. Guest addressability for a given DMA request is limited to the minimum of the value reported through this field and the adjusted guest address width of the corresponding page-table structure. (Adjusted guest address widths supported by hardware are reported through the SAGAW field). Implementations are recommended to support MGAW at least equal to the physical addressability (host address width) of the platform.	26h	RO
15:13	RSVD	Reserved.	0h	RO
12:8	SAGAW	This 5-bit field indicates the supported adjusted guest address widths (which in turn represents the levels of page-table walks for the 4KB base page size) supported by the hardware implementation. A value of 1 in any of these bits indicates the corresponding adjusted guest address width is supported. The adjusted guest address widths corresponding to various bit positions within this field are: 0: 30-bit AGAW (2-level page table) 1: 39-bit AGAW (3-level page table) 2: 48-bit AGAW (4-level page table) 3: 57-bit AGAW (5-level page table) 4: 64-bit AGAW (6-level page table) Software must ensure that the adjusted guest address width used to setup the page tables is one of the supported guest address widths reported in this field.	04h	RO
7	CM	0: Not-present and erroneous entries are not cached in any of the renmapping caches. Invalidations are not required for modifications to individual not present or invalid entries. However, any modifications that result in decreasing the effective permissions or partial permission increases require invalidations for them to be effective. 1: Not-present and erroneous mappings may be cached in the remapping caches. Any software updates to the remapping structures (including updates to "not-present" or erroneous entries) require explicit invalidation. Hardware implementations of this architecture must support a value of 0 in this field.	0h	RO
6	PHMR	0: Indicates protected high-memory region is not supported. 1: Indicates protected high-memory region is supported.	1h	RO

continued...

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RO; ROV	
Size: 64	Default Value: 01C0000C40660462h		Address Offset: 8h	
Bit Range	Acronym	Description	Default	Access
5	PLMR	0: Indicates protected low-memory region is not supported. 1: Indicates protected low-memory region is supported.	1h	RO
4	RWBF	0: Indicates no write-buffer flushing is needed to ensure changes to memory-resident structures are visible to hardware. 1: Indicates software must explicitly flush the write buffers to ensure updates made to memory-resident remapping structures are visible to hardware.	0h	RO
3	AFL	0: Indicates advanced fault logging is not supported. Only primary fault logging is supported. 1: Indicates advanced fault logging is supported.	0h	RO
2:0	ND	000b: Hardware supports 4-bit domain-ids with support for up to 16 domains. 001b: Hardware supports 6-bit domain-ids with support for up to 64 domains. 010b: Hardware supports 8-bit domain-ids with support for up to 256 domains. 011b: Hardware supports 10-bit domain-ids with support for up to 1024 domains. 100b: Hardware supports 12-bit domain-ids with support for up to 4K domains. 100b: Hardware supports 14-bit domain-ids with support for up to 16K domains. 110b: Hardware supports 16-bit domain-ids with support for up to 64K domains. 111b: Reserved.	2h	RO

5.3.3 ECAP—Extended Capability Register

Register to report remapping hardware extended capabilities

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RO; ROV	
Size: 64	Default Value: 0000007E3FF0505Eh		Address Offset: 10h	
Bit Range	Acronym	Description	Default	Access
63:40	RSVD	Reserved.	000000h	RO
39:35	PSS	This field reports the PASID size supported by the remapping hardware for requestswith- PASID. A value of N in this field indicates hardware supports PASID field of N+1 bits (For example, value of 7 in this field, indicates 8-bit PASIDs are supported). Requests-with-PASID with PASID value beyond the limit specified by this field are treated as error by the remapping hardware. This field is valid only when PASID field is reported as Set.	0Fh	RO
34	EAFS	0: Hardware does not support the extended-accessed (EA) bit in first-level paging-structure entries. 1: Hardware supports the extendedaccessed (EA) bit in first-level paging-structure entries. This field is valid only when PASID field is reported as Set.	1h	ROV

continued...

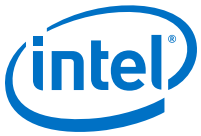

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RO; ROV	
Size: 64	**Default Value: 0000007E3FF0505Eh**		**Address Offset: 10h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
33	NWFS	0: Hardware ignores the "No Write" (NW) flag in Device-TLB translationrequests, and behaves as if NW is always 0. 1: Hardware supports the "No Write" (NW) flag in Device-TLB translationrequests. This field is valid only when Device-TLB support (DT) field is reported as Set.	1h	ROV
32	POT	0: Hardware does not support PASID-only Translation Type in extended-context-entries 1: Hardware supports PASID-only Translation Type in extended-context-entries	0h	RO
31	SRS	0: H/W does not support requests-with-PASID seeking supervisor privilege 1: H/W supports requests-with-PASID seeking supervisor privilege	0h	RO
30	ERS	0: H/W does not support requests seeking execute permission 1: H/W supports requests seeking execute permission	0h	RO
29	PRS	0: Hardware does not support Page Requests 1: Hardware supports Page Requests	1h	ROV
28	PASID	0: Hardware does not support process address space IDs. 1: Hardware supports Process Address Space IDs.	1h	ROV
27	DIS	0: Hardware does not support deferred invalidations of IOTLB and Device-TLB. 1: Hardware supports deferred invalidations of IOTLB and Device-TLB.	1h	ROV
26	NEST	0: Hardware does not support nested translations. 1: Hardware supports nested translations.	1h	ROV
25	MTS	0: Hardware does not support Memory Type 1: Hardware supports Memory Type	1h	ROV
24	ECS	0: Hardware does not support extended-root-entries and Extended Context-Entries 1: Hardware supports extended-root-entries and Extended Context-Entries	1h	ROV
23:20	MHMV	The value in this field indicates the maximum supported value for the Handle Mask (HM) field in the interrupt entry cache invalidation descriptor (iec_inv_dsc). This field is valid only when the IR field in Extended Capability register is reported as Set.	Fh	RO
19:18	RSVD	Reserved.	0h	RO
17:8	IRO	This field specifies the offset to the IOTLB registers relative to the register base address of this remapping hardware unit. If the register base address is X, and the value reported in this field is Y, the address for the first IOTLB invalidation register is calculated as X+(16*Y).	050h	RO
7	SC	0: Hardware does not support 1-setting of the SNP field in the page-table entries. 1: Hardware supports the 1-setting of the SNP field in the page-table entries.	0h	RO

continued...

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RO; ROV	
Size: 64	Default Value: 0000007E3FF0505Eh		Address Offset: 10h	
Bit Range	Acronym	Description	Default	Access
6	PT	0: Hardware does not support pass-through translation type in context entries. 1: Hardware supports pass-through translation type in context entries.	1h	ROV
5	RSVD	Reserved.	0h	RO
4	EIM	0: On Intel®64 platforms, hardware supports only 8-bit APIC-IDs (xAPIC mode). 1: On Intel®64 platforms, hardware supports 32-bit APIC-IDs (x2APIC mode). This field is valid only on Intel®64 platforms reporting Interrupt Remapping support (IR field Set).	1h	ROV
3	IR	0: Hardware does not support interrupt remapping. 1: Hardware supports interrupt remapping. Implementations reporting this field as Set must also support Queued Invalidation (QI).	1h	ROV
2	DT	0: Hardware does not support device-IOTLBs. 1: Hardware supports Device-IOTLBs. Implementations reporting this field as Set must also support Queued Invalidation (QI).	1h	ROV
1	QI	0: Hardware does not support queued invalidations. 1: Hardware supports queued invalidations.	1h	ROV
0	C	This field indicates if hardware access to the root, context, page-table and interrupt-remap structures are coherent (snooped) or not. 0: Indicates hardware accesses to remapping structures are non-coherent. 1: Indicates hardware accesses to remapping structures are coherent. Hardware access to advanced fault log and invalidation queue are always coherent.	0h	RO

5.3.4 GCMD—Global Command Register

Register to control remapping hardware. If multiple control fields in this register need to be modified, software must serialize the modifications through multiple writes to this register.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: WO; RO	
Size: 32	Default Value: 00000000h		Address Offset: 18h	
Bit Range	Acronym	Description	Default	Access
31	TE	Software writes to this field to request hardware to enable/disable DMA-remapping: 0: Disable DMA remapping 1: Enable DMA remapping Hardware reports the status of the translation enable operation through the TES field in the Global Status register. There may be active DMA requests in the platform when software updates this field. Hardware must enable or disable remapping logic only at determinstic transaction boundaries, so that any in-flight transaction is either	0h	WO

continued...

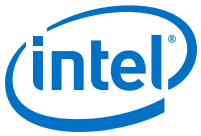

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: WO; RO	
Size: 32	**Default Value: 00000000h**		**Address Offset: 18h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		subject to remapping or not at all. Hardware implementations supporting DMA draining must drain any in-flight DMA read/write requests queued within the Root-Complex before completing the translation enable command and reflecting the status of the command through the TES field in the Global Status register. The value returned on a read of this field is undefined.		
30	SRTP	Software sets this field to set/update the root-entry table pointer used by hardware. The root-entry table pointer is specified through the Root-entry Table Address (RTA_REG) register. Hardware reports the status of the "Set Root Table Pointer" operation through the RTPS field in the Global Status register. The "Set Root Table Pointer" operation must be performed before enabling or re-enabling (after disabling) DMA remapping through the TE field. After a "Set Root Table Pointer" operation, software must globally invalidate the context cache and then globally invalidate of IOTLB. This is required to ensure hardware uses only the remapping structures referenced by the new root table pointer, and not stale cached entries. While DMA remapping hardware is active, software may update the root table pointer through this field. However, to ensure valid in-flight DMA requests are deterministically remapped, software must ensure that the structures referenced by the new root table pointer are programmed to provide the same remapping results as the structures referenced by the previous root-table pointer. Clearing this bit has no effect. The value returned on read of this field is undefined.	0h	WO
29	SFL	This field is valid only for implementations supporting advanced fault logging. Software sets this field to request hardware to set/update the fault-log pointer used by hardware. The fault-log pointer is specified through Advanced Fault Log register. Hardware reports the status of the 'Set Fault Log' operation through the FLS field in the Global Status register. The fault log pointer must be set before enabling advanced fault logging (through EAFL field). Once advanced fault logging is enabled, the fault log pointer may be updated through this field while DMA remapping is active. Clearing this bit has no effect. The value returned on read of this field is undefined.	0h	RO
28	EAFL	This field is valid only for implementations supporting advanced fault logging. Software writes to this field to request hardware to enable or disable advanced fault logging: 0: Disable advanced fault logging. In this case, translation faults are reported through the Fault Recording registers. 1: Enable use of memory-resident fault log. When enabled, translation faults are recorded in the memory-resident log. The fault log pointer must be set in hardware (through the SFL field) before enabling advaned fault logging. Hardware reports the status of the advaned fault logging enable operation through the AFLS field in the Global Status register. The value returned on read of this field is undefined.	0h	RO

continued...

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: WO; RO	
Size: 32	Default Value: 00000000h		Address Offset: 18h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
27	WBF	This bit is valid only for implementations requiring write buffer flushing. Software sets this field to request that hardware flush the Root-Complex internal write buffers. This is done to ensure any updates to the memory-resident remapping structures are not held in any internal write posting buffers. Hardware reports the status of the write buffer flushing operation through the WBFS field in the Global Status register. Cleraing this bit has no effect. The value returned on a read of this field is undefined.	0h	RO
26	QIE	This field is valid only for implementations supporting queued invalidations. Software writes to this field to enable or disable queued invalidations. 0: Disable queued invalidations. 1: Enable use of queued invalidations. Hardware reports the status of queued invalidation enable operation through QIES field in the Global Status register. The value returned on a read of this field is undefined.	0h	WO
25	IRE	This field is valid only for implementations supporting interrupt remapping. 0: Disable interrupt-remapping hardware 1: Enable interrupt-remapping hardware Hardware reports the status of the interrupt remapping enable operation through the IRES field in the Global Status register. There may be active interrupt requests in the platform when software updates this field. Hardware must enable or disable interrupt-remapping logic only at deterministic transaction boundaries, so that any in-flight interrupts are either subject to remapping or not at all. Hardware implementations must drain any in-flight interrupts requests queued in the Root-Complex before completing the interrupt-remapping enable command and reflecting the status of the command through the IRES field in the Global Status register. The value returned on a read of this field is undefined.	0h	WO
24	SIRTP	This field is valid only for implementations supporting interrupt-remapping. Software sets this field to set/update the interrupt remapping table pointer used by hardware. The interrupt remapping table pointer is specified through the Interrupt Remapping Table Address (IRTA_REG) register. Hardware reports the status of the 'Set Interrupt Remap Table Pointer' operation through the IRTPS field in the Global Status register. The 'Set Interrupt Remap Table Pointer' operation must be performed before enabling or re-enabling (after disabling) interrupt-remapping hardware through the IRE field. After a 'Set Interrupt Remap Table Pointer' operation, software must globally invalidate the interrupt entry cache. This is required to ensure hardware uses only the interrupt-remapping entries referenced by the new interrupt remap table pointer, and not any stale cached entries. While interrupt remapping is active, software may update the interrupt remapping table pointer through this field. However, to ensure valid in-flight interrupt requests are deterministically remapped, software must ensure that the	0h	WO

continued...

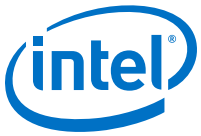

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: WO; RO	
Size: 32	**Default Value: 00000000h**		**Address Offset: 18h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		structures referenced by the new interrupt remap table pointer are programmed to provide the same remapping results as the structures referenced by the previous interrupt remap table pointer. Clearing this bit has no effect. The value returned on a read of this field is undefined.		
23	CFI	This field is valid only for Intel®64 implementations supporting interrupt-remapping. Software writes to this field to enable or disable Compatibility Format interrupts on Intel®64 platforms. The value in this field is effective only when interrupt-remapping is enabled and Extended Interrupt Mode (x2APIC mode) is not enabled. 0: Block Compatibility format interrupts. 1: Process Compatibility format interrupts as pass-through (bypass interrupt remapping). Hardware reports the status of updating this field through the CFIS field in the Global Status register. The value returned on a read of this field is undefined.	0h	WO
22:0	RSVD	Reserved.	000000h	RO

5.3.5 GSTS—Global Status Register

Register to report general remapping hardware status.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RO_V; ROV; RO	
Size: 32	**Default Value: 00000000h**		**Address Offset: 1Ch**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31	TES	This field indicates the status of DMA-remapping hardware. 0: DMA-remapping hardware is not enabled 1: DMA-remapping hardware is enabled	0h	ROV
30	RTPS	This field indicates the status of the root- table pointer in hardware. This field is cleared by hardware when software sets the SRTP field in the Global Command register. This field is set by hardware when hardware completes the 'Set Root Table Pointer' operation using the value provided in the Root-Entry Table Address register.	0h	RO_V
29	FLS	This field: - Is cleared by hardware when software Sets the SFL field in the Global Command register. - Is Set by hardware whn hardware completes the 'Set Fault Log Pointer' operation using the value provided in the Advanced Fault Log register.	0h	RO
28	AFLS	This field is valid only for implementations supporting advanced fault logging. It indicates the advanced fault logging status: 0: Advanced Fault Logging is not enabled. 1: Advanced Fault Logging is enabled.	0h	RO

continued...

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RO_V; ROV; RO	
Size: 32	Default Value: 00000000h		Address Offset: 1Ch	
Bit Range	Acronym	Description	Default	Access
27	WBFS	This field is valid only for implementations requiring write buffer flushing. This field indicates the status of the write buffer flush command. It is: - Set by hardware when software sets the WBF field in the Global Command register. - Cleared by hardware when hardware completes the write buffer flushing operation.	0h	RO
26	QIES	This field indicates queued invalidation enable status. 0: queued invalidation is not enabled 1: queued invalidation is enabled	0h	RO_V
25	IRES	This field indicates the status of Interrupt-remapping hardware. 0: Interrupt-remapping hardware is not enabled 1: Interrupt-remapping hardware is enabled	0h	ROV
24	IRTPS	This field indicates the status of the interrupt remapping table pointer in hardware. This field is cleared by hardware when software sets the SIRTP field in the Global Command register. This field is Set by hardware when hardware completes the set interrupt remap table pointer operation using the value provided in the Interrupt Remapping Table Address register.	0h	RO_V
23	CFIS	This field indicates the status of Compatibility format interrupts on Intel®64 implementations supporting interrupt-remapping. The value reported in this field is applicable only when interrupt-remapping is enabled and Extended Interrupt Mode (x2APIC mode) is not enabled. 0: Compatibility format interrupts are blocked. 1: Compatibility format interrupts are processed as pass-through (bypassing interrupt remapping).	0h	RO_V
22:0	RSVD	Reserved.	000000h	RO

5.3.6 RTADDR—Root-Entry Table Address Register

Register providing the base address of root-entry table.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW_V; RW	
Size: 64	Default Value: 0000000000000000h		Address Offset: 20h	
Bit Range	Acronym	Description	Default	Access
63:39	RSVD	Reserved.	0000000h	RO
38:12	RTA	This register points to base of page aligned, 4KB-sized root-entry table in system memory. Hardware ignores and not implements bits 63:HAW, where HAW is the host address width. Software specifies the base address of the root-entry table through this register, and programs it in hardware through	0000000h	RW

continued...

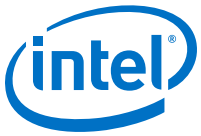

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW_V; RW	
Size: 64	**Default Value: 0000000000000000h**		**Address Offset: 20h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		the SRTP field in the Global Command register. Reads of this register returns value that was last programmed to it.		
11	RTT	This field specifies the type of root-table referenced by the Root Table Address (RTA) field; 0: Root Table / 1: Extended Root Table	0h	RW_V
10:0	RSVD	Reserved.	000h	RO

5.3.7 CCMD—Context Command Register

Register to manage context cache. The act of writing the uppermost byte of the CCMD_REG with the ICC field Set causes the hardware to perform the context-cache invalidation.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW; ROV; RW_V	
Size: 64	**Default Value: 0800000000000000h**		**Address Offset: 28h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63	ICC	Software requests invalidation of context-cache by setting this field. Software must also set the requested invalidation granularity by programming the CIRG field. Software must read back and check the ICC field is Clear to confirm the invalidation is complete. Software must not update this register when this field is set. Hardware clears the ICC field to indicate the invalidation request is complete. Hardware also indicates the granularity at which the invalidation operation was performed through the CAIG field. Software must submit a context-cache invalidation request through this field only when there are no invalidation requests pending at this remapping hardware unit. Since information from the context-cache may be used by hardware to tag IOTLB entries, software must perform domain-selective (or global) invalidation of IOTLB after the context cache invalidation has completed. Hardware implementations reporting write-buffer flushing requirement (RWBF=1 in Capability register) must implicitly perform a write buffer flush before invalidating the context cache.	0h	RW_V
62:61	CIRG	Software provides the requested invalidation granularity through this field when setting the ICC field: 00: Reserved. 01: Global Invalidation request. 10: Domain-selective invalidation request. The target domain-id must be specified in the DID field. 11: Device-selective invalidation request. The target source-id(s) must be specified through the SID and FM fields, and the domain-id (that was programmed in the context-entry for these device(s)) must be provided in the DID field. Hardware implementations may process an invalidation request by performing invalidation at a coarser granularity than requested. Hardware indicates completion of the	0h	RW

continued...

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW; ROV; RW_V	
Size: 64	Default Value: 0800000000000000h		Address Offset: 28h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		invalidation request by clearing the ICC field. At this time, hardware also indicates the granularity at which the actual invalidation was performed through the CAIG field.		
60:59	CAIG	Hardware reports the granularity at which an invalidation request was processed through the CAIG field at the time of reporting invalidation completion (by clearing the ICC field). The following are the encodings for this field: 00: Reserved. 01: Global Invalidation performed. This could be in response to a global, domain-selective or device-selective invalidation request. 10: Domain-selective invalidation performed using the domain-id specified by software in the DID field. This could be in response to a domain-selective or device-selective invalidation request. 11: Device-selective invalidation performed using the source-id and domain-id specified by software in the SID and FM fields. This can only be in response to a device-selective invalidation request.	1h	ROV
58:34	RSVD	Reserved.	0000000h	RO
33:32	FM	Software may use the Function Mask to perform device-selective invalidations on behalf of devices supporting PCI Express Phantom Functions. This field specifies which bits of the function number portion (least significant three bits) of the SID field to mask when performing device-selective invalidations. The following encodings are defined for this field: 00: No bits in the SID field masked. 01: Mask most significant bit of function number in the SID field. 10: Mask two most significant bit of function number in the SID field. 11: Mask all three bits of function number in the SID field. The context-entries corresponding to all the source-ids specified through the FM and SID fields must have to the domain-id specified in the DID field.	0h	RW
31:16	SID	Indicates the source-id of the device whose corresponding context-entry needs to be selectively invalidated. This field along with the FM field must be programmed by software for device-selective invalidation requests.	0000h	RW
15:8	RSVD	Reserved.	00h	RO
7:0	DID	Indicates the id of the domain whose context-entries need to be selectively invalidated. This field must be programmed by software for both domain-selective and device-selective invalidation requests. The Capability register reports the domain-id width supported by hardware. Software must ensure that the value written to this field is within this limit. Hardware may ignore and not implement bits15:N, where N is the supported domain-id width reported in the Capability register.	00h	RW

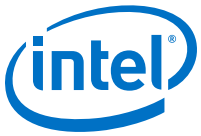

5.3.8 FSTS—Fault Status Register

Register indicating the various error status.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW1CS; ROSV; RO	
Size: 32	**Default Value: 00000000h**		**Address Offset: 34h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:8	FRI	This field is valid only when the PPF field is Set. The FRI field indicates the index (from base) of the fault recording register to which the first pending fault was recorded when the PPF field was Set by hardware. The value read from this field is undefined when the PPF field is clear.	00h	RO
7	PRO	Hardware detected a Page Request Overflow error. Hardware implementations not supporting the Page Request Queue implement this bit as RsvdZ.	0h	RW1CS
6	ITE	Hardware detected a Device-IOTLB invalidation completion time-out. At this time, a fault event may be generated based on the programming of the Fault Event Control register. Hardware implementations not supporting device Device-IOTLBs implement this bit as RsvdZ.	0h	RO
5	ICE	Hardware received an unexpected or invalid Device-IOTLB invalidation completion. This could be due to either an invalid ITag or invalid source-id in an invalidation completion response. At this time, a fault event may be generated based on the programming of hte Fault Event Control register. Hardware implementations not supporting Device-IOTLBs implement this bit as RsvdZ.	0h	RO
4	IQE	Hardware detected an error associated with the invalidation queue. This could be due to either a hardware error while fetching a descriptor from the invalidation queue, or hardware detecting an erroneous or invalid descriptor in the invalidation queue. At this time, a fault event may be generated based on the programming of the Fault Event Control register. Hardware implementations not supporting queued invalidations implement this bit as RsvdZ.	0h	RW1CS
3	APF	When this field is Clear, hardware sets this field when the first fault record (at index 0) is written to a fault log. At this time, a fault event is generated based on the programming of the Fault Event Control register. Software writing 1 to this field clears it. Hardware implementations not supporting advanced fault logging implement this bit as RsvdZ.	0h	RO
2	AFO	Hardware sets this field to indicate advanced fault log overflow condition. At this time, a fault event is generated based on the programming of the Fault Event Control register.	0h	RO

continued...

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW1CS; ROSV; RO	
Size: 32	Default Value: 00000000h		Address Offset: 34h	
Bit Range	Acronym	Description	Default	Access
		Software writing 1 to this field clears it. Hardware implementations not supporting advanced fault logging implement this bit as RsvdZ.		
1	PPF	This field indicates if there are one or more pending faults logged in the fault recording registers. Hardware computes this field as the logical OR of Fault (F) fields across all the fault recording registers of this remapping hardware unit. 0: No pending faults in any of the fault recording registers 1: One or more fault recording registers has pending faults. The FRI field is updated by hardware whenever the PPF field is set by hardware. Also, depending on the programming of Fault Event Control register, a fault event is generated when hardware sets this field.	0h	ROSV
0	PFO	Hardware sets this field to indicate overflow of fault recording registers. Software writing 1 clears this field. When this field is Set, hardware does not record any new faults until software clears this field.	0h	RW1CS

5.3.9 FECTL—Fault Event Control Register

Register specifying the fault event interrupt message control bits.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: ROV; RW	
Size: 32	Default Value: 80000000h		Address Offset: 38h	
Bit Range	Acronym	Description	Default	Access
31	IM	0: No masking of interrupt. When an interrupt condition is detected, hardware issues an interrupt message (using the Fault Event Data and Fault Event Address register values). 1: This is the value on reset. Software may mask interrupt message generation by setting this field. Hardware is prohibited from sending the interrupt message when this field is set.	1h	RW
30	IP	Hardware sets the IP field whenever it detects an interrupt condition, which is defined as: When primary fault logging is active, an interrupt condition occurs when hardware records a fault through one of the Fault Recording registers and sets the PPF field in Fault Status register. When advanced fault logging is active, an interrupt condition occurs when hardware records a fault in the first fault record (at index 0) of the current fault log and sets the APF field in the Fault Status register. Hardware detected error associated with the Invalidation Queue, setting the IQE field in the Fault Status register. Hardware detected invalid Device-IOTLB invalidation completion, setting the ICE field in the Fault Status register. Hardware detected Device-IOTLB invalidation completion time-out, setting the ITE field in the Fault Status register. If any of the status fields in the Fault Status register was already Set at the time of setting any of these fields, it is not treated as a new interrupt condition.	0h	ROV

continued...

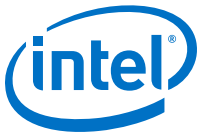

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: ROV; RW	
Size: 32	**Default Value: 80000000h**		**Address Offset: 38h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		The IP field is kept set by hardware while the interrupt message is held pending. The interrupt message could be held pending due to interrupt mask (IM field) being Set or other transient hardware conditions. The IP field is cleared by hardware as soon as the interrupt message pending condition is serviced. This could be due to either: Hardware issuing the interrupt message due to either change in the transient hardware condition that caused interrupt message to be held pending, or due to software clearing the IM field.. Software servicing all the pending interrupt status fields in the Fault Status register as follows: - When primary fault logging is active, software clearing the Fault (F) field in all the Fault Recording registers with faults, causing the PPF field in Fault Status register to be evaluated as clear. - Software clearing other status fields in the Fault Status register by writing back the value read from the respective fields.		
29:0	RSVD	Reserved.	00000000h	RO

5.3.10 FEDATA—Fault Event Data Register

Register specifying the interrupt message data

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW	
Size: 32	**Default Value: 00000000h**		**Address Offset: 3Ch**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	EIMD	This field is valid only for implementations supporting 32-bit interrupt data fields. Hardware implementations supporting only 16-bit interrupt data may treat this field as RsvdZ.	0000h	RW
15:0	IMD	Data value in the interrupt request.	0000h	RW

5.3.11 FEADDR—Fault Event Address Register

Register specifying the interrupt message address.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW	
Size: 32	**Default Value: 00000000h**		**Address Offset: 40h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:2	MA	When fault events are enabled, the contents of this register specify the DWORD-aligned address (bits 31:2) for the interrupt request.	00000000h	RW
1:0	RSVD	Reserved.	0h	RO

5.3.12 FEUADDR—Fault Event Upper Address Register

Register specifying the interrupt message upper address.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 44h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:0	MUA	Hardware implementations supporting Extended Interrupt Mode are required to implement this register. Hardware implementations not supporting Extended Interrupt Mode may treat this field as RsvdZ.	00000000h	RW

5.3.13 AFLOG—Advanced Fault Log Register

Register to specify the base address of the memory-resident fault-log region. This register is treated as RsvdZ for implementations not supporting advanced translation fault logging (AFL field reported as 0 in the Capability register).

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RO	
Size: 64	Default Value: 0000000000000000h		Address Offset: 58h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:12	FLA	This field specifies the base of 4KB aligned fault-log region in system memory. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address width. Software specifies the base address and size of the fault log region through this register, and programs it in hardware through the SFL field in the Global Command register. When implemented, reads of this field return the value that was last programmed to it.	0000000000000h	RO
11:9	FLS	This field specifies the size of the fault log region pointed by the FLA field. The size of the fault log region is 2^X * 4KB, where X is the value programmed in this register. When implemented, reads of this field return the value that was last programmed to it.	0h	RO
8:0	RSVD	Reserved.	000h	RO

5.3.14 PMEN—Protected Memory Enable Register

Register to enable the DMA-protected memory regions setup through the PLMBASE, PLMLIMT, PHMBASE, PHMLIMIT registers. This register is always treated as RO for implementations not supporting protected memory regions (PLMR and PHMR fields reported as Clear in the Capability register).
Protected memory regions may be used by software to securely initialize remapping structures in memory. To avoid impact to legacy BIOS usage of memory, software is recommended to not overlap protected memory regions with any reserved memory regions of the platform reported through the Reserved Memory Region Reporting (RMRR) structures.

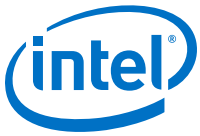

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: ROV; RW	
Size: 32	**Default Value: 00000000h**		**Address Offset: 64h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31	EPM	This field controls DMA accesses to the protected low-memory and protected high-memory regions. 0: Protected memory regions are disabled. 1: Protected memory regions are enabled. DMA requests accessing protected memory regions are handled as follows: - When DMA remapping is not enabled, all DMA requests accessing protected memory regions are blocked. - When DMA remapping is enabled: - DMA requests processed as pass-through (Translation Type value of 10b in Context-Entry) and accessing the protected memory regions are blocked. - DMA requests with translated address (AT=10b) and accessing the protected memory regions are blocked. - DMA requests that are subject to address remapping, and accessing the protected memory regions may or may not be blocked by hardware. For such requests, software must not depend on hardware protection of the protected memory regions, and instead program the DMA-remapping page-tables to not allow DMA to protected memory regions. Remapping hardware access to the remapping structures are not subject to protected memory region checks. DMA requests blocked due to protected memory region violation are not recorded or reported as remapping faults. Hardware reports the status of the protected memory enable/disable operation through the PRS field in this register. Hardware implementations supporting DMA draining must drain any in-flight translated DMA requests queued within the Root-Complex before indicating the protected memory region as enabled through the PRS field.	0h	RW
30:1	RSVD	Reserved.	00000000h	RO
0	PRS	This field indicates the status of protected memory region(s): 0: Protected memory region(s) disabled. 1: Protected memory region(s) enabled.	0h	ROV

5.3.15 PLMBASE—Protected Low-Memory Base Register

Register to set up the base address of DMA-protected low-memory region below 4GB. This register must be set up before enabling protected memory through PMEN_REG, and must not be updated when protected memory regions are enabled. This register is always treated as RO for implementations not supporting protected low memory region (PLMR field reported as Clear in the Capability register).
The alignment of the protected low memory region base depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1s to this register, and finding the most significant zero bit position with 0 in the value read back from the register. Bits N:0 of this register is decoded by hardware as all 0s.
Software must setup the protected low memory region below 4GB.
Software must not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 68h	
Bit Range	Acronym	Description	Default	Access
31:20	PLMB	This register specifies the base of protected low-memory region in system memory.	000h	RW
19:0	RSVD	Reserved.	00000h	RO

5.3.16 PLMLIMIT—Protected Low-Memory Limit Register

Register to set up the limit address of DMA-protected low-memory region below 4GB. This register must be set up before enabling protected memory through PMEN_REG, and must not be updated when protected memory regions are enabled.
This register is always treated as RO for implementations not supporting protected low memory region (PLMR field reported as Clear in the Capability register).
The alignment of the protected low memory region limit depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1's to this register, and finding most significant zero bit position with 0 in the value read back from the register. Bits N:0 of the limit register is decoded by hardware as all 1s.
The Protected low-memory base and limit registers functions as follows:
- Programming the protected low-memory base and limit registers with the same value
  in bits 31:(N+1) specifies a protected low-memory region of size 2^(N+1) bytes.
- Programming the protected low-memory limit register with a value less than the protected low-memory base register disables the protected low-memory region.

Software must not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 6Ch	
Bit Range	Acronym	Description	Default	Access
31:20	PLML	This register specifies the last host physical address of the DMA-protected low-memory region in system memory.	000h	RW
19:0	RSVD	Reserved.	00000h	RO

5.3.17 PHMBASE—Protected High-Memory Base Register

Register to set up the base address of DMA-protected high-memory region. This register must be set up before enabling protected memory through PMEN_REG, and must not be updated when protected memory regions are enabled.
This register is always treated as RO for implementations not supporting protected high memory region (PHMR field reported as Clear in the Capability register).
The alignment of the protected high memory region base depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1's to this register, and finding most significant zero bit position below host address width (HAW) in the value read back from the register. Bits N:0 of this register are decoded by hardware as all 0s.
Software may setup the protected high memory region either above or below 4GB.
Software must not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

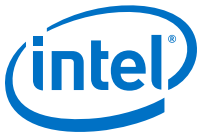

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW	
Size: 64	Default Value: 0000000000000000h		Address Offset: 70h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:20	PHMB	This register specifies the base of protected (high) memory region in system memory. Hardware ignores, and does not implement, bits 63:HAW, where HAW is the host address width.	00000h	RW
19:0	RSVD	Reserved.	00000h	RO

5.3.18 PHMLIMIT—Protected High-Memory Limit Register

Register to set up the limit address of DMA-protected high-memory region. This register must be set up before enabling protected memory through PMEN_REG, and must not be updated when protected memory regions are enabled.
This register is always treated as RO for implementations not supporting protected high memory region (PHMR field reported as Clear in the Capability register).
The alignment of the protected high memory region limit depends on the number of reserved bits (N:0) of this register. Software may determine the value of N by writing all 1's to this register, and finding most significant zero bit position below host address width (HAW) in the value read back from the register. Bits N:0 of the limit register is decoded by hardware as all 1s.
The protected high-memory base & limit registers functions as follows.
- Programming the protected low-memory base and limit registers with the same value
  in bits HAW:(N+1) specifies a protected low-memory region of size 2^(N+1) bytes.
- Programming the protected high-memory limit register with a value less than the protected high-memory base register disables the protected high-memory region.
Software must not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW	
Size: 64	Default Value: 0000000000000000h		Address Offset: 78h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:20	PHML	This register specifies the last host physical address of the DMA-protected high-memory region in system memory. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address width.	00000h	RW
19:0	RSVD	Reserved.	00000h	RO

5.3.19 IQH—Invalidation Queue Head Register

Register indicating the invalidation queue head. This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: ROV	
Size: 64	Default Value: 0000000000000000h		Address Offset: 80h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:19	RSVD	Reserved.	000000000000h	RO
18:4	QH	Specifies the offset (128-bit aligned) to the invalidation queue for the command that will be fetched next by hardware. Hardware resets this field to 0 whenever the queued invalidation is disabled (QIES field Clear in the Global Status register).	0000h	ROV
3:0	RSVD	Reserved.	0h	RO

5.3.20 IQT—Invalidation Queue Tail Register

Register indicating the invalidation tail head. This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW_L	
Size: 64	Default Value: 0000000000000000h		Address Offset: 88h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:19	RSVD	Reserved.	000000000000h	RO
18:4	QT	Specifies the offset (128-bit aligned) to the invalidation queue for the command that will be written next by software.	0000h	RW_L
3:0	RSVD	Reserved.	0h	RO

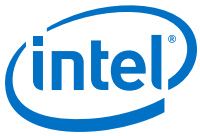

5.3.21 IQA—Invalidation Queue Address Register

Register to configure the base address and size of the invalidation queue. This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW_L	
Size: 64	Default Value: 0000000000000000h		Address Offset: 90h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:12	IQA	This field points to the base of 4KB aligned invalidation request queue. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address width. Reads of this field return the value that was last programmed to it.	0000000h	RW_L
11:3	RSVD	Reserved.	000h	RO
2:0	QS	This field specifies the size of the invalidation request queue. A value of X in this field indicates an invalidation request queue of (2^X) 4KB pages. The number of entries in the invalidation queue is 2^(X + 8).	0h	RW_L

5.3.22 ICS—Invalidation Completion Status Register

Register to report completion status of invalidation wait descriptor with Interrupt Flag (IF) Set.
This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW1CS	
Size: 32	Default Value: 00000000h		Address Offset: 9Ch	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:1	RSVD	Reserved.	00000000h	RO
0	IWC	Indicates completion of Invalidation Wait Descriptor with Interrupt Flag (IF) field Set. Hardware implementations not supporting queued invalidations implement this field as RsvdZ.	0h	RW1CS

5.3.23 IECTL—Invalidation Event Control Register

Register specifying the invalidation event interrupt control bits.
This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: ROV; RW_L	
Size: 32	Default Value: 80000000h		Address Offset: A0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31	IM	0: No masking of interrupt. When a invalidation event condition is detected, hardware issues an interrupt message (using the Invalidation Event Data & Invalidation	1h	RW_L

continued...

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: ROV; RW_L	
Size: 32	Default Value: 80000000h		Address Offset: A0h	
Bit Range	Acronym	Description	Default	Access
		Event Address register values). 1: This is the value on reset. Software may mask interrupt message generation by setting this field. Hardware is prohibited from sending the interrupt message when this field is Set.		
30	IP	Hardware sets the IP field whenever it detects an interrupt condition. Interrupt condition is defined as: - An Invalidation Wait Descriptor with Interrupt Flag (IF) field Set completed, setting the IWC field in the Invalidation Completion Status register. - If the IWC field in the Invalidation Completion Status register was already Set at the time of setting this field, it is not treated as a new interrupt condition. The IP field is kept Set by hardware while the interrupt message is held pending. The interrupt message could be held pending due to interrupt mask (IM field) being Set, or due to other transient hardware conditions. The IP field is cleared by hardware as soon as the interrupt message pending condition is serviced. This could be due to either: - Hardware issuing the interrupt message due to either change in the transient hardware condition that caused interrupt message to be held pending or due to software clearing the IM field. - Software servicing the IWC field in the Invalidation Completion Status register.	0h	ROV
29:0	RSVD	Reserved.	00000000h	RO

5.3.24 IEDATA—Invalidation Event Data Register

Register specifying the Invalidation Event interrupt message data.
This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW_L	
Size: 32	Default Value: 00000000h		Address Offset: A4h	
Bit Range	Acronym	Description	Default	Access
31:16	EIMD	This field is valid only for implementations supporting 32-bit interrupt data fields. Hardware implementations supporting only 16-bit interrupt data treat this field as Rsvd.	0000h	RW_L
15:0	IMD	Data value in the interrupt request.	0000h	RW_L

5.3.25 IEADDR—Invalidation Event Address Register

Register specifying the Invalidation Event Interrupt message address.
This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

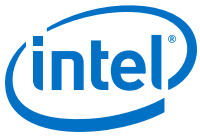

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW_L	
Size: 32	Default Value: 00000000h		Address Offset: A8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:2	MA	When fault events are enabled, the contents of this register specify the DWORD-aligned address (bits 31:2) for the interrupt request.	00000000h	RW_L
1:0	RSVD	Reserved.	0h	RO

5.3.26 IEUADDR—Invalidation Event Upper Address Register

Register specifying the Invalidation Event interrupt message upper address.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW_L	
Size: 32	Default Value: 00000000h		Address Offset: ACh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:0	MUA	Hardware implementations supporting Queued Invalidations and Extended Interrupt Mode are required to implement this register. Hardware implementations not supporting Queued Invalidations or Extended Interrupt Mode may treat this field as RsvdZ.	00000000h	RW_L

5.3.27 IRTA—Interrupt Remapping Table Address Register

Register providing the base address of Interrupt remapping table. This register is treated as RsvdZ by implementations reporting Interrupt Remapping (IR) as not supported in the Extended Capability register.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW_L; ROV	
Size: 64	Default Value: 0000000000000000h		Address Offset: B8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:12	IRTA	This field points to the base of 4KB aligned interrupt remapping table. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address width. Reads of this field returns value that was last programmed to it.	0000000h	RW_L
11	EIME	This field is used by hardware on Intel®64 platforms as follows: 0: xAPIC mode is active. Hardware interprets only low 8-bits of Destination-ID field in the IRTEs. The high 24-bits of the Destination-ID field are treated as reserved. 1: x2APIC mode is active. Hardware interprets all 32-bits of Destination-ID field in the IRTEs.	0h	ROV

continued...

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW_L; ROV	
Size: 64	Default Value: 0000000000000000h		Address Offset: B8h	
Bit Range	Acronym	Description	Default	Access
		This field is implemented as RsvdZ on implementations reporting Extended Interrupt Mode (EIM) field as Clear in Extended Capability register.		
10:4	RSVD	Reserved.	00h	RO
3:0	S	This field specifies the size of the interrupt remapping table. The number of entries in the interrupt remapping table is 2^(X+1), where X is the value programmed in this field.	0h	RW_L

5.3.28 FRCDL—Fault Recording Low Register

Register to record fault information when primary fault logging is active. Hardware reports the number and location of fault recording registers through the Capability register. This register is relevant only for primary fault logging.
This register is sticky and can be cleared only through power good reset or by software clearing the RW1C fields by writing a 1.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: ROSV	
Size: 64	Default Value: 0000000000000000h		Address Offset: 400h	
Bit Range	Acronym	Description	Default	Access
63:12	FI	When the Fault Reason (FR) field indicates one of the DMA-remapping fault conditions, bits 63:12 of this field contain the page address in the faulted DMA request. Hardware treats bits 63:N as reserved (0), where N is the maximum guest address width (MGAW) supported. When the Fault Reason (FR) field indicates one of the interrupt-remapping fault conditions, bits 63:48 of this field indicate the interrupt_index computed for the faulted interrupt request, and bits 47:12 are cleared. This field is relevant only when the F field is Set.	0000000000000h	ROSV
11:0	RSVD	Reserved.	000h	RO

5.3.29 FRCDH—Fault Recording High Register

Register to record fault information when primary fault logging is active. Hardware reports the number and location of fault recording registers through the Capability register. This register is relevant only for primary fault logging.
This register is sticky and can be cleared only through power good reset or by software clearing the RW1C fields by writing a 1.

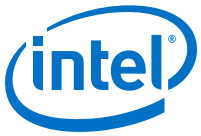

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: ROSV; ROV; RW1CS	
Size: 64	**Default Value: 0000000000000000h**		**Address Offset: 408h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63	F	Hardware sets this field to indicate a fault is logged in this Fault Recording register. The F field is set by hardware after the details of the fault is recorded in other fields. When this field is Set, hardware may collapse additional faults from the same source-id (SID). Software writes the value read from this field to Clear it.	0h	RW1CS
62	T	Type of the faulted request: 0: Write request 1: Read request or AtomicOp request This field is relevant only when the F field is Set, and when the fault reason (FR) indicates one of the DMA-remapping fault conditions.	0h	ROSV
61:60	AT	This field captures the AT field from the faulted DMA request. Hardware implementations not supporting Device-IOTLBs (DI field Clear in Extended Capability register) treat this field as RsvdZ. When supported, this field is valid only when the F field is Set, and when the fault reason (FR) indicates one of the DMA-remapping fault conditions.	0h	ROV
59:40	PN	PASID value in the faulted request. This field is relevant only when the PP field is set. Hardware implementations not supporting PASID (PASID field Clear in Extended Capability register) implement this field as RsvdZ.	00000h	ROSV
39:32	FR	Reason for the fault. This field is relevant only when the F field is set.	00h	ROSV
31	PP	When set, indicates the faulted request has a PASID tag. The value of the PASID field is reported in the PASID Value (PV) field. This field is relevant only when the F field is Set, and when the fault reason (FR) indicates one of the non-recoverable address translation fault conditions. Hardware implementations not supporting PASID (PASID field Clear in Extended Capability register) implement this field as RsvdZ.	0h	ROSV
30	EXE	When set, indicates Execute permission was requested by the faulted read request. This field is relevant only when the PP field and T field are both Set. Hardware implementations not supporting PASID (PASID field Clear in Extended Capability register) implement this field as RsvdZ.	0h	ROSV
29	PRIV	When set, indicates Supervisor privilege was requested by the faulted request. This field is relevant only when the PP field is Set. Hardware implementations not supporting PASID (PASID field Clear in Extended Capability register) implement this field as RsvdZ.	0h	ROSV
28:16	RSVD	Reserved.	0000h	RO
15:0	SID	Requester-id associated with the fault condition. This field is relevant only when the F field is set.	0000h	ROSV

5.3.30 IVA—Invalidate Address Register

Register to provide the DMA address whose corresponding IOTLB entry needs to be invalidated through the corresponding IOTLB Invalidate register. This register is a write-only register.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW	
Size: 64	Default Value: 0000000000000000h		Address Offset: 500h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:12	ADDR	Software provides the DMA address that needs to be page-selectively invalidated. To make a page-selective invalidation request to hardware, software must first write the appropriate fields in this register, and then issue the appropriate page-selective invalidate command through the IOTLB_REG. Hardware ignores bits 63 : N, where N is the maximum guest address width (MGAW) supported.	0000000h	RW
11:7	RSVD	Reserved.	00h	RO
6	IH	The field provides hint to hardware about preserving or flushing the non-leaf (page-directory) entries that may be cached in hardware: 0: Software may have modified both leaf and non-leaf page-table entries corresponding to mappings specified in the ADDR and AM fields. On a page-selective invalidation request, hardware must flush both the cached leaf and non-leaf page-table entries corresponding tot he mappings specified by ADDR and AM fields. 1: Software has not modified any non-leaf page-table entries corresponding to mappings specified in the ADDR and AM fields. On a page-selective invalidation request, hardware may preserve the cached non-leaf page-table entries corresponding to mappings specified by ADDR and AM fields.	0h	RW
5:0	AM	The value in this field specifies the number of low order bits of the ADDR field that must be masked for the invalidation operation. This field enables software to request invalidation of contiguous mappings for size-aligned regions. For example: Mask Value / ADDR bits masked / Pages invalidated 0 None 1 1 12 2 2 13:12 4 3 14:12 8 4 15:12 16 ... . When invalidating mappings for super-pages, software must specify the appropriate mask value. For example, when invalidating mapping for a 2MB page, software must specify an address mask value of at least 9. Hardware implementations report the maximum supported mask value through the Capability register.	00h	RW

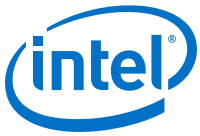

5.3.31 IOTLB—IOTLB Invalidate Register

Register to invalidate IOTLB. The act of writing the upper byte of the IOTLB_REG with IVT field Set causes the hardware to perform the IOTLB invalidation.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW; ROV; RW_V	
Size: 64	**Default Value: 0200000000000000h**		**Address Offset: 508h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63	IVT	Software requests IOTLB invalidation by setting this field. Software must also set the requested invalidation granularity by programming the IIRG field. Hardware clears the IVT field to indicate the invalidation request is complete. Hardware also indicates the granularity at which the invalidation operation was performed through the IAIG field. Software must not submit another invalidation request through this register while the IVT field is Set, nor update the associated Invalidate Address register. Software must not submit IOTLB invalidation requests when there is a context-cache invalidation request pending at this remapping hardware unit. Hardware implementations reporting write-buffer flushing requirement (RWBF=1 in Capability register) must implicitly perform a write buffer flushing before invalidating the IOTLB.	0h	RW_V
62	RSVD	Reserved.	0h	RO
61:60	IIRG	When requesting hardware to invalidate the IOTLB (by setting the IVT field), software writes the requested invalidation granularity through this field. The following are the encodings for the field. 00: Reserved. 01: Global invalidation request. 10: Domain-selective invalidation request. The target domain-id must be specified in the DID field. 11: Page-selective invalidation request. The target address, mask and invalidation hint must be specified in the Invalidate Address register, and the domain-id must be provided in the DID field. Hardware implementations may process an invalidation request by performing invalidation at a coarser granularity than requested. Hardware indicates completion of the invalidation request by clearing the IVT field. At this time, the granularity at which actual invalidation was performed is reported through the IAIG field	0h	RW
59	RSVD	Reserved.	0h	RO
58:57	IAIG	Hardware reports the granularity at which an invalidation request was processed through this field when reporting invalidation completion (by clearing the IVT field). The following are the encodings for this field. 00: Reserved. This indicates hardware detected an incorrect invalidation request and ignored the request. Examples of incorrect invalidation requests include detecting an unsupported address mask value in Invalidate Address register for page-selective invalidation requests. 01: Global Invalidation performed. This could be in response to a global, domain-selective, or page-selective invalidation request. 10: Domain-selective invalidation performed using the domain-id specified by software in the DID field. This could be in response to a domain-selective or a page-	1h	ROV

continued...

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW; ROV; RW_V	
Size: 64	Default Value: 0200000000000000h		Address Offset: 508h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		selective invalidation request. 11: Domain-page-selective invalidation performed using the address, mask and hint specified by software in the Invalidate Address register and domain-id specified in DID field. This can be in response to a page-selective invalidation request.		
56:50	RSVD	Reserved.	00h	RO
49	DR	This field is ignored by hardware if the DRD field is reported as clear in the Capability register. When the DRD field is reported as Set in the Capability register, the following encodings are supported for this field: 0: Hardware may complete the IOTLB invalidation without draining any translated DMA read requests. 1: Hardware must drain DMA read requests.	0h	RW
48	DW	This field is ignored by hardware if the DWD field is reported as Clear in the Capability register. When the DWD field is reported as Set in the Capability register, the following encodings are supported for this field: 0: Hardware may complete the IOTLB invalidation without draining DMA write requests. 1: Hardware must drain relevant translated DMA write requests.	0h	RW
47:40	RSVD	Reserved.	00h	RO
39:32	DID	Indicates the ID of the domain whose IOTLB entries need to be selectively invalidated. This field must be programmed by software for domain-selective and page-selective invalidation requests. The Capability register reports the domain-id width supported by hardware. Software must ensure that the value written to this field is within this limit. Hardware ignores and not implements bits 47:(32+N), where N is the supported domain-id width reported in the Capability register.	00h	RW
31:0	RSVD	Reserved.	00000000h	RO

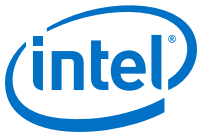

5.3.32 ARCHDIS—DMA Remap Engine Policy Control

This register contains all architectural disables and defeatures for the graphics DMA remap engine.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW_L; RO_KFW; RO; RW_KL	
Size: 32	**Default Value: 00000001h**		**Address Offset: FF0h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31	DMAR_LCKDN	This register bit protects all the DMA remap engine specific policy configuration registers. Once this bit is set by software all the DMA remap engine registers within the range 0xF00 to 0xFFC will be read-only. This bit can only be clear through platform reset.	0h	RW_KL
30	DMA_RSRV_CTL	This bit indicates whether Reserved Bit checking is supported or not (i.e. support for Fault Reason 0xA, 0xB, or 0xC). 0 - HW supports reserved field checking in root, context and page translation structures. 1 - HW ignores reserved field checking in root, context, and page translation structures.	0h	RW_L
29:16	RSVD	Reserved.	0000h	RO
15	NWFSCAPDIS	This bit allows hiding the NWFS Capability. 0: ECAP_REG[NWFS] is determined by its own default value. 1: ECAP_REG[NWFS] is set to 0b.	0h	RW_L
14	MTSCAPDIS	This bit allows hiding the MTS Capability. 0: ECAP_REG[MTS] is determined by its own default value. 1: ECAP_REG[MTS] is set to 0b.	0h	RW_L
13	EAFSCAPDIS	This bit allows hiding the EAFS Capability. 0: ECAP_REG[EAFS] is determined by its own default value. 1: ECAP_REG[EAFS] is set to 0b.	0h	RW_L
12	FL64KPCAPCTRL	This bit allows enabling/disabling the FL64KP Capability. 0: ECAP_REG[FL64KP] is set to 0b to disable first-level 64k pages capability. 1: ECAP_REG[FL64KP] is set to 1b to enable first-level 64k pages capability.	0h	RW_L
11	DTCAPDIS	This bit allows hiding the Device TLB Capability. 0: ECAP_REG[DT] is determined by its own default value. 1: ECAP_REG[DT] is set to 0b.	0h	RW_L
10	PASIDCAPDIS	This bit allows hiding the PASID Capability. 0: ECAP_REG[PASID] is determined by its own default value. 1: ECAP_REG[PASID] is set to 0b.	0h	RW_L
9	ECSCAPDIS	This bit allows hiding the Extended Context Capability. 0: ECAP_REG[ECS] is determined by its own default value. 1: ECAP_REG[ECS] is set to 0b. Additionally hardware will prevent writing of '1' to RTADDR_REG.b[11].	0h	RW_L

continued...

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RW_L; RO_KFW; RO; RW_KL	
Size: 32	Default Value: 00000001h		Address Offset: FF0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
8	SCCAPDIS	This bit allows hiding the Snoop Control Capability. 0: ECAP_REG[SC] is determined by its own default value. 1: ECAP_REG[SC] is set to 0b.	0h	RO
7	PTCAPDIS	This bit allows hiding the Pass Through Capability. 0: ECAP_REG[PT] is determined by its own default value. 1: ECAP_REG[PT] is set to 0b.	0h	RW_L
6	IRCAPDIS	This bit allows hiding the Interrupt Remapping Capability. 0: ECAP_REG[IR] is determined by its own default value. 1: ECAP_REG[IR] is set to 0b.	0h	RO_KFW
5	QICAPDIS	This bit allows hiding the Queued Invalidation Capability. 0: ECAP_REG[QI] is determined by its own default value. 1: ECAP_REG[QI] is set to 0b.	0h	RO_KFW
4	NESTCAPDIS	This bit allows hiding the Nested Translation Capability. 0: CAP_REG[NEST] is determined by its own default value. 1: CAP_REG[NEST] is set to 0b.	0h	RW_L
3	DISCAPDIS	This bit allows hiding the Deferred Invalidation Support Capability. 0: CAP_REG[DIS] is determined by its own default value. 1: CAP_REG[DIS] is set to 0b.	0h	RW_L
2	PRSCAPDIS	This bit allows hiding the Page Request Capability. 0: CAP_REG[PRS] is determined by its own default value. 1: CAP_REG[PRS] is set to 0b.	0h	RW_L
1	FL1GPCAPDIS	This bit allows hiding the First Level 1G Page Capability. 0: CAP_REG[FL1GP] is determined by its own default value. 1: CAP_REG[FL1GP] is set to 0b.	0h	RW_L
0	SLLPSCAPCTRL	This bit allows enabling/disabling the Super Page Capability. 0: CAP_REG[SLLPS] is set to 0x0 to disable superpages. 1: CAP_REG[SLLPS] is set to 0x3 to enable superpages. When SLLPSCAPCTRL is set to 0, CAP_REG[SLLPS]=0. If software ignores it and sets up Super Pages then IMPH will generate VT-d fault.	1h	RW_L

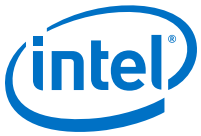

5.3.33 UARCHDIS—DMA Remap Engine Policy Control

This register contains all microarchitectural disables and defeatures for the graphics DMA remap engine.

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RO; RW_L	
Size: 32	Default Value: 00000000h		Address Offset: FF4h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:23	RSVD	Reserved.	000h	RO
22	NO_TLBLKUP_PEND	When this bit is set, all entries which which hit to pending on another request's TLB allocation in the default engine are not allowed to look up peer aperture TLBs for a following graphics walk. They must do all page walks (including root and context) in the IGD engine.	0h	RW_L
21	IQ_COH_DIS	When this bit is set to 1b, read requests from the Invalidation Queue are done in a non-coherent manner (no snoops are generated).	0h	RW_L
20	L3_HIT2PEND_DIS	When set, this bit forces a lookup which matches an L3 TLB entry in PEND state to be treated as a miss without allocation.	0h	RW_L
19	L2_HIT2PEND_DIS	When set, this bit forces a lookup which matches an L2 TLB entry in PEND state to be treated as a miss without allocation.	0h	RO
18	L1_HIT2PEND_DIS	When set, this bit forces a lookup which matches an L1 TLB entry in PEND state to be treated as a miss without allocation.	0h	RW_L
17	L0_HIT2PEND_DIS	When set, this bit forces a lookup which matches an L0 TLB entry in PEND state to be treated as a miss without allocation.	0h	RW_L
16	CC_HIT2PEND_DIS	When set, this bit forces a lookup which matches a context cache entry in PEND state to be treated as a miss without allocation.	0h	RW_L
15	L3DIS	1: L3 TLB is disabled, and each GPA request that looks up the L3 will result in a miss. 0: Normal mode (default). L3 is enabled.	0h	RW_L
14	L2DIS	1: L2 TLB is disabled, and each GPA request that looks up the L2 will result in a miss. 0: Normal mode (default). L2 is enabled.	0h	RO
13	L1DIS	1: L1 TLB is disabled, and each GPA request that looks up the L1 will result in a miss. 0: Normal mode (default). L1 is enabled.	0h	RW_L
12	L0DIS	1: L0 TLB is disabled, and each GPA request that looks up the L0 will result in a miss. 0: Normal mode (default). L0 is enabled.	0h	RW_L
11	CCDIS	1: Context Cache is disabled. Each GPA request results in a miss and will request a root walk. 0: Normal mode (default). Context Cache is enabled.	0h	RW_L

continued...

B/D/F/Type: 0/0/0/MEM/GFXVTBAR			Access: RO; RW_L	
Size: 32	Default Value: 00000000h		Address Offset: FF4h	
Bit Range	Acronym	Description	Default	Access
10:2	RSVD	Reserved.	000h	RO
1	GLBIOTLBINV	This bit controls the IOTLB Invalidation behaviour of the DMA remap engine. When this bit is set, any type of IOTLB Invalidation will be promoted to Global IOTLB Invalidation. This promotion applies to both register-based invalidation and queued invalidation.	0h	RO
0	GLBCTXTINV	This bit controls the Context Invalidation behaviour of the DMA remap engine. When this bit is set, any type of Context Invalidation will be promoted to Global Context Invalidation. This promotion applies to both register-based invalidation and queued invalidation.	0h	RO

5.4 PXPEPBAR Registers Summary

Offset	Register ID—Description	Default Value	Access
14	EPVC0RCTL—EP VC 0 Resource Control on page 330	800000FFh	RO; RW

5.4.1 EPVC0RCTL—EP VC 0 Resource Control

Controls the resources associated with Egress Port Virtual Channel 0.

B/D/F/Type: 0/0/0/MEM/PXPEPBAR			Access: RO; RW	
Size: 32	Default Value: 800000FFh		Address Offset: 14h	
Bit Range	Acronym	Description	Default	Access
31	VC0E	VC0 Enable: For VC0 this is hardwired to 1 and read only as VC0 can never be disabled.	1h	RO
30:27	RSVD	Reserved.	0h	RO
26:24	VC0ID	VC0 ID: Assigns a VC ID to the VC resource. For VC0 this is hardwired to 0 and read only.	0h	RO
23:20	RSVD	Reserved.	0h	RO
19:17	PAS	Port Arbitration Select: This field configures the VC resource to provide a particular Port Arbitration service. The value of 0h corresponds to the bit position of the only asserted bit in the Port Arbitration Capability field.	0h	RW
16:8	RSVD	Reserved.	000h	RO
7:1	TCVC0M	TC/VC0 Map: Indicates the TCs (Traffic Classes) that are mapped to the VC resource. Bit locations within this field correspond to TC values. For example, when bit 7 is set in this field, TC7 is mapped to this VC resource. When more than one bit in this field is set, it indicates that multiple TCs are mapped to the VC resource. In order to remove one or more TCs from the TC/VC Map of an enabled VC, software must ensure that no new or outstanding transactions with the TC labels are targeted at the given Link.	7Fh	RW
0	TC0VC0M	TC0/VC0 Map: Traffic Class 0 is always routed to VC0.	1h	RO

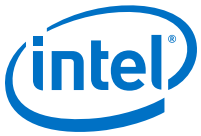

5.5 VC0PREMAP Registers Summary

Offset	Register ID—Description	Default Value	Access
0	VER—Version Register on page 332	00000010h	RO
8	CAP—Capability Register on page 332	00D2008C20660462h	RO; ROV
10	ECAP—Extended Capability Register on page 335	0000000000F010DAh	RO; ROV
18	GCMD—Global Command Register on page 337	00000000h	WO; RO
1C	GSTS—Global Status Register on page 340	00000000h	RO_V; ROV; RO
20	RTADDR—Root-Entry Table Address Register on page 341	0000000000000000h	RO; RW
28	CCMD—Context Command Register on page 342	0000000000000000h	RW; ROV; RW_V
34	FSTS—Fault Status Register on page 343	00000000h	RW1CS; ROSV; RO
38	FECTL—Fault Event Control Register on page 345	80000000h	ROV; RW
3C	FEDATA—Fault Event Data Register on page 346	00000000h	RW
40	FEADDR—Fault Event Address Register on page 346	00000000h	RW
44	FEUADDR—Fault Event Upper Address Register on page 346	00000000h	RW
58	AFLOG—Advanced Fault Log Register on page 346	0000000000000000h	RO
64	PMEN—Protected Memory Enable Register on page 347	00000000h	ROV; RW
68	PLMBASE—Protected Low-Memory Base Register on page 348	00000000h	RW
6C	PLMLIMIT—Protected Low-Memory Limit Register on page 348	00000000h	RW
70	PHMBASE—Protected High-Memory Base Register on page 349	0000000000000000h	RW
78	PHMLIMIT—Protected High-Memory Limit Register on page 349	0000000000000000h	RW
80	IQH—Invalidation Queue Head Register on page 350	0000000000000000h	ROV
88	IQT—Invalidation Queue Tail Register on page 350	0000000000000000h	RW_L
90	IQA—Invalidation Queue Address Register on page 351	0000000000000000h	RW_L
9C	ICS—Invalidation Completion Status Register on page 351	00000000h	RW1CS
A0	IECTL—Invalidation Event Control Register on page 351	80000000h	ROV; RW_L
A4	IEDATA—Invalidation Event Data Register on page 352	00000000h	RW_L
A8	IEADDR—Invalidation Event Address Register on page 352	00000000h	RW_L
AC	IEUADDR—Invalidation Event Upper Address Register on page 353	00000000h	RW_L
B8	IRTA—Interrupt Remapping Table Address Register on page 353	0000000000000000h	RW_L; ROV
100	IVA—Invalidate Address Register on page 354	0000000000000000h	RW
108	IOTLB—IOTLB Invalidate Register on page 355	0000000000000000h	RW; ROV; RW_V
200	FRCDL—Fault Recording Low Register on page 357	0000000000000000h	ROSV
208	FRCDH—Fault Recording High Register on page 357	0000000000000000h	ROSV; RO; RW1CS

5.5.1 VER—Version Register

Register to report the architecture version supported. Backward compatibility for the architecture is maintained with new revision numbers, allowing software to load remapping hardware drivers written for prior architecture versions.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RO	
Size: 32	Default Value: 00000010h		Address Offset: 0h	
Bit Range	Acronym	Description	Default	Access
31:8	RSVD	Reserved.	000000h	RO
7:4	MAJOR	Indicates supported architecture version.	1h	RO
3:0	MINOR	Indicates supported architecture minor version.	0h	RO

5.5.2 CAP—Capability Register

Register to report general remapping hardware capabilities

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RO; ROV	
Size: 64	Default Value: 00D2008C20660462h		Address Offset: 8h	
Bit Range	Acronym	Description	Default	Access
63:59	RSVD	Reserved.	00h	RO
58	SL64KP	A value of 1 in this field indicates 64-KByte page size is supported for second-level translation.	0h	RO
57	FL64KP	A value of 1 in this field indicates 64-KByte page size is supported for first-level translation.	0h	RO
56	FL1GP	A value of 1 in this field indicates 1-GByte page size is supported for first-level translation.	0h	RO
55	DRD	0: Hardware does not support draining of DMA read requests. 1: Hardware supports draining of DMA read requests.	1h	RO
54	DWD	0: Hardware does not support draining of DMA write requests. 1: Hardware supports draining of DMA write requests.	1h	RO
53:48	MAMV	The value in this field indicates the maximum supported value for the Address Mask (AM) field in the Invalidation Address register (IVA_REG) and IOTLB Invalidation Descriptor (iotlb_inv_dsc). This field is valid only when the PSI field in Capability register is reported as Set.	12h	RO
47:40	NFR	Number of fault recording registers is computed as N+1, where N is the value reported in this field. Implementations must support at least one fault recording register (NFR = 0) for each remapping hardware unit in the platform. The maximum number of fault recording registers per remapping hardware unit is 256.	00h	RO
39	PSI	0: Hardware supports only domain and global invalidates for IOTLB 1: Hardware supports page selective, domain and global invalidates for IOTLB.	1h	ROV

continued...

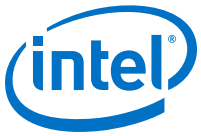

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RO; ROV	
Size: 64	**Default Value: 00D2008C20660462h**		**Address Offset: 8h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		Hardware implementations reporting this field as set are recommended to support a Maximum Address Mask Value (MAMV) value of at least 9.		
38	RSVD	Reserved.	0h	RO
37:34	SLLPS	This field indicates the super page sizes supported by hardware. A value of 1 in any of these bits indicates the corresponding super-page size is supported. The super-page sizes corresponding to various bit positions within this field are: 0: 21-bit offset to page frame (2MB) 1: 30-bit offset to page frame (1GB) 2: 39-bit offset to page frame (512GB) 3: 48-bit offset to page frame (1TB) Hardware implementations supporting a specific super-page size must support all smaller super-page sizes, i.e. only valid values for this field are 0000b, 0001b, 0011b, 0111b, 1111b.	3h	ROV
33:24	FRO	This field specifies the location to the first fault recording register relative to the register base address of this remapping hardware unit. If the register base address is X, and the value reported in this field is Y, the address for the first fault recording register is calculated as X+(16*Y).	020h	RO
23	RSVD	Reserved.	0h	RO
22	ZLR	0: Indicates the remapping hardware unit blocks (and treats as fault) zero length DMA read requests to write-only pages. 1: Indicates the remapping hardware unit supports zero length DMA read requests to write-only pages. DMA remapping hardware implementations are recommended to report ZLR field as Set.	1h	RO
21:16	MGAW	This field indicates the maximum DMA virtual addressability supported by remapping hardware. The Maximum Guest Address Width (MGAW) is computed as (N +1), where N is the value reported in this field. For example, a hardware implementation supporting 48-bit MGAW reports a value of 47 (101111b) in this field. If the value in this field is X, untranslated and translated DMA requests to addresses above 2^(x+1)-1 are always blocked by hardware. Translations requests to address above 2^(x+1)-1 from allowed devices return a null Translation Completion Data Entry with R=W=0. Guest addressability for a given DMA request is limited to the minimum of the value reported through this field and the adjusted guest address width of the corresponding page-table structure. (Adjusted guest address widths supported by hardware are reported through the SAGAW field). Implementations are recommended to support MGAW at least equal to the physical addressability (host address width) of the platform.	26h	RO
15:13	RSVD	Reserved.	0h	RO

continued...

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RO; ROV	
Size: 64	**Default Value: 00D2008C20660462h**		**Address Offset: 8h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
12:8	SAGAW	This 5-bit field indicates the supported adjusted guest address widths (which in turn represents the levels of page-table walks for the 4KB base page size) supported by the hardware implementation. A value of 1 in any of these bits indicates the corresponding adjusted guest address width is supported. The adjusted guest address widths corresponding to various bit positions within this field are: 0: 30-bit AGAW (2-level page table) 1: 39-bit AGAW (3-level page table) 2: 48-bit AGAW (4-level page table) 3: 57-bit AGAW (5-level page table) 4: 64-bit AGAW (6-level page table) Software must ensure that the adjusted guest address width used to setup the page tables is one of the supported guest address widths reported in this field.	04h	RO
7	CM	0: Not-present and erroneous entries are not cached in any of the renmapping caches. Invalidations are not required for modifications to individual not present or invalid entries. However, any modifications that result in decreasing the effective permissions or partial permission increases require invalidations for them to be effective. 1: Not-present and erroneous mappings may be cached in the remapping caches. Any software updates to the remapping structures (including updates to "not-present" or erroneous entries) require explicit invalidation. Hardware implementations of this architecture must support a value of 0 in this field.	0h	RO
6	PHMR	0: Indicates protected high-memory region is not supported. 1: Indicates protected high-memory region is supported.	1h	RO
5	PLMR	0: Indicates protected low-memory region is not supported. 1: Indicates protected low-memory region is supported.	1h	RO
4	RWBF	0: Indicates no write-buffer flushing is needed to ensure changes to memory-resident structures are visible to hardware. 1: Indicates software must explicitly flush the write buffers to ensure updates made to memory-resident remapping structures are visible to hardware.	0h	RO
3	AFL	0: Indicates advanced fault logging is not supported. Only primary fault logging is supported. 1: Indicates advanced fault logging is supported.	0h	RO
2:0	ND	000b: Hardware supports 4-bit domain-ids with support for up to 16 domains. 001b: Hardware supports 6-bit domain-ids with support for up to 64 domains. 010b: Hardware supports 8-bit domain-ids with support for up to 256 domains. 011b: Hardware supports 10-bit domain-ids with support for up to 1024 domains. 100b: Hardware supports 12-bit domain-ids with support for up to 4K domains. 100b: Hardware supports 14-bit domain-ids with support for up to 16K domains.	2h	RO

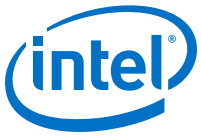

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RO; ROV	
Size: 64	Default Value: 00D2008C20660462h		Address Offset: 8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		110b: Hardware supports 16-bit domain-ids with support for up to 64K domains. 111b: Reserved.		

5.5.3 ECAP—Extended Capability Register

Register to report remapping hardware extended capabilities

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RO; ROV	
Size: 64	Default Value: 0000000000F010DAh		Address Offset: 10h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:40	RSVD	Reserved.	000000h	RO
39:35	PSS	This field reports the PASID size supported by the remapping hardware for requestswith- PASID. A value of N in this field indicates hardware supports PASID field of N+1 bits (For example, value of 7 in this field, indicates 8-bit PASIDs are supported). Requests-with-PASID with PASID value beyond the limit specified by this field are treated as error by the remapping hardware. This field is valid only when PASID field is reported as Set.	00h	RO
34	EAFS	0: Hardware does not support the extended-accessed (EA) bit in first-level paging-structure entries. 1: Hardware supports the extendedaccessed (EA) bit in first-level paging-structure entries. This field is valid only when PASID field is reported as Set.	0h	RO
33	NWFS	0: Hardware ignores the "No Write" (NW) flag in Device-TLB translationrequests, and behaves as if NW is always 0. 1: Hardware supports the "No Write" (NW) flag in Device-TLB translationrequests. This field is valid only when Device-TLB support (DT) field is reported as Set.	0h	RO
32	POT	0: Hardware does not support PASID-only Translation Type in extended-context-entries 1: Hardware supports PASID-only Translation Type in extended-context-entries	0h	RO
31	SRS	0: H/W does not support requests-with-PASID seeking supervisor privilege 1: H/W supports requests-with-PASID seeking supervisor privilege	0h	RO
30	ERS	0: H/W does not support requests seeking execute permission 1: H/W supports requests seeking execute permission	0h	RO
29	PRS	0: Hardware does not support Page Requests 1: Hardware supports Page Requests	0h	RO
28	PASID	0: Hardware does not support process address space IDs. 1: Hardware supports Process Address Space IDs.	0h	RO

continued...

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RO; ROV	
Size: 64	Default Value: 0000000000F010DAh		Address Offset: 10h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
27	DIS	0: Hardware does not support deferred invalidations of IOTLB and Device-TLB. 1: Hardware supports deferred invalidations of IOTLB and Device-TLB.	0h	RO
26	NEST	0: Hardware does not support nested translations. 1: Hardware supports nested translations.	0h	RO
25	MTS	0: Hardware does not support Memory Type 1: Hardware supports Memory Type	0h	RO
24	ECS	0: Hardware does not support extended-root-entries and Extended Context-Entries 1: Hardware supports extended-root-entries and Extended Context-Entries	0h	RO
23:20	MHMV	The value in this field indicates the maximum supported value for the Handle Mask (HM) field in the interrupt entry cache invalidation descriptor (iec_inv_dsc). This field is valid only when the IR field in Extended Capability register is reported as Set.	Fh	RO
19:18	RSVD	Reserved.	0h	RO
17:8	IRO	This field specifies the offset to the IOTLB registers relative to the register base address of this remapping hardware unit. If the register base address is X, and the value reported in this field is Y, the address for the first IOTLB invalidation register is calculated as X+(16*Y).	010h	RO
7	SC	0: Hardware does not support 1-setting of the SNP field in the page-table entries. 1: Hardware supports the 1-setting of the SNP field in the page-table entries.	1h	ROV
6	PT	0: Hardware does not support pass-through translation type in context entries. 1: Hardware supports pass-through translation type in context entries.	1h	ROV
5	RSVD	Reserved.	0h	RO
4	EIM	0: On Intel®64 platforms, hardware supports only 8-bit APIC-IDs (xAPIC mode). 1: On Intel®64 platforms, hardware supports 32-bit APIC-IDs (x2APIC mode). This field is valid only on Intel®64 platforms reporting Interrupt Remapping support (IR field Set).	1h	ROV
3	IR	0: Hardware does not support interrupt remapping. 1: Hardware supports interrupt remapping. Implementations reporting this field as Set must also support Queued Invalidation (QI).	1h	ROV

continued...

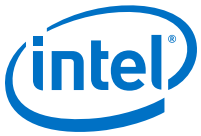

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RO; ROV	
Size: 64	**Default Value: 0000000000F010DAh**		**Address Offset: 10h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
2	DT	0: Hardware does not support device-IOTLBs. 1: Hardware supports Device-IOTLBs. Implementations reporting this field as Set must also support Queued Invalidation (QI).	0h	RO
1	QI	0: Hardware does not support queued invalidations. 1: Hardware supports queued invalidations.	1h	ROV
0	C	This field indicates if hardware access to the root, context, page-table and interrupt-remap structures are coherent (snooped) or not. 0: Indicates hardware accesses to remapping structures are non-coherent. 1: Indicates hardware accesses to remapping structures are coherent. Hardware access to advanced fault log and invalidation queue are always coherent.	0h	RO

5.5.4 GCMD—Global Command Register

Register to control remapping hardware. If multiple control fields in this register need to be modified, software must serialize the modifications through multiple writes to this register.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: WO; RO	
Size: 32	**Default Value: 00000000h**		**Address Offset: 18h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31	TE	Software writes to this field to request hardware to enable/ disable DMA-remapping: 0: Disable DMA remapping 1: Enable DMA remapping Hardware reports the status of the translation enable operation through the TES field in the Global Status register. There may be active DMA requests in the platform when software updates this field. Hardware must enable or disable remapping logic only at determinstic transaction boundaries, so that any in-flight transaction is either subject to remapping or not at all. Hardware implementations supporting DMA draining must drain any in-flight DMA read/write requests queued within the Root-Complex before completing the translation enable command and reflecting the status of the command through the TES field in the Global Status register. The value returned on a read of this field is undefined.	0h	WO
30	SRTP	Software sets this field to set/update the root-entry table pointer used by hardware. The root-entry table pointer is specified through the Root-entry Table Address (RTA_REG) register. Hardware reports the status of the "Set Root Table Pointer" operation through the RTPS field in the Global Status register. The "Set Root Table Pointer" operation must be performed before enabling or re-enabling (after disabling) DMA remapping through the TE field. After a "Set Root Table Pointer" operation, software must	0h	WO

continued...

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: WO; RO	
Size: 32	Default Value: 00000000h		Address Offset: 18h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		globally invalidate the context cache and then globally invalidate of IOTLB. This is required to ensure hardware uses only the remapping structures referenced by the new root table pointer, and not stale cached entries. While DMA remapping hardware is active, software may update the root table pointer through this field. However, to ensure valid in-flight DMA requests are deterministically remapped, software must ensure that the structures referenced by the new root table pointer are programmed to provide the same remapping results as the structures referenced by the previous root-table pointer. Clearing this bit has no effect. The value returned on read of this field is undefined.		
29	SFL	This field is valid only for implementations supporting advanced fault logging. Software sets this field to request hardware to set/update the fault-log pointer used by hardware. The fault-log pointer is specified through Advanced Fault Log register. Hardware reports the status of the 'Set Fault Log' operation through the FLS field in the Global Status register. The fault log pointer must be set before enabling advanced fault logging (through EAFL field). Once advanced fault logging is enabled, the fault log pointer may be updated through this field while DMA remapping is active. Clearing this bit has no effect. The value returned on read of this field is undefined.	0h	RO
28	EAFL	This field is valid only for implementations supporting advanced fault logging. Software writes to this field to request hardware to enable or disable advanced fault logging: 0: Disable advanced fault logging. In this case, translation faults are reported through the Fault Recording registers. 1: Enable use of memory-resident fault log. When enabled, translation faults are recorded in the memory-resident log. The fault log pointer must be set in hardware (through the SFL field) before enabling advaned fault logging. Hardware reports the status of the advaned fault logging enable operation through the AFLS field in the Global Status register. The value returned on read of this field is undefined.	0h	RO
27	WBF	This bit is valid only for implementations requiring write buffer flushing. Software sets this field to request that hardware flush the Root-Complex internal write buffers. This is done to ensure any updates to the memory-resident remapping structures are not held in any internal write posting buffers. Hardware reports the status of the write buffer flushing operation through the WBFS field in the Global Status register. Cleraing this bit has no effect. The value returned on a read of this field is undefined.	0h	RO
26	QIE	This field is valid only for implementations supporting queued invalidations. Software writes to this field to enable or disable queued invalidations. 0: Disable queued invalidations.	0h	WO

continued...

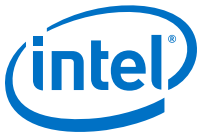

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: WO; RO	
Size: 32	**Default Value: 00000000h**		**Address Offset: 18h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		1: Enable use of queued invalidations. Hardware reports the status of queued invalidation enable operation through QIES field in the Global Status register. The value returned on a read of this field is undefined.		
25	IRE	This field is valid only for implementations supporting interrupt remapping. 0: Disable interrupt-remapping hardware 1: Enable interrupt-remapping hardware Hardware reports the status of the interrupt remapping enable operation through the IRES field in the Global Status register. There may be active interrupt requests in the platform when software updates this field. Hardware must enable or disable interrupt-remapping logic only at deterministic transaction boundaries, so that any in-flight interrupts are either subject to remapping or not at all. Hardware implementations must drain any in-flight interrupts requests queued in the Root-Complex before completing the interrupt-remapping enable command and reflecting the status of the command through the IRES field in the Global Status register. The value returned on a read of this field is undefined.	0h	WO
24	SIRTP	This field is valid only for implementations supporting interrupt-remapping. Software sets this field to set/update the interrupt remapping table pointer used by hardware. The interrupt remapping table pointer is specified through the Interrupt Remapping Table Address (IRTA_REG) register. Hardware reports the status of the 'Set Interrupt Remap Table Pointer' operation through the IRTPS field in the Global Status register. The 'Set Interrupt Remap Table Pointer' operation must be performed before enabling or re-enabling (after disabling) interrupt-remapping hardware through the IRE field. After a 'Set Interrupt Remap Table Pointer' operation, software must globally invalidate the interrupt entry cache. This is required to ensure hardware uses only the interrupt-remapping entries referenced by the new interrupt remap table pointer, and not any stale cached entries. While interrupt remapping is active, software may update the interrupt remapping table pointer through this field. However, to ensure valid in-flight interrupt requests are deterministically remapped, software must ensure that the structures referenced by the new interrupt remap table pointer are programmed to provide the same remapping results as the structures referenced by the previous interrupt remap table pointer. Clearing this bit has no effect. The value returned on a read of this field is undefined.	0h	WO
23	CFI	This field is valid only for Intel®64 implementations supporting interrupt-remapping. Software writes to this field to enable or disable Compatibility Format interrupts on Intel®64 platforms. The value in this field is effective only when interrupt-remapping is enabled and Extended Interrupt Mode (x2APIC mode) is not enabled. 0: Block Compatibility format interrupts. 1: Process Compatibility format interrupts as pass-through (bypass interrupt remapping).	0h	WO

continued...

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: WO; RO	
Size: 32	Default Value: 00000000h		Address Offset: 18h	
Bit Range	Acronym	Description	Default	Access
		Hardware reports the status of updating this field through the CFIS field in the Global Status register. The value returned on a read of this field is undefined.		
22:0	RSVD	Reserved.	000000h	RO

5.5.5 GSTS—Global Status Register

Register to report general remapping hardware status.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RO_V; ROV; RO	
Size: 32	Default Value: 00000000h		Address Offset: 1Ch	
Bit Range	Acronym	Description	Default	Access
31	TES	This field indicates the status of DMA-remapping hardware. 0: DMA-remapping hardware is not enabled 1: DMA-remapping hardware is enabled	0h	ROV
30	RTPS	This field indicates the status of the root- table pointer in hardware. This field is cleared by hardware when software sets the SRTP field in the Global Command register. This field is set by hardware when hardware completes the 'Set Root Table Pointer' operation using the value provided in the Root-Entry Table Address register.	0h	RO_V
29	FLS	This field: - Is cleared by hardware when software Sets the SFL field in the Global Command register. - Is Set by hardware whn hardware completes the 'Set Fault Log Pointer' operation using the value provided in the Advanced Fault Log register.	0h	RO
28	AFLS	This field is valid only for implementations supporting advanced fault logging. It indicates the advanced fault logging status: 0: Advanced Fault Logging is not enabled. 1: Advanced Fault Logging is enabled.	0h	RO
27	WBFS	This field is valid only for implementations requiring write buffer flushing. This field indicates the status of the write buffer flush command. It is: - Set by hardware when software sets the WBF field in the Global Command register. - Cleared by hardware when hardware completes the write buffer flushing operation.	0h	RO
26	QIES	This field indicates queued invalidation enable status. 0: queued invalidation is not enabled 1: queued invalidation is enabled	0h	RO_V
25	IRES	This field indicates the status of Interrupt-remapping hardware. 0: Interrupt-remapping hardware is not enabled 1: Interrupt-remapping hardware is enabled	0h	ROV

continued...

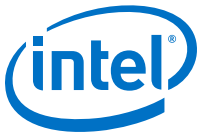

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RO_V; ROV; RO	
Size: 32	**Default Value: 00000000h**		**Address Offset: 1Ch**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
24	IRTPS	This field indicates the status of the interrupt remapping table pointer in hardware. This field is cleared by hardware when software sets the SIRTP field in the Global Command register. This field is Set by hardware when hardware completes the set interrupt remap table pointer operation using the value provided in the Interrupt Remapping Table Address register.	0h	RO_V
23	CFIS	This field indicates the status of Compatibility format interrupts on Intel®64 implementations supporting interrupt-remapping. The value reported in this field is applicable only when interrupt-remapping is enabled and Extended Interrupt Mode (x2APIC mode) is not enabled. 0: Compatibility format interrupts are blocked. 1: Compatibility format interrupts are processed as pass-through (bypassing interrupt remapping).	0h	RO_V
22:0	RSVD	Reserved.	000000h	RO

5.5.6 RTADDR—Root-Entry Table Address Register

Register providing the base address of root-entry table.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RO; RW	
Size: 64	**Default Value: 0000000000000000h**		**Address Offset: 20h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:12	RTA	This register points to base of page aligned, 4KB-sized root-entry table in system memory. Hardware ignores and not implements bits 63:HAW, where HAW is the host address width. Software specifies the base address of the root-entry table through this register, and programs it in hardware through the SRTP field in the Global Command register. Reads of this register returns value that was last programmed to it.	0000000h	RW
11	RTT	PLACEHOLDER: This field specifies the type of root-table referenced by the Root Table Address (RTA) field; 0: Root Table / 1: Extended Root Table	0h	RO
10:0	RSVD	Reserved.	000h	RO

5.5.7 CCMD—Context Command Register

Register to manage context cache. The act of writing the uppermost byte of the CCMD_REG with the ICC field Set causes the hardware to perform the context-cache invalidation.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW; ROV; RW_V	
Size: 64	Default Value: 0000000000000000h		Address Offset: 28h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63	ICC	Software requests invalidation of context-cache by setting this field. Software must also set the requested invalidation granularity by programming the CIRG field. Software must read back and check the ICC field is Clear to confirm the invalidation is complete. Software must not update this register when this field is set. Hardware clears the ICC field to indicate the invalidation request is complete. Hardware also indicates the granularity at which the invalidation operation was performed through the CAIG field. Software must submit a context-cache invalidation request through this field only when there are no invalidation requests pending at this remapping hardware unit. Since information from the context-cache may be used by hardware to tag IOTLB entries, software must perform domain-selective (or global) invalidation of IOTLB after the context cache invalidation has completed. Hardware implementations reporting write-buffer flushing requirement (RWBF=1 in Capability register) must implicitly perform a write buffer flush before invalidating the context cache.	0h	RW_V
62:61	CIRG	Software provides the requested invalidation granularity through this field when setting the ICC field: 00: Reserved. 01: Global Invalidation request. 10: Domain-selective invalidation request. The target domain-id must be specified in the DID field. 11: Device-selective invalidation request. The target source-id(s) must be specified through the SID and FM fields, and the domain-id (that was programmed in the context-entry for these device(s)) must be provided in the DID field. Hardware implementations may process an invalidation request by performing invalidation at a coarser granularity than requested. Hardware indicates completion of the invalidation request by clearing the ICC field. At this time, hardware also indicates the granularity at which the actual invalidation was performed through the CAIG field.	0h	RW
60:59	CAIG	Hardware reports the granularity at which an invalidation request was processed through the CAIG field at the time of reporting invalidation completion (by clearing the ICC field). The following are the encodings for this field: 00: Reserved. 01: Global Invalidation performed. This could be in response to a global, domain-selective or device-selective invalidation request. 10: Domain-selective invalidation performed using the domain-id specified by software in the DID field. This could be in response to a domain-selective or device-selective invalidation request. 11: Device-selective invalidation performed using the	0h	ROV

continued...

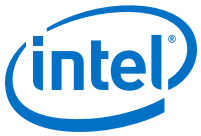

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW; ROV; RW_V	
Size: 64	**Default Value: 0000000000000000h**		**Address Offset: 28h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		source-id and domain-id specified by software in the SID and FM fields. This can only be in response to a device-selective invalidation request.		
58:34	RSVD	Reserved.	0000000h	RO
33:32	FM	Software may use the Function Mask to perform device-selective invalidations on behalf of devices supporting PCI Express Phantom Functions. This field specifies which bits of the function number portion (least significant three bits) of the SID field to mask when performing device-selective invalidations. The following encodings are defined for this field: 00: No bits in the SID field masked. 01: Mask most significant bit of function number in the SID field. 10: Mask two most significant bit of function number in the SID field. 11: Mask all three bits of function number in the SID field. The context-entries corresponding to all the source-ids specified through the FM and SID fields must have to the domain-id specified in the DID field.	0h	RW
31:16	SID	Indicates the source-id of the device whose corresponding context-entry needs to be selectively invalidated. This field along with the FM field must be programmed by software for device-selective invalidation requests.	0000h	RW
15:8	RSVD	Reserved.	00h	RO
7:0	DID	Indicates the id of the domain whose context-entries need to be selectively invalidated. This field must be programmed by software for both domain-selective and device-selective invalidation requests. The Capability register reports the domain-id width supported by hardware. Software must ensure that the value written to this field is within this limit. Hardware may ignore and not implement bits15:N, where N is the supported domain-id width reported in the Capability register.	00h	RW

5.5.8 FSTS—Fault Status Register

Register indicating the various error status.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW1CS; ROSV; RO	
Size: 32	**Default Value: 00000000h**		**Address Offset: 34h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:16	RSVD	Reserved.	0000h	RO
15:8	FRI	This field is valid only when the PPF field is Set. The FRI field indicates the index (from base) of the fault recording register to which the first pending fault was recorded when the PPF field was Set by hardware. The value read from this field is undefined when the PPF field is clear.	00h	RO

continued...

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW1CS; ROSV; RO	
Size: 32	Default Value: 00000000h		Address Offset: 34h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
7	PRO	Hardware detected a Page Request Overflow error. Hardware implementations not supporting the Page Request Queue implement this bit as RsvdZ.	0h	RO
6	ITE	Hardware detected a Device-IOTLB invalidation completion time-out. At this time, a fault event may be generated based on the programming of the Fault Event Control register. Hardware implementations not supporting device Device-IOTLBs implement this bit as RsvdZ.	0h	RO
5	ICE	Hardware received an unexpected or invalid Device-IOTLB invalidation completion. This could be due to either an invalid ITag or invalid source-id in an invalidation completion response. At this time, a fault event may be generated based on the programming of hte Fault Event Control register. Hardware implementations not supporting Device-IOTLBs implement this bit as RsvdZ.	0h	RO
4	IQE	Hardware detected an error associated with the invalidation queue. This could be due to either a hardware error while fetching a descriptor from the invalidation queue, or hardware detecting an erroneous or invalid descriptor in the invalidation queue. At this time, a fault event may be generated based on the programming of the Fault Event Control register. Hardware implementations not supporting queued invalidations implement this bit as RsvdZ.	0h	RW1CS
3	APF	When this field is Clear, hardware sets this field when the first fault record (at index 0) is written to a fault log. At this time, a fault event is generated based on the programming of the Fault Event Control register. Software writing 1 to this field clears it. Hardware implementations not supporting advanced fault logging implement this bit as RsvdZ.	0h	RO
2	AFO	Hardware sets this field to indicate advanced fault log overflow condition. At this time, a fault event is generated based on the programming of the Fault Event Control register. Software writing 1 to this field clears it. Hardware implementations not supporting advanced fault logging implement this bit as RsvdZ.	0h	RO
1	PPF	This field indicates if there are one or more pending faults logged in the fault recording registers. Hardware computes this field as the logical OR of Fault (F) fields across all the fault recording registers of this remapping hardware unit. 0: No pending faults in any of the fault recording registers 1: One or more fault recording registers has pending faults. The FRI field is updated by hardware whenever the PPF field is set by hardware. Also, depending on the programming of Fault Event Control register, a fault event is generated when hardware sets this field.	0h	ROSV
0	PFO	Hardware sets this field to indicate overflow of fault recording registers. Software writing 1 clears this field. When this field is Set, hardware does not record any new faults until software clears this field.	0h	RW1CS

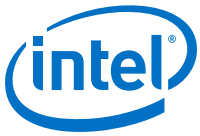

5.5.9 FECTL—Fault Event Control Register

Register specifying the fault event interrupt message control bits.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: ROV; RW	
Size: 32	**Default Value: 80000000h**		**Address Offset: 38h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31	IM	0: No masking of interrupt. When an interrupt condition is detected, hardware issues an interrupt message (using the Fault Event Data and Fault Event Address register values). 1: This is the value on reset. Software may mask interrupt message generation by setting this field. Hardware is prohibited from sending the interrupt message when this field is set.	1h	RW
30	IP	Hardware sets the IP field whenever it detects an interrupt condition, which is defined as: When primary fault logging is active, an interrupt condition occurs when hardware records a fault through one of the Fault Recording registers and sets the PPF field in Fault Status register. When advanced fault logging is active, an interrupt condition occurs when hardware records a fault in the first fault record (at index 0) of the current fault log and sets the APF field in the Fault Status register. Hardware detected error associated with the Invalidation Queue, setting the IQE field in the Fault Status register. Hardware detected invalid Device-IOTLB invalidation completion, setting the ICE field in the Fault Status register. Hardware detected Device-IOTLB invalidation completion time-out, setting the ITE field in the Fault Status register. If any of the status fields in the Fault Status register was already Set at the time of setting any of these fields, it is not treated as a new interrupt condition. The IP field is kept set by hardware while the interrupt message is held pending. The interrupt message could be held pending due to interrupt mask (IM field) being Set or other transient hardware conditions. The IP field is cleared by hardware as soon as the interrupt message pending condition is serviced. This could be due to either: Hardware issuing the interrupt message due to either change in the transient hardware condition that caused interrupt message to be held pending, or due to software clearing the IM field.. Software servicing all the pending interrupt status fields in the Fault Status register as follows: - When primary fault logging is active, software clearing the Fault (F) field in all the Fault Recording registers with faults, causing the PPF field in Fault Status register to be evaluated as clear. - Software clearing other status fields in the Fault Status register by writing back the value read from the respective fields.	0h	ROV
29:0	RSVD	Reserved.	00000000h	RO

5.5.10 FEDATA—Fault Event Data Register

Register specifying the interrupt message data

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 3Ch	
Bit Range	Acronym	Description	Default	Access
31:16	EIMD	This field is valid only for implementations supporting 32-bit interrupt data fields. Hardware implementations supporting only 16-bit interrupt data may treat this field as RsvdZ.	0000h	RW
15:0	IMD	Data value in the interrupt request.	0000h	RW

5.5.11 FEADDR—Fault Event Address Register

Register specifying the interrupt message address.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 40h	
Bit Range	Acronym	Description	Default	Access
31:2	MA	When fault events are enabled, the contents of this register specify the DWORD-aligned address (bits 31:2) for the interrupt request.	00000000h	RW
1:0	RSVD	Reserved.	0h	RO

5.5.12 FEUADDR—Fault Event Upper Address Register

Register specifying the interrupt message upper address.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 44h	
Bit Range	Acronym	Description	Default	Access
31:0	MUA	Hardware implementations supporting Extended Interrupt Mode are required to implement this register. Hardware implementations not supporting Extended Interrupt Mode may treat this field as RsvdZ.	00000000h	RW

5.5.13 AFLOG—Advanced Fault Log Register

Register to specify the base address of the memory-resident fault-log region. This register is treated as RsvdZ for implementations not supporting advanced translation fault logging (AFL field reported as 0 in the Capability register).

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RO	
Size: 64	Default Value: 0000000000000000h		Address Offset: 58h	
Bit Range	Acronym	Description	Default	Access
63:12	FLA	This field specifies the base of 4KB aligned fault-log region in system memory. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address	0000000000000h	RO

continued...

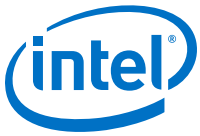

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RO	
Size: 64	**Default Value: 0000000000000000h**		**Address Offset: 58h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		width. Software specifies the base address and size of the fault log region through this register, and programs it in hardware through the SFL field in the Global Command register. When implemented, reads of this field return the value that was last programmed to it.		
11:9	FLS	This field specifies the size of the fault log region pointed by the FLA field. The size of the fault log region is 2^X * 4KB, where X is the value programmed in this register. When implemented, reads of this field return the value that was last programmed to it.	0h	RO
8:0	RSVD	Reserved.	000h	RO

5.5.14 PMEN—Protected Memory Enable Register

Register to enable the DMA-protected memory regions setup through the PLMBASE, PLMLIMT, PHMBASE, PHMLIMIT registers. This register is always treated as RO for implementations not supporting protected memory regions (PLMR and PHMR fields reported as Clear in the Capability register).
Protected memory regions may be used by software to securely initialize remapping structures in memory. To avoid impact to legacy BIOS usage of memory, software is recommended to not overlap protected memory regions with any reserved memory regions of the platform reported through the Reserved Memory Region Reporting (RMRR) structures.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: ROV; RW	
Size: 32	**Default Value: 00000000h**		**Address Offset: 64h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31	EPM	This field controls DMA accesses to the protected low-memory and protected high-memory regions. 0: Protected memory regions are disabled. 1: Protected memory regions are enabled. DMA requests accessing protected memory regions are handled as follows: - When DMA remapping is not enabled, all DMA requests accessing protected memory regions are blocked. - When DMA remapping is enabled: - DMA requests processed as pass-through (Translation Type value of 10b in Context-Entry) and accessing the protected memory regions are blocked. - DMA requests with translated address (AT=10b) and accessing the protected memory regions are blocked. - DMA requests that are subject to address remapping, and accessing the protected memory regions may or may not be blocked by hardware. For such requests, software must not depend on hardware protection of the protected memory regions, and instead program the DMA-remapping page-tables to not allow DMA to protected memory regions. Remapping hardware access to the remapping structures are not subject to protected memory region checks. DMA requests blocked due to protected memory region violation are not recorded or reported as remapping faults. Hardware reports the status of the protected memory	0h	RW

continued...

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: ROV; RW	
Size: 32	Default Value: 00000000h		Address Offset: 64h	
Bit Range	Acronym	Description	Default	Access
		enable/disable operation through the PRS field in this register. Hardware implementations supporting DMA draining must drain any in-flight translated DMA requests queued within the Root-Complex before indicating the protected memory region as enabled through the PRS field.		
30:1	RSVD	Reserved.	00000000h	RO
0	PRS	This field indicates the status of protected memory region(s): 0: Protected memory region(s) disabled. 1: Protected memory region(s) enabled.	0h	ROV

5.5.15 PLMBASE—Protected Low-Memory Base Register

Register to set up the base address of DMA-protected low-memory region below 4GB. This register must be set up before enabling protected memory through PMEN_REG, and must not be updated when protected memory regions are enabled.
This register is always treated as RO for implementations not supporting protected low memory region (PLMR field reported as Clear in the Capability register).
The alignment of the protected low memory region base depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1s to this register, and finding the most significant zero bit position with 0 in the value read back from the register. Bits N:0 of this register is decoded by hardware as all 0s.
Software must setup the protected low memory region below 4GB.
Software must not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 68h	
Bit Range	Acronym	Description	Default	Access
31:20	PLMB	This register specifies the base of protected low-memory region in system memory.	000h	RW
19:0	RSVD	Reserved.	00000h	RO

5.5.16 PLMLIMIT—Protected Low-Memory Limit Register

Register to set up the limit address of DMA-protected low-memory region below 4GB. This register must be set up before enabling protected memory through PMEN_REG, and must not be updated when protected memory regions are enabled.
This register is always treated as RO for implementations not supporting protected low memory region (PLMR field reported as Clear in the Capability register).
The alignment of the protected low memory region limit depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1's to this register, and finding most significant zero bit position with 0 in the value read back from the register. Bits N:0 of the limit register is decoded by hardware as all 1s.
The Protected low-memory base and limit registers functions as follows:
- Programming the protected low-memory base and limit registers with the same value
  in bits 31:(N+1) specifies a protected low-memory region of size 2^(N+1) bytes.

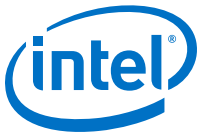

- Programming the protected low-memory limit register with a value less than the protected low-memory base register disables the protected low-memory region.

Software must not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW	
Size: 32	Default Value: 00000000h		Address Offset: 6Ch	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:20	PLML	This register specifies the last host physical address of the DMA-protected low-memory region in system memory.	000h	RW
19:0	RSVD	Reserved.	00000h	RO

5.5.17 PHMBASE—Protected High-Memory Base Register

Register to set up the base address of DMA-protected high-memory region. This register must be set up before enabling protected memory through PMEN_REG, and must not be updated when protected memory regions are enabled.

This register is always treated as RO for implementations not supporting protected high memory region (PHMR field reported as Clear in the Capability register).

The alignment of the protected high memory region base depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1's to this register, and finding most significant zero bit position below host address width (HAW) in the value read back from the register. Bits N:0 of this register are decoded by hardware as all 0s.

Software may setup the protected high memory region either above or below 4GB.

Software must not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW	
Size: 64	Default Value: 0000000000000000h		Address Offset: 70h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:20	PHMB	This register specifies the base of protected (high) memory region in system memory. Hardware ignores, and does not implement, bits 63:HAW, where HAW is the host address width.	00000h	RW
19:0	RSVD	Reserved.	00000h	RO

5.5.18 PHMLIMIT—Protected High-Memory Limit Register

Register to set up the limit address of DMA-protected high-memory region. This register must be set up before enabling protected memory through PMEN_REG, and must not be updated when protected memory regions are enabled.

This register is always treated as RO for implementations not supporting protected high memory region (PHMR field reported as Clear in the Capability register).

The alignment of the protected high memory region limit depends on the number of reserved bits (N:0) of this register. Software may determine the value of N by writing all 1's to this register, and finding most significant zero bit position below host address width (HAW) in the value read back from the register. Bits N:0 of the limit register is decoded by hardware as all 1s.

The protected high-memory base & limit registers functions as follows.

- Programming the protected low-memory base and limit registers with the same value
  in bits HAW:(N+1) specifies a protected low-memory region of size 2^(N+1) bytes.
- Programming the protected high-memory limit register with a value less than the protected high-memory base register disables the protected high-memory region. Software must not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW	
Size: 64	Default Value: 0000000000000000h		Address Offset: 78h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:20	PHML	This register specifies the last host physical address of the DMA-protected high-memory region in system memory. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address width.	00000h	RW
19:0	RSVD	Reserved.	00000h	RO

5.5.19 IQH—Invalidation Queue Head Register

Register indicating the invalidation queue head. This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: ROV	
Size: 64	Default Value: 0000000000000000h		Address Offset: 80h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:19	RSVD	Reserved.	000000000000h	RO
18:4	QH	Specifies the offset (128-bit aligned) to the invalidation queue for the command that will be fetched next by hardware. Hardware resets this field to 0 whenever the queued invalidation is disabled (QIES field Clear in the Global Status register).	0000h	ROV
3:0	RSVD	Reserved.	0h	RO

5.5.20 IQT—Invalidation Queue Tail Register

Register indicating the invalidation tail head. This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW_L	
Size: 64	Default Value: 0000000000000000h		Address Offset: 88h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:19	RSVD	Reserved.	000000000000h	RO
18:4	QT	Specifies the offset (128-bit aligned) to the invalidation queue for the command that will be written next by software.	0000h	RW_L
3:0	RSVD	Reserved.	0h	RO

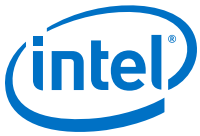

5.5.21 IQA—Invalidation Queue Address Register

Register to configure the base address and size of the invalidation queue. This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW_L	
Size: 64	Default Value: 0000000000000000h		Address Offset: 90h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:12	IQA	This field points to the base of 4KB aligned invalidation request queue. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address width. Reads of this field return the value that was last programmed to it.	0000000h	RW_L
11:3	RSVD	Reserved.	000h	RO
2:0	QS	This field specifies the size of the invalidation request queue. A value of X in this field indicates an invalidation request queue of (2^X) 4KB pages. The number of entries in the invalidation queue is 2^(X + 8).	0h	RW_L

5.5.22 ICS—Invalidation Completion Status Register

Register to report completion status of invalidation wait descriptor with Interrupt Flag (IF) Set.
This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW1CS	
Size: 32	Default Value: 00000000h		Address Offset: 9Ch	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:1	RSVD	Reserved.	00000000h	RO
0	IWC	Indicates completion of Invalidation Wait Descriptor with Interrupt Flag (IF) field Set. Hardware implementations not supporting queued invalidations implement this field as RsvdZ.	0h	RW1CS

5.5.23 IECTL—Invalidation Event Control Register

Register specifying the invalidation event interrupt control bits.
This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: ROV; RW_L	
Size: 32	Default Value: 80000000h		Address Offset: A0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31	IM	0: No masking of interrupt. When a invalidation event condition is detected, hardware issues an interrupt message (using the Invalidation Event Data & Invalidation	1h	RW_L

continued...

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: ROV; RW_L	
Size: 32	Default Value: 80000000h		Address Offset: A0h	
Bit Range	Acronym	Description	Default	Access
		Event Address register values). 1: This is the value on reset. Software may mask interrupt message generation by setting this field. Hardware is prohibited from sending the interrupt message when this field is Set.		
30	IP	Hardware sets the IP field whenever it detects an interrupt condition. Interrupt condition is defined as: - An Invalidation Wait Descriptor with Interrupt Flag (IF) field Set completed, setting the IWC field in the Invalidation Completion Status register. - If the IWC field in the Invalidation Completion Status register was already Set at the time of setting this field, it is not treated as a new interrupt condition. The IP field is kept Set by hardware while the interrupt message is held pending. The interrupt message could be held pending due to interrupt mask (IM field) being Set, or due to other transient hardware conditions. The IP field is cleared by hardware as soon as the interrupt message pending condition is serviced. This could be due to either: - Hardware issuing the interrupt message due to either change in the transient hardware condition that caused interrupt message to be held pending or due to software clearing the IM field. - Software servicing the IWC field in the Invalidation Completion Status register.	0h	ROV
29:0	RSVD	Reserved.	00000000h	RO

5.5.24 IEDATA—Invalidation Event Data Register

Register specifying the Invalidation Event interrupt message data.
This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW_L	
Size: 32	Default Value: 00000000h		Address Offset: A4h	
Bit Range	Acronym	Description	Default	Access
31:16	EIMD	This field is valid only for implementations supporting 32-bit interrupt data fields. Hardware implementations supporting only 16-bit interrupt data treat this field as Rsvd.	0000h	RW_L
15:0	IMD	Data value in the interrupt request.	0000h	RW_L

5.5.25 IEADDR—Invalidation Event Address Register

Register specifying the Invalidation Event Interrupt message address.
This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

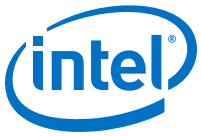

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW_L	
Size: 32	Default Value: 00000000h		Address Offset: A8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:2	MA	When fault events are enabled, the contents of this register specify the DWORD-aligned address (bits 31:2) for the interrupt request.	00000000h	RW_L
1:0	RSVD	Reserved.	0h	RO

5.5.26 IEUADDR—Invalidation Event Upper Address Register

Register specifying the Invalidation Event interrupt message upper address.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW_L	
Size: 32	Default Value: 00000000h		Address Offset: ACh	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:0	MUA	Hardware implementations supporting Queued Invalidations and Extended Interrupt Mode are required to implement this register. Hardware implementations not supporting Queued Invalidations or Extended Interrupt Mode may treat this field as RsvdZ.	00000000h	RW_L

5.5.27 IRTA—Interrupt Remapping Table Address Register

Register providing the base address of Interrupt remapping table. This register is treated as RsvdZ by implementations reporting Interrupt Remapping (IR) as not supported in the Extended Capability register.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW_L; ROV	
Size: 64	Default Value: 0000000000000000h		Address Offset: B8h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:12	IRTA	This field points to the base of 4KB aligned interrupt remapping table. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address width. Reads of this field returns value that was last programmed to it.	0000000h	RW_L
11	EIME	This field is used by hardware on Intel®64 platforms as follows: 0: xAPIC mode is active. Hardware interprets only low 8-bits of Destination-ID field in the IRTEs. The high 24-bits of the Destination-ID field are treated as reserved. 1: x2APIC mode is active. Hardware interprets all 32-bits of Destination-ID field in the IRTEs.	0h	ROV

continued...

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW_L; ROV	
Size: 64	Default Value: 0000000000000000h		Address Offset: B8h	
Bit Range	Acronym	Description	Default	Access
		This field is implemented as RsvdZ on implementations reporting Extended Interrupt Mode (EIM) field as Clear in Extended Capability register.		
10:4	RSVD	Reserved.	00h	RO
3:0	S	This field specifies the size of the interrupt remapping table. The number of entries in the interrupt remapping table is 2^(X+1), where X is the value programmed in this field.	0h	RW_L

5.5.28 IVA—Invalidate Address Register

Register to provide the DMA address whose corresponding IOTLB entry needs to be invalidated through the corresponding IOTLB Invalidate register. This register is a write-only register.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW	
Size: 64	Default Value: 0000000000000000h		Address Offset: 100h	
Bit Range	Acronym	Description	Default	Access
63:39	RSVD	Reserved.	0000000h	RO
38:12	ADDR	Software provides the DMA address that needs to be page-selectively invalidated. To make a page-selective invalidation request to hardware, software must first write the appropriate fields in this register, and then issue the appropriate page-selective invalidate command through the IOTLB_REG. Hardware ignores bits 63 : N, where N is the maximum guest address width (MGAW) supported.	0000000h	RW
11:7	RSVD	Reserved.	00h	RO
6	IH	The field provides hint to hardware about preserving or flushing the non-leaf (page-directory) entries that may be cached in hardware: 0: Software may have modified both leaf and non-leaf page-table entries corresponding to mappings specified in the ADDR and AM fields. On a page-selective invalidation request, hardware must flush both the cached leaf and non-leaf page-table entries corresponding tot he mappings specified by ADDR and AM fields. 1: Software has not modified any non-leaf page-table entries corresponding to mappings specified in the ADDR and AM fields. On a page-selective invalidation request, hardware may preserve the cached non-leaf page-table entries corresponding to mappings specified by ADDR and AM fields.	0h	RW
5:0	AM	The value in this field specifies the number of low order bits of the ADDR field that must be masked for the invalidation operation. This field enables software to request invalidation of contiguous mappings for size-aligned regions. For example: Mask Value ADDR bits masked Pages invalidated 0 None 1 1 12 2	00h	RW

continued...

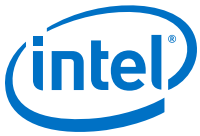

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW	
Size: 64	**Default Value: 0000000000000000h**		**Address Offset: 100h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		2 13:12 4 3 14:12 8 4 15:12 16 ... . When invalidating mappings for super-pages, software must specify the appropriate mask value. For example, when invalidating mapping for a 2MB page, software must specify an address mask value of at least 9. Hardware implementations report the maximum supported mask value through the Capability register.		

5.5.29 IOTLB—IOTLB Invalidate Register

Register to invalidate IOTLB. The act of writing the upper byte of the IOTLB_REG with IVT field Set causes the hardware to perform the IOTLB invalidation.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW; ROV; RW_V	
Size: 64	**Default Value: 0000000000000000h**		**Address Offset: 108h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63	IVT	Software requests IOTLB invalidation by setting this field. Software must also set the requested invalidation granularity by programming the IIRG field. Hardware clears the IVT field to indicate the invalidation request is complete. Hardware also indicates the granularity at which the invalidation operation was performed through the IAIG field. Software must not submit another invalidation request through this register while the IVT field is Set, nor update the associated Invalidate Address register. Software must not submit IOTLB invalidation requests when there is a context-cache invalidation request pending at this remapping hardware unit. Hardware implementations reporting write-buffer flushing requirement (RWBF=1 in Capability register) must implicitly perform a write buffer flushing before invalidating the IOTLB.	0h	RW_V
62	RSVD	Reserved.	0h	RO
61:60	IIRG	When requesting hardware to invalidate the IOTLB (by setting the IVT field), software writes the requested invalidation granularity through this field. The following are the encodings for the field. 00: Reserved. 01: Global invalidation request. 10: Domain-selective invalidation request. The target domain-id must be specified in the DID field. 11: Page-selective invalidation request. The target address, mask and invalidation hint must be specified in the Invalidate Address register, and the domain-id must be provided in the DID field. Hardware implementations may process an invalidation request by performing invalidation at a coarser granularity than requested. Hardware indicates completion of the invalidation request by clearing the IVT field. At this time, the granularity at which actual invalidation was performed is reported through the IAIG field	0h	RW

continued...

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: RW; ROV; RW_V	
Size: 64	**Default Value: 0000000000000000h**		**Address Offset: 108h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
59	RSVD	Reserved.	0h	RO
58:57	IAIG	Hardware reports the granularity at which an invalidation request was processed through this field when reporting invalidation completion (by clearing the IVT field). The following are the encodings for this field. 00: Reserved. This indicates hardware detected an incorrect invalidation request and ignored the request. Examples of incorrect invalidation requests include detecting an unsupported address mask value in Invalidate Address register for page-selective invalidation requests. 01: Global Invalidation performed. This could be in response to a global, domain-selective, or page-selective invalidation request. 10: Domain-selective invalidation performed using the domain-id specified by software in the DID field. This could be in response to a domain-selective or a page-selective invalidation request. 11: Domain-page-selective invalidation performed using the address, mask and hint specified by software in the Invalidate Address register and domain-id specified in DID field. This can be in response to a page-selective invalidation request.	0h	ROV
56:50	RSVD	Reserved.	00h	RO
49	DR	This field is ignored by hardware if the DRD field is reported as clear in the Capability register. When the DRD field is reported as Set in the Capability register, the following encodings are supported for this field: 0: Hardware may complete the IOTLB invalidation without draining any translated DMA read requests. 1: Hardware must drain DMA read requests.	0h	RW
48	DW	This field is ignored by hardware if the DWD field is reported as Clear in the Capability register. When the DWD field is reported as Set in the Capability register, the following encodings are supported for this field: 0: Hardware may complete the IOTLB invalidation without draining DMA write requests. 1: Hardware must drain relevant translated DMA write requests.	0h	RW
47:40	RSVD	Reserved.	00h	RO
39:32	DID	Indicates the ID of the domain whose IOTLB entries need to be selectively invalidated. This field must be programmed by software for domain-selective and page-selective invalidation requests. The Capability register reports the domain-id width supported by hardware. Software must ensure that the value written to this field is within this limit. Hardware ignores and not implements bits 47:(32+N), where N is the supported domain-id width reported in the Capability register.	00h	RW
31:0	RSVD	Reserved.	00000000h	RO

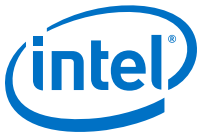

5.5.30 FRCDL—Fault Recording Low Register

Register to record fault information when primary fault logging is active. Hardware reports the number and location of fault recording registers through the Capability register. This register is relevant only for primary fault logging.
This register is sticky and can be cleared only through power good reset or by software clearing the RW1C fields by writing a 1.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: ROSV	
Size: 64	Default Value: 0000000000000000h		Address Offset: 200h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:12	FI	When the Fault Reason (FR) field indicates one of the DMA-remapping fault conditions, bits 63:12 of this field contain the page address in the faulted DMA request. Hardware treats bits 63:N as reserved (0), where N is the maximum guest address width (MGAW) supported. When the Fault Reason (FR) field indicates one of the interrupt-remapping fault conditions, bits 63:48 of this field indicate the interrupt_index computed for the faulted interrupt request, and bits 47:12 are cleared. This field is relevant only when the F field is Set.	0000000000000h	ROSV
11:0	RSVD	Reserved.	000h	RO

5.5.31 FRCDH—Fault Recording High Register

Register to record fault information when primary fault logging is active. Hardware reports the number and location of fault recording registers through the Capability register. This register is relevant only for primary fault logging.
This register is sticky and can be cleared only through power good reset or by software clearing the RW1C fields by writing a 1.

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: ROSV; RO; RW1CS	
Size: 64	Default Value: 0000000000000000h		Address Offset: 208h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63	F	Hardware sets this field to indicate a fault is logged in this Fault Recording register. The F field is set by hardware after the details of the fault is recorded in other fields. When this field is Set, hardware may collapse additional faults from the same source-id (SID). Software writes the value read from this field to Clear it.	0h	RW1CS
62	T	Type of the faulted request: 0: Write request 1: Read request or AtomicOp request This field is relevant only when the F field is Set, and when the fault reason (FR) indicates one of the DMA-remapping fault conditions.	0h	ROSV
61:60	AT	This field captures the AT field from the faulted DMA request. Hardware implementations not supporting Device-IOTLBs (DI field Clear in Extended Capability register) treat this field as RsvdZ. When supported, this field is valid only when the F field is Set, and when the fault reason (FR) indicates one of the DMA-remapping fault conditions.	0h	RO

continued...

B/D/F/Type: 0/0/0/MEM/VTDPC0BAR			Access: ROSV; RO; RW1CS	
Size: 64	Default Value: 0000000000000000h		Address Offset: 208h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
59:40	PN	PASID value in the faulted request. This field is relevant only when the PP field is set. Hardware implementations not supporting PASID (PASID field Clear in Extended Capability register) implement this field as RsvdZ.	00000h	RO
39:32	FR	Reason for the fault. This field is relevant only when the F field is set.	00h	ROSV
31	PP	When set, indicates the faulted request has a PASID tag. The value of the PASID field is reported in the PASID Value (PV) field. This field is relevant only when the F field is Set, and when the fault reason (FR) indicates one of the non-recoverable address translation fault conditions. Hardware implementations not supporting PASID (PASID field Clear in Extended Capability register) implement this field as RsvdZ.	0h	RO
30	EXE	When set, indicates Execute permission was requested by the faulted read request. This field is relevant only when the PP field and T field are both Set. Hardware implementations not supporting PASID (PASID field Clear in Extended Capability register) implement this field as RsvdZ.	0h	RO
29	PRIV	When set, indicates Supervisor privilege was requested by the faulted request. This field is relevant only when the PP field is Set. Hardware implementations not supporting PASID (PASID field Clear in Extended Capability register) implement this field as RsvdZ.	0h	RO
28:16	RSVD	Reserved.	0000h	RO
15:0	SID	Requester-id associated with the fault condition. This field is relevant only when the F field is set.	0000h	ROSV

5.6 MEM GTTMMADR Registers Summary

Offset	Register ID—Description	Default Value	Access
108000	MTOLUD—Top of Low Usable DRAM on page 359	00100000h	RO_V
108080	MTOUUD—Top of Upper Usable DRAM on page 359	0000000000000000h	RO_V
1080C0	MBDSM—Base Data of Stolen Memory on page 360	00000000h	RO_V
108100	MBGSM—Base of GTT stolen Memory on page 361	00100000h	RO_V
108180	MPMEN—Protected Memory Enable Register on page 361	00000000h	RO_V
1081C0	MPLMBASE—Protected Low-Memory Base Register on page 362	00000000h	RO_V
108200	MPLMLIMIT—Protected Low-Memory Limit Register on page 362	00000000h	RO_V
108240	MPHMBASE—Protected High-Memory Base Register on page 363	0000000000000000h	RO_V
108280	MPHMLIMIT—Protected High-Memory Limit Register on page 363	0000000000000000h	RO_V
1082C0	MPAVPC—Protected Audio Video Path Control on page 364	00000000h	RO_V
108300	MGCMD—Global Command Register on page 364	00000000h	RO_V; WO; RO
138158	PRIMARY—PRIMARY_PLANE_TURBO_POWER_POLICY on page 367	00000000h	RW
13815C	SECONDARY—SECONDARY_PLANE_TURBO_POWER_POLICY on page 368	00000010h	RW

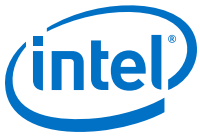

5.6.1 MTOLUD—Top of Low Usable DRAM

This 32 bit register defines the Top of Low Usable DRAM. TSEG, GTT Graphics memory and Graphics Stolen Memory are within the DRAM space defined. From the top, the Host optionally claims 1 to 64MBs of DRAM for internal graphics if enabled, 1or 2MB of DRAM for GTT Graphics Stolen Memory (if enabled) and 1, 2, or 8 MB of DRAM for TSEG if enabled.

Programming Example:
C1DRB3 is set to 4GB
TSEG is enabled and TSEG size is set to 1MB
Internal Graphics is enabled, and Graphics Mode Select is set to 32MB
GTT Graphics Stolen Memory Size set to 2MB
BIOS knows the OS requires 1G of PCI space.
BIOS also knows the range from 0_FEC0_0000h to 0_FFFF_FFFFh is not usable by the system. This 20MB range at the very top of addressable memory space is lost to APIC and Intel TXT.

According to the above equation, TOLUD is originally calculated to: 4GB = 1_0000_0000h

The system memory requirements are: 4GB (max addressable space) - 1GB (pci space) - 35MB (lost memory) = 3GB - 35MB (minimum granularity) = 0_ECB0_0000h

Since 0_ECB0_0000h (PCI and other system requirements) is less than 1_0000_0000h, TOLUD should be programmed to ECBh.
These bits are Intel TXT lockable.

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RO_V	
Size: 32	Default Value: 00100000h		Address Offset: 108000h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:20	TOLUD	This register contains bits 31 to 20 of an address one byte above the maximum DRAM memory below 4G that is usable by the operating system. Address bits 31 down to 20 programmed to 01h implies a minimum memory size of 1MB. Configuration software must set this value to the smaller of the following 2 choices: maximum amount memory in the system minus ME stolen memory plus one byte or the minimum address allocated for PCI memory. Address bits 19:0 are assumed to be 0_0000h for the purposes of address comparison. The Host interface positively decodes an address towards DRAM if the incoming address is less than the value programmed in this register. The Top of Low Usable DRAM is the lowest address above both Graphics Stolen memory and TSEG. BIOS determines the base of Graphics Stolen Memory by subtracting the Graphics Stolen Memory Size from TOLUD and further decrements by TSEG size to determine base of TSEG. All the Bits in this register are locked in Intel TXT mode. This register must be 1MB aligned when reclaim is enabled.	001h	RO_V
19:1	RSVD	Reserved.	00000h	RO
0	LOCK	This bit will lock all writeable settings in this register, including itself.	0h	RO_V

5.6.2 MTOUUD—Top of Upper Usable DRAM

This 64 bit register defines the Top of Upper Usable DRAM.
Configuration software must set this value to TOM minus all ME stolen memory if reclaim is disabled. If reclaim is enabled, this value must be set to reclaim limit +

1byte, 1MB aligned, since reclaim limit is 1MB aligned. Address bits 19:0 are assumed to be 000_0000h for the purposes of address comparison. The Host interface positively decodes an address towards DRAM if the incoming address is less than the value programmed in this register and greater than or equal to 4GB.
BIOS Restriction: Minimum value for TOUUD is 4GB.
These bits are Intel TXT lockable.

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RO_V	
Size: 64	Default Value: 0000000000000000h		Address Offset: 108080h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:20	TOUUD	This register contains bits 38 to 20 of an address one byte above the maximum DRAM memory above 4G that is usable by the operating system. Configuration software must set this value to TOM minus all ME stolen memory if reclaim is disabled. If reclaim is enabled, this value must be set to reclaim limit 1MB aligned since reclaim limit + 1byte is 1MB aligned. Address bits 19:0 are assumed to be 000_0000h for the purposes of address comparison. The Host interface positively decodes an address towards DRAM if the incoming address is less than the value programmed in this register and greater than 4GB. All the bits in this register are locked in Intel TXT mode.	00000h	RO_V
19:1	RSVD	Reserved.	00000h	RO
0	LOCK	This bit will lock all writeable settings in this register, including itself.	0h	RO_V

5.6.3 MBDSM—Base Data of Stolen Memory

This register contains the base address of graphics data stolen DRAM memory. BIOS determines the base of graphics data stolen memory by subtracting the graphics data stolen memory size (PCI Device 0 offset 52 bits 7:4) from TOLUD (PCI Device 0 offset BC bits 31:20).

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RO_V	
Size: 32	Default Value: 00000000h		Address Offset: 1080C0h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:20	BDSM	This register contains bits 31 to 20 of the base address of stolen DRAM memory. BIOS determines the base of graphics stolen memory by subtracting the graphics stolen memory size (PCI Device 0 offset 50 bits 15:8) from TOLUD (PCI Device 0 offset BC bits 31:20).	000h	RO_V
19:1	RSVD	Reserved.	00000h	RO
0	LOCK	This bit will lock all writeable settings in this register, including itself.	0h	RO_V

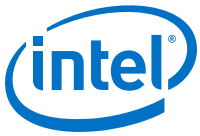

5.6.4 MBGSM—Base of GTT stolen Memory

This register contains the base address of stolen DRAM memory for the GTT. BIOS determines the base of GTT stolen memory by subtracting the GTT graphics stolen memory size (PCI Device 0 offset 52 bits 9:8) from the Graphics Base of Data Stolen Memory (PCI Device 0 offset B0 bits 31:20).

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RO_V	
Size: 32	**Default Value: 00100000h**		**Address Offset: 108100h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:20	BGSM	This register contains the base address of stolen DRAM memory for the GTT. BIOS determines the base of GTT stolen memory by subtracting the GTT graphics stolen memory size (PCI Device 0 offset 50 bits 7:6) from the Graphics Base of Data Stolen Memory (PCI Device 0 offset B0 bits 31:20).	001h	RO_V
19:1	RSVD	Reserved.	00000h	RO
0	LOCK	This bit will lock all writeable settings in this register, including itself.	0h	RO_V

5.6.5 MPMEN—Protected Memory Enable Register

Register to enable the DMA-protected memory regions setup through the PLMBASE, PLMLIMT, PHMBASE, PHMLIMIT registers. This register is always treated as RO for implementations not supporting protected memory regions (PLMR and PHMR fields reported as Clear in the Capability register).
Protected memory regions may be used by software to securely initialize remapping structures in memory. To avoid impact to legacy BIOS usage of memory, software is recommended to not overlap protected memory regions with any reserved memory regions of the platform reported through the Reserved Memory Region Reporting (RMRR) structures.

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RO_V	
Size: 32	**Default Value: 00000000h**		**Address Offset: 108180h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31	EPM	This field controls DMA accesses to the protected low-memory and protected high-memory regions. 0: Protected memory regions are disabled. 1: Protected memory regions are enabled. DMA requests accessing protected memory regions are handled as follows: - When DMA remapping is not enabled, all DMA requests accessing protected memory regions are blocked. - When DMA remapping is enabled: - DMA requests processed as pass-through (Translation Type value of 10b in Context-Entry) and accessing the protected memory regions are blocked. - DMA requests with translated address (AT=10b) and accessing the protected memory regions are blocked. - DMA requests that are subject to address remapping, and accessing the protected memory regions may or may not be blocked by hardware. For such requests, software must not depend on hardware protection of the protected memory regions, and instead program the DMA-remapping page-tables to not allow DMA to protected memory regions.	0h	RO_V

continued...

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RO_V	
Size: 32	Default Value: 00000000h		Address Offset: 108180h	
Bit Range	Acronym	Description	Default	Access
		Remapping hardware access to the remapping structures are not subject to protected memory region checks. DMA requests blocked due to protected memory region violation are not recorded or reported as remapping faults. Hardware reports the status of the protected memory enable/disable operation through the PRS field in this register. Hardware implementations supporting DMA draining must drain any in-flight translated DMA requests queued within the Root-Complex before indicating the protected memory region as enabled through the PRS field.		
30:1	RSVD	Reserved.	00000000h	RO
0	PRS	This field indicates the status of protected memory region(s): 0: Protected memory region(s) disabled. 1: Protected memory region(s) enabled.	0h	RO_V

5.6.6 MPLMBASE—Protected Low-Memory Base Register

Register to set up the base address of DMA-protected low-memory region below 4GB. This register must be set up before enabling protected memory through PMEN_REG, and must not be updated when protected memory regions are enabled.
This register is always treated as RO for implementations not supporting protected low memory region (PLMR field reported as Clear in the Capability register).
The alignment of the protected low memory region base depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1s to this register, and finding the most significant zero bit position with 0 in the value read back from the register. Bits N:0 of this register is decoded by hardware as all 0s.
Software must setup the protected low memory region below 4GB.
Software must not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RO_V	
Size: 32	Default Value: 00000000h		Address Offset: 1081C0h	
Bit Range	Acronym	Description	Default	Access
31:20	PLMB	This register specifies the base of protected low-memory region in system memory.	000h	RO_V
19:0	RSVD	Reserved.	00000h	RO

5.6.7 MPLMLIMIT—Protected Low-Memory Limit Register

Register to set up the limit address of DMA-protected low-memory region below 4GB. This register must be set up before enabling protected memory through PMEN_REG, and must not be updated when protected memory regions are enabled.
This register is always treated as RO for implementations not supporting protected low memory region (PLMR field reported as Clear in the Capability register).
The alignment of the protected low memory region limit depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1's to this register, and finding most significant zero bit position with 0 in the value read back from the register. Bits N:0 of the limit register is decoded by hardware as all 1s.
The Protected low-memory base and limit registers functions as follows:

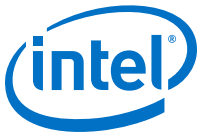

- Programming the protected low-memory base and limit registers with the same value
  in bits 31:(N+1) specifies a protected low-memory region of size 2^(N+1) bytes.
- Programming the protected low-memory limit register with a value less than the protected low-memory base register disables the protected low-memory region.

Software must not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RO_V	
Size: 32	Default Value: 00000000h		Address Offset: 108200h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:20	PLML	This register specifies the last host physical address of the DMA-protected low-memory region in system memory.	000h	RO_V
19:0	RSVD	Reserved.	00000h	RO

5.6.8 MPHMBASE—Protected High-Memory Base Register

Register to set up the base address of DMA-protected high-memory region. This register must be set up before enabling protected memory through PMEN_REG, and must not be updated when protected memory regions are enabled.
This register is always treated as RO for implementations not supporting protected high memory region (PHMR field reported as Clear in the Capability register).
The alignment of the protected high memory region base depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1's to this register, and finding most significant zero bit position below host address width (HAW) in the value read back from the register. Bits N:0 of this register are decoded by hardware as all 0s.
Software may setup the protected high memory region either above or below 4GB.
Software must not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RO_V	
Size: 64	Default Value: 0000000000000000h		Address Offset: 108240h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
63:39	RSVD	Reserved.	0000000h	RO
38:20	PHMB	This register specifies the base of protected (high) memory region in system memory. Hardware ignores, and does not implement, bits 63:HAW, where HAW is the host address width.	00000h	RO_V
19:0	RSVD	Reserved.	00000h	RO

5.6.9 MPHMLIMIT—Protected High-Memory Limit Register

Register to set up the limit address of DMA-protected high-memory region. This register must be set up before enabling protected memory through PMEN_REG, and must not be updated when protected memory regions are enabled.
This register is always treated as RO for implementations not supporting protected high memory region (PHMR field reported as Clear in the Capability register).
The alignment of the protected high memory region limit depends on the number of reserved bits (N:0) of this register. Software may determine the value of N by writing all 1's to this register, and finding most significant zero bit position below host

address width (HAW) in the value read back from the register. Bits N:0 of the limit register is decoded by hardware as all 1s.
The protected high-memory base & limit registers functions as follows.
- Programming the protected low-memory base and limit registers with the same value
  in bits HAW:(N+1) specifies a protected low-memory region of size 2^(N+1) bytes.
- Programming the protected high-memory limit register with a value less than the protected high-memory base register disables the protected high-memory region.
Software must not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RO_V	
Size: 64	Default Value: 0000000000000000h		Address Offset: 108280h	
Bit Range	Acronym	Description	Default	Access
63:39	RSVD	Reserved.	0000000h	RO
38:20	PHML	This register specifies the last host physical address of the DMA-protected high-memory region in system memory. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address width.	00000h	RO_V
19:0	RSVD	Reserved.	00000h	RO

5.6.10 MPAVPC—Protected Audio Video Path Control

All the bits in this register are locked by Intel TXT. When locked the RW bits are RO.

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RO_V	
Size: 32	Default Value: 00000000h		Address Offset: 1082C0h	
Bit Range	Acronym	Description	Default	Access
31:3	RSVD	Reserved.	00000000h	RO
2	PAVPLCK	This bit locks all writeable contents in this register when set (including itself). Only a hardware reset can unlock the register again. This lock bit needs to be set only if PAVP is enabled (bit 1 of this register is asserted).	0h	RO_V
1:0	RSVD	Reserved.	0h	RO

5.6.11 MGCMD—Global Command Register

Register to control remapping hardware. If multiple control fields in this register need to be modified, software must serialize the modifications through multiple writes to this register.

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RO_V; WO; RO	
Size: 32	Default Value: 00000000h		Address Offset: 108300h	
Bit Range	Acronym	Description	Default	Access
31	TE	Software writes to this field to request hardware to enable/ disable DMA-remapping: 0: Disable DMA remapping 1: Enable DMA remapping	0h	RO_V

continued...

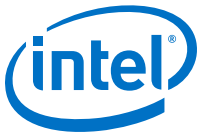

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RO_V; WO; RO	
Size: 32	**Default Value: 00000000h**		**Address Offset: 108300h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		Hardware reports the status of the translation enable operation through the TES field in the Global Status register. There may be active DMA requests in the platform when software updates this field. Hardware must enable or disable remapping logic only at determinstic transaction boundaries, so that any in-flight transaction is either subject to remapping or not at all. Hardware implementations supporting DMA draining must drain any in-flight DMA read/write requests queued within the Root-Complex before completing the translation enable command and reflecting the status of the command through the TES field in the Global Status register. The value returned on a read of this field is undefined.		
30	SRTP	Software sets this field to set/update the root-entry table pointer used by hardware. The root-entry table pointer is specified through the Root-entry Table Address (RTA_REG) register. Hardware reports the status of the "Set Root Table Pointer" operation through the RTPS field in the Global Status register. The "Set Root Table Pointer" operation must be performed before enabling or re-enabling (after disabling) DMA remapping through the TE field. After a "Set Root Table Pointer" operation, software must globally invalidate the context cache and then globally invalidate of IOTLB. This is required to ensure hardware uses only the remapping structures referenced by the new root table pointer, and not stale cached entries. While DMA remapping hardware is active, software may update the root table pointer through this field. However, to ensure valid in-flight DMA requests are deterministically remapped, software must ensure that the structures referenced by the new root table pointer are programmed to provide the same remapping results as the structures referenced by the previous root-table pointer. Clearing this bit has no effect. The value returned on read of this field is undefined.	0h	WO
29	SFL	This field is valid only for implementations supporting advanced fault logging. Software sets this field to request hardware to set/update the fault-log pointer used by hardware. The fault-log pointer is specified through Advanced Fault Log register. Hardware reports the status of the 'Set Fault Log' operation through the FLS field in the Global Status register. The fault log pointer must be set before enabling advanced fault logging (through EAFL field). Once advanced fault logging is enabled, the fault log pointer may be updated through this field while DMA remapping is active. Clearing this bit has no effect. The value returned on read of this field is undefined.	0h	RO
28	EAFL	This field is valid only for implementations supporting advanced fault logging. Software writes to this field to request hardware to enable or disable advanced fault logging: 0: Disable advanced fault logging. In this case, translation faults are reported through the Fault Recording registers.	0h	RO

continued...

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RO_V; WO; RO	
Size: 32	Default Value: 00000000h		Address Offset: 108300h	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		1: Enable use of memory-resident fault log. When enabled, translation faults are recorded in the memory-resident log. The fault log pointer must be set in hardware (through the SFL field) before enabling advaned fault logging. Hardware reports the status of the advaned fault logging enable operation through the AFLS field in the Global Status register. The value returned on read of this field is undefined.		
27	WBF	This bit is valid only for implementations requiring write buffer flushing. Software sets this field to request that hardware flush the Root-Complex internal write buffers. This is done to ensure any updates to the memory-resident remapping structures are not held in any internal write posting buffers. Hardware reports the status of the write buffer flushing operation through the WBFS field in the Global Status register. Cleraing this bit has no effect. The value returned on a read of this field is undefined.	0h	RO
26	QIE	This field is valid only for implementations supporting queued invalidations. Software writes to this field to enable or disable queued invalidations. 0: Disable queued invalidations. 1: Enable use of queued invalidations. Hardware reports the status of queued invalidation enable operation through QIES field in the Global Status register. The value returned on a read of this field is undefined.	0h	RO_V
25	IRE	This field is valid only for implementations supporting interrupt remapping. 0: Disable interrupt-remapping hardware 1: Enable interrupt-remapping hardware Hardware reports the status of the interrupt remapping enable operation through the IRES field in the Global Status register. There may be active interrupt requests in the platform when software updates this field. Hardware must enable or disable interrupt-remapping logic only at deterministic transaction boundaries, so that any in-flight interrupts are either subject to remapping or not at all. Hardware implementations must drain any in-flight interrupts requests queued in the Root-Complex before completing the interrupt-remapping enable command and reflecting the status of the command through the IRES field in the Global Status register. The value returned on a read of this field is undefined.	0h	RO_V
24	SIRTP	This field is valid only for implementations supporting interrupt-remapping. Software sets this field to set/update the interrupt remapping table pointer used by hardware. The interrupt remapping table pointer is specified through the Interrupt Remapping Table Address (IRTA_REG) register. Hardware reports the status of the 'Set Interrupt Remap Table Pointer' operation through the IRTPS field in the Global Status register. The 'Set Interrupt Remap Table Pointer' operation must be performed before enabling or re-enabling (after disabling)	0h	WO

continued...

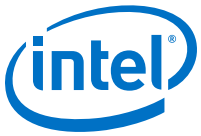

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RO_V; WO; RO	
Size: 32	**Default Value: 00000000h**		**Address Offset: 108300h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
		interrupt-remapping hardware through the IRE field. After a 'Set Interrupt Remap Table Pointer' operation, software must globally invalidate the interrupt entry cache. This is required to ensure hardware uses only the interrupt-remapping entries referenced by the new interrupt remap table pointer, and not any stale cached entries. While interrupt remapping is active, software may update the interrupt remapping table pointer through this field. However, to ensure valid in-flight interrupt requests are deterministically remapped, software must ensure that the structures referenced by the new interrupt remap table pointer are programmed to provide the same remapping results as the structures referenced by the previous interrupt remap table pointer. Clearing this bit has no effect. The value returned on a read of this field is undefined.		
23	CFI	This field is valid only for Intel®64 implementations supporting interrupt-remapping. Software writes to this field to enable or disable Compatibility Format interrupts on Intel®64 platforms. The value in this field is effective only when interrupt-remapping is enabled and Extended Interrupt Mode (x2APIC mode) is not enabled. 0: Block Compatibility format interrupts. 1: Process Compatibility format interrupts as pass-through (bypass interrupt remapping). Hardware reports the status of updating this field through the CFIS field in the Global Status register. The value returned on a read of this field is undefined.	0h	RO_V
22:0	RSVD	Reserved.	000000h	RO

5.6.12 PRIMARY—PRIMARY_PLANE_TURBO_POWER_POLICY

The PRIMARY_PLANE_TURBO_POWER_POLICY and SECONDARY_PLANE_TURBO_POWER_POLICY are used together to balance the power budget betwen the two power planes.
The power plane with the higher policy will get a higher priority. The default values for these registers give a higher priority to the secondary power plane.

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RW	
Size: 32	**Default Value: 00000000h**		**Address Offset: 138158h**	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:5	RSVD	Reserved.	0000000h	RO
4:0	PRIPTP	Priority Level. A higher number implies a higher priority.	00h	RW

5.6.13 SECONDARY—SECONDARY_PLANE_TURBO_POWER_POLICY

The PRIMARY_PLANE_TURBO_POWER_POLICY and SECONDARY_PLANE_TURBO_POWER_POLICY are used together to balance the power budget betwen the two power planes.
The power plane with the higher policy will get a higher priority. The default values for these registers give a higher priority to the secondary power plane.

B/D/F/Type: 0/2/0/MEM/GTTMMADR			Access: RW	
Size: 32	Default Value: 00000010h		Address Offset: 13815Ch	
Bit Range	**Acronym**	**Description**	**Default**	**Access**
31:5	RSVD	Reserved.	0000000h	RO
4:0	SECPTP	Priority Level. A higher number implies a higher priority.	10h	RW