

# Intel® Core™ i7-600, i5-500, i5-400 and i3-300 Mobile Processor Series

Datasheet — Volume Two

---

*This is volume 2 of 2. Refer to Document Number 322812 for Volume 1  
November 2010*



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

64-bit computing on Intel architecture requires a computer system with a processor, chipset, BIOS, operating system, device drivers and applications enabled for Intel® 64 architecture. Processors will not operate (including 32-bit operation) without an Intel® 64 architecture-enabled BIOS. Performance will vary depending on your hardware and software configurations. Consult with your system vendor for more information

Intel® Core™ i7-600, i5-500, i5-400 and i3-300 Mobile Processor Series 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.

Δ Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See [http://www.intel.com/products/processor\\_number](http://www.intel.com/products/processor_number) for details.

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

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

Copyright © 2010, Intel Corporation. All rights reserved.



# Contents

---

<b>1</b>	<b>Processor Configuration Registers</b> .....	12
1.1	Register Terminology .....	12
1.2	System Address Map.....	14
1.2.1	Legacy Address Range .....	17
1.2.2	Main Memory Address Range (1 MB - TOLUD) .....	19
1.2.3	Main Memory Address Space (4 GB to TOUUD).....	27
1.2.4	PCI Express* Configuration Address Space .....	34
1.2.5	PCI Express Graphics Attach (PEG) .....	34
1.2.6	Graphics Memory Address Ranges .....	35
1.2.7	System Management Mode (SMM) .....	36
1.2.8	SMM and VGA Access through GTT TLB .....	37
1.2.9	I/O Address Space.....	37
1.3	Configuration Process and Registers.....	39
1.3.1	Platform Configuration Structure .....	39
1.4	Configuration Mechanisms .....	40
1.4.1	Standard PCI Configuration Mechanism .....	40
1.4.2	PCI Express Enhanced Configuration Mechanism.....	40
1.5	Routing Configuration Accesses .....	42
1.5.1	Internal Device Configuration Accesses.....	43
1.5.2	Bridge Related Configuration Accesses .....	44
1.6	Processor Register Introduction .....	45
1.7	I/O Mapped Registers .....	46
1.8	PCI Device 0.....	46
1.8.1	VID - Vendor Identification .....	48
1.8.2	DID - Device Identification.....	48
1.8.3	PCICMD - PCI Command .....	49
1.8.4	PCISTS - PCI Status .....	51
1.8.5	RID - Revision Identification .....	52
1.8.6	CC - Class Code .....	53
1.8.7	MLT - Master Latency Timer .....	54
1.8.8	HDR - Header Type.....	54
1.8.9	SVID - Subsystem Vendor Identification .....	54
1.8.10	SID - Subsystem Identification.....	55
1.8.11	CAPPTR - Capabilities Pointer .....	55
1.8.12	PXPEPBAR - PCI Express Egress Port Base Address .....	56
1.8.13	MCHBAR - Processor Memory Mapped Register Range Base.....	57
1.8.14	GGC - Processor Graphics Control Register.....	58
1.8.15	DEVEN - Device Enable .....	62
1.8.16	DMIBAR - Root Complex Register Range Base Address .....	63
1.8.17	LAC - Legacy Access Control .....	64
1.8.18	REMAPBASE - Remap Base Address Register .....	66
1.8.19	REMAPLIMIT - Remap Limit Address Register .....	66
1.8.20	TOM - Top of Memory .....	67
1.8.21	TOUUD - Top of Upper Usable DRAM .....	68
1.8.22	GBSM - Graphics Base of Stolen Memory .....	69
1.8.23	BGSM - Base of GTT Stolen Memory.....	70
1.8.24	TSEGMB - TSEG Memory Base.....	71
1.8.25	TOLUD - Top of Low Usable DRAM .....	71



1.8.26	PBFC - Primary Buffer Flush Control	73
1.8.27	SBFC - Secondary Buffer Flush Control	73
1.8.28	ERRSTS - Error Status	74
1.8.29	ERRCMD - Error Command	75
1.8.30	SMICMD - SMI Command	76
1.8.31	SCICMD - SCI Command	77
1.8.32	SKPD - Scratchpad Data	77
1.8.33	CAPID0 - Capability Identifier	78
1.9	Device 0 MCHBAR DRAM Controls	78
1.9.1	CSZMAP - Channel Size Mapping	80
1.9.2	CHDECMISC - Channel Decode Misc	81
1.9.3	CODRB0 - Channel 0 DRAM Rank Boundary Address 0	83
1.9.4	CODRB1 - Channel 0 DRAM Rank Boundary Address 1	84
1.9.5	CODRB2 - Channel 0 DRAM Rank Boundary Address 2	84
1.9.6	CODRB3 - Channel 0 DRAM Rank Boundary Address 3	85
1.9.7	CODRA01 - Channel 0 DRAM Rank 0,1 Attribute	85
1.9.8	CODRA23 - Channel 0 DRAM Rank 2,3 Attribute	87
1.9.9	COCYTRKPCHG - Channel 0 CYCTRK PCHG	88
1.9.10	COCYTRKACT - Channel 0 CYCTRK ACT	89
1.9.11	COCYTRKWR - Channel 0 CYCTRK WR	90
1.9.12	COCYTRKRD - Channel 0 CYCTRK READ	91
1.9.13	COCYTRKREFR - Channel 0 CYCTRK REFR	92
1.9.14	COREFRCTRL - Channel 0 DRAM Refresh Control	92
1.9.15	COCKECTRL - Channel 0 CKE Control	95
1.9.16	COODTCTRL - Channel 0 ODT Control	96
1.9.17	CODTC - Channel 0 DRAM Throttling Control	97
1.9.18	CODTPEW - Channel 0 DRAM Rank Throttling Passive Event	98
1.9.19	CODTAEW - Channel 0 DRAM Rank Throttling Active Event	99
1.9.20	C1DRB0 - Channel 1 DRAM Rank Boundary Address 0	99
1.9.21	C1DRB1 - Channel 1 DRAM Rank Boundary Address 1	100
1.9.22	C1DRB2 - Channel 1 DRAM Rank Boundary Address 2	100
1.9.23	C1DRB3 - Channel 1 DRAM Rank Boundary Address 3	101
1.9.24	C1DRA01 - Channel 1 DRAM Rank 0,1 Attributes	101
1.9.25	C1DRA23 - Channel 1 DRAM Rank 2,3 Attributes	102
1.9.26	C1CYTRKPCHG - Channel 1 CYCTRK PCHG	102
1.9.27	C1CYTRKACT - Channel 1 CYCTRK ACT	103
1.9.28	C1CYTRKWR - Channel 1 CYCTRK WR	104
1.9.29	C1CYTRKRD - Channel 1 CYCTRK READ	105
1.9.30	C1CKECTRL - Channel 1 CKE Control	106
1.9.31	C1REFRCTRL - Channel 1 DRAM Refresh Control	107
1.9.32	C1ODTCTRL - Channel 1 ODT Control	110
1.9.33	C1DTC - Channel 1 DRAM Throttling Control	111
1.9.34	C1DTPEW - Channel 1 DRAM Rank Throttling Passive Event	112
1.9.35	C1DTAEW - Channel 1 DRAM Rank Throttling Active Event	113
1.9.36	DDRMPLL1 - DDR PLL BIOS	114
1.10	Device 0 MCHBAR Thermal Management Controls	114
1.10.1	TSC1 - Thermal Sensor Control 1	116
1.10.2	TSS1 - Thermal Sensor Status 1	117
1.10.3	TR1 - Thermometer Read 1	118
1.10.4	TOF1 - Thermometer Offset 1	119
1.10.5	RTR1 - Relative Thermometer Read 1	119
1.10.6	TSTTPA1 - Thermal Sensor Temperature Trip Point A1	120
1.10.7	TSTTPB1 - Thermal Sensor Temperature Trip Point B1	121



1.10.8	HWTHROTCTRL1 - Hardware Throttle Control 1	122
1.10.9	TIS1 - Thermal Interrupt Status 1	123
1.10.10	TERATE - Thermometer Mode Enable and Rate	125
1.10.11	TERRCMD - Thermal Error Command	127
1.10.12	TSMICMD - Thermal SMI Command	128
1.10.13	TSCICMD - Thermal SCI Command	129
1.10.14	TINTRCMD - Thermal INTR Command	130
1.10.15	EXTTSCS - External Thermal Sensor Control and Status	131
1.10.16	MCHTSWDT - Memory Controller Thermal Sensor Watch Dog Timer	133
1.10.17	MEMTDPCTW - Memory TDP Controller Registers	133
1.10.18	MTDPCCRWTWHOTTH - Memory TDP Controller Combined RD/WR Thermal Weight Hot Thresholds	134
1.10.19	MTDPCCRWTWHOTTH2 - Memory TDP Controller Combined RD/WR Thermal Weight Hot Thresholds 2	136
1.10.20	MTDPCCRWTWHOTTH3 - Memory TDP Controller Combined RD/WR Thermal Weight Hot Thresholds 3	137
1.10.21	MTDPCCRWTWHOTTH4 - Memory TDP Controller Combined RD/WR Thermal Weight Hot Thresholds 4	138
1.10.22	MTDPCHOTTHINT - Memory TDP Controller Hot Throttled Intervals	139
1.10.23	MTDPCHOTTHINT2 - Memory TDP Controller Hot Throttled Intervals 2	140
1.10.24	MTDPCTLAUXTNTINT - Memory TDP Controller Aux and Throttle-NonThrottle Intervals	141
1.10.25	MTDPCMISC - Memory TDP Controller Miscellaneous Control	141
1.10.26	TSFUSE - Thermal Sensor Fuses	142
1.11	MCHBAR Render Thermal Throttling Controls	143
1.11.1	THERMSTCTL - Render Thermal State Control	143
1.11.2	RSTDBYCTL - Render Standby State Control	144
1.11.3	VIDCTL - VID Control	145
1.11.4	WDTMRTS - Watchdog Timer For Thermal Sensor Trip	145
1.11.5	WDTSTPSZ - Watchdog Timer Based Px Step Size	145
1.12	Device 0 MCHBAR ACPI Power Management Controls	146
1.12.1	C3C6ET - C3/C6 EntryTimers	146
1.12.2	SLFRCS - Self-Refresh Channel Status	147
1.12.3	DSLFRCS - PM Memory Subsystem	147
1.12.4	PMCFG - Power Management Configuration	148
1.13	PCI Device1	148
1.13.1	VID1 - Vendor Identification	151
1.13.2	DID1 - Device Identification	151
1.13.3	PCICMD1 - PCI Command	152
1.13.4	PCISTS1 - PCI Status	154
1.13.5	RID1 - Revision Identification	156
1.13.6	CC1 - Class Code	157
1.13.7	CL1 - Cache Line Size	157
1.13.8	HDR1 - Header Type	158
1.13.9	PBUSN1 - Primary Bus Number	158
1.13.10	SBUSN1 - Secondary Bus Number	159
1.13.11	SUBUSN1 - Subordinate Bus Number	159
1.13.12	IOBASE1 - I/O Base Address	160
1.13.13	IOLIMIT1 - I/O Limit Address	160
1.13.14	SSTS1 - Secondary Status	161
1.13.15	MBASE1 - Memory Base Address	162
1.13.16	MLIMIT1 - Memory Limit Address	162
1.13.17	PMBASE1 - Prefetchable Memory Base Address	163



1.13.18	PMLIMIT1 - Prefetchable Memory Limit Address	164
1.13.19	PMBASEU1 - Prefetchable Memory Base Address Upper	165
1.13.20	PMLIMITU1 - Prefetchable Memory Limit Address Upper	165
1.13.21	CAPPTR1 - Capabilities Pointer	166
1.13.22	INTRLINE1 - Interrupt Line	167
1.13.23	INTRPIN1 - Interrupt Pin	167
1.13.24	BCTRL1 - Bridge Control	168
1.13.25	PM_CAPID1 - Power Management Capabilities	170
1.13.26	PM_CS1 - Power Management Control/Status	171
1.13.27	SS_CAPID - Subsystem ID and Vendor ID Capabilities	173
1.13.28	SS - Subsystem ID and Subsystem Vendor ID	173
1.13.29	MSI_CAPID - Message Signaled Interrupts Capability ID	174
1.13.30	MC - Message Control	174
1.13.31	MA - Message Address	175
1.13.32	MD - Message Data	176
1.13.33	PEG_CAPL - PCI Express-G Capability List	176
1.13.34	PEG_CAP - PCI Express-G Capabilities	177
1.13.35	DCAP - Device Capabilities	178
1.13.36	DCTL - Device Control	179
1.13.37	DSTS - Device Status	180
1.13.38	LCAP - Link Capabilities	181
1.13.39	LCTL - Link Control	185
1.13.40	LSTS - Link Status	187
1.13.41	SLOTCAP - Slot Capabilities	189
1.13.42	SLOTCTL - Slot Control	191
1.13.43	SLOTSTS - Slot Status	193
1.13.44	RCTL - Root Control	196
1.13.45	RSTS - Root Status	197
1.13.46	LCTL2 - Link Control 2	198
1.13.47	LSTS2 - Link Status 2	200
1.13.48	PEGLC - PCI Express-G Legacy Control	201
1.14	PCI Device1 - Extended Configuration	202
1.14.1	VCECH - Virtual Channel Enhanced Capability Header	202
1.14.2	PVCCAP1 - Port VC Capability Register 1	203
1.14.3	PVCCAP2 - Port VC Capability Register 2	204
1.14.4	PVCCCTL - Port VC Control	204
1.14.5	VCORCAP - VC0 Resource Capability	205
1.14.6	VCORCTL - VC0 Resource Control	206
1.14.7	VCORSTS - VC0 Resource Status	207
1.14.8	PEGSSTS - PCI Express-G Sequence Status	208
1.14.9	PEGTXDEMPSEL - PEG Transmit De-Emphasis Select Register	208
1.15	DMIBAR	211
1.15.1	DMIVCECH - DMI Virtual Channel Enhanced Capability	212
1.15.2	DMIPVCCAP1 - DMI Port VC Capability Register 1	213
1.15.3	DMIPVCCAP2 - DMI Port VC Capability Register 2	213
1.15.4	DMIPVCCTL - DMI Port VC Control	214
1.15.5	DMIVCORCAP - DMI VC0 Resource Capability	214
1.15.6	DMIVCORCTL0 - DMI VC0 Resource Control	215
1.15.7	DMIVCORSTS - DMI VC0 Resource Status	216
1.15.8	DMIVC1RCAP - DMI VC1 Resource Capability	217
1.15.9	DMIVC1RCTL1 - DMI VC1 Resource Control	218
1.15.10	DMIVC1RSTS - DMI VC1 Resource Status	219
1.15.11	DMILCAP - DMI Link Capabilities	220



1.15.12	DMILCTL - DMI Link Control	221
1.15.13	DMILSTS - DMI Link Status	222
1.16	PCI Device 2 Function 0	222
1.16.1	VID2 - Vendor Identification	224
1.16.2	DID2 - Device Identification	225
1.16.3	PCICMD2 - PCI Command	225
1.16.4	PCISTS2 - PCI Status	226
1.16.5	RID2 - Revision Identification	228
1.16.6	CC - Class Code	229
1.16.7	CLS - Cache Line Size	230
1.16.8	MLT2 - Master Latency Timer	230
1.16.9	HDR2 - Header Type	231
1.16.10	GTMMADR - Graphics Translation Table, Memory Mapped Range Address	231
1.16.11	GMADR - Graphics Memory Range Address	232
1.16.12	IOBAR - I/O Base Address	233
1.16.13	SVID2 - Subsystem Vendor Identification	234
1.16.14	SID2 - Subsystem Identification	234
1.16.15	ROMADR - Video BIOS ROM Base Address	235
1.16.16	CAPPOINT - Capabilities Pointer	235
1.16.17	INTRLINE - Interrupt Line	236
1.16.18	INTRPIN - Interrupt Pin	236
1.16.19	MINGNT - Minimum Grant	236
1.16.20	MAXLAT - Maximum Latency	237
1.16.21	GGCTL - Graphics Enhanced Intel® SpeedStep Technology Capability	237
1.16.22	MGGC - Processor Graphics Control Register	239
1.16.23	DEVEN - Device Enable	241
1.16.24	SSRW - Software Scratch Read Write	242
1.16.25	BSM - Base of Stolen Memory	242
1.16.26	HSRW - Hardware Scratch Read Write	242
1.16.27	MSAC - Multi Size Aperture Control	243
1.16.28	MC - Message Control	243
1.16.29	MA - Message Address	245
1.16.30	MD - Message Data	245
1.16.31	PMCAPIID - Power Management Capabilities ID	246
1.16.32	PMCAP - Power Management Capabilities	246
1.16.33	PMCS - Power Management Control/Status	247
1.16.34	SWSMI - Software SMI	248
1.16.35	GSE - Graphics System Event Register	249
1.16.36	ASLS - ASL Storage	250
1.17	Device 2 IO	250
1.17.1	Index - MMIO Address Register	251
1.17.2	Data - MMIO Data Register	251
1.18	DMI and PEG VC0/VCp Remap Registers	252
1.18.1	VER_REG - Version Register	254
1.18.2	CAP_REG - Capability Register	255
1.18.3	ECAP_REG - Extended Capability Register	258
1.18.4	GCMD_REG - Global Command Register	261
1.18.5	GSTS_REG - Global Status Register	265
1.18.6	RTADDR_REG - Root-Entry Table Address Register	267
1.18.7	CCMD_REG - Context Command Register	268
1.18.8	FSTS_REG - Fault Status Register	271
1.18.9	FECTL_REG - Fault Event Control Register	273
1.18.10	FEDATA_REG - Fault Event Data Register	275



1.18.11	FEADDR_REG - Fault Event Address Register	275
1.18.12	FEUADDR_REG - Fault Event Upper Address Register	276
1.18.13	AFLOG_REG - Advanced Fault Log Register	276
1.18.14	PMEM_REG - Protected Memory Enable Register	278
1.18.15	PLMBASE_REG - Protected Low-Memory Base Register	279
1.18.16	PLMLIMIT_REG - Protected Low-Memory Limit Register	280
1.18.17	PHMBASE_REG - Protected High-Memory Base Register	281
1.18.18	PHMLIMIT_REG - Protected High-Memory Limit Register	282
1.18.19	IQH_REG - Invalidation Queue Head Register	283
1.18.20	IQT_REG - Invalidation Queue Tail Register	283
1.18.21	IOA_REG - Invalidation Queue Address Register	284
1.18.22	ICS_REG - Invalidation Completion Status Register	284
1.18.23	IECTL_REG - Invalidation Event Control Register	285
1.18.24	IEDATA_REG - Invalidation Event Data Register	286
1.18.25	IEADDR_REG - Invalidation Event Address Register	286
1.18.26	IEUADDR_REG - Invalidation Event Upper Address Register	287
1.18.27	IRTA_REG - Interrupt Remapping Table Address Register	287
1.18.28	IVA_REG - Invalidate Address Register	288
1.18.29	IOTLB_REG - IOTLB Invalidate Register	290
1.18.30	FRCD_REG - Fault Recording Registers	293
1.18.31	VTCMPLRESR - VT-d Completion Resource Dedication	294
1.18.32	VTFTCHARBCTL - VC0/VCp Intel VT-d Fetch Arbiter Control	297
1.18.33	PEGVTCMPLRESR - PEG VT-d Completion Resource Dedication	298
1.19	DMI VC1 REMAP Registers	300
1.19.1	VER_REG - Version Register	302
1.19.2	CAP_REG - Capability Register	302
1.19.3	ECAP_REG - Extended Capability Register	306
1.19.4	GCMD_REG - Global Command Register	308
1.19.5	GSTS_REG - Global Status Register	312
1.19.6	RTADDR_REG - Root-Entry Table Address Register	314
1.19.7	CCMD_REG - Context Command Register	315
1.19.8	FSTS_REG - Fault Status Register	318
1.19.9	FECTL_REG - Fault Event Control Register	320
1.19.10	FEDATA_REG - Fault Event Data Register	322
1.19.11	FEADDR_REG - Fault Event Address Register	322
1.19.12	FEUADDR_REG - Fault Event Upper Address Register	323
1.19.13	AFLOG_REG - Advanced Fault Log Register	323
1.19.14	PMEN_REG - Protected Memory Enable Register	324
1.19.15	PLMBASE_REG - Protected Low-Memory Base Register	326
1.19.16	PLMLIMIT_REG - Protected Low-Memory Limit Register	327
1.19.17	PHMBASE_REG - Protected High-Memory Base Register	328
1.19.18	PHMLIMIT_REG - Protected High-Memory Limit Register	329
1.19.19	IQH_REG - Invalidation Queue Head Register	330
1.19.20	IQT_REG - Invalidation Queue Tail Register	330
1.19.21	IOA_REG - Invalidation Queue Address Register	331
1.19.22	ICS_REG - Invalidation Completion Status Register	332
1.19.23	IECTL_REG - Invalidation Event Control Register	332
1.19.24	IEDATA_REG - Invalidation Event Data Register	333
1.19.25	IEADDR_REG - Invalidation Event Address Register	334
1.19.26	IEUADDR_REG - Invalidation Event Upper Address Register	334
1.19.27	IRTA_REG - Interrupt Remapping Table Address Register	335
1.19.28	IVA_REG - Invalidate Address Register	336
1.19.29	IOTLB_REG - IOTLB Invalidate Register	338





1.19.30	FRCD_REG - Fault Recording Registers	341
1.19.31	VTPOLICY - DMA Remap Engine Policy Control	342
1.20	GFXVTBAR	344
1.20.1	VER_REG - Version Register	346
1.20.2	CAP_REG - Capability Register	346
1.20.3	ECAP_REG - Extended Capability Register	351
1.20.4	GCMD_REG - Global Command Register	353
1.20.5	GSTS_REG - Global Status Register	357
1.20.6	RTADDR_REG - Root-Entry Table Address Register	359
1.20.7	CCMD_REG - Context Command Register	360
1.20.8	FSTS_REG - Fault Status Register	363
1.20.9	FECTL_REG - Fault Event Control Register	365
1.20.10	FEDATA_REG - Fault Event Data Register	367
1.20.11	FEADDR_REG - Fault Event Address Register	367
1.20.12	FEUADDR_REG - Fault Event Upper Address Register	368
1.20.13	AFLOG_REG - Advanced Fault Log Register	368
1.20.14	PMEN_REG - Protected Memory Enable Register	369
1.20.15	PLMBASE_REG - Protected Low Memory Base Register	370
1.20.16	PLMLIMIT_REG - Protected Low Memory Limit Register	371
1.20.17	PHMBASE_REG - Protected High Memory Base Register	372
1.20.18	PHMLIMIT_REG - Protected High Memory Limit Register	373
1.20.19	IQH_REG - Invalidation Queue Head	374
1.20.20	IQT_REG - Invalidation Queue Tail	374
1.20.21	IQA_REG - Invalidation Queue Address	375
1.20.22	ICS_REG - Invalidation Completion Status	376
1.20.23	IECTL_REG - Invalidation Completion Event Control	376
1.20.24	IEDATA_REG - Invalidation Completion Event Data	377
1.20.25	IEADDR_REG - Invalidation Completion Event Address	378
1.20.26	IEUADDR_REG - Invalidation Completion Event Upper Address	378
1.20.27	IRTA_REG - Interrupt Remapping Table Address	379
1.20.28	IVA_REG - Invalidate Address Register	380
1.20.29	IOTLB_REG - IOTLB Invalidate Register	382
1.20.30	FRCD_REG - Fault Recording Registers	384
1.21	Intel® Trusted Execution Technology (Intel® TXT) Specific Registers	386
1.21.1	TXT.DID - TXT Device ID Register	386
1.21.2	TXT.DPR - DMA Protected Range	387
1.21.3	TXT.PUBLIC.KEY.LOWER - TXT Processor Public Key Hash Lower Half	388
1.21.4	TXT.PUBLIC.KEY.UPPER - TXT Processor Public Key Hash Upper Half	388



## Figures

Figure 1	System Address Range .....	16
Figure 2	Microsoft MS-DOS* Legacy Address Range .....	17
Figure 3	Main Memory Address Range .....	20
Figure 4	PCI Memory Address Range .....	25
Figure 5	Less Than 4 GB of Physical Memory (No Remap) .....	29
Figure 6	Greater than 4 GB of Physical Memory .....	30
Figure 7	4 GB or Less of Physical Memory .....	31
Figure 8	More Than 4 GB, Remap Enabled .....	33
Figure 9	Memory Map to PCI Express Device Configuration Space .....	41
Figure 10	Processor Configuration Cycle Flow Chart .....	43

## Tables

Table 1	SMM Regions .....	36
Table 2	Device Number Assignment for Internal Processor Devices .....	39
Table 3	Device 0 Function 0 Register Summary .....	47
Table 4	Device 0 MCHBAR DRAM Controls Summary .....	78



## Revision History

---

Revision Number	Description	Revision Date
001	<ul style="list-style-type: none"><li>Initial Release</li></ul>	January 2010
002	<ul style="list-style-type: none"><li>Updated Figure 8 for better clarification</li></ul>	November 2010

§



# 1 Processor Configuration Registers

This is volume 2 of the *Intel Core i7-600, i5-500, i5-400 and i3-300 Mobile Processor Series Datasheet*. Throughout this document, the Intel Core i7-600, i5-500, i5-400 and i3-300 Mobile Processor Series may be referred to as simply the processor. This document provides register information for the processor.

## 1.1 Register Terminology

The following table shows the register-related terminology that is used in this document.

(Sheet 1 of 3)

Item	Description
<b>RO</b>	<b>Read Only bit(s)</b> . Writes to these bits have no effect. These are static values only.
<b>RO-V</b>	<b>Read Only/Volatile bit(s)</b> . Writes to these bits have no effect. These are status bits only. The value to be read may change based on internal events.
<b>RO-V-S</b>	<b>Read Only/Volatile/Sticky bit(s)</b> . Writes to these bits have no effect. These are status bits only. The value to be read may change based on internal events. Bits are not returned to their default values by “warm” reset, but is reset with a cold/complete reset (for PCI Express* related bits a cold reset is “Power Good Reset” as defined in the <i>PCI Express Base Specification</i> ).
<b>AF</b>	<b>Atomic Flag bit(s)</b> . The first time the bit is read with an enabled byte, it returns the value 0, but a side-effect of the read is that the value changes to 1. Any subsequent reads with enabled bytes return a 1 until a 1 is written to the bit. When the bit is read, but the byte is not enabled, the state of the bit does not change, and the value returned is irrelevant, but will match the state of the bit.  When a 0 is written to the bit, there is no effect. When a 1 is written to the bit, its value becomes 0, until the next byte-enabled read. When the bit is written, but the byte is not enabled, there is no effect.  Conceptually, this is “Read to Set, Write 1 to Clear”
<b>RW</b>	<b>Read/Write bit(s)</b> . These bits can be read and written by software. Hardware may only change the state of this bit by reset.
<b>RW1C</b>	<b>Read/Write 1 to Clear bit(s)</b> . These bits can be read. Internal events may set this bit. A software write of 1 clears (sets to 0) the corresponding bit(s) and a write of 0 has no effect.



(Sheet 2 of 3)

Item	Description
<b>RW1C-L-S</b>	<b>Read/Write 1 to Clear/Lockable/Sticky bit(s).</b> These bits can be read. Internal events may set this bit. A software write of 1 clears (sets to 0) the corresponding bit(s) and a write of 0 has no effect. Bits are not cleared by "warm" reset, but is reset with a cold/complete reset (for PCI Express related bits a cold reset is "Power Good Reset" as defined in the PCI Express Base spec). Additionally there is a Key bit (which is marked RW-K or RW-L-K) that, when set, prohibits this bit field from being writable (bit field becomes Read Only/Volatile).
<b>RW1C-S</b>	<b>Read/Write 1 to Clear/Sticky bit(s).</b> These bits can be read. Internal events may set this bit. A software write of 1 clears (sets to 0) the corresponding bit(s) and a write of 0 has no effect. Bits are not cleared by "warm" reset, but is reset with a cold/complete reset (for PCI Express related bits a cold reset is "Power Good Reset" as defined in the PCI Express Base spec).
<b>RW-K</b>	<b>Read/Write/Key bit(s).</b> These bits can be read and written by software. Additionally this bit, when set, prohibits some other target bit field from being writable (bit fields become Read Only).
<b>RW-L</b>	<b>Read/Write/Lockable bit(s).</b> These bits can be read and written by software. Additionally there is a Key bit (which is marked RW-K or RW-L-K) that, when set, prohibits this bit field from being writable (bit field becomes Read Only).
<b>RW-L-K</b>	<b>Read/Write/Lockable/Key bit(s).</b> These bits can be read and written by software. This bit, when set, prohibits some other bit field(s) from being writable (bit fields become Read Only). Additionally there is a Key bit (which is marked RW-K or RW-L-K) that, when set, prohibits this bit field from being writable (bit field becomes Read Only). Conceptually, this may be a cascaded lock, or it may be self-locking when in its non-default state. When self-locking, it differs from RW-O in that writing back the default value will not set the lock.
<b>RW-V</b>	<b>Write/Volatile bit(s).</b> These bits can be read and written by software. Hardware may set or clear the bit based on internal events, possibly sooner than any subsequent software read could retrieve the value written.
<b>RW-V-L</b>	<b>Read/Write/Volatile/Lockable bit(s).</b> These bits can be read and written by software. Hardware may set or clear the bit based upon internal events, possibly sooner than any subsequent software read could retrieve the value written. Additionally there is a bit (which is marked RW-K or RW-L-K) that, when set, prohibits this bit field from being writable (bit field becomes Read Only).
<b>RW-V-L-S</b>	<b>Read/Write/Volatile/Lockable/Sticky bit(s).</b> These bits can be read and written by software. Hardware may set or clear the bit based upon internal events, possibly sooner than any subsequent software read could retrieve the value written. Additionally there is a bit (which is marked RW-K or RW-L-K) that, when set, prohibits this bit field from being writable (bit field becomes Read Only). These bits return to their default values on cold reset.
<b>RW-S</b>	<b>Read/Write/Sticky bit(s).</b> These bits can be read and written by software. Bits are not returned to their default values by "warm" reset, but will return to default values with a cold/complete reset (for PCI Express related bits a cold reset is "Power Good Reset" as defined in the PCI Express spec).



(Sheet 3 of 3)

Item	Description
<b>RW-O</b>	<b>Read/Write Once bit(s).</b> Reads prior to the first write return the default value. The first write after warm reset stores any value written. Any subsequent write to this bit field is ignored. All subsequent reads return the first value written. The value returns to default on warm reset. If there are multiple RW-O or RW-O-S fields within a DWORD, they should be written all at once (atomically) to avoid capturing an incorrect value.
<b>RW-O-S</b>	<b>Read/Write Once/Sticky bit(s).</b> Reads prior to the first write return the default value. The first write after cold reset stores any value written. Any subsequent write to this bit field is ignored. All subsequent reads return the first value written. The value returns to default on cold reset. If there are multiple RW-O or RW-O-S fields within a DWORD, they should be written all at once (atomically) to avoid capturing an incorrect value.
<b>W</b>	<b>Write-only.</b> These bits may be written by software, but will always return zeros when read. They are used for write side-effects. Any data written to these registers cannot be retrieved.
<b>W1C</b>	<b>Write 1 to Clear-only.</b> These bits may be cleared by software by writing a 1. Writing a 0 has no effect. The state of the bits cannot be read directly. The states of such bits are tracked outside the CPU and all read transactions to the address of such bits are routed to the other agent. Write transactions to these bits go to both agents.

## 1.2 System Address Map

**Note:** The processor is a multi-chip package (MCP) and basically consists of the CPU and the north bridge chipset, i.e., GMCH combined together in a single package. Hence this section will make reference to CPU as well as GMCH address mapping.

The processor supports 64 GB (36 bit) of addressable memory space and 64 KB+3 of addressable I/O space. The CPU performs decoding that historically occurred within the GMCH. Specifically, the GMCH address decoding for CPU initiated PAM, 15 M-16 M ISA hole, SMM CSEG/TSEG, PCIexBAR, and DRAM accesses will occur within the CPU and the GMCH has no direct knowledge. In addition, the Intel® Management Engine (Intel® ME) will move to the PCH, so Intel ME associated register ranges have been removed from the Graphics Controller. This section focuses on how the memory space is partitioned and what the separate memory regions are used for. I/O address space has simpler mapping and is explained near the end of this section.

The processor supports PEG (PCI Express Graphics) port upper prefetchable base/limit registers. This allows the PEG unit to claim IO accesses above 32 bit, complying with the PCI Express Spec. Addressing of greater than 4 GB is allowed on either the DMI Interface or PCI Express interface. The MCP supports a maximum of 16 GB of DRAM. No DRAM memory is accessible above 16 GB. DRAM capacity is limited by the number of address pins available.



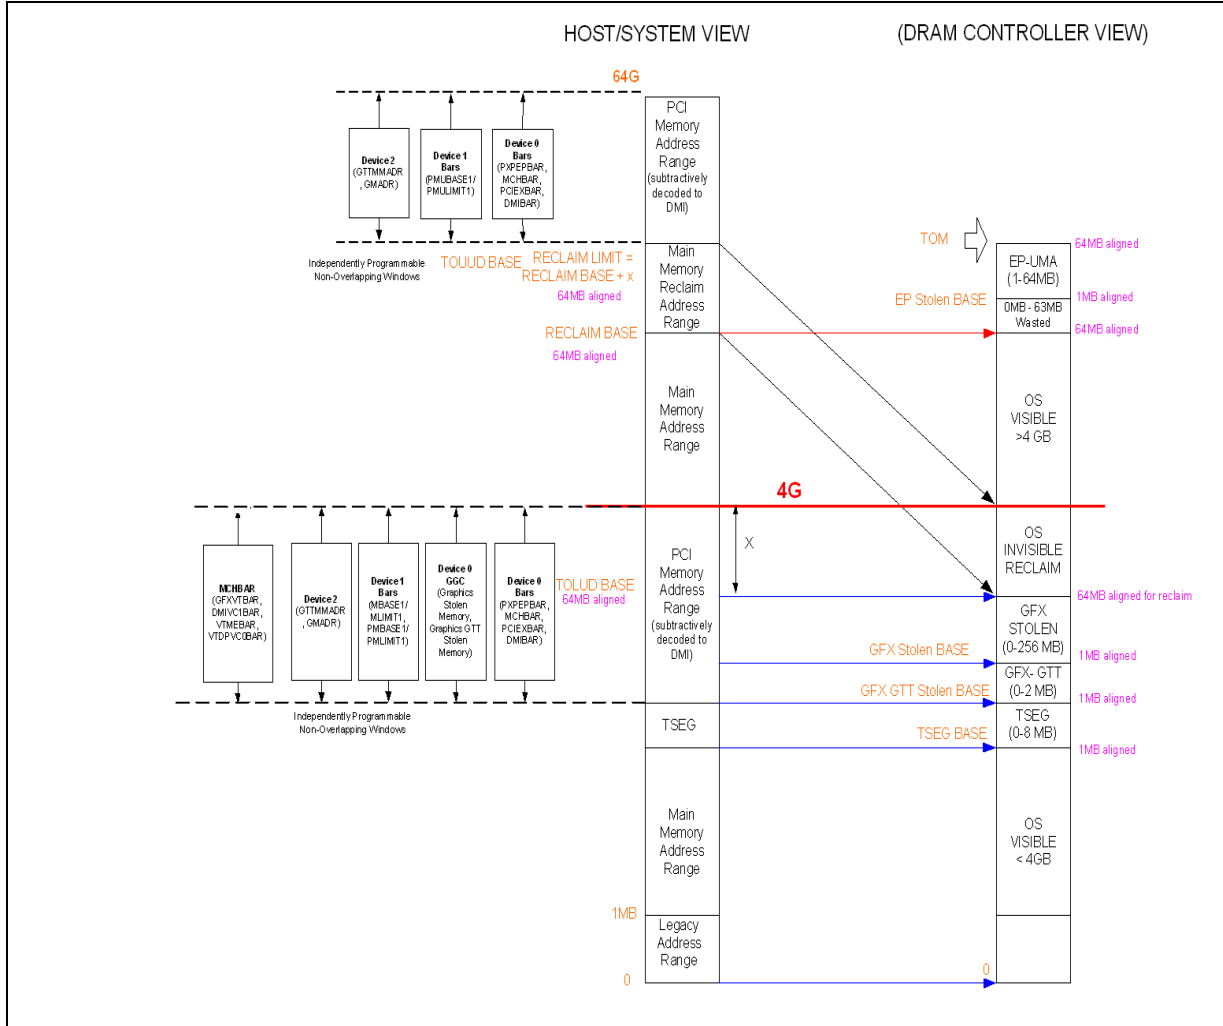
When running in internal graphics mode, tileX/tileY/linear reads/writes to GMADR range are supported. Write accesses to GMADR linear regions are supported from both DMI and PEG. GMADR write accesses to tileX and tileY regions (defined via fence registers) are not supported from DMI or the PEG port. GMADR read accesses are not supported from either DMI or PEG.

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 GMCH 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.

Figure 1 represents system memory address map in a simplified form.



Figure 1. System Address Range







### 1.2.1 Legacy Address Range

This area is divided into the following address regions:

- 0 – 640 KB – Microsoft MS-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

**Figure 2. Microsoft MS-DOS\* Legacy Address Range**

000F_FFFFh	System BIOS (Upper) 64KB	1MB
000F_0000h		960KB
000E_FFFFh	Extended System BIOS (Lower) 64KB (16KBx4)	896KB
000E_0000h		768KB
000D_FFFFh	Expansion Area 128KB (16KBx8)	640KB
000C_0000h		640KB
000B_FFFFh	Legacy Video Area (SMM Memory) 128KB	640KB
000A_0000h		640KB
0009_FFFFh	DOS Area	640KB
0000_0000h		640KB



### 1.2.1.1 DOS Range (0000\_0000h – 0009\_FFFFh)

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

### 1.2.1.2 Legacy Video Area (000A\_0000h-000B\_FFFFh)

The legacy 128-KB VGA memory range, frame buffer, (000A\_0000h – 000B\_FFFFh) can be mapped to IGD (Device 2), to PCI Express (Device 1), and/or to the DMI Interface. The appropriate mapping depends on which devices are enabled and the programming of the VGA steering bits. Based on the VGA steering bits, priority for VGA mapping is constant. The GMCH always decodes internally mapped devices first. Internal to the GMCH, decode priority is:

1. IGD
2. PCI Express
3. DMI Interface (subtractive)

Non-SMM-mode CPU accesses to this range are considered to be to the Video Buffer Area as described above. The CPU will route these accesses on the non-coherent (NCS or NCB) channels.

The GMCH 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 & MDAP). This region is also the default for SMM space.

#### 1.2.1.2.1 Compatible SMRAM Address Range (000A\_0000h-000B\_FFFFh)

Unlike FSB platforms, the GMCH sees no SMM indication with CPU accesses. When compatible SMM space is enabled, SMM-mode CPU accesses to this range route to physical system DRAM at 000A\_0000h - 000B\_FFFFh. The CPU performs the decode and routes the access to physical system DRAM.

PCI Express and DMI originated cycles to enabled SMM space are not allowed and are considered to be to the Video Buffer Area, if IGD is not enabled as the VGA device. DMI initiated writes cycles are attempted as peer writes cycles to a VGA enabled PCIe port.

#### 1.2.1.2.2 Monochrome Adapter (MDA) Range (000B\_0000h-000B\_7FFFh)

Legacy support requires the ability to have a second graphics controller (monochrome) in the system. Accesses in the standard VGA range are forwarded to IGD, PCI-Express, or the DMI Interface (depending on configuration bits). Since the monochrome adapter may be mapped to any of these devices, the GMCH must decode cycles in the MDA range (000B\_0000h - 000B\_7FFFh) and forward either to IGD, PCI-Express, or the DMI Interface. This capability is controlled by a VGA steering bits and the legacy configuration bit (MDAP bit). In addition to the memory range B0000h to B7FFFh, the GMCH decodes IO cycles at 3B4h, 3B5h, 3B8h, 3B9h, 3BAh and 3BFh and forwards them to the either IGD, PCI-Express, and/or the DMI Interface.



### 1.2.1.3 PAM (000C\_0000h-000F\_FFFFh)

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. The CPU documentation will now contain the registers and decode rules/restrictions.

The PAM registers have moved to the CPU. For the PAM register details, refer to CPU documentation.

- ISA Expansion Area (000C\_0000h-000D\_FFFFh)
- Extended System BIOS Area (000E\_0000h-000E\_FFFFh)
- System BIOS Area (000F\_0000h-000F\_FFFFh)

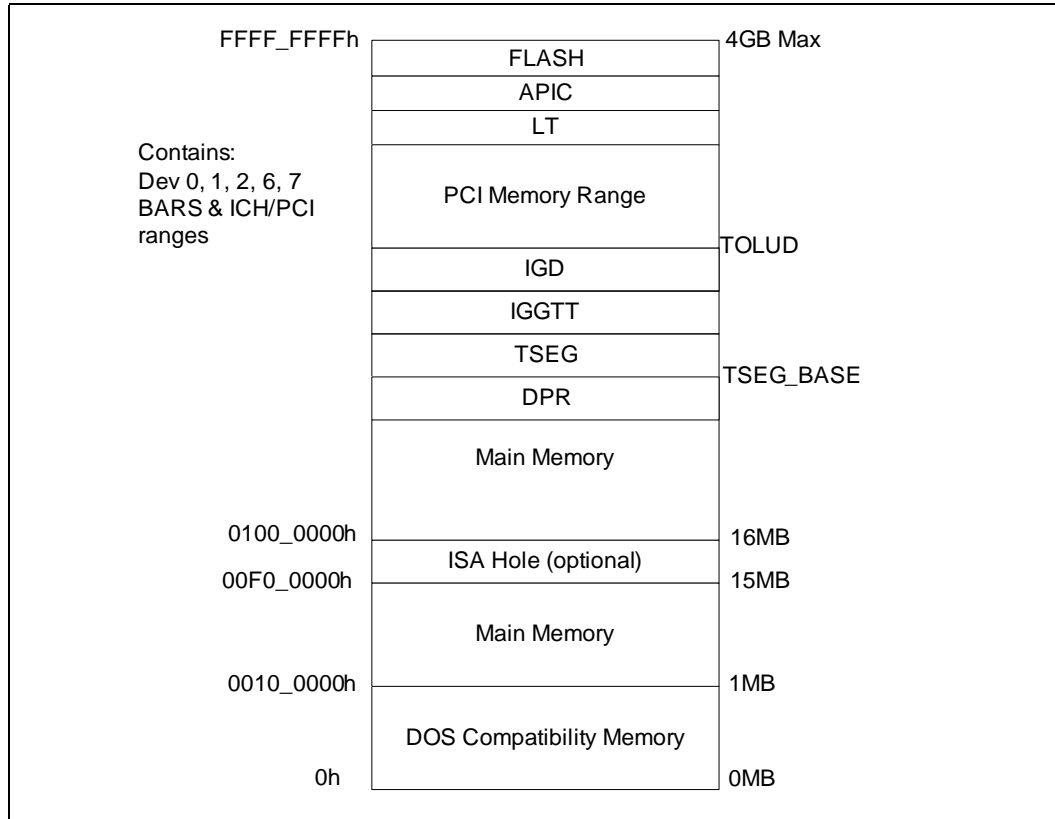
The CPU contains the PAM registers and the GMCH has no knowledge of the register programming. The CPU decodes the request and routes to the appropriate destination (DRAM or DMI) by sending the request on HOM or NCS/NCB.

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 is routed to 000C\_0000h. Writes will have the byte enables de-asserted.

## 1.2.2 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 GMCH (as programmed in the TOLUD register). The CPU will route all addresses within this range as HOM accesses which is forwarded by the GMCH to the DRAM unless it falls into the optional TSEG, optional ISA Hole, or optional IGD stolen VGA memory.

Figure 3. Main Memory Address Range



### 1.2.2.1 ISA Hole (15 MB-16 MB)

This register moved to the CPU. As such, the CPU performs the necessary decode and routes the request appropriately. Specifically, if no hole is created, the CPU will route the request to DRAM (HOM channel). If a hole is created, the CPU will route the request on NCS/NCB, since the request does not target DRAM.

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

### 1.2.2.2 TSEG

The TSEG register moved from the GMCH to the CPU. The GMCH will have no direct knowledge of the TSEG size. For CPU initiated transactions, the CPU will perform necessary decode and route appropriately on HOM (to DRAM) or NCS/NCB.

TSEG is below IGD stolen memory, which is at the Top of Low Usable physical memory (TOLUD). When SMM is enabled, the maximum amount of memory available to the system is equal to the amount of physical DRAM minus the value in the TSEG register. BIOS will calculate and program a register, so the GMCH has knowledge of where (TOLUD)-(GFX stolen)-(GFX GTT stolen)-(TSEG) is located. This is indicated by the TSEG\_BASE register.



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

When the extended SMRAM space is enabled, CPU accesses to the TSEG range without SMM attribute or without WB attribute are handled by the CPU as invalid accesses. Refer to the CPU documentation for how the CPU handles these accesses.

Non-CPU originated accesses are not allowed to SMM space. PCI-Express, DMI, and Internal Graphics originated cycle to enabled SMM space are handled as invalid cycle type with reads and writes to location 0 and byte enables turned off for writes.

### 1.2.2.3 Protected Memory Range (PMR) – (Programmable)

For robust and secure launch of the MVMM, the MVMM code and private data needs 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® Trusted Execution Technology (Intel® TxT), and is optional for non-Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d) 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.
- **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, 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.



#### 1.2.2.4 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 CPU streams.

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:

1. 0 to (TSEG\_BASE – DPR size – 1) for DMA traffic
2. TSEG\_BASE
3. 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 could use the previous decode until the new decode can guarantee PV. No flushing of cycles is required. On a clock by clock basis proper decode with the previous or new decode needs to be guaranteed.

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

#### 1.2.2.5 Pre-allocated Memory

Voids of physical addresses that are not accessible as general system memory and reside within 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.

#### 1.2.2.6 GFX Stolen Spaces

##### 1.2.2.6.1 GTT Stolen Space (GSM)

GSM is allocated to store the GFX translation table entries, depending on Intel VT-d support it may be divided into two sections.

##### 1.2.2.6.2 Global GTT Stolen Space (GGSM)

GGSM always exists regardless of Intel VT-d as long as internal GFX is enabled. This space is allocated to store accesses as page table entries are getting updated through virtual GTTMMADR range. Hardware is responsible to map PTEs into this physical space.



**Note:** Direct accesses to GGSM are not allowed, only hardware translations and fetches can be directed to GGSM.

#### 1.2.2.6.3 Shadow GTT Stolen Space (SGSM)

Shadow GSM is only used once internal GFX and Intel VT-d translations are enabled. The purpose of shadow GSM is to provide a physical space to hardware, where Intel VT-d translation for PTE updates can be made on the fly and re-written back into physical memory.

#### 1.2.2.7 Intel® Management Engine (Intel® ME) UMA

Intel ME can be allocated UMA memory. The Intel ME memory is “stolen” from the top of the Host address map. Intel ME stolen memory base is calculated by subtracting the amount of memory stolen by the Manageability Engine from TOM.

**Note:** Only Intel ME can access this space; it is not accessible by or coherent with any CPU side accesses.

#### 1.2.2.8 PCI Memory Address Range (TOLUD - 4 GB)

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

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

1. Addresses decoded to the PCI Express Memory Window defined by the MBASE1, MLIMIT1, registers are mapped to PCI Express.
2. Addresses decoded to the PCI Express prefetchable Memory Window defined by the PMBASE1, PMLIMIT1, registers are mapped to PCI Express.

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

1. Addresses decode to the internal graphics translation window (GMADR)
2. Addresses decode to the Internal graphics translation table or IGD registers. (GTTMMADR)



With Intel VT-d enabling configuration, there are exceptions to this rule:

1. Addresses decoded to the memory mapped window to Graphics Intel VT-d remap engine registers (GFXVTBAR)
2. Addresses decoded to the memory mapped window to DMI VC1 Intel VT-d remap engine registers (DMIVC1BAR)
3. Addresses decoded to the memory mapped window to PEG/DMI/Intel ME VC0 Intel VT-d remap engine registers (VTDPVC0BAR)
4. TCm accesses (to Intel ME stolen memory) from PCH do not go through Intel VT-d remap engines.

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

**Note:** 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 4. PCI Memory Address Range

FFFF_FFFFh	High BIOS	4GB
FFE0_0000h	DMI Interface (subtractive decode)	4GB - 2MB
FEF0_0000h	MSI Interrupts	4GB - 17MB
FEE0_0000h	DMI Interface (subtractive decode)	4GB - 18MB
FED0_0000h	Local (CPU) APIC	4GB - 19MB
FEC8_0000h	I/O APIC	
FEC0_0000h	DMI Interface (subtractive decode)	4GB - 20MB
F000_0000h	PCI Express Configuration Space	4GB - 256MB
E000_0000h	DMI Interface (subtractive decode)	4GB - 512MB
		<i>Possible address range/size (not guaranteed)</i>
		<i>BARs, Internal Graphics ranges, PCI Express Port, CHAPADR could be here.</i>
		TOLUD



#### 1.2.2.9 APIC Configuration Space (FECO\_0000h-FECF\_FFFFh)

This range is reserved for APIC configuration space. The I/O APIC(s) usually reside in the PCH portion of the chip-set, 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 (FECO\_0000h to FEC7\_FFFFh) are always forwarded to DMI.

The GMCH optionally supports additional I/O APICs behind the PCI Express "Graphics" port. When enabled via the PCI Express Configuration register (Device 1 Offset 200h) the PCI Express port(s) will positively decode a subset of the APIC configuration space. Specifically,

- Device 1 can be enabled to claim FECC\_0000h thru FECF\_FFFFh.

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 GMCH, and would be disabled in typical Desktop systems. When disabled, any access within entire APIC Configuration space (FECO\_0000h to FECF\_FFFFh) is forwarded to DMI.

#### 1.2.2.10 MSI Interrupt Memory Space (FEE0\_0000-FEEF\_FFFF)

Any PCI Express or DMI device may issue a Memory Write to 0FEE<sub>x</sub>xxxxh. This Memory Write cycle does not go to DRAM. The GMCH will forward this Memory Write along with the data to the CPU as an Interrupt Message Transaction.

This interrupt message is delivered to the CPU as an IntPhysical or IntLogical message.

#### 1.2.2.11 High BIOS Area

For security reasons, the GMCH will now positively decode this range to DMI. This positive decode will guarantee any overlapping ranges is ignored.

The top 2 MB (FEE0\_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 CPU begins execution from the High BIOS after reset. This region is positively decoded to DMI Interface so that the upper subset of this region aliases to 16 MB–256 KB range. The actual address space required for the BIOS is less than 2 MB, but the minimum CPU MTRR range for this region is 2 MB, so that full 2 MB must be considered.



### 1.2.3 Main Memory Address Space (4 GB to TOUUD)

The processor supports 36-bit addressing. The maximum main memory size supported is 16GB 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 is handled similar to other GMCHs.

Upstream read and write accesses above 36-bit addressing is 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).

Intel ME stolen size register reflects the total amount of physical memory stolen by the Manageability Engine. Intel ME stolen memory is located at the top of physical memory. Intel ME stolen memory base is calculated by subtracting the amount of memory stolen by the Manageability Engine 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 Manageability Engine's stolen size. If remap is enabled, then it will reflect the remap limit. Note, when there is more than 4 GB of DRAM and reclaim is enabled, the reclaim base is the same as TOM minus Intel ME stolen memory size to the nearest 64-MB alignment (shown in case 2 below).

#### **Top of Low Usable DRAM (TOLUD)**

TOLUD register is restricted to 4 GB memory (A[31:20]), but the processor can support up to 16 GB, limited by DRAM pins. For physical memory greater than 4 GB, the TOUUD register helps identify the address range in between the 4-GB boundary and the top of physical memory. This identifies memory that can be directly accessed (including remap address calculation) which is useful for memory access indication and early path indication. When remap is enabled, TOLUD must be 64-MB aligned, but when remap is disabled, 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 and program this register, so the GMCH has knowledge of where (TOLUD)-(Gfx stolen)-(Gfx GTT stolen)-(TSEG) is located. IO blocks use this minus DPR for upstream DRAM decode.



### 1.2.3.1 Programming Model

The memory boundaries of interest are:

- Bottom of logical address remap window defined by the REMAPBASE register, which is calculated and loaded by BIOS.
- Top of logical address remap window defined by the REMAPLIMIT register, which is calculated and loaded by BIOS.
- Bottom of physical remap memory defined by the existing TOLUD register.
- Top of physical remap memory, which is implicitly defined by either 4 GB or TOM minus Manageability Engine stolen size.

Determine the following Mapping steps:

1. TOM
2. TOM minus Intel ME stolen size
3. MMIO allocation
4. TOLUD
5. GFX stolen base
6. GFX GTT stolen base
7. TSEG base
8. Remap base/limit
9. TOUUD


The following diagrams show the four possible general cases of remapping.

- Case #1: Less than 4 GB of Physical Memory, no remap
- Case #2: Greater than 4 GB of Physical Memory
- Case #3: 4 GB or Less of Physical Memory
- Case #4: Greater than 4 GB of Physical Memory, remap



1.2.3.1.1 Case #1: Less Than 4 GB of Physical Memory (No Remap)

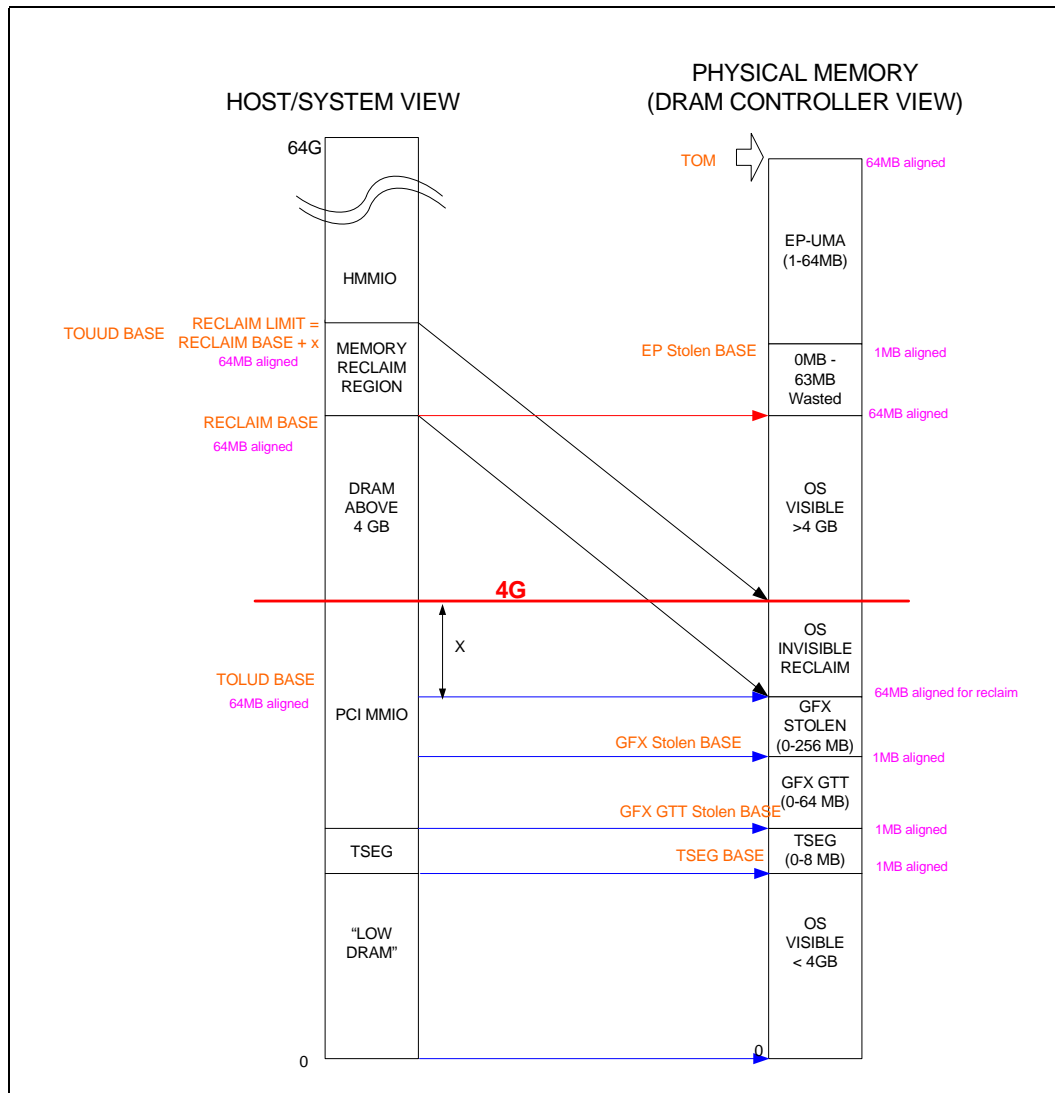
Figure 5. Less Than 4 GB of Physical Memory (No Remap)



- Populated Physical Memory = 2 GB
- Address Space allocated to memory mapped IO = 1 GB
- Remapped Physical Memory = 0 GB
- TOM – 020h (2 GB)
- Intel ME stolen size – 00001b (1 MB)
- TOUUD – 07FFh (2 GB minus 1 MB) (1 MB aligned)
- TOLUD – 01F00h (2GB minus 64 MB)
- REMAPBASE – 3FFh (64 GB – 1 boundary, default)
- REMAPLIMIT – 000h (0-GB boundary, default)

1.2.3.1.2 Case #2: Greater than 4 GB of Physical Memory

Figure 6. Greater than 4 GB of Physical Memory



In this case, the amount of memory remapped is the range between TOLUD and 4 GB. This physical memory is mapped to the logical address range defined between the REMAPBASE and the REMAPLIMIT registers.

**Example: 5 GB of Physical Memory, with 1 GB allocated to Memory Mapped IO:**

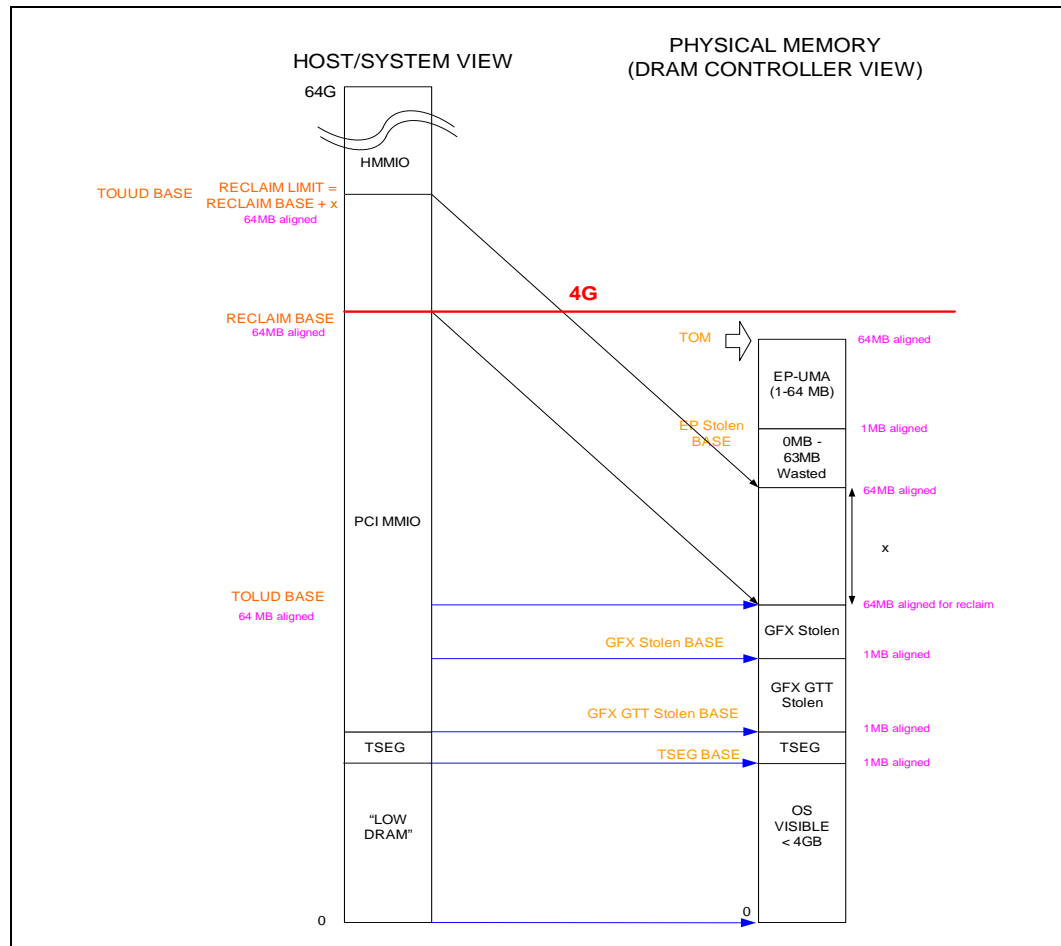
- Populated Physical Memory = 5 GB
- Address Space allocated to memory mapped IO = 1 GB
- Remapped Physical Memory = 1GB
- TOM – 050h (5 GB)
- Intel ME stolen size – 00000b (0 MB)



- TOUUD – 1800h (6 GB) (1 MB aligned)
- TOLUD – 06000h (3 GB) (64 MB aligned because remap is enabled and the remap register has 64MB granularity)
- REMAPBASE – 050h (5 GB)
- REMAPLIMIT – 05Fh (6 GB – 1 boundary)

1.2.3.1.3 Case #3: 4 GB or Less of Physical Memory

Figure 7. 4 GB or Less of Physical Memory



In this case the amount of memory remapped is the range between TOLUD and TOM minus the Intel ME stolen memory. This physical memory is mapped to the logical address range defined between the REMAPBASE and the REMAPLIMIT registers.



**Example: 3 GB of Physical Memory, with 2 GB allocated to Memory Mapped IO:**

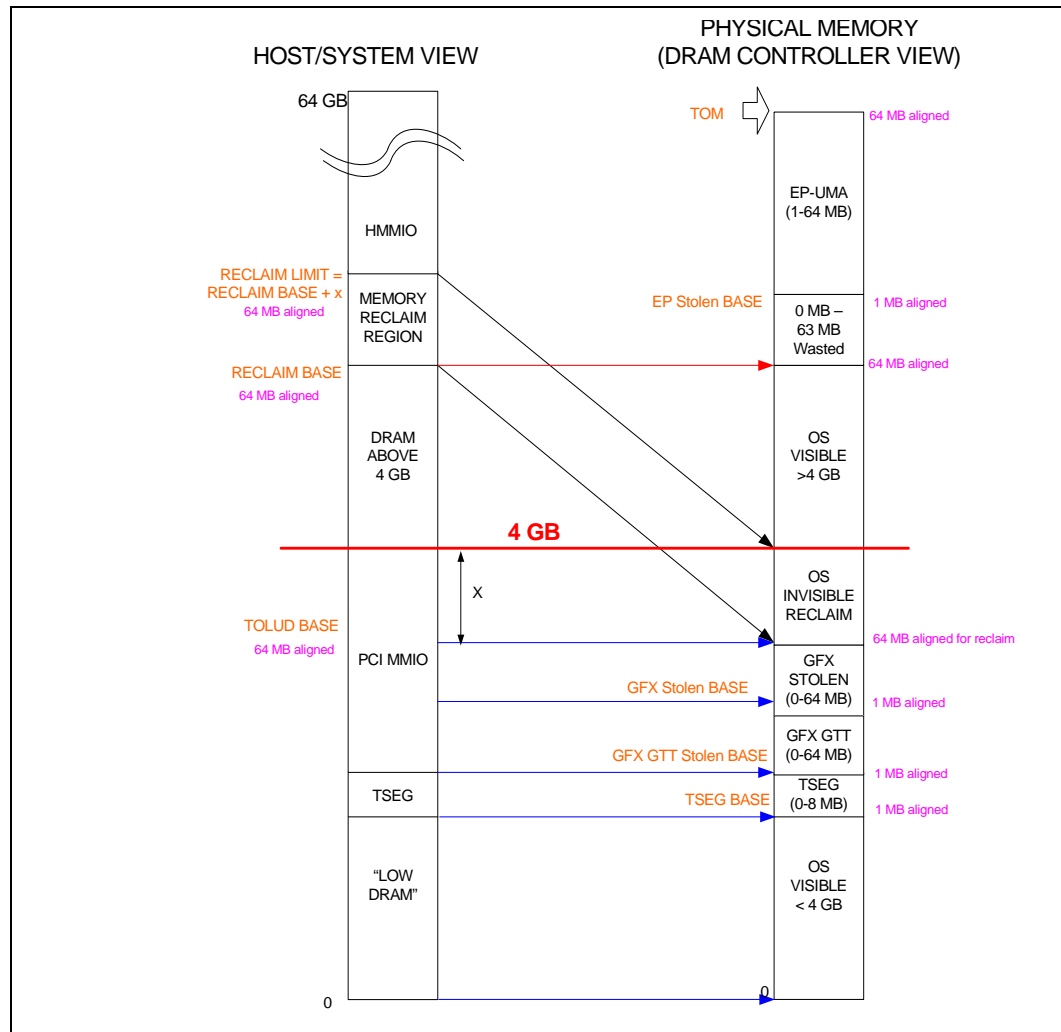
- Populated Physical Memory = 3 GB
- Address Space allocated to memory mapped IO = 2 GB
- Remapped Physical Memory = 1 GB
- TOM – 030h (3 GB)
- Intel ME stolen size – 00000b (0 MB)
- TOUUD – 1400h (5 GB) (1-MB aligned)
- TOLUD – 02000h (2 GB) (64-MB aligned because remap is enabled and the remap register has 64MB granularity)
- REMAPBASE – 040h (4 GB)
- REMAPLIMIT – 04Fh (5 GB – 1 boundary)





1.2.3.1.4 Case #4: Greater Than 4 GB of Physical Memory, Remap

Figure 8. More Than 4 GB, Remap Enabled



In this case the amount of memory remapped is the range between TOLUD and 4 GB. This physical memory is mapped to the logical address range defined between the REMAPBASE and the REMAPLIMIT registers.

**Example: 5 GB of Physical Memory, with 1 GB allocated to Memory Mapped IO:**

- Populated Physical Memory = 5 GB
- Address Space allocated to memory mapped IO = 1 GB
- Remapped Physical Memory = 1 GB
- TOM – 050h (5 GB)
- Intel ME stolen size – 00000b (0 MB)
- TOUUD – 17FFh (6GB-1 MB) (1 MB aligned)



- TOLUD – 06000h (3 GB) (64 MB aligned because remap is enabled and the remap register has 64MB granularity)
- REMAPBASE – 050h (5 GB)
- REMAPLIMIT – 05Fh (6 GB – 1 boundary)

#### 1.2.4 PCI Express\* Configuration Address Space

PCIEXBAR has moved to the CPU. The CPU now detects memory accesses targeting PCIEXBAR. BIOS must assign this address range such that it will not conflict with any other address ranges.

#### 1.2.5 PCI Express Graphics Attach (PEG)

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

- The first range is controlled via the Memory Base Register (MBASE) and Memory Limit Register (MLIMIT) registers.
- The second range is controlled via 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 GMCH 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 GMCH positively decodes memory accesses to PCI Express memory address space as defined by the following equations:

$$\text{Memory\_Base\_Address} \leq \text{Address} \leq \text{Memory\_Limit\_Address}$$

$$\text{Prefetchable\_Memory\_Base\_Address} \leq \text{Address} \leq \text{Prefetchable\_Memory\_Limit\_Address}$$

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.



Note that the GMCH memory range registers described above are used to allocate memory address space for any PCI Express devices sitting on PCI Express that require such a window.

The PCICMD1 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.

For the processor, the upper PMUBASE/PMULIMIT registers have been implemented for PCI Express Spec compliance. The processor locates MMIO space above 4 GB using these registers.

### 1.2.6 Graphics Memory Address Ranges

The GMCH can be programmed to direct memory accesses to IGD when addresses are within any of five ranges specified via registers in GMCH Device 2 configuration space.

1. The Graphics Memory Aperture Base Register (GMADR) is used to access graphics memory allocated via the graphics translation table.
2. The Graphics Translation Table Base Register (GTTADR) is used to access the translation table and graphics control registers. This is part of 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 steal any physical DRAM memory space.

Alternatively, these ranges can reside above 4 GB, similar to other BARs which 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.



### 1.2.6.1 IOBAR Mapped Access to Device 2 MMIO Space

Device 2, integrated graphics device, contains an IOBAR register. If Device 2 is enabled, then 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. An IO write to this port loads the offset of the MMIO register or offset into the GTT that needs to be accessed. An IO Read returns the current value of this register. See IOBAR rules for detailed information.

**MMIO\_Data:** MMIO\_DATA is a 32-bit register. An IO write to this port is re-directed to the MMIO register pointed to by the MMIO-index register. An IO read to this port is re-directed to the MMIO register pointed to by the MMIO-index register. See IOBAR rules for detailed information.

The result of accesses through IOBAR can be:

1. Accesses directed to the GTT table. (i.e., route to DRAM)
2. Accesses to internal graphics registers with the GMCH (i.e. route to internal configuration bus)
3. Accesses to internal graphics display registers now located within the PCH. (i.e., 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 IO registers which are mapped through the MMIO space. VGA registers must be accessed directly through the dedicated VGA IO ports.

### 1.2.7 System Management Mode (SMM)

The CPU handles all SMM mode transaction routing. The GMCH has no direct knowledge of SMM mode. The GMCH will never allow IO devices access to CSEG/TSEG/HSEG ranges. Refer to the CPU EAS for full register details, behaviors, and restrictions.

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

Table 1. 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



### 1.2.8 SMM and VGA Access through GTT TLB

Accesses through GTT TLB address translation SMM DRAM space are not allowed. Writes is routed to Memory address 000C\_0000h with byte enables de-asserted and reads is 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 GMADR range will not be snooped. Only PCI Express and DMI accesses to GMADR linear range (defined via 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 enabled SMM DRAM space, the request is 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 will have no address translation concerns. PCI Express and DMI Interface reads to GMADR is 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 they are not in SMM (actually, to ensure they are not above base of TSEG or 640 K -1 M. Thus, they are invalid and go to address 000C\_0000h, but that isn't specific to PCI Express or DMI; it applies to CPU or internal graphics engines.

### 1.2.9 I/O Address Space

The GMCH generates either DMI Interface or PCI Express bus cycles for all CPU I/O accesses that it does not claim. The GMCH no longer contains the two internal registers in the CPU I/O space, Configuration Address Register (CONFIG\_ADDRESS) and the Configuration Data Register (CONFIG\_DATA).

The CPU allows 64K+3 bytes to be addressed within the I/O space. The GMCH propagates the CPU I/O address without any translation on to the destination bus and therefore provides addressability for 64K+3 byte locations. Note that 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 CPU 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 IO decode and the associated control is explained later.



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 GMCH 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 request will route as a read to Memory address 000C\_0000h so a completion is naturally generated (whether the original request was a read or write). The transaction will complete with an UR completion status.

CPU I/O reads that lie within 8-byte boundaries but cross 4-byte boundaries are issued from the CPU as 1 transaction. The GMCH will break this into 2 separate transactions. I/O writes that lie within 8-byte boundaries but cross 4-byte boundaries is split into two transactions by the CPU.

### 1.2.9.1 PCI Express I/O Address Mapping

The GMCH can be programmed to direct non-memory (I/O) accesses to the PCI Express bus interface when CPU initiated I/O cycle addresses are within the PCI Express I/O address range. This range is controlled via the I/O Base Address (IOBASE) and I/O Limit Address (IOLIMIT) registers in GMCH Device 1 or Device 6 (if a second PEG port is enabled) 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 GMCH assumes that 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 4-KB boundary and produces a size granularity of 4 KB.

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

$$\text{I/O\_Base\_Address} \leq \text{CPU I/O Cycle Address} \leq \text{I/O\_Limit\_Address}$$

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 GMCH also forwards accesses to the Legacy VGA I/O ranges according to the settings in the Device 1 configuration registers BCTRL (VGA Enable) and PCICMD1 (IOAE1), 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, the GMCH will decode legacy monochrome IO 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.

**Note:** The GMCH Device 1 I/O address range registers defined above are used for all I/O space allocation for any devices requiring such a window on PCI-Express. The PCICMD1 register can disable the routing of I/O cycles to PCI-Express.



## 1.3 Configuration Process and Registers

### 1.3.1 Platform Configuration Structure

The DMI physically connects the processor and the Intel PCH; so, from a configuration standpoint, the DMI is logically PCI Bus 0. As a result, all devices internal to the processor and the Intel PCH appear to be on PCI Bus 0.

**Note:** The PCH internal LAN controller does not appear on Bus 0 - it appears on the external PCI bus (whose number is configurable).

The system's primary PCI expansion bus is physically attached to the Intel PCH and, from a configuration perspective, appears to be a hierarchical PCI bus behind a PCI-to-PCI bridge and therefore has a programmable PCI Bus number. The PCI Express Graphics Attach appears to system software to be a real PCI bus behind a PCI-to-PCI bridge that is a device resident on PCI Bus 0.

**Note:** A physical PCI bus 0 does not exist. DMI and the internal devices in the processor and Intel PCH logically constitute PCI Bus 0 to configuration software. This is shown in the following figure.

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

- **Device 0** Host Bridge/DRAM Controller. Logically this 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 appears as a “virtual” PCI-to-PCI bridge residing on PCI Bus 0 and is compliant with PCI Express Base Specification. Device 1 contains the standard PCI-to-PCI bridge registers and the standard PCI Express/PCI configuration registers (including the PCI Express memory address mapping). It also contains Isochronous and Virtual Channel controls in the PCI Express extended configuration space.
- **Device 2** Internal Graphics Device. Logically, this appears as an APCI device residing on PCI Bus 0. Physically, Device 2 contains the configurations registers for 3D, 2D and display functions.

**Table 2. Device Number Assignment for Internal Processor Devices**

Processor Function	Device Number
Host Bridge/DRAM Controller	Device 0
Host-to-PCI Express Bridge (virtual P2P)	Device 1
Internal Graphics Device	Device 2



## 1.4 Configuration Mechanisms

The processor is the originator of configuration cycles. Internal to the processor transactions received through both configuration mechanisms are translated to the same format.

### 1.4.1 Standard PCI Configuration Mechanism

The following is the mechanism for translating processor I/O bus cycles to configuration cycles.

The PCI specification defines a slot based “configuration space” that allows each device to contain up to eight functions with each function containing up to 256, 8-bit configuration registers. The PCI specification defines two bus cycles to access the PCI configuration space: Configuration Read and Configuration Write. Memory and I/O spaces are supported directly by the processor. Configuration space is supported by a mapping mechanism implemented within the processor.

The configuration access mechanism makes use of the CONFIG\_ADDRESS Register (at I/O address 0CF8h through 0CFBh) and CONFIG\_DATA Register (at I/O address 0CFCh through 0CFFh). To reference a configuration register, a dword I/O write cycle is used to place a value into CONFIG\_ADDRESS that specifies the PCI bus, the device on that bus, the function within the device and a specific configuration register of the device function being accessed. CONFIG\_ADDRESS[31] must be 1 to enable a configuration cycle. CONFIG\_DATA then becomes a window into the four bytes of configuration space specified by the contents of CONFIG\_ADDRESS. Any read or write to CONFIG\_DATA will result in the processor translating the CONFIG\_ADDRESS into the appropriate configuration cycle.


The processor is responsible for translating and routing the processor's I/O accesses to the CONFIG\_ADDRESS and CONFIG\_DATA registers to internal processor configuration registers, DMI or PCI Express.

### 1.4.2 PCI Express Enhanced Configuration Mechanism

PCI Express extends the configuration space to 4096 bytes per device/function as compared to 256 bytes allowed by the latest *PCI Local Bus Specification*. PCI Express configuration space is divided into a PCI 3.0 compatible region, which consists of the first 256B of a logical device's configuration space and a PCI Express extended region which consists of the remaining configuration space.

The PCI compatible region can be accessed using either the Standard PCI Configuration Mechanism or using the PCI Express Enhanced Configuration Mechanism described in this section. The extended configuration registers may only be accessed using the PCI Express Enhanced Configuration Mechanism. To maintain compatibility with PCI configuration addressing mechanisms, system software must access the extended configuration space using 32-bit operations (32-bit aligned) only. These 32-bit operations include byte enables allowing only appropriate bytes within the dword to be



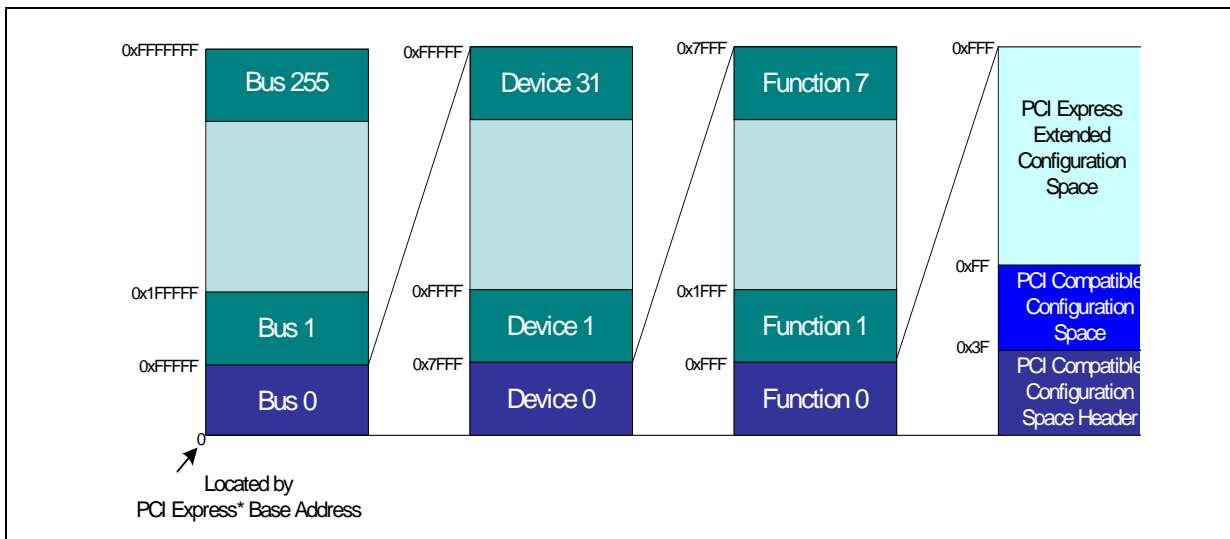


accessed. Locked transactions to the PCI Express memory mapped configuration address space are not supported. All changes made using either access mechanism are equivalent.

The PCI Express Enhanced Configuration Mechanism utilizes a flat memory-mapped address space to access device configuration registers. This address space is reported by the system firmware to the operating system. There is a register, PCIEXBAR, that defines the base address for the block of addresses below 4 GB for the configuration space associated with busses, devices and functions that are potentially a part of the PCI Express root complex hierarchy. In the PCIEXBAR register there are controls to limit the size of this reserved memory mapped space. 256 MB is the amount of address space required to reserve space for every bus, device, and function that could possibly exist. Options for 128 MB and 64 MB exist in order to free up those addresses for other uses. In these cases, the number of busses and all of their associated devices and functions are limited to 128 or 64 busses, respectively.

The PCI Express Configuration Transaction Header includes an additional four bits (ExtendedRegisterAddress[3:0]) between the Function Number and Register Address fields to provide indexing into the 4 KB of configuration space allocated to each potential device. For PCI Compatible Configuration Requests, the Extended Register Address field must be all zeros.

**Figure 9. Memory Map to PCI Express Device Configuration Space**



Just the same as with PCI devices, each device is selected based on decoded address information that is provided as a part of the address portion of Configuration Request packets. A PCI Express device will decode all address information fields (Bus, Device, Function and extended address numbers) to provide access to the correct register.



To access this space (step 1 is done only once by BIOS):

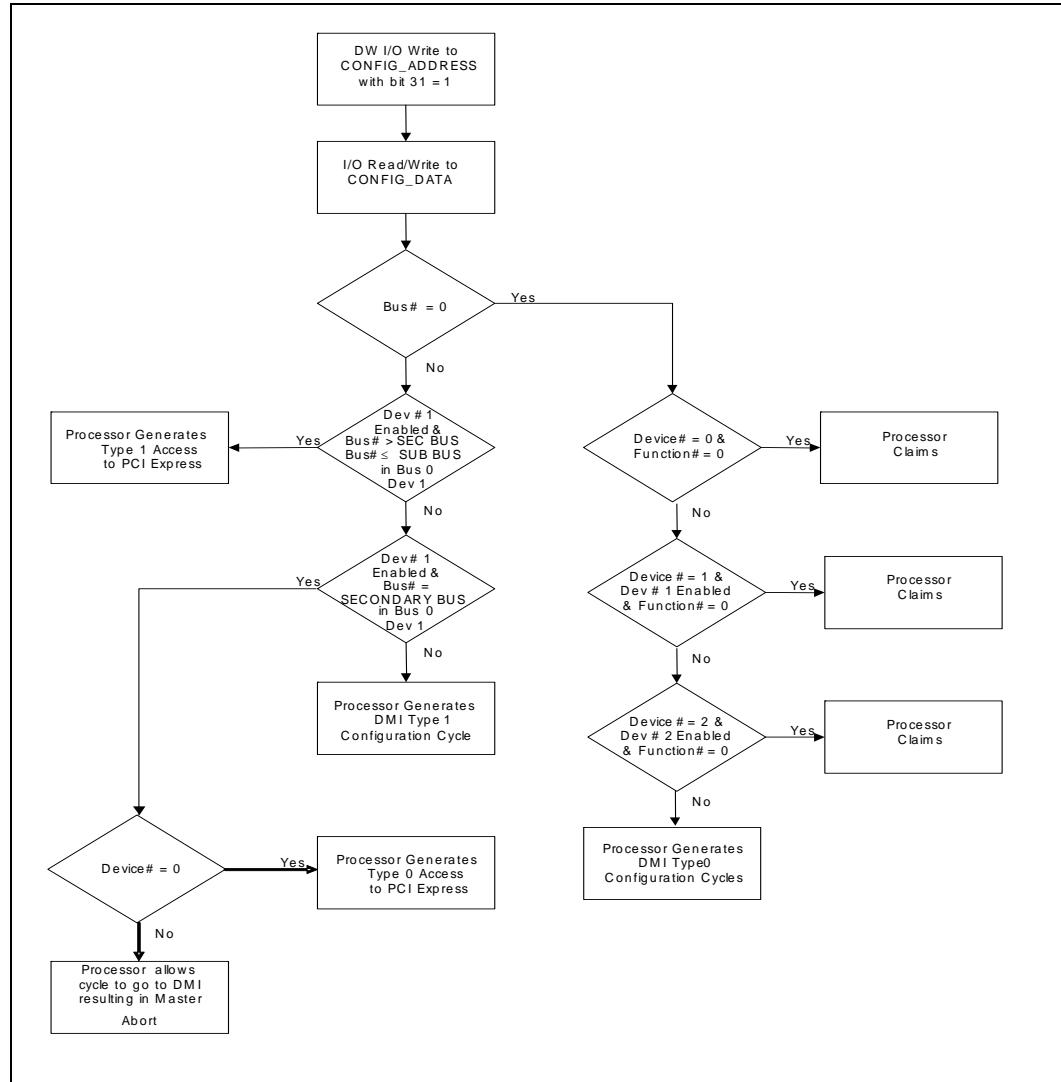
1. Write to CSR address 0x01050 to enable the PCI Express enhanced configuration mechanism by writing 1 to Bit 0 of the GQ1\_CR\_PCIEBAR register. Allocate either 256, 128, or 64 busses to PCI Express by writing "000", "111", or "110" respectively to Bits 3:1. Pick a naturally aligned base address for mapping the configuration space onto memory space using 1 MB per bus number and write that base address into Bits 39:20.
2. Calculate the host address of the register you wish to set using (PCI Express base + (bus number \* 1 MB) + (device number \* 32 KB) + (function number \* 4 KB) + (1 B \* offset within the function) = host address)
3. Use a memory write or memory read cycle to the calculated host address to write or read that register.

## 1.5 Routing Configuration Accesses

The processor supports two PCI related interfaces: DMI and PCI Express. The processor is responsible for routing PCI and PCI Express configuration cycles to the appropriate device that is an integrated part of the processor or to one of these two interfaces. Configuration cycles to the PCH internal devices and Primary PCI (including downstream devices) are routed to the PCH via DMI. Configuration cycles to both the PCI Express Graphics PCI compatibility configuration space and the PCI Express Graphics extended configuration space are routed to the PCI Express Graphics port device or associated link.



Figure 10. Processor Configuration Cycle Flow Chart



### 1.5.1 Internal Device Configuration Accesses

The processor decodes the Bus Number (Bits 23:16) and the Device Number fields of the CONFIG\_ADDRESS register. If the Bus Number field of CONFIG\_ADDRESS is 0 the configuration cycle is targeting a PCI Bus 0 device.

If the targeted PCI Bus 0 device exists in the processor and is not disabled, the configuration cycle is claimed by the appropriate device.



## 1.5.2 Bridge Related Configuration Accesses

Configuration accesses on PCI Express or DMI are PCI Express Configuration TLPs.

- Bus Number [7:0] is Header Byte 8 [7:0]
- Device Number [4:0] is Header Byte 9 [7:3]
- Function Number [2:0] is Header Byte 9 [2:0]

Special fields for this type of TLP:

- Extended Register Number [3:0] is Header Byte 10 [3:0]
- Register Number [5:0] is Header Byte 11 [7:2]

See the *PCI Express Local Base Specification* for more information on both the PCI 3.0-compatible and PCI Express Enhanced Configuration Mechanism and transaction rules.

### 1.5.2.1 PCI Express Configuration Accesses

When the Bus Number of a Type 1 Standard PCI Configuration cycle or PCI Express Enhanced Configuration access matches the Device 1 Secondary Bus Number a PCI Express Type 0 Configuration TLP is generated on the PCI Express link targeting the device directly on the opposite side of the link. This should be Device 0 on the bus number assigned to the PCI Express link (likely Bus 1).

The device on other side of link must be Device 0. The processor will Master Abort any Type 0 Configuration access to a non-zero Device number. If there is to be more than one device on that side of the link there must be a bridge implemented in the downstream device.

When the Bus Number of a Type 1 Standard PCI Configuration cycle or PCI Express Enhanced Configuration access is within the claimed range (between the upper bound of the bridge device's Subordinate Bus Number register and the lower bound of the bridge device's Secondary Bus Number register) but doesn't match the Device 1 Secondary Bus Number, a PCI Express Type 1 Configuration TLP is generated on the secondary side of the PCI Express link.

PCI Express Configuration Writes:

- Internally the processor will translate writes to PCI Express extended configuration space to configuration writes on the backbone.
- Posted writes to extended space are non-posted on the PCI Express or DMI (i.e., translated to config writes)

### 1.5.2.2 DMI Configuration Accesses

Accesses to disabled processor internal devices, bus numbers not claimed by the Host-PCI Express bridge, or PCI Bus 0 devices not part of the processor will subtractively decode to the PCH and consequently be forwarded over the DMI via a PCI Express configuration TLP.



If the Bus Number is zero, the processor will generate a Type 0 Configuration Cycle TLP on DMI. If the Bus Number is non-zero, and falls outside the range claimed by the Host-PCI Express bridge, the processor will generate a Type 1 Configuration Cycle TLP on DMI.

The PCH routes configurations accesses in a manner similar to the processor. The PCH decodes the configuration TLP and generates a corresponding configuration access. Accesses targeting a device on PCI Bus 0 may be claimed by an internal device. The PCH compares the non-zero Bus Number with the Secondary Bus Number and Subordinate Bus Number registers of its PCI-to-PCI bridges to determine if the configuration access is meant for Primary PCI, or some other downstream PCI bus or PCI Express link.

Configuration accesses that are forwarded to the PCH, but remain unclaimed by any device or bridge will result in a master abort.

## 1.6 Processor Register Introduction

The processor contains two sets of software accessible registers, accessed via the Host Processor I/O address space: Control registers and internal configuration registers.

- Control registers are I/O mapped into the processor I/O space, which control access to PCI and PCI Express configuration space (see section entitled I/O Mapped Registers).
- Internal configuration registers residing within the processor are partitioned into three logical device register sets ("logical" since they reside within a single physical device). The first register set is dedicated to Host Bridge functionality (i.e., DRAM configuration, other chip-set operating parameters and optional features). The second register block is dedicated to Host-PCI Express Bridge functions (controls PCI Express interface configurations and operating parameters). The third register block is for the internal graphics functions.

The processor internal registers (I/O Mapped, Configuration and PCI Express Extended Configuration registers) are accessible by the Host processor. The registers that reside within the lower 256 bytes of each device can be accessed as Byte, Word (16 bit), or Dword (32 bit) quantities, with the exception of CONFIG\_ADDRESS, which can only be accessed as a Dword. All multi-byte numeric fields use "little-endian" ordering (i.e., lower addresses contain the least significant parts of the field). Registers which reside in bytes 256 through 4095 of each device may only be accessed using memory mapped transactions in Dword (32 bit) quantities.

Some of the processor registers described in this section contain reserved bits. These bits are labeled "Reserved". Software must deal correctly with fields that are reserved. On reads, software must use appropriate masks to extract the defined bits and not rely on reserved bits being any particular value. On writes, software must ensure that the values of reserved bit positions are preserved. That is, the values of reserved bit positions must first be read, merged with the new values for other bit positions and then written back. Note the software does not need to perform read, merge, and write operation for the Configuration Address Register.



In addition to reserved bits within a register, the processor contains address locations in the configuration space of the Host Bridge entity that are marked either “Reserved” or “Intel Reserved”. The processor responds to accesses to Reserved address locations by completing the host cycle. When a Reserved register location is read, a zero value is returned. (Reserved registers can be 8, 16, or 32 bits in size). Writes to Reserved registers have no effect on the processor. Registers that are marked as Intel Reserved must not be modified by system software. Writes to Intel Reserved registers may cause system failure. Reads from Intel Reserved registers may return a non-zero value.

Upon a Full Reset, the processor sets its entire set of internal configuration registers to predetermined default states. Some register values at reset are determined by external strapping options. The default state represents the minimum functionality feature set required to successfully bringing up the system. Hence, it does not represent the optimal system configuration. It is the responsibility of the system initialization software (usually BIOS) to properly determine the DRAM configurations, operating parameters and optional system features that are applicable, and to program the processor registers accordingly.

## 1.7 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.

## 1.8 PCI Device 0

The DRAM Controller registers are in Device 0 (D0), Function 0 (F0).

**Warning:** Address locations that are not listed are considered Intel Reserved registers locations. Reads to Reserved registers may return non-zero values. Writes to reserved locations may cause system failures.

All registers that are defined in the latest *PCI Local Bus Specification*, but are not necessary or implemented in this component are simply not included in this document. The reserved/un-implemented space in the PCI configuration header space is not documented as such in this summary.



Table 3. Device 0 Function 0 Register Summary (Sheet 1 of 2)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Vendor Identification	VID	0	1	8086h	RO
Device Identification	DID	2	3	0044h	RO
PCI Command	PCICMD	4	5	0006h	RO; RW
PCI Status	PCISTS	6	7	0090h	RWC; RO
Revision Identification	RID	8	8	12h	RO
Class Code	CC	9	B	060000h	RO
Master Latency Timer	MLT	D	D	00h	RO
Header Type	HDR	E	E	00h	RO
Subsystem Vendor Identification	SVID	2C	2D	0000h	RW-O
Subsystem Identification	SID	2E	2F	0000h	RW-O
Capabilities Pointer	CAPPTR	34	34	E0h	RO
PCI Express Egress Port Base Address	PXPEPBAR	40	47	0000000000000000h	RW-L; RO
Processor Memory Mapped Register Range Base	MCHBAR	48	4F	0000000000000000h	RW-L; RO
Processor Graphics Control Register	GGC	52	53	0030h	RW-L; RO
Device Enable	DEVEN	54	57	0000210Bh	RW-L; RO; RW
Root Complex Register Range Base Address	DMIBAR	68	6F	0000000000000000h	RW-L; RO
Legacy Access Control	LAC	97	97	00h	RW
Remap Base Address Register	REMAPBASE	98	99	03FFh	RO; RW-L
Remap Limit Address Register	REMAPLIMIT	9A	9B	0000h	RO; RW-L
Top of Memory	TOM	A0	A1	0001h	RO; RW-L
Top of Upper Usable DRAM	TOUUD	A2	A3	0000h	RW-L
Graphics Base of Stolen Memory	GBSM	A4	A7	00000000h	RW-L; RO
Base of GTT stolen Memory	BGSM	A8	AB	00000000h	RW-L; RO
TSEG Memory Base	TSEGMB	AC	AF	00000000h	RO; RW-L
Top of Low Usable DRAM	TOLUD	B0	B1	0010h	RW-L; RO
Primary Buffer Flush Control	PBFC	C0	C3	00000000h	RO; W
Secondary Buffer Flush Control	SBFC	C4	C7	00000000h	RO; W
Error Status	ERRSTS	C8	C9	0000h	RO; RWC-S
Error Command	ERRCMD	CA	C8	0000h	RO; RW
SMI Command	SMICMD	CC	CD	0000h	RO; RW



**Table 3. Device 0 Function 0 Register Summary (Sheet 2 of 2)**

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
SCI Command	SCICMD	CE	CF	0000h	RO; RW
Scratchpad Data	SKPD	DC	DF	00000000h	RW
Capability Identifier	CAPID0	E0	EB	00000002000 00000010C0 009	RO

### 1.8.1 VID - Vendor Identification

B/D/F/Type: 0/0/0/PCI  
 Address Offset: 0-1h  
 Default Value: 8086h  
 Access: RO  
 Size: 16 bits

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

Bit	Access	Default Value	Description
15:0	RO	8086h	<b>Vendor Identification Number (VID)</b> PCI standard identification for Intel.

### 1.8.2 DID - Device Identification

B/D/F/Type: 0/0/0/PCI  
 Address Offset: 2-3h  
 Default Value: 0044h  
 Access: RO  
 Size: 16 bits

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

Bit	Access	Default Value	Description
15:0	RO	0044h	<b>Device Identification Number (DID)</b> Identifier assigned to the processor core/primary PCI device.





### 1.8.3 PCICMD - PCI Command

B/D/F/Type: 0/0/0/PCI  
 Address Offset: 4-5h  
 Default Value: 0006h  
 Access: RO; RW  
 Size: 16 bits

Since processor Device 0 does not physically reside on PCI\_A many of the bits are not implemented.

(Sheet 1 of 2)

Bit	Access	Default Value	Description						
15:10	RO	00h	<b>Reserved</b>						
9	RO	0b	<p><b>Fast Back-to-Back Enable (FB2B)</b>                      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 hard wired to 0. Writes to this bit position have no effect.</p>						
8	RW	0b	<p><b>SERR Enable (SERRE)</b>                      This bit is a global enable bit for Device 0 SERR messaging. The processor does not have an SERR signal. The processor communicates the SERR condition by sending an SERR message over DMI to the PCH.                      0 = The SERR message is not generated by the processor for Device 0.                      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.                      This bit only controls SERR messaging for Device 0. Device 1 has its own SERRE bits to control error reporting for error conditions occurring in that device. The control bits are used in a logical OR manner to enable the SERR DMI message mechanism.</p> <table border="1" data-bbox="787 1396 1274 1522"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0b</td> <td>Device 0 SERR disabled</td> </tr> <tr> <td>1b</td> <td>Device 0 SERR enabled</td> </tr> </tbody> </table>	Encoding	Description	0b	Device 0 SERR disabled	1b	Device 0 SERR enabled
Encoding	Description								
0b	Device 0 SERR disabled								
1b	Device 0 SERR enabled								
7	RO	0b	<p><b>Address/Data Stepping Enable (ADSTEP)</b> Address/data stepping is not implemented in the processor, and this bit is hard wired to 0. Writes to this bit position have no effect.</p>						



(Sheet 2 of 2)

Bit	Access	Default Value	Description						
6	RW	0b	<p><b>Parity Error Enable (PERRE)</b>            Controls whether or not the Master Data Parity Error bit in the PCI Status register can be set.            0 = Master Data Parity Error bit in PCI Status register <b>cannot</b> be set.            1 = Master Data Parity Error bit in PCI Status register CAN be set.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0b</td> <td>Master Data Parity Error <b>cannot</b> be set</td> </tr> <tr> <td>1b</td> <td>Master Data Parity Error can be set</td> </tr> </tbody> </table>	Encoding	Description	0b	Master Data Parity Error <b>cannot</b> be set	1b	Master Data Parity Error can be set
Encoding	Description								
0b	Master Data Parity Error <b>cannot</b> be set								
1b	Master Data Parity Error can be set								
5	RO	0b	<p><b>VGA Palette Snoop Enable (VGASNOOP)</b>            The processor does not implement this bit and it is hard wired to a 0. Writes to this bit position have no effect.</p>						
4	RO	0b	<p><b>Memory Write and Invalidate Enable (MWIE)</b>            The processor will never issue memory write and invalidate commands. This bit is therefore hard wired to 0. Writes to this bit position will have no effect.</p>						
3	RO	0b	<p><b>Special Cycle Enable (SCE)</b>            The processor does not implement this bit and it is hard wired to a 0. Writes to this bit position have no effect.</p>						
2	RO	1b	<p><b>Bus Master Enable (BME)</b>            The processor is always enabled as a master on the backbone. This bit is hard wired to a "1". Writes to this bit position have no effect.</p>						
1	RO	1b	<p><b>Memory Access Enable (MAE)</b>            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 hard wired to 1. Writes to this bit position have no effect.</p>						
0	RO	0b	<p><b>I/O Access Enable (IOAE)</b>            This bit is not implemented in the processor and is hard wired to a 0. Writes to this bit position have no effect.</p>						



### 1.8.4 PCISTS - PCI Status

B/D/F/Type: 0/0/0/PCI  
 Address Offset: 6-7h  
 Default Value: 0090h  
 Access: RWC; RO  
 Size: 16 bits

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

(Sheet 1 of 2)

Bit	Access	Default Value	Description
15	RWC	0b	<b>Detected Parity Error (DPE)</b> This bit is set when this Device receives a Poisoned TLP.
14	RWC	0b	<b>Signaled System Error (SSE)</b> This bit is set to 1 when the processor 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.
13	RWC	0b	<b>Received Master Abort Status (RMAS)</b> 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.
12	RWC	0b	<b>Received Target Abort Status (RTAS)</b> 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.
11	RO	0b	<b>Signaled Target Abort Status (STAS)</b> The processor will not generate a Target Abort DMI completion packet or Special Cycle. This bit is not implemented in the processor and is hard wired to a 0. Writes to this bit position have no effect.
10:9	RO	00b	<b>DEVSEL Timing (DEVT)</b> These bits are hard wired 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 processor.
8	RWC	0b	<b>Master Data Parity Error Detected (DPD)</b> 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.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
7	RO	1b	<b>Fast Back-to-Back (FB2B)</b> This bit is hard wired 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 processor.
6	RO	0b	<i>Reserved</i>
5	RO	0b	<b>66-MHz Capable (66MC)</b> Does not apply to PCI Express. Must be hard wired to 0.
4	RO	1b	<b>Capability List (CLIST)</b> This bit is hard wired 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.
3	RO	0b	<i>Reserved</i>
2:0	RO	000b	<i>Reserved</i>

### 1.8.5 RID - Revision Identification

B/D/F/Type: 0/0/0/PCI  
Address Offset: 8h  
Default Value: 12h  
Access: RO  
Size: 8 bits

This register contains the revision number of the processor. The Revision ID (RID) is a traditional 8-bit Read Only (RO) register located at offset 08h in the standard PCI header of every PCI/PCI Express compatible device and function.

Following reset, the SRID is returned when the RID is read at offset 08h. The SRID value reflects the actual product stepping. To select the CRID value, BIOS/configuration software writes a key value of 69h to Bus 0, Device 0, Function 0 (DMI device) of the CPU's RID register at offset 08h. This causes the CRID to be returned when the RID is read at offset 08h.

#### Stepping Revision ID (SRID)

This register contains the revision number of the CPU.

The SRID is a 8-bit hardwired value assigned by Intel, based on product's stepping. The SRID is not a directly addressable PCI register. The SRID value is reflected through the RID register when appropriately addressed.



### Compatible Revision ID (CRID)

The CRID is an 8-bit hardwired value assigned by Intel during manufacturing process. Normally, the value assigned as the CRID will be identical to the SRID value of a previous stepping of the product with which the new product is deemed “compatible”.

The CRID is not a directly addressable PCI register. The CRID value is reflected through the RID register when appropriately addressed.

Bit	Access	Default Value	Description
7:0	RO	10h	<b>Revision Identification Number (RID)</b> This is an 8-bit value that indicates the revision identification number for the processor Device 0. For the C-2 Stepping, these values are: SRID = 12h CRID = 02h

### 1.8.6 CC - Class Code

B/D/F/Type: 0/0/0/PCI  
 Address Offset: 9-Bh  
 Default Value: 060000h  
 Access: RO  
 Size: 24 bits

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

Bit	Access	Default Value	Description
23:16	RO	06h	<b>Base Class Code (BCC)</b> This is an 8-bit value that indicates the base class code for the processor. This code has the value 06h, indicating a Bridge device.
15:8	RO	00h	<b>Sub-Class Code (SUBCC)</b> This is an 8-bit value that indicates the category of Bridge into which the processor falls. The code is 00h indicating a Host Bridge.
7:0	RO	00h	<b>Programming Interface (PI)</b> 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.



### 1.8.7 MLT - Master Latency Timer

B/D/F/Type: 0/0/0/PCI  
Address Offset: Dh  
Default Value: 00h  
Access: RO  
Size: 8 bits

Device 0 in the processor is not a PCI master. Therefore this register is not implemented.

Bit	Access	Default Value	Description
7:0	RO	00h	Reserved

### 1.8.8 HDR - Header Type

B/D/F/Type: 0/0/0/PCI  
Address Offset: Eh  
Default Value: 00h  
Access: RO  
Size: 8 bits

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

Bit	Access	Default Value	Description
7:0	RO	00h	PCI Header (HDR) This field always returns 0 to indicate that the processor is a single function device with standard header layout. Reads and writes to this location have no effect.

### 1.8.9 SVID - Subsystem Vendor Identification

B/D/F/Type: 0/0/0/PCI  
Address Offset: 2C-2Dh  
Default Value: 0000h  
Access: RW-O  
Size: 16 bits

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

Bit	Access	Default Value	Description
15:0	RW-O	0000h	<b>Subsystem Vendor ID (SUBVID)</b> 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.



### 1.8.10 SID - Subsystem Identification

B/D/F/Type: 0/0/0/PCI  
 Address Offset: 2E-2Fh  
 Default Value: 0000h  
 Access: RW-O  
 Size: 16 bits

This value is used to identify a particular subsystem.

Bit	Access	Default Value	Description
15:0	RW-O	0000h	<b>Subsystem ID (SUBID)</b> This field should be programmed during BIOS initialization. After it has been written once, it becomes read only.

### 1.8.11 CAPPTR - Capabilities Pointer

B/D/F/Type: 0/0/0/PCI  
 Address Offset: 34h  
 Default Value: E0h  
 Access: RO  
 Size: 8 bits

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

Bit	Access	Default Value	Description
7:0	RO	E0h	<b>Capabilities Pointer (CAPPTR)</b> Pointer to the offset of the first capability ID register block. In this case the first capability is the product-specific Capability Identifier (CAPIDO).



### 1.8.12 PXPEPBAR - PCI Express Egress Port Base Address

B/D/F/Type: 0/0/0/PCI  
 Address Offset: 40-47h  
 Default Value: 0000000000000000h  
 Access: RW-L; RO  
 Size: 64 bits

This is the base address for the PCI Express Egress Port MMIO Configuration space. There is no physical memory within this 4-KB window that can be addressed. The 4 KB reserved by this register does not alias to any PCI 3.0-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.

Bit	Access	Default Value	Description						
63:36	RO	0000000h	<i>Reserved</i>						
35:12	RW-L	000000h	<p><b>PCI Express Egress Port MMIO Base Address (PXPEPBAR)</b></p> <p>This field corresponds to Bits 35:12 of the base address PCI Express Egress Port MMIO configuration space. BIOS will program this register resulting in a base address for a 4-KB block of contiguous memory address space. This register ensures that a naturally aligned 4-KB space is allocated within the first 64 GB of addressable memory space. System Software uses this base address to program the processor MMIO register set. All the bits in this register are locked in Intel VT-d mode.</p>						
11:1	RO	000h	<i>Reserved</i>						
0	RW-L	0b	<p><b>PXPEPBAR Enable (PXPEPBAREN)</b></p> <p>0 = PXPEPBAR is disabled and does not claim any memory            1 = PXPEPBAR memory mapped accesses are claimed and decoded appropriately</p> <p>This register is locked by Intel VT-d.</p> <table border="1" data-bbox="760 1417 1242 1549"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0b</td> <td>PXPEPBAR is disabled</td> </tr> <tr> <td>1b</td> <td>PXPEPBAR is enabled</td> </tr> </tbody> </table>	Encoding	Description	0b	PXPEPBAR is disabled	1b	PXPEPBAR is enabled
Encoding	Description								
0b	PXPEPBAR is disabled								
1b	PXPEPBAR is enabled								





### 1.8.13 MCHBAR - Processor Memory Mapped Register Range Base

B/D/F/Type: 0/0/0/PCI  
 Address Offset: 48-4Fh  
 Default Value: 0000000000000000h  
 Access: RW-L; RO  
 Size: 64 bits

This is the base address for the processor Memory Mapped Configuration space. There is no physical memory within this 16-KB window that can be addressed. The 16-KB reserved by this register does not alias to any PCI 3.0-compliant memory mapped space. On reset, the processor MMIO Memory Mapped Configuration space is disabled and must be enabled by writing a 1 to MCHBAREN [Device 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. The 16KB space reserved by the MCHBAR register is not accessible during Intel TXT mode of operation or if the Intel® Management Engine (Intel® ME) security lock is asserted (MESMLCK.ME\_SM\_lock at PCI Device 0, Function 0, Offset F4h) except for the following offset ranges.

- 02B8h to 02BFh: Channel 0 Throttle Counter Status Registers
- 06B8h to 06BFh: Channel 1 Throttle Counter Status Registers
- 0CD0h to 0CFFh: Thermal Sensor Control Registers
- 3000h to 3FFFh: Unlocked registers for future expansion

(Sheet 1 of 2)

Bit	Access	Default Value	Description
63:36	RO	0000000h	<i>Reserved</i>
35:14	RW-L	000000h	<b>Processor Memory Mapped Base Address (MCHBAR)</b> This field corresponds to Bits 35:14 of the base address processor Memory Mapped configuration space. BIOS will program this register resulting in a base address for a 16-KB block of contiguous memory address space. This register ensures that a naturally aligned 16-KB space is allocated within the first 64 GB of addressable memory space. System Software uses this base address to program the processor Memory Mapped register set. All the bits in this register are locked in Intel VT-d mode.
13:1	RO	0000h	<i>Reserved</i>



(Sheet 2 of 2)

Bit	Access	Default Value	Description						
0	RW-L	0b	<b>MCHBAR Enable (MCHBAREN)</b> 0 = MCHBAR is disabled and does not claim any memory 1 = MCHBAR memory mapped accesses are claimed and decoded appropriately This register is locked by Intel TXT. <table border="1"><thead><tr><th>Encoding</th><th>Description</th></tr></thead><tbody><tr><td>0b</td><td>MCHBAR is disabled</td></tr><tr><td>1b</td><td>MCHBAR is enabled</td></tr></tbody></table>	Encoding	Description	0b	MCHBAR is disabled	1b	MCHBAR is enabled
Encoding	Description								
0b	MCHBAR is disabled								
1b	MCHBAR is enabled								

### 1.8.14 GGC - Processor Graphics Control Register

B/D/F/Type: 0/0/0/PCI  
Address Offset: 52-53h  
Default Value: 0030h  
Access: RW-L; RO  
Size: 16 bits  
All the bits in this register are Intel TXT lockable.

(Sheet 1 of 4)

Bit	Access	Default Value	Description
15:12	RO	0h	<i>Reserved</i>




(Sheet 2 of 4)

Bit	Access	Default Value	Description														
11:8	RW-L	0h	<p><b>GTT Graphics Memory Size (GGMS)</b></p> <p>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.</p> <p>GSM is assumed to be a contiguous physical DRAM space with DSM, and BIOS needs to allocate a contiguous memory chunk. Hardware will drive the base of GSM from DSM only using the GSM size programmed in the register.</p> <p>0h: No memory pre-allocated. GTT cycles (Mem and IO) are not claimed.</p> <p>1h: No Intel® Virtualization Technology (Intel® VT-d) mode, 1 MB of memory pre-allocated for GTT.</p> <p>3h: No Intel VT-d mode, 2 MB of memory pre-allocated for GTT.</p> <p>9h: Intel VT-d mode, 2 MB of memory pre-allocated for 1 MB of Global GTT and 1 MB for Shadow GTT.</p> <p>Ah: Intel VT-d mode, 3 MB of memory pre-allocated for 1.5 MB of Global GTT and 1.5 MB for Shadow GTT.</p> <p>Bh: Intel VT-d mode, 4 MB of memory pre-allocated for 2 MB of Global GTT and 2 MB for Shadow GTT.</p> <p>All unspecified encoding of this register field are reserved, hardware functionality is not guaranteed if used. This register is locked and becomes Read Only when the D_LCK bit in the SMRAM register is set.</p> <table border="1" data-bbox="773 1163 1268 1507"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>3h</td> <td>No Intel VT -d mode, 2 MB</td> </tr> <tr> <td>1h</td> <td>No Intel VT -d mode, 1 MB</td> </tr> <tr> <td>0h</td> <td>No memory preallocated</td> </tr> <tr> <td>9h</td> <td>Intel VT -d mode, 2 MB</td> </tr> <tr> <td>Ah</td> <td>Intel VT -d mode, 3 MB</td> </tr> <tr> <td>Bh</td> <td>Intel VT -d mode, 4 MB</td> </tr> </tbody> </table>	Encoding	Description	3h	No Intel VT -d mode, 2 MB	1h	No Intel VT -d mode, 1 MB	0h	No memory preallocated	9h	Intel VT -d mode, 2 MB	Ah	Intel VT -d mode, 3 MB	Bh	Intel VT -d mode, 4 MB
Encoding	Description																
3h	No Intel VT -d mode, 2 MB																
1h	No Intel VT -d mode, 1 MB																
0h	No memory preallocated																
9h	Intel VT -d mode, 2 MB																
Ah	Intel VT -d mode, 3 MB																
Bh	Intel VT -d mode, 4 MB																



(Sheet 3 of 4)

Bit	Access	Default Value	Description																						
7:4	RW-L	3h	<p><b>Graphics Mode Select (GMS)</b></p> <p>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.</p> <p>0h: No memory pre-allocated. 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.</p> <p>1h-4h: Reserved.</p> <p>5h-Dh: DVMT (UMA) mode, memory pre-allocated for frame buffer, in quantities as shown in the Encoding table.</p> <p>Eh-Fh: Reserved.</p> <p><b>Note:</b> This register is locked and becomes Read Only when the D_LCK bit in the SMRAM register is set. This register is also Intel VT-d lockable.</p> <p>Hardware does not clear or set any of these bits automatically based on IGD being disabled/enabled.</p> <p>BIOS Requirement: BIOS must not set this field to 0h if IVD (Bit 1 of this register) is 0.</p> <table border="1" data-bbox="750 1020 1247 1465"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>No memory pre-allocated</td> </tr> <tr> <td>5h</td> <td>32 MB</td> </tr> <tr> <td>6h</td> <td>48 MB</td> </tr> <tr> <td>7h</td> <td>64 MB</td> </tr> <tr> <td>8h</td> <td>128 MB</td> </tr> <tr> <td>9h</td> <td>256 MB</td> </tr> <tr> <td>Ah</td> <td>96 MB</td> </tr> <tr> <td>Bh</td> <td>160 MB</td> </tr> <tr> <td>Ch</td> <td>224 MB</td> </tr> <tr> <td>Dh</td> <td>352 MB</td> </tr> </tbody> </table>	Encoding	Description	0h	No memory pre-allocated	5h	32 MB	6h	48 MB	7h	64 MB	8h	128 MB	9h	256 MB	Ah	96 MB	Bh	160 MB	Ch	224 MB	Dh	352 MB
Encoding	Description																								
0h	No memory pre-allocated																								
5h	32 MB																								
6h	48 MB																								
7h	64 MB																								
8h	128 MB																								
9h	256 MB																								
Ah	96 MB																								
Bh	160 MB																								
Ch	224 MB																								
Dh	352 MB																								
3:2	RO	00b	Reserved																						



(Sheet 4 of 4)

Bit	Access	Default Value	Description						
1	RW-L	0b	<p><b>IGD VGA Disable (IVD)</b></p> <p>0 = Enable. Device 2 (IGD) claims VGA memory and IO cycles, the Sub-Class Code within Device 2 Class Code register is 00.</p> <p>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.</p> <p>BIOS Requirement: BIOS must not set this bit to 0 if the GMS field (Bits 6:4 of this register) pre-allocates no memory. This bit MUST be set to 1 if Device 2 is disabled via register (DEVEN[3] = 0).</p> <p>This register is locked by Intel VT-d.</p> <table border="1" data-bbox="776 779 1256 909"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0b</td> <td>Enable</td> </tr> <tr> <td>1b</td> <td>Disable</td> </tr> </tbody> </table>	Encoding	Description	0b	Enable	1b	Disable
Encoding	Description								
0b	Enable								
1b	Disable								
0	RO	0b	<i>Reserved</i>						



### 1.8.15 DEVEN - Device Enable

B/D/F/Type: 0/0/0/PCI  
 Address Offset: 54-57h  
 Default Value: 0000010Bh  
 Access: RW-L; RO; RW  
 Size: 32 bits  
 BIOS Optimal Default 000000h

Allows for enabling/disabling of PCI devices and functions that are within the processor. 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.

Bit	Access	Default Value	Description
31:15	RO	0h	<i>Reserved</i>
14:13	RW-L	0b	<i>Reserved</i>
12:12	RO	0h	<i>Reserved</i>
11	RW-L	0b	<i>Reserved</i>
10	RW-L	0b	<i>Reserved</i>
9:9	RO	0h	<i>Reserved</i>
8	RW-L	1b	<i>Reserved</i>
7:4	RO	0h	<i>Reserved</i>
3	RW-L	1b	<b>Internal Graphics Engine Function 0 (D2F0EN)</b> 0 = Bus 0 Device 2 Function 0 is disabled and hidden 1 = Bus 0 Device 2 Function 0 is enabled and visible
2:2	RO	0h	<i>Reserved</i>
1	RW-L	1b	<b>PCI Express Port (D1EN)</b> 0 = Bus 0 Device 1 Function 0 is disabled and hidden. 1 = Bus 0 Device 1 Function 0 is enabled and visible.
0	RO	1b	<b>Host Bridge (D0EN)</b> Bus 0 Device 0 Function 0 may not be disabled and is therefore hard wired to 1.



### 1.8.16 DMIBAR - Root Complex Register Range Base Address

B/D/F/Type: 0/0/0/PCI  
 Address Offset: 68-6Fh  
 Default Value: 0000000000000000h  
 Access: RW-L; RO  
 Size: 64 bits

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 processor. There is no physical memory within this 4-KB window that can be addressed. The 4-KB reserved by this register does not alias to any PCI 3.0-compliant memory mapped space. On reset, the Root Complex configuration space is disabled and must be enabled by writing a 1 to DMIBAREN [Device 0, Offset 68h, Bit 0] All the bits in this register are locked in Intel TXT mode.

Bit	Access	Default Value	Description
63:36	RO	0000000h	<i>Reserved (DMIBAR_rsv)</i>
35:12	RW-L	000000h	<b>DMI Base Address (DMIBAR)</b> This field corresponds to Bits 35:12 of the base address DMI configuration space. BIOS will program this register resulting in a base address for a 4-KB block of contiguous memory address space. This register ensures that a naturally aligned 4-KB space is allocated within the first 64 GB 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 VT-d mode.
11:1	RO	000h	<i>Reserved</i>
0	RW-L	0b	<b>DMIBAR Enable (DMIBAREN)</b> 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 VT-d.

Encoding	Description
0b	DMIBAR disabled
1b	DMIBAR enabled



### 1.8.17 LAC - Legacy Access Control

B/D/F/Type: 0/0/0/PCI  
Address Offset: 97h  
Default Value: 00h  
Access: RW  
Size: 8 bits  
BIOS Optimal Default 00h

This 8-bit register controls steering of MDA cycles.

There can only be at most one MDA device in the system. BIOS must not program bits 1:0 to 11b.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
7:2	RO	0h	<i>Reserved</i>
1	RW	0b	<i>Reserved</i>





(Sheet 2 of 2)

Bit	Access	Default Value	Description																					
0	RW	0b	<p><b>PEG0 MDA Present (MDAPO)</b></p> <p>This bit works with the VGA Enable bits in the BCTRL register of Device 1 to control the routing of CPU initiated transactions targeting MDA compatible I/O and memory address ranges. This bit should not be set if device 1's VGA Enable bit is not set.</p> <p>If Device 1's VGA enable bit is not set, then accesses to I/O address range x3BCh-x3BFh remain on the backbone.</p> <p>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 if the address is within the corresponding IOBASE and IOLIMIT, otherwise they remain on the backbone.</p> <p>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)</p> <p>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.</p> <p>The following table shows the behavior for all combinations of MDA and VGA:</p> <table border="1"> <thead> <tr> <th>VGAEN</th> <th>MDAP</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>All References to MDA and VGA space are not claimed by Device 1</td> </tr> <tr> <td>0</td> <td>1</td> <td>Illegal combination</td> </tr> <tr> <td>1</td> <td>0</td> <td>All VGA and MDA references are routed to an external graphics device attached to PCI Express Device 1</td> </tr> <tr> <td>1</td> <td>1</td> <td>All VGA references are routed to an external graphics device attached to PCI Express Device 1. MDA references are not claimed by Device 1.</td> </tr> </tbody> </table> <p>VGA and MDA memory cycles can only be routed across PEG0 when MAE (PCICMD1[1]) is set. VGA and MDA I/O cycles can only be routed across PEG0 if IOAE (PCICMD1[0]) is set.</p> <table border="1"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0b</td> <td>No MDA</td> </tr> <tr> <td>1b</td> <td>MDA Present</td> </tr> </tbody> </table>	VGAEN	MDAP	Description	0	0	All References to MDA and VGA space are not claimed by Device 1	0	1	Illegal combination	1	0	All VGA and MDA references are routed to an external graphics device attached to PCI Express Device 1	1	1	All VGA references are routed to an external graphics device attached to PCI Express Device 1. MDA references are not claimed by Device 1.	Encoding	Description	0b	No MDA	1b	MDA Present
VGAEN	MDAP	Description																						
0	0	All References to MDA and VGA space are not claimed by Device 1																						
0	1	Illegal combination																						
1	0	All VGA and MDA references are routed to an external graphics device attached to PCI Express Device 1																						
1	1	All VGA references are routed to an external graphics device attached to PCI Express Device 1. MDA references are not claimed by Device 1.																						
Encoding	Description																							
0b	No MDA																							
1b	MDA Present																							



### 1.8.18 REMAPBASE - Remap Base Address Register

B/D/F/Type: 0/0/0/PCI  
Address Offset: 98-99h  
Default Value: 03FFh  
Access: RO; RW-L  
Size: 16 bits

Bit	Access	Default Value	Description
15:10	RO	000000b	<i>Reserved</i>
9:0	RW-L	3FFh	<b>Remap Base Address [35 26] (REMAPBASE)</b> 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[25:0] of the Remap Base Address are assumed to be 0's. Thus the bottom of the defined memory range is aligned to a 64-MB 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 VT-d lockable or Intel ME stolen Memory lockable.

### 1.8.19 REMAPLIMIT - Remap Limit Address Register

B/D/F/Type: 0/0/0/PCI  
Address Offset: 9A-9Bh  
Default Value: 0000h  
Access: RO; RW-L  
Size: 16 bits

Bit	Access	Default Value	Description
15:10	RO	000000b	<i>Reserved</i>
9:0	RW-L	000h	<b>Remap Limit Address [35 26] (REMAPLMT)</b> 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[25:0] of the remap limit address are assumed to be F. Thus the top of the defined range is one byte less than a 64-MB boundary. 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 VT-d lockable or Intel ME stolen Memory lockable.



### 1.8.20 TOM - Top of Memory

B/D/F/Type: 0/0/0/PCI  
 Address Offset: A0-A1h  
 Default Value: 0001h  
 Access: RO; RW-L  
 Size: 16 bits

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

Bit	Access	Default Value	Description
15:10	RO	00h	<i>Reserved</i>
9:0	RW-L	001h	<p><b>Top of Memory (TOM)</b></p> <p>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 35:26 (64-MB granularity). Bits 25:0 are assumed to be 0. All the bits in this register are locked in Intel VT-d mode.</p> <p>MCH determines the base of EP stolen memory by subtracting the EP stolen memory size from TOM.</p>



### 1.8.21 TOUUD - Top of Upper Usable DRAM

B/D/F/Type: 0/0/0/PCI  
Address Offset: A2-A3h  
Default Value: 0000h  
Access: RW-L  
Size: 16 bits

This 16-bit register defines the Top of Upper Usable DRAM. Configuration software must set this value to TOM minus all EP stolen memory aligned down to a 64-MB boundary. If reclaim is enabled, this value must be set to reclaim limit 64-MB aligned, since reclaim limit + 1 byte is 64-MB 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 4 GB.

These bits are Intel TXT lockable.

Bit	Access	Default Value	Description
15:0	RW-L	0000h	<b>TOUUD (TOUUD)</b> This register contains Bits 35:20 of an address one byte above the maximum DRAM memory above 4 GB that is usable by the operating system. Configuration software must set this value to TOM minus all EP stolen memory if reclaim is disabled. If reclaim is enabled, this value must be set to reclaim limit 64-MB aligned since reclaim limit + 1 byte is 64-MB 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 4 GB. All the bits in this register are locked in Intel VT-d mode.



### 1.8.22 GBSM - Graphics Base of Stolen Memory

B/D/F/Type: 0/0/0/PCI  
 Address Offset: A4-A7h  
 Default Value: 00000000h  
 Access: RW-L; RO  
 Size: 32 bits

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 B0 Bits 15:04).

This register is locked and becomes Read Only when CMD.LOCK.MEMCONFIG is received or when ME\_SM\_LOCK is set to 1.

Bit	Access	Default Value	Description
31:20	RW-L	000h	<b>Graphics Base of Stolen Memory (GBSM)</b> This register contains Bits 31:20 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 52 Bits 6:4) from TOLUD (PCI Device 0 Offset B0 Bits 15:04).
19:0	RO	00000h	<i>Reserved</i>



### 1.8.23 BGSM - Base of GTT Stolen Memory

B/D/F/Type: 0/0/0/PCI  
Address Offset: A8-ABh  
Default Value: 00000000h  
Access: RW-L; RO  
Size: 32 bits

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 where TOLUD (PCI Device 0 Offset B0 Bits 15:04) would be if there were no memory reserved for TSEG, internal graphics, or the GTT for internal graphics.

This register is locked and becomes Read Only when CMD.LOCK. MEMCONFIG is received or when ME\_SM\_LOCK is set to 1.

Bit	Access	Default Value	Description
31:20	RW-L	000h	<b>Graphics Base of Stolen Memory (GBSM)</b> This register contains Bits 31:20 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 52 Bits 9:8) from the graphics stolen memory base (PCI Device 0 Offset A4 Bits 31:20).
19:0	RO	00000h	<i>Reserved</i>



### 1.8.24 TSEGMB - TSEG Memory Base

B/D/F/Type: 0/0/0/PCI  
 Address Offset: AC-AFh  
 Default Value: 00000000h  
 Access: RO; RW-L  
 Size: 32 bits

This register contains the base address of TSEG DRAM memory. BIOS determines the base of TSEG memory which must be at or below graphics GTT stolen base (PCI Device 0 Offset A8 Bits 31:20).

Once D\_LCK has been set, these bits becomes read only.

Bit	Access	Default Value	Description
31:20	RW-L	000h	<b>TSEG Memory Base (TSEGMB)</b> This register contains Bits 31:20 of the base address of TSEG DRAM memory. BIOS determines the base of TSEG memory by subtracting the TSEG size (located within the CPU register space) from graphics GTT stolen base (PCI Device 0 Offset A8 Bits 31:20). Once D_LCK has been set, these bits become read only.
19:0	RO	00000h	<i>Reserved</i>

### 1.8.25 TOLUD - Top of Low Usable DRAM

B/D/F/Type: 0/0/0/PCI  
 Address Offset: B0-B1h  
 Default Value: 0010h  
 Access: RW-L; RO  
 Size: 16 bits

This 16-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, processor optionally claims 1 to 64 MBs of DRAM for internal graphics if enabled, 1 or 2 MB 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 4 GB
- TSEG is enabled and TSEG size is set to 1 MB
- Internal Graphics is enabled, and Graphics Mode Select is set to 32 MB
- GTT Graphics Stolen Memory Size set to 2 MB
- BIOS knows the OS requires 1 GB of PCI space.
- BIOS also knows the range from 0\_FEC0\_0000h to 0\_FFFF\_FFFFh is not usable by the system. This 20-MB 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: 4 GB = 1\_0000\_0000h

The system memory requirements are: 4 GB (max addressable space) – 1 GB (PCI space) – 35 MB (lost memory) = 3 GB – 35 MB (minimum granularity) = 0\_ECBO\_0000h

Since 0\_ECBO\_0000h (PCI and other system requirements) is less than 1\_0000\_0000h, TOLUD should be programmed to ECBh.

These bits are Intel TXT lockable.

Bit	Access	Default Value	Description
15:4	RW-L	001h	<p><b>Top of Low Usable DRAM (TOLUD)</b></p> <p>This register contains Bits 31:20 of an address one byte above the maximum DRAM memory below 4 GB that is usable by the operating system. Address Bits 31 down to 20 programmed to 01h implies a minimum memory size of 1 MB.</p> <p>Configuration software must set this value to the smaller of the following two choices: Maximum amount memory in the system minus Intel ME stolen memory plus one byte or the minimum address allocated for PCI memory.</p> <p>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.</p> <p>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 VT-d mode.</p> <p>This register must be 64-MB aligned when reclaim is enabled.</p>
3:0	RO	0000b	<i>Reserved</i>





### 1.8.26 PBFC - Primary Buffer Flush Control

B/D/F/Type: 0/0/0/PCI  
 Address Offset: C0-C3h  
 Default Value: 00000000h  
 Access: RO; W  
 Size: 32 bits

Bit	Access	Default Value	Description
31:1	RO	00000000h	<i>Reserved</i> These bits must remain reserved to avoid a potential hang condition.
0	W	0b	<b>Primary CWB Flush Control (PCWBFLSH)</b> A CPU write to this bit flushes the PCWB of all writes. The data associated with the write to this register is discarded. Note: The write completion indication can be returned before the flush has finished

### 1.8.27 SBFC - Secondary Buffer Flush Control

B/D/F/Type: 0/0/0/PCI  
 Address Offset: C4-C7h  
 Default Value: 00000000h  
 Access: RO; W  
 Size: 32 bits

Bit	Access	Default Value	Description
31:1	RO	00000000h	<i>Reserved</i> These bits must remain reserved to avoid a potential hang condition.
0	W	0b	<b>Secondary CWB Flush Control (SCWBFLSH)</b> A CPU write to this bit flushes the SCWB of all writes. The data associated with the write to this register is discarded.



### 1.8.28 ERRSTS - Error Status

B/D/F/Type: 0/0/0/PCI  
 Address Offset: C8-C9h  
 Default Value: 0000h  
 Access: RO; RWC-S  
 Size: 16 bits  
 BIOS Optimal Default 0h

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.

Bit	Access	Default Value	Description
15:13	RO	000b	<i>Reserved</i>
12	RWC-S	0b	<b>Device 2 Software Generated Event for SMI (GSGESMI)</b> This indicates the source of the SMI was a Device 2 Software Event.
11	RWC-S	0b	<b>Memory Controller Thermal Sensor Event for SMI / SCI / SERR (GTSE)</b> Indicates that a Memory controller Thermal Sensor trip has occurred and an SMI, SCI or SERR has been generated. The status bit is set only if a message is sent based on thermal event enables in Error command, SMI command and SCI command registers. A trip point can generate one of SMI, SCI, or SERR interrupts (two or more per event is illegal). Multiple trip points can generate the same interrupt, if software chooses this mode, subsequent trips may be lost. If this bit is already set, then an interrupt message will not be sent on a new thermal sensor event.
10	RO	0b	<i>Reserved</i>
9	RWC-S	0b	<b>LOCK to non-DRAM Memory Flag (LCKF)</b> When this bit is set to 1, the memory controller has detected a lock operation to memory space that did not map into DRAM.
8	RO	0b	<i>Reserved</i>
7	RWC-S	0b	<b>DRAM Throttle Flag (DTF)</b> 0 = Software has cleared this flag since the most recent throttling event. 1 = Indicates that a DRAM Throttling condition occurred.
6:2	RO	00h	<i>Reserved</i>
1	RWC-S	0b	<i>Reserved</i>
0	RWC-S	0b	<i>Reserved</i>



### 1.8.29 ERRCMD - Error Command

B/D/F/Type: 0/0/0/PCI  
 Address Offset: CA-CBh  
 Default Value: 0000h  
 Access: RO; RW  
 Size: 16 bits

This register controls the memory controller responses to various system errors. Since the processor does not have an SERRB signal, SERR messages are passed from the Processor to the PCH over DMI.

When a bit in this register is set, a SERR message is 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.

Bit	Access	Default Value	Description
15:12	RO	0h	<i>Reserved</i>
11	RW	0b	<b>SERR on Memory Controller Thermal Sensor Event (TSESERR)</b> 0 = Reporting of this condition via SERR messaging is disabled. 1 = The Memory controller generates a DMI SERR special cycle when Bit 11 of the ERRSTS is set. The SERR must not be enabled at the same time as the SMI for the same thermal sensor event.
10	RO	0b	<i>Reserved</i>
9	RW	0b	<b>SERR on LOCK to non-DRAM Memory (LCKERR)</b> 0 = Reporting of this condition via SERR messaging is disabled 1 = The memory controller will generate a DMI SERR special cycle whenever a CPU lock cycle is detected that does not hit DRAM.
8	RW	0b	<i>Reserved</i>
7	RW	0b	<b>SERR on DRAM Throttle Condition (ERR)</b> 0 = Reporting of this condition via SERR messaging is disabled. 1 = The memory controller generates a DMI SERR special cycle when a DRAM Read or Write Throttle condition occurs.
6:2	RO	00h	<i>Reserved</i>
1	RW	0b	<i>Reserved</i>
0	RW	0b	<i>Reserved</i>



### 1.8.30 SMICMD - SMI Command

B/D/F/Type: 0/0/0/PCI  
Address Offset: CC-CDh  
Default Value: 0000h  
Access: RO; RW  
Size: 16 bits

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.

Bit	Access	Default Value	Description
15:12	RO	0h	<i>Reserved</i>
11	RW	0b	<b>SMI on GMCH Thermal Sensor Trip (TSTSMI)</b> 1 = A SMI DMI special cycle is generated by GMCH when the thermal sensor trip requires an SMI. A thermal sensor trip point cannot generate more than one special cycle. 0 = Reporting of this condition via SMI messaging is disabled.
10:2	RO	000h	<i>Reserved</i>
1	RW	0b	<b>SMI on Multiple-Bit DRAM ECC Error (DMESMI)</b> 1 = The GMCH 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.
0	RW	0b	<b>SMI on Single-bit ECC Error (DSESMI)</b> 1 = The GMCH 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.



### 1.8.31 SCICMD - SCI Command

B/D/F/Type: 0/0/0/PCI  
 Address Offset: CE-CFh  
 Default Value: 0000h  
 Access: RO; RW  
 Size: 16 bits

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.

Bit	Access	Default Value	Description
15:12	RO	0h	<i>Reserved</i>
11	RW	0b	<b>SCI on Processor Thermal Sensor Trip (TSTSCI)</b> 0 = Reporting of this condition via SCI messaging is disabled. 1 = A SCI DMI special cycle is generated by processor when the thermal sensor trip requires an SCI. A thermal sensor trip point cannot generate more than one special cycle.
10:2	RO	000h	<i>Reserved</i>
1	RW	0b	<i>Reserved</i>
0	RW	0b	<i>Reserved</i>

### 1.8.32 SKPD - Scratchpad Data

B/D/F/Type: 0/0/0/PCI  
 Address Offset: DC-DFh  
 Default Value: 00000000h  
 Access: RW  
 Size: 32 bits

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

Bit	Access	Default Value	Description
31:0	RW	00000000h	<b>Scratchpad Data (SKPD)</b> 1 DWORD of data storage.



### 1.8.33 CAPID0 - Capability Identifier

B/D/F/Type: 0/0/0/PCI  
 Address Offset: E0-EBh  
 Default Value: 0000000200000000010C0009h  
 Access: RO  
 Size: 96 bits

Control of bits in this register are for capability SKU differentiation.

Bit	Access	Default Value	Description
95:0	RO		Reserved

## 1.9 Device 0 MCHBAR DRAM Controls

Table 4. Device 0 MCHBAR DRAM Controls Summary (Sheet 1 of 2)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Channel Size Mapping	CSZMAP	100	107	0000000000000000h	RW-L
Channel Decode Misc.	CHDECMISC	111	111	00h	RW-L; RO
Channel 0 DRAM Rank Boundary Address 0	CODRBO	200	201	0000h	RO; RW-L
Channel 0 DRAM Rank Boundary Address 1	CODRB1	202	203	0000h	RO; RW-L
Channel 0 DRAM Rank Boundary Address 2	CODRB2	204	205	0000h	RO; RW-L
Channel 0 DRAM Rank Boundary Address 3	CODRB3	206	207	0000h	RW-L; RO
Channel 0 DRAM Rank 0,1 Attribute	CODRA01	208	209	0000h	RW-L
Channel 0 DRAM Rank 2,3 Attribute	CODRA23	20A	20B	0000h	RW-L
Channel 0 CYCTRK PCHG	COCYCTRKPCHG	250	251	0000h	RO; RW
Channel 0 CYCTRK ACT	COCYCTRKACT	252	255	00000000h	RO; RW
Channel 0 CYCTRK WR	COCYCTRKWR	256	257	0000h	RW
Channel 0 CYCTRK READ	COCYCTRKRDR	258	25A	000000h	RW; RO
Channel 0 CYCTRK REFR	COCYCTRKREFR	25B	25C	0000h	RO; RW
Channel 0 Refresh Control	COREFCTRL	2B9	2BA	0618h	RO; RW



Table 4. Device 0 MCHBAR DRAM Controls Summary (Sheet 2 of 2)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Channel 0 CKE Control	C0CKECTRL	260	263	000009FEh	RW-L; RO; RW
Channel 0 ODT Control	C0ODTCTRL	29C	29F	00000000h	RW; RO
Channel 0 DRAM Throttling Control	C0DTC	2B4	2B7	00000000h	RO; RW-L-K; RW-L
Channel 0 DRAM Rank Throttling Passive Event	C0DTPEW	2A8	2AB	00000000h	RW-L
Channel 0 DRAM Rank Throttling Active Event	C0DTAEW	2AC	2B3	000000000000000h	RO; RW-L
Channel 1 DRAM Rank Boundary Address 0	C1DRB0	600	601	0000h	RW-L; RO
Channel 1 DRAM Rank Boundary Address 1	C1DRB1	602	603	0000h	RW-L; RO
Channel 1 DRAM Rank Boundary Address 2	C1DRB2	604	605	0000h	RW-L; RO
Channel 1 DRAM Rank Boundary Address 3	C1DRB3	606	607	0000h	RW-L; RO
Channel 1 DRAM Rank 0,1 Attributes	C1DRA01	608	609	0000h	RW-L
Channel 1 DRAM Rank 2,3 Attributes	C1DRA23	60A	60B	0000h	RW-L
Channel 1 CYCTRK PCHG	C1CYCTRKPCHG	650	651	0000h	RW; RO
Channel 1 CYCTRK ACT	C1CYCTRKACT	652	655	00000000h	RW; RO
Channel 1 CYCTRK WR	C1CYCTRKWR	656	657	0000h	RW
Channel 1 CYCTRK READ	C1CYCTRKRD	658	65A	000000h	RW; RO
Channel 1 CKE Control	C1CKECTRL	660	663	0000080h	RW; RW-L; RO
Channel 1 CYCTRK REFR	C1CYCTRKREFR	65B	65C	0000h	RO; RW
Channel 1 ODT Control	C1ODTCTRL	69C	69F	00000000h	RW; RO
Channel 1 DRAM Throttling Control	C1DTC	6B4	6B7	00000000h	RO; RW-L-K; RW-L
Channel 1 DRAM Rank Throttling Passive Event	C1DTPEW	6A8	6AB	00000000h	RW-L
Channel 1 DRAM Rank Throttling Active Event	C1DTAEW	6AC	6B3	000000000000000h	RO; RW-L



### 1.9.1 CSZMAP - Channel Size Mapping

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 100-107h  
 Default Value: 0000000000000000h  
 Access: RW-L  
 Size: 64 bits  
 BIOS Optimal Default 0000h

This register indicates the total memory which is mapped to Interleaved and Asymmetric operation respectively (1MB granularity) used for Channel address decode.

Bit	Access	Default Value	Description
63:48	RO	0h	<i>Reserved</i>
47:32	RW-L	0000h	<b>2 Channel Size (2CHSZ)</b> This register indicates the total memory which is mapped to 2-channel operation (1-MB granularity) This register is locked by Intel ME stolen Memory lock and may also be forced to 0000h by the Performance Dual Channel Disable fuse.
31:16	RW-L	0000h	<b>1 Channel Size (1CHSZ)</b> This register indicates the total memory which is mapped to 1-channel operation (1-MB granularity) This register is locked by Intel ME stolen Memory lock.
15:0	RW-L	0000h	<b>Channel 0 Single Channel Size (COCSIZE)</b> This register indicates the quantity of memory physically in channel 0 which is mapped to 1-channel operation (1-MB granularity).





### 1.9.2 CHDECMISC - Channel Decode Misc.

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 111h  
 Default Value: 00h  
 Access: RW-L; RO  
 Size: 8 bits  
 BIOS Optimal Default 0h  
 Enhanced addressing configuration bits.

(Sheet 1 of 2)

Bit	Access	Default Value	Description	
7	RW-L	0b	<b>Enhanced Address for SO-DIMM Select (ENHDIMMSEL)</b> This bit can be set when enhanced mode of addressing for ranks is enabled and all four ranks are populated with equal amount of memory. 0 = Use Standard methods for SO-DIMM Select. 1 = Use Enhanced Address as SO-DIMM Select. This register is locked by Intel ME stolen Memory lock.	
			<b>Encoding</b>	<b>Description</b>
			0b	Standard methods for SO-DIMM Select
			1b	Enhanced Address as SO-DIMM Select
6:5	RW-L	00b	<b>Enhanced Mode Select (ENHMODESEL)</b> 00: Swap Enabled for Bank Selects and Rank Selects 01: XOR Enabled for Bank Selects and Rank Selects 10: Swap Enabled for Bank Selects only 11: XOR Enabled for Bank Select only This register is locked by Intel ME stolen Memory lock.	
			<b>Encoding</b>	<b>Description</b>
			00b	Swap Enabled for Bank Selects and Rank Selects
			01b	XOR Enabled for Bank Selects and Rank Selects
			10b	Swap Enabled for Bank Selects only
11b	XOR Enabled for Bank Select only			
4	RO	0b	Reserved	



(Sheet 2 of 2)

Bit	Access	Default Value	Description	
3	RW-L	0b	<b>Channel 1 Enhanced Mode (CH1_ENHMODE)</b> This bit indicates that enhanced addressing mode of operation is enabled for channel 1. Enhanced addressing mode of operation should be enabled only when both the channels are equally populated with same size and same type of DRAM memory. An added restriction is that the number of ranks/channel has to be 1, 2 or 4. <b>Note:</b> If any of the channels is in enhanced mode, the other channel should also be in enhanced mode. This register is locked by Intel ME stolen Memory lock.	
			<b>Encoding</b>	<b>Description</b>
			0b	Standard addressing
			1b	Enhanced addressing
2	RW-L	0b	<b>Channel 0 Enhanced Mode (CH0_ENHMODE)</b> This bit indicates that enhanced addressing mode of operation is enabled for Channel 0. Enhanced addressing mode of operation should be enabled only when both the channels are equally populated with same size and same type of DRAM memory. An added restriction is that the number of ranks/channel has to be 1, 2 or 4. <b>Note:</b> If any of the two channels is in enhanced mode, the other channel should also be in enhanced mode. This register is locked by Intel ME stolen Memory lock.	
			<b>Encoding</b>	<b>Description</b>
			0b	Standard addressing
			1b	Enhanced addressing
1:0	RO	0h	<i>Reserved</i>	



### 1.9.3 CODRBO - Channel 0 DRAM Rank Boundary Address 0

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 200-201h  
 Default Value: 0000h  
 Access: RW-L; RO  
 Size: 16 bits

The DRAM Rank Boundary Registers define the upper boundary address of each DRAM rank with a granularity of 64 MB. Each rank has its own single-word DRB register. These registers are used to determine which chip select is active for a given address. Channel and rank map:

- Ch0 Rank0: 200h
- Ch0 Rank1: 202h
- Ch0 Rank2: 204h
- Ch0 Rank3: 206h
- Ch1 Rank0: 600h
- Ch1 Rank1: 602h
- Ch1 Rank2: 604h
- Ch1 Rank3: 606h

Programming guide:

If Channel 0 is empty, all of the CODRBs are programmed with 00h.  
 CODRB0 = Total memory in Channel 0 Rank 0 (in 64-MB increments)  
 CODRB1 = Total memory in Channel 0 Rank 0 + Channel 0 Rank 1 (in 64-MB increments) and so on.  
 If Channel 1 is empty, all of the C1DRBs are programmed with 00h.  
 C1DRB0 = Total memory in Channel 1 Rank 0 (in 64-MB increments)  
 C1DRB1 = Total memory in Channel 1 Rank 0 + Channel 1 Rank 1 (in 64-MB increments) and so on.

Bit	Access	Default Value	Description
15:10	RO	00h	<i>Reserved</i>
9:0	RW-L	000h	<b>Channel 0 DRAM Rank Boundary Address 0 (CODRBA0)</b> This register defines the DRAM rank boundary for Rank 0 of Channel 0 (64-MB granularity) = R0 R0 = Total Rank 0 memory size/64 MB R1 = Total Rank 1 memory size/64 MB R2 = Total Rank 2 memory size/64 MB R3 = Total Rank 3 memory size/64 MB This register is locked by Intel ME stolen Memory lock.



### 1.9.4 CODRB1 - Channel 0 DRAM Rank Boundary Address 1

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 202-203h  
 Default Value: 0000h  
 Access: RW-L; RO  
 Size: 16 bits  
 See CODRB0.

Bit	Access	Default Value	Description
15:10	RO	00h	<i>Reserved</i>
9:0	RW-L	000h	<b>Channel 0 DRAM Rank Boundary Address 1 (CODRBA1)</b> This register defines the DRAM rank boundary for Rank 1 of Channel 0 (64-MB granularity) = (R1 + R0) R0 = Total Rank 0 memory size/64 MB R1 = Total Rank 1 memory size/64 MB R2 = Total Rank 2 memory size/64 MB R3 = Total Rank 3 memory size/64 MB This register is locked by Intel ME stolen Memory lock.

### 1.9.5 CODRB2 - Channel 0 DRAM Rank Boundary Address 2

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 204-205h  
 Default Value: 0000h  
 Access: RO; RW-L  
 Size: 16 bits  
 See CODRB0.

Bit	Access	Default Value	Description
15:10	RO	00h	<i>Reserved</i>
9:0	RW-L	000h	<b>Channel 0 DRAM Rank Boundary Address 2 (CODRBA2)</b> This register defines the DRAM rank boundary for Rank 2 of Channel 0 (64-MB granularity) = (R2 + R1 + R0) R0 = Total Rank 0 memory size/64 MB R1 = Total Rank 1 memory size/64 MB R2 = Total Rank 2 memory size/64 MB R3 = Total Rank 3 memory size/64 MB This register is locked by Intel ME stolen Memory lock.



### 1.9.6 CODRB3 - Channel 0 DRAM Rank Boundary Address 3

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 206-207h  
 Default Value: 0000h  
 Access: RO; RW-L  
 Size: 16 bits  
 See CODRB0.

Bit	Access	Default Value	Description
15:10	RO	00h	Reserved
9:0	RW-L	000h	<b>Channel 0 DRAM Rank Boundary Address 3 (CODRBA3)</b> This register defines the DRAM rank boundary for Rank 3 of Channel 0 (64 MB granularity) = (R3 + R2 + R1 + R0) R0 = Total Rank 0 memory size/64 MB R1 = Total Rank 1 memory size/64 MB R2 = Total Rank 2 memory size/64 MB R3 = Total Rank 3 memory size/64 MB This register is locked by Intel ME stolen Memory lock.

### 1.9.7 CODRA01 - Channel 0 DRAM Rank 0,1 Attribute

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 208-209h  
 Default Value: 0000h  
 Access: RW-L  
 Size: 16 bits

The DRAM Rank Attribute Registers define the page sizes/number of banks to be used when accessing different ranks. These registers should be left with their default value (all zeros) for any rank that is unpopulated, as determined by the corresponding CxDRB registers. Each byte of information in the CxDRA registers describes the page size of a pair of ranks. Channel and rank map:

- Ch0 Rank 0, 1: 208h-209h
- Ch0 Rank 2, 3: 20Ah-20Bh
- Ch1 Rank 0, 1: 608h - 609h
- Ch1 Rank 2, 3: 60Ah - 60Bh

DRA[7:0] = "00" means cfg0, DRA[7:0] = "01" means cfg1.... DRA[7:0] = "09" means cfg9 and so on.



DRA		Address Usage					Rank Capacity	Page Size
Cfg	Tech	Depth in rows	Device Width	Row Bits	Column Bits	Bank Bits		
00h through 85h	Reserved							
86h	1 Gb	128 M	8	14	10	3	1 GB	8 KB
87h	1 Gb	64 M	16	13	10	3	512 MB	8 KB
88h	2 Gb	256 M	8	15	10	3	2 GB	8 KB
89h	2 Gb	128 M	16	14	10	3	1 GB	8 KB
8Ah	Reserved							
8Bh	Reserved							
8Ch through FFh	Reserved							

Bit	Access	Default Value	Description
15:8	RW-L	00h	<b>Channel 0 DRAM Rank-1 Attributes (CODRA1)</b> This register defines DRAM pagesize/number-of-banks for Rank 1 for given channel. See table in register description for programming. This register is locked by Intel ME stolen Memory lock.
7:0	RW-L	00h	<b>Channel 0 DRAM Rank-0 Attributes (CODRA0)</b> This register defines DRAM page size/number-of-banks for Rank 0 for given channel. See table in register description for programming. This register is locked by Intel ME stolen Memory lock.



### 1.9.8 CODRA23 - Channel 0 DRAM Rank 2,3 Attribute

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 20A-20Bh  
 Default Value: 0000h  
 Access: RW-L  
 Size: 16 bits  
 See CODRA01.

Bit	Access	Default Value	Description
15:8	RW-L	00h	<b>Channel 0 DRAM Rank-3 Attributes (CODRA3)</b> This register defines DRAM pagesize/number-of-banks for Rank 3 for given channel. See table in register description for programming. This register is locked by Intel ME stolen Memory lock.
7:0	RW-L	00h	<b>Channel 0 DRAM Rank-2 Attributes (CODRA2)</b> This register defines DRAM pagesize/number-of-banks for Rank 2 for given channel. See table in register description for programming. This register is locked by Intel ME stolen Memory lock.



### 1.9.9 COCYCTRKPCHG - Channel 0 CYCTRK PCHG

B/D/F/Type: 0/0/0/MCHBAR  
Address Offset: 250-251h  
Default Value: 0000h  
Access: RO; RW  
Size: 16 bits

Channel 0 CYCTRK Precharge Registers.

Bit	Access	Default Value	Description
15:11	RO	00h	<i>Reserved</i>
10:6	RW	00h	<b>Write To Precharge Delay (C0sd_cr_wr_pchg)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the WRITE and PRE commands to the same rank-bank. Corresponds to the tWR parameter in the DDR3 Specification.
5:2	RW	0h	<b>Read To Precharge Delay (C0sd_cr_rd_pchg)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the READ and PRE commands to the same rank-bank.
1:0	RW	00b	<b>Precharge To Precharge Delay (C0sd_cr_pchg_pchg)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between two PRE commands to the same rank.





### 1.9.10 COCYCTRKACT - Channel 0 CYCTRK ACT

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 252-255h  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits

Channel 0 CYCTRK Activate Registers.

Bit	Access	Default Value	Description
31:28	RO	00b	Reserved
27:22	RW	00h	<b>Activate Window Count (C0sd_cr_act_windowcnt)</b> This configuration register indicates the window duration (in DRAM clocks) during which the controller counts the number of activate commands which are launched to a particular rank. If the number of activate commands launched within this window is greater than 4, then a check is implemented to block launch of further activates to this rank for the rest of the duration of this window.
21	RW	0b	<b>Max Activate Check (C0sd_cr_maxact_dischk)</b> This configuration register enables the check which ensures that there are no more than four activates to a particular rank in a given window.
20:17	RW	0h	<b>Activate to Activate Delay (C0sd_cr_act_act)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between two ACT commands to the same rank. Corresponds to the tRRD parameter in the DDR3 Specification.
16:13	RW	0h	<b>Precharge to Activate Delay (C0sd_cr_pre_act)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the PRE and ACT commands to the same rank-bank. Corresponds to the tRP parameter in the DDR3 Specification.
12:9	RW	0h	<b>Precharge All to Activate Delay (C0sd_cr_preall_act)</b> From the launch of a precharge-all command wait for this many memory bus clocks before launching an activate command. Corresponds to the tPALL_RP parameter.
8:0	RW	000h	<b>Refresh to Activate Delay (C0sd_cr_rfsh_act)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between REF and ACT commands to the same rank. Corresponds to the tRFC parameter in the DDR3 Specification.



### 1.9.11 COCYCTRKWR - Channel 0 CYCTRK WR

B/D/F/Type: 0/0/0/MCHBAR  
Address Offset: 256-257h  
Default Value: 0000h  
Access: RW  
Size: 16 bits

Channel 0 CYCTRK WR Registers.

Bit	Access	Default Value	Description
15:12	RW	0h	<b>Activate-to-Write Delay (C0sd_cr_act_wr)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the ACT and WRITE commands to the same rank-bank. Corresponds to the tRCD_wr parameter DDR3 specification.
11:8	RW	0h	<b>Same Rank Write-to-Write Delay (C0sd_cr_wrsr_wr)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between two WRITE commands to the same rank.
7:4	RW	0h	<b>Different Rank Write-to-Write Delay (C0sd_cr_wrdr_wr)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between two WRITE commands to different ranks. Corresponds to the tWR_WR parameter in the DDR3 specification.
3:0	RW	0h	<b>Read-to-Write Delay (C0sd_cr_rd_wr)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the READ and WRITE commands. Corresponds to the tRD_WR parameter.



### 1.9.12 COCYCTRKR - Channel 0 CYCTRK READ

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 258-25Ah  
 Default Value: 000000h  
 Access: RO; RW  
 Size: 24 bits

Channel 0 CYCTRK RD Registers.

Bit	Access	Default Value	Description
23:21	RO	000b	<i>Reserved</i>
20:17	RW	0h	<b>Minimum Activate-to-Read Delay (COsd_cr_act_rd)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the ACT and READ commands to the same rank-bank. Corresponds to tRCD_rd parameter in the DDR3 specification.
16:12	RW	00h	<b>Same Rank Write-to-Read Delayed (COsd_cr_wrsr_rd)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the WRITE and READ commands to the same rank. Corresponds to the tWTR parameter in the DDR3 specification.
11:8	RW	0h	<b>Different Ranks Write-to-Read Delayed (COsd_cr_wrdr_rd)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the WRITE and READ commands to different ranks. Corresponds to the tWR_RD parameter DDR3 specification.
7:4	RW	0h	<b>Same Rank Read-to-Read Delayed (COsd_cr_rdsr_rd)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between two READ commands to the same rank.
3:0	RW	0h	<b>Different Ranks Read-to-Read Delayed (COsd_cr_rddr_rd)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between two READ commands to different ranks. Corresponds to the tRD_RD parameter.



### 1.9.13 COCYCTRKREFR - Channel 0 CYCTRK REFR

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 25B-25Ch  
 Default Value: 0000h  
 Access: RO; RW  
 Size: 16 bits

Channel 0 CYCTRK Refresh Registers.

Bit	Access	Default Value	Description
15:13	RO	000b	<i>Reserved</i>
12:9	RW	0h	<b>Same Rank Precharge All to Refresh Delay (COsd_cr_pchgall_rfsh)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the PRE-ALL and REF commands to the same rank.
8:0	RW	000h	<b>Same Rank Refresh to Refresh Delay (COsd_cr_rfsh_rfsh)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between two REF commands to the same rank.

### 1.9.14 COREFRCTRL - Channel 0 DRAM Refresh Control

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 269-26Eh  
 Default Value: 241830000C30h  
 Access: RO; RW  
 Size: 48 bits

Settings to configure the DRAM refresh controller.

(Sheet 1 of 3)

Bit	Access	Default Value	Description
47	RO	0b	<i>Reserved</i>
46:44	RW	010b	<b>Initial Refresh Count (INITREFCNT)</b> Initial Refresh Count Value
43:38	RW	10h	<b>Direct Rcomp Quiet Window (DIRQUIET)</b> This configuration setting indicates the amount of refresh_tick events to wait before the service of Rcomp request in non-default mode of independent rank refresh.
37:32	RW	18h	<b>Indirect Rcomp Quiet Window (INDIRQUIET)</b> This configuration setting indicates the amount of refresh_tick events to wait before the service of Rcomp request in non-default mode of independent rank refresh.



(Sheet 2 of 3)

Bit	Access	Default Value	Description								
31:27	RW	06h	<p><b>Rcomp Wait (RCOMPWAIT)</b> This configuration setting indicates the amount of refresh_tick events to wait before the service of Rcomp request in non-default mode of independent rank refresh.</p>								
26	RW	0b	<i>Reserved</i>								
25	RW	0b	<p><b>Refresh Counter Enable (REFCNTEN)</b> This bit is used to enable the refresh counter to count during times that DRAM is not in self-refresh, but refreshes are not enabled. Such a condition may occur due to need to reprogram SO-DIMMs following DRAM controller switch. This bit has no effect when Refresh is enabled (i.e., there is no mode where Refresh is enabled but the counter does not run) so, in conjunction with Bit 23 REFEN, the modes are:</p> <table border="1" data-bbox="776 829 1276 1073"> <thead> <tr> <th>REFEN: REFCNTEN</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0:0</td> <td>Normal refresh disable</td> </tr> <tr> <td>0:1</td> <td>Refresh disabled, but counter is accumulating refreshes.</td> </tr> <tr> <td>1:X</td> <td>Normal refresh enable</td> </tr> </tbody> </table>	REFEN: REFCNTEN	Description	0:0	Normal refresh disable	0:1	Refresh disabled, but counter is accumulating refreshes.	1:X	Normal refresh enable
REFEN: REFCNTEN	Description										
0:0	Normal refresh disable										
0:1	Refresh disabled, but counter is accumulating refreshes.										
1:X	Normal refresh enable										
24	RW	0b	<p><b>All Rank Refresh (ALLRKREF)</b> This configuration bit enables (by default) that all the ranks are refreshed in a staggered/atomic fashion. If set, the ranks are refreshed in an independent fashion. 0 = Ranks are refreshed atomically staggered 1 = Ranks are refreshed independently</p>								
23	RW	0b	<p><b>Refresh Enable (REFEN)</b> Refresh is enabled. 0 = Disabled 1 = Enabled</p>								
22	RW	0b	<p><b>DDR Initialization Done (INITDONE)</b> Indicates that DDR initialization is complete.</p>								
21:20	RW	00b	<p><b>DRAM Refresh Hysteresis (REFHYSTERISIS)</b> Hysteresis level - Useful for dref_high watermark cases. The dref_high flag is set when the dref_high watermark level is exceeded, and is cleared when the refresh count is less than the hysteresis level. This bit should be set to a value less than the high watermark level.</p> <p>00: 3 01: 4 10: 5 11: 6</p>								



(Sheet 3 of 3)

Bit	Access	Default Value	Description
19:18	RW	00b	<b>DRAM Refresh Panic Watermark (REFPANICWMM)</b> When the refresh count exceeds this level, a refresh request is launched to the scheduler and the dref_panic flag is set. 00: 5 01: 6 10: 7 11: 8
17:16	RW	00b	<b>DRAM Refresh High Watermark (REFHIGHWM)</b> When the refresh count exceeds this level, a refresh request is launched to the scheduler and the dref_high flag is set. 00: 3 01: 4 10: 5 11: 6
15:14	RW	00b	<b>DRAM Refresh Low Watermark (REFLOWWM)</b> When the refresh count exceeds this level, a refresh request is launched to the scheduler and the dref_low flag is set. 00: 1 01: 2 10: 3 11: 4
13:0	RW	0C30h	<b>Refresh Counter Time Out Value (REFTIMEOUT)</b> At various frequencies this results in the following values: 400 MHz -> C30 hex 533 MHz -> 104B hex 666 MHz -> 1450 hex



### 1.9.15 COCKECTRL - Channel 0 CKE Control

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 260-263h  
 Default Value: 000009FEh  
 Access: RW-L; RO; RW  
 Size: 32 bits

CKE controls for Channel 0.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31:28	RO	0h	<i>Reserved</i>
27	RW	0b	<b>Start the Self-refresh Exit Sequence (sd0_cr_srcstart)</b>  This configuration register indicates the request to start the self-refresh exit sequence.
26:24	RW	000b	<b>CKE Pulse Width Requirement in High Phase (sd0_cr_cke_pw_hl_safe)</b>  This configuration register indicates CKE pulse width requirement in high phase. Corresponds to the tCKE (high) parameter in the DDR3 specification.
23	RW-L	0b	<b>Rank 3 Population (sd0_cr_rankpop3)</b> 1: Rank 3 populated. 0: Rank 3 not populated. This register is locked by Intel ME stolen Memory lock.
22	RW-L	0b	<b>Rank 2 Population (sd0_cr_rankpop2)</b> 1: Rank 2 populated. 0: Rank 2 not populated. This register is locked by Intel ME stolen Memory lock.
21	RW-L	0b	<b>Rank 1 Population (sd0_cr_rankpop1)</b> 1: Rank 1 populated. 0: Rank 1 not populated. This register is locked by Intel ME stolen Memory lock.
20	RW-L	0b	<b>Rank 0 Population (sd0_cr_rankpop0)</b> 1: Rank 0 populated. 0: Rank 0 not populated. This register is locked by Intel ME stolen Memory lock.
19:17	RW	000b	<b>CKE Pulse Width Requirement in Low Phase (sd0_cr_cke_pw_lh_safe)</b>  This configuration register indicates CKE pulse width requirement in low phase. Corresponds to the tCKE (low) parameter in the DDR3 specification.
16:14	RO	000b	<i>Reserved</i>



(Sheet 2 of 2)

Bit	Access	Default Value	Description
13:10	RW	2h	<b>Minimum Power down Exit to Non-Read Command Spacing (sd0_cr_txp)</b> This configuration register indicates the minimum number of clocks to wait following assertion of CKE before issuing a non-read command. Ah-Fh: Reserved. 2h-9h: 2-9clocks. 0h-1h: Reserved.
9:1	RW	0FFh	<b>Self Refresh Exit Count (sd0_cr_slfrfsh_exit_cnt)</b> This configuration register indicates the Self refresh exit count. (Program to 255). Corresponds to the tXSNR/tXSRD parameters in the DDR3 Specification.
0	RW	0b	<b>Only 1 DIMM Populated (sd0_cr_singledimmpop)</b> 0: There is more than one DIMM or SO-DIMM in this channel 1: There is only one DIMM or SO-DIMM in this channel

### 1.9.16 COODTCTRL - Channel 0 ODT Control

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 29C-29Fh  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits

Bit	Access	Default Value	Description
31:12	RO	00000h	<i>Reserved</i>
11:8	RW	0h	<b>DRAM ODT for Read Commands (sd0_cr_odt_duration_rd)</b> Specifies the duration to assert DRAM ODT for Read Commands. The Async value should be used when the Dynamic Powerdown bit is set. Else use the Sync value.
7:4	RW	0h	<b>DRAM ODT for Write Commands (sd0_cr_odt_duration_wr)</b> Specifies the duration to assert DRAM ODT for Write Commands. The Async value should be used when the Dynamic Powerdown bit is set. Else use the Sync value.
3:0	RW	0h	<b>MCH ODT for Read Commands (sd0_cr_mchodt_duration)</b> Specifies the duration to assert MCH ODT for Read Commands.





### 1.9.17 CODTC - Channel 0 DRAM Throttling Control

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 2B4-2B7h  
 Default Value: 00000000h  
 Access: RO; RW-L-K; RW-L  
 Size: 32 bits

Programmable Event weights are input into the averaging filter. Each Event weight is an normalized 8-bit value that the BIOS must program. The BIOS must account for burst length and 1N/2N rule considerations. It is also possible for BIOS to take into account loading variations of memory caused as a function of memory types and population of ranks.

Bit	Access	Default Value	Description
31:24	RO	00h	<i>Reserved</i>
23	RW-L-K	0b	<b>DRAM Throttle Lock (DTLOCK)</b> This bit secures the DRAM throttling control registers DT*EW and DTC. Once a 1 is written to this bit, all of these configuration register bits become read-only.
22	RW-L	0b	<i>Reserved</i>
21	RW-L	0b	<b>DRAM Bandwidth Based Throttling Enable (DBBTE)</b> 0 = Bandwidth Threshold (WAB) is not used for throttling. 1 = Bandwidth Threshold (WAB) is used for throttling.  If both Bandwidth based and thermal sensor based throttling modes are on and the thermal sensor trips, weighted average WAT is used for throttling.
20	RW-L	0b	<i>Reserved</i>
19	RO	0b	<i>Reserved</i>
18:16	RW-L	000b	<i>Reserved</i>
15:8	RW-L	00h	<b>Weighted Average Bandwidth Limit (WAB)</b> Average weighted bandwidth allowed per clock during for bandwidth based throttling. The memory controller does not allow any transactions to proceed on the System Memory bus if the output of the filter equals or exceeds this value.
7:0	RW-L	00h	<b>Weighted Average Thermal Limit (WAT)</b> Average weighted bandwidth allowed per clock during for thermal sensor enabled throttling. The memory controller does not allow any transactions to proceed on the System Memory bus if the output of the filter equals or exceeds this value.



### 1.9.18 CODTPEW - Channel 0 DRAM Rank Throttling Passive Event

B/D/F/Type: 0/0/0/MCHBAR  
Address Offset: 2A8-2ABh  
Default Value: 00000000h  
Access: RW-L  
Size: 32 bits

Programmable Event weights are input into the averaging filter. Each Event weight is an normalized 8-bit value that the BIOS must program. The BIOS must account for burst length and 1N/2N rule considerations. It is also possible for BIOS to take into account loading variations of memory caused as a function of memory types and population of ranks. GMCH implements 4 independent filters, one per rank. All bits in this register can be locked by the DTLOCK bit in the CODTC register.

Bit	Access	Default Value	Description
31:24	RW-L	00h	<i>Reserved</i>
23:16	RW-L	00h	<i>Reserved</i>
15:8	RW-L	00h	<i>Reserved</i>
7:0	RW-L	00h	<i>Reserved</i>



### 1.9.19 CODTAEW - Channel 0 DRAM Rank Throttling Active Event

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 2AC-2B3h  
 Default Value: 0000000000000000h  
 Access: RO; RW-L  
 Size: 64 bits

Programmable Event weights are input into the averaging filter. Each Event weight is an normalized 8-bit value that the BIOS must program. The BIOS must account for burst length and 1N/2N rule considerations. It is also possible for BIOS to take into account loading variations of memory caused as a function of memory types and population of ranks. GMCH implements four independent filters, one per rank. During a given clock, GMCH asserts a command to the DRAM (via CSB assertion). Based on the command type, one of the weights specified in this register is added to the appropriate weight specified in CODTPEW and input to the filter. All bits in this register can be locked by the DTLOCK bit in the CODTC register.

Bit	Access	Default Value	Description
63:48	RO	0000h	Reserved
47:40	RW-L	00h	Reserved
39:32	RW-L	00h	Reserved
31:24	RW-L	00h	Reserved
23:16	RW-L	00h	Reserved
15:8	RW-L	00h	Reserved
7:0	RW-L	00h	Reserved

### 1.9.20 C1DRB0 - Channel 1 DRAM Rank Boundary Address 0

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 600-601h  
 Default Value: 0000h  
 Access: RW-L; RO  
 Size: 16 bits

The operation of this register is detailed in the description for register CODRBO.

Bit	Access	Default Value	Description
15:10	RO	000000b	Reserved
9:0	RW-L	000h	<b>Channel 1 DRAM Rank Boundary Address 0 (C1DRBA0)</b> See CODRBO. This register is locked by Intel ME stolen Memory lock.



### 1.9.21 C1DRB1 - Channel 1 DRAM Rank Boundary Address 1

B/D/F/Type: 0/0/0/MCHBAR  
Address Offset: 602-603h  
Default Value: 0000h  
Access: RO; RW-L  
Size: 16 bits

The operation of this register is detailed in the description for register C0DRB0.

Bit	Access	Default Value	Description
15:10	RO	000000b	<i>Reserved</i>
9:0	RW-L	000h	<b>Channel 1 DRAM Rank Boundary Address 1 (C1DRBA1)</b>  See C0DRB1. This register is locked by Intel ME stolen Memory lock.

### 1.9.22 C1DRB2 - Channel 1 DRAM Rank Boundary Address 2

B/D/F/Type: 0/0/0/MCHBAR  
Address Offset: 604-605h  
Default Value: 0000h  
Access: RW-L; RO  
Size: 16 bits

The operation of this register is detailed in the description for register C0DRB0.

Bit	Access	Default Value	Description
15:10	RO	000000b	<i>Reserved</i>
9:0	RW-L	000h	<b>Channel 1 DRAM Rank Boundary Address 2 (C1DRBA2)</b>  See C0DRB2. This register is locked by Intel ME stolen Memory lock.



### 1.9.23 C1DRB3 - Channel 1 DRAM Rank Boundary Address 3

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 606-607h  
 Default Value: 0000h  
 Access: RW-L; RO  
 Size: 16 bits

The operation of this register is detailed in the description for register CODRB0.

Bit	Access	Default Value	Description
15:10	RO	000000b	<i>Reserved</i>
9:0	RW-L	000h	<b>Channel 1 DRAM Rank Boundary Address 3 (C1DRBA3)</b>  <b>See CODRB3.</b> This register is locked by Intel ME stolen Memory lock.

### 1.9.24 C1DRA01 - Channel 1 DRAM Rank 0,1 Attributes

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 608-609h  
 Default Value: 0000h  
 Access: RW-L  
 Size: 16 bits

The operation of this register is detailed in the description for register CODRA01.

Bit	Access	Default Value	Description
15:8	RW-L	00h	<b>Channel 1 DRAM Rank-1 Attributes (C1DRA1)</b> <b>See CODRA1.</b> This register is locked by Intel ME stolen Memory lock.
7:0	RW-L	00h	<b>Channel 1 DRAM Rank-0 Attributes (C1DRA0)</b> <b>See CODRA0.</b> This register is locked by Intel ME stolen Memory lock.



### 1.9.25 C1DRA23 - Channel 1 DRAM Rank 2,3 Attributes

B/D/F/Type: 0/0/0/MCHBAR  
Address Offset: 60A-60Bh  
Default Value: 0000h  
Access: RW-L  
Size: 16 bits

The operation of this register is detailed in the description for register CODRA01.

Bit	Access	Default Value	Description
15:8	RW-L	00h	<b>Channel 1 DRAM Rank-3 Attributes (C1DRA3)</b> See CODRA3. This register is locked by Intel ME stolen Memory lock.
7:0	RW-L	00h	<b>Channel 1 DRAM Rank-2 Attributes (C1DRA2)</b> See CODRA2. This register is locked by Intel ME stolen Memory lock.

### 1.9.26 C1CYCTRPCHG - Channel 1 CYCTRK PCHG

B/D/F/Type: 0/0/0/MCHBAR  
Address Offset: 650-651h  
Default Value: 0000h  
Access: RO; RW  
Size: 16 bits

Channel 1 CYCTRK Precharge.

Bit	Access	Default Value	Description
15:11	RO	00000b	<i>Reserved</i>
10:6	RW	00000b	<b>Write To PRE Delayed (C1sd_cr_wr_pchg)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the WRITE and PRE commands to the same rank-bank. Corresponds to tWR at DDR Spec.
5:2	RW	0000b	<b>READ To PRE Delayed (C1sd_cr_rd_pchg)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the READ and PRE commands to the same rank-bank
1:0	RW	00b	<b>PRE To PRE Delayed (C1sd_cr_pchg_pchg)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between two PRE commands to the same rank.



### 1.9.27 C1CYCTRKACT - Channel 1 CYCTRK ACT

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 652-655h  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits  
 Channel 1 CYCTRK ACT.

Bit	Access	Default Value	Description
31:28	RO	0h	<i>Reserved</i>
27:22	RW	000000b	<b>ACT Window Count (C1sd_cr_act_windowcnt)</b> This configuration register indicates the window duration (in DRAM clocks) during which the controller counts the number of activate commands which are launched to a particular rank. If the number of activate commands launched within this window is greater than 4, then a check is implemented to block launch of further activates to this rank for the rest of the duration of this window.
21	RW	0b	<b>Max ACT Check (C1sd_cr_maxact_dischk)</b> This configuration register enables the check which ensures that there are no more than four activates to a particular rank in a given window.
20:17	RW	0000b	<b>ACT to ACT Delayed (C1sd_cr_act_act)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between two ACT commands to the same rank. Corresponds to tRRD at DDR Spec.
16:13	RW	0000b	<b>PRE to ACT Delayed (C1sd_cr_pre_act)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the PRE-ALL and ACT commands to the same rank. Corresponds to tRP at DDR Spec.
12:9	RW	0h	<b>ALLPRE-to-ACT Delay (C1sd_cr_preall_act)</b> From the launch of a prechargeall command wait for these number of clocks before launching a activate command. Corresponds to tPALL_RP.
8:0	RW	00000000 0b	<b>REF-to-ACT Delayed (C1sd_cr_rfsh_act)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between REF and ACT commands to the same rank. Corresponds to tRFC at DDR Spec.



### 1.9.28 C1CYCTRKWR - Channel 1 CYCTRK WR

B/D/F/Type: 0/0/0/MCHBAR  
Address Offset: 656-657h  
Default Value: 0000h  
Access: RW  
Size: 16 bits  
Channel 1 CYCTRK WR

Bit	Access	Default Value	Description
15:12	RW	0h	<b>ACT-to-Write Delay (C1sd_cr_act_wr)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the ACT and WRITE commands to the same rank-bank. Corresponds to tRCD_wr at DDR Spec.
11:8	RW	0h	<b>Same Rank Write-to-Write Delayed (C1sd_cr_wrsr_wr)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between two WRITE commands to the same rank.
7:4	RW	0h	<b>Different Rank Write-to-Write Delay (C1sd_cr_wrd_r_wr)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between two WRITE commands to different ranks. Corresponds to tWR_WR at DDR Spec.
3:0	RW	0h	<b>READ-to-WRITE Delay (C1sd_cr_rd_wr)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the READ and WRITE commands. Corresponds to tRD_WR.





### 1.9.29 C1CYCTRKRD - Channel 1 CYCTRK READ

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 658-65Ah  
 Default Value: 000000h  
 Access: RO; RW  
 Size: 24 bits

Channel 1 CYCTRK READ

Bit	Access	Default Value	Description
23:21	RO	0h	<i>Reserved</i>
20:17	RW	0h	<b>Min ACT-to-READ Delayed (C1sd_cr_act_rd)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the ACT and READ commands to the same rank-bank Corresponds to tRCD_rd at DDR Spec.
16:12	RW	00000b	<b>Same Rank Write-to-READ Delayed (C1sd_cr_wrsr_rd)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the WRITE and READ commands to the same rank. Corresponds to tWTR at DDR Spec.
11:8	RW	0000b	<b>Different Ranks Write-to-READ Delayed (C1sd_cr_wrd_r_rd)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between the WRITE and READ commands to different ranks. Corresponds to tWR_RD at DDR Spec.
7:4	RW	0000b	<b>Same Rank Read-to-Read Delayed (C1sd_cr_rdsr_rd)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between two READ commands to the same rank.
3:0	RW	0000b	<b>Different Ranks Read-to-Read Delayed (C1sd_cr_rddr_rd)</b> This configuration register indicates the minimum allowed spacing (in DRAM clocks) between two READ commands to different ranks. Corresponds to tRD_RD.



### 1.9.30 C1CKECTRL - Channel 1 CKE Control

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 660-663h  
 Default Value: 00000800h  
 Access: RW; RW-L; RO  
 Size: 32 bits  
 Channel 1 CKE Controls

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31:28	RO	0h	<i>Reserved</i>
27	RW	0b	<b>start the self-refresh exit sequence (sd1_cr_srcstart)</b> This configuration register indicates the request to start the self-refresh exit sequence
26:24	RW	000b	<b>CKE pulse width requirement in high phase (sd1_cr_cke_pw_hl_safe)</b> This configuration register indicates CKE pulse width requirement in high phase. Corresponds to tCKE (high) at DDR Spec.
23	RW-L	0b	<b>Rank 3 Population (sd1_cr_rankpop3)</b> 0: Rank 3 not populated 1: Rank 3 populated This register is locked by Intel ME stolen Memory lock.
22	RW-L	0b	<b>Rank 2 Population (sd1_cr_rankpop2)</b> 0: Rank 2 not populated 1: Rank 2 populated This register is locked by Intel ME stolen Memory lock.
21	RW-L	0b	<b>Rank 1 Population (sd1_cr_rankpop1)</b> 0: Rank 1 not populated 1: Rank 1 populated This register is locked by Intel ME stolen Memory lock.
20	RW-L	0b	<b>Rank 0 Population (sd1_cr_rankpop0)</b> Rank 0 not populated Rank 0 populated This register is locked by Intel ME stolen Memory lock.
19:17	RW	000b	<b>CKE pulse width requirement in low phase (sd1_cr_cke_pw_lh_safe)</b> This configuration register indicates CKE pulse width requirement in low phase. Corresponds to tCKE (low) at DDR Spec.
16:14	RO	000b	<i>Reserved</i>



(Sheet 2 of 2)

Bit	Access	Default Value	Description
13:10	RW	0010b	<b>Minimum Powerdown Exit to Non-Read command spacing (sd1_cr_txp)</b> This configuration register indicates the minimum number of clocks to wait following assertion of CKE before issuing a non-read command. 1010-1111=Reserved 0010-1001=2-9 clocks 0000-0001=Reserved.
9:1	RW	00000000 0b	<b>Self Refresh Exit Count (sd1_cr_slfrsh_exit_cnt)</b> This configuration register indicates the Self refresh exit count. (Program to 255) Corresponds to tXSNR/tXSRD at DDR Spec.
0	RW	0b	<b>Indicates only 1 SO-DIMM populated (sd1_cr_singledimmpop)</b> This configuration register indicates that only 1 SO_DIMM is populated.

### 1.9.31 C1REFRCTRL - Channel 1 DRAM Refresh Control

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 669-66Eh  
 Default Value: 241830000C30h  
 Access: RO; RW  
 Size: 48 bits  
 Settings to configure the DRAM refresh controller.

(Sheet 1 of 3)

Bit	Access	Default Value	Description
47	RO	0b	<i>Reserved</i>
46:44	RW	010b	<b>Initial Refresh Count (INITREFCNT)</b> Specifies Initial Refresh Count Value
43:38	RW	10h	<b>Direct Rcomp Quiet Window (DIRQUIET)</b> This configuration setting indicates the amount of refresh_tick events to wait before the service of Rcomp request in non-default mode of independent rank refresh.
37:32	RW	18h	<b>Indirect Rcomp Quiet Window (INDIRQUIET)</b> This configuration setting indicates the amount of refresh_tick events to wait before the service of Rcomp request in non-default mode of independent rank refresh.
31:27	RW	00110b	<b>Rcomp Wait (RCOMPWAIT)</b> This configuration setting indicates the amount of refresh_tick events to wait before the service of Rcomp request in non-default mode of independent rank refresh.
26	RW	0b	<i>Reserved</i>



(Sheet 2 of 3)

Bit	Access	Default Value	Description								
25	RW	0b	<p><b>Refresh Counter Enable (REFCNTEN)</b></p> <p>This bit is used to enable the refresh counter to count during times that DRAM is not in self-refresh, but refreshes are not enabled. Such a condition may occur due to need to reprogram SO-DIMMs following DRAM controller switch.</p> <p>This bit has no effect when Refresh is enabled (i.e. there is no mode where Refresh is enabled but the counter does not run) So, in conjunction with Bit 23 REFEN, the modes are:</p> <table border="1"> <thead> <tr> <th>REFEN: REFCNTEN</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0:0</td> <td>Normal refresh disable</td> </tr> <tr> <td>0:1</td> <td>Refresh disabled, but counter is accumulating refreshes</td> </tr> <tr> <td>1:X</td> <td>Normal refresh enable</td> </tr> </tbody> </table>	REFEN: REFCNTEN	Description	0:0	Normal refresh disable	0:1	Refresh disabled, but counter is accumulating refreshes	1:X	Normal refresh enable
REFEN: REFCNTEN	Description										
0:0	Normal refresh disable										
0:1	Refresh disabled, but counter is accumulating refreshes										
1:X	Normal refresh enable										
24	RW	0b	<p><b>All Rank Refresh (ALLRKREF)</b></p> <p>This configuration bit enables (by default) that all the ranks are refreshed in a staggered/atomic fashion. If set, the ranks are refreshed in an independent fashion.</p>								
23	RW	0b	<p><b>Refresh Enable (REFEN)</b></p> <p>Refresh is enabled.</p> <p>0 = Disabled 1 = Enabled</p>								
22	RW	0b	<p><b>DDR Initialization Done (INITDONE)</b></p> <p>Indicates that DDR initialization is complete.</p>								
21:20	RW	00b	<p><b>DRAM Refresh Hysteresis (REFHYSTERISIS)</b></p> <p>Hysteresis level - Useful for dref_high watermark cases. The dref_high flag is set when the dref_high watermark level is exceeded, and is cleared when the refresh count is less than the hysteresis level. This bit should be set to a value less than the high watermark level.</p> <p>00: 3 01: 4 10: 5 11: 6</p>								
19:18	RW	00b	<p><b>DRAM Refresh Panic Watermark (REFPANICWM)</b></p> <p>When the refresh count exceeds this level, a refresh request is launched to the scheduler and the dref_panic flag is set.</p> <p>00: 5 01: 6 10: 7 11: 8</p>								



(Sheet 3 of 3)

Bit	Access	Default Value	Description
17:16	RW	00b	<b>DRAM Refresh High Watermark (REFHIGHWM)</b> When the refresh count exceeds this level, a refresh request is launched to the scheduler and the dref_high flag is set.  00: 3 01: 4 10: 5 11: 6
15:14	RW	00b	<b>DRAM Refresh Low Watermark (REFLOWWM)</b> When the refresh count exceeds this level, a refresh request is launched to the scheduler and the dref_low flag is set.  00: 1 01: 2 10: 3 11: 4
13:0	RW	00110000 110000b	<b>Refresh Counter Time Out Value (REFTIMEOUT)</b> At various mclk freq's this results in the following values:  400 MHz -> C30 hex 533 MHz -> 104B hex 666 MHz -> 1450 hex



### 1.9.32 C1ODTCTRL - Channel 1 ODT Control

B/D/F/Type: 0/0/0/MCHBAR  
Address Offset: 69C-69Fh  
Default Value: 00000000h  
Access: RO; RW  
Size: 32 bits  
ODT controls.

Bit	Access	Default Value	Description
31:12	RO	00000h	<i>Reserved</i>
11:8	RW	0h	<b>DRAM ODT for Read Commands</b> Specifies the duration to assert DRAM ODT for Read Commands. The Async value should be used when the Dynamic Powerdown bit is set. Else use the Sync value.
7:4	RW	0h	<b>DRAM ODT for Write Commands</b> Specifies the duration to assert DRAM ODT for Write Commands. The Async value should be used when the Dynamic Powerdown bit is set. Else use the Sync value.
3:0	RW	0h	<b>MCH ODT for Read Commands</b> Specifies the duration to assert MCH ODT for Read Commands.



### 1.9.33 C1DTC - Channel 1 DRAM Throttling Control

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 6B4-6B7h  
 Default Value: 00000000h  
 Access: RO; RW-L-K; RW-L  
 Size: 32 bits

Programmable Event weights are input into the averaging filter. Each Event weight is a normalized 8-bit value that the BIOS must program. The BIOS must account for burst length and 1N/2N rule considerations. It is also possible for BIOS to take into account loading variations of memory caused as a function of memory types and population of ranks.

Bit	Access	Default Value	Description
31:24	RO	00h	<i>Reserved</i>
23	RW-L-K	0b	<b>DRAM Throttle Lock (DTLOCK)</b> This bit secures the DRAM throttling control registers DT*EW and DTC. Once a 1 is written to this bit, all of these configuration register bits become read-only.
22	RW-L	0b	<i>Reserved</i>
21	RW-L	0b	<b>DRAM Bandwidth Based Throttling Enable (DBBTE)</b> 0 = Bandwidth Threshold (WAB) is not used for throttling. 1 = Bandwidth Threshold (WAB) is used for throttling. If both Bandwidth based and thermal sensor based throttling modes are on and the thermal sensor trips, weighted average WAT is used for throttling.
20	RW-L	0b	<b>DRAM Thermal Sensor Trip Enable (DTSTE)</b> 0 = Memory controller throttling is not initiated when the memory controller thermal sensor trips. 1 = Memory controller throttling is initiated when the memory controller thermal sensor trips and the Filter output is equal to or exceeds thermal threshold WAT.
19	RO	0b	<i>Reserved</i>
18:16	RW-L	000b	<i>Reserved</i>
15:8	RW-L	00h	<b>Weighted Average Bandwidth Limit (WAB)</b> Average weighted bandwidth allowed per clock during for bandwidth based throttling. The memory controller does not allow any transactions to proceed on the System Memory bus if the output of the filter equals or exceeds this value.
7:0	RW-L	00h	<b>Weighted Average Thermal Limit (WAT)</b> Average weighted bandwidth allowed per clock during for thermal sensor enabled throttling. The memory controller does not allow any transactions to proceed on the System Memory bus if the output of the filter equals or exceeds this value.



### 1.9.34 C1DTPEW - Channel 1 DRAM Rank Throttling Passive Event

B/D/F/Type: 0/0/0/MCHBAR  
Address Offset: 6A8-6ABh  
Default Value: 00000000h  
Access: RW-L  
Size: 32 bits

Programmable Event weights are input into the averaging filter. Each Event weight is an normalized 8-bit value that the BIOS must program. The BIOS must account for burst length and 1N/2N rule considerations. It is also possible for BIOS to take into account loading variations of memory caused as a function of memory types and population of ranks. GMCH implements 4 independent filters, one per rank. All bits in this register can be locked by the DTLOCK bit in the C1DTC register.

Bit	Access	Default Value	Description
31:24	RW-L	00h	<i>Reserved</i>
23:16	RW-L	00h	<i>Reserved</i>
15:8	RW-L	00h	<i>Reserved</i>
7:0	RW-L	00h	<i>Reserved</i>





### 1.9.35 C1DTAEW - Channel 1 DRAM Rank Throttling Active Event

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 6AC-6B3h  
 Default Value: 0000000000000000h  
 Access: RO; RW-L  
 Size: 64 bits

Programmable Event weights are input into the averaging filter. Each Event weight is an normalized 8-bit value that the BIOS must program. The BIOS must account for burst length and 1N/2N rule considerations. It is also possible for BIOS to take into account loading variations of memory caused as a function of memory types and population of ranks. GMCH implements four independent filters, one per rank. During a given clock, GMCH asserts a command to the DRAM (via CSB assertion). Based on the command type, one of the weights specified in this register is added to the appropriate weight specified in C1DTPEW and input to the filter. All bits in this register can be locked by the DTLOCK bit in the C1DTC register.

Bit	Access	Default Value	Description
63:48	RO	0000h	<i>Reserved</i>
47:40	RW-L	00h	<i>Reserved</i>
39:32	RW-L	00h	<i>Reserved</i>
31:24	RW-L	00h	<i>Reserved</i>
23:16	RW-L	00h	<i>Reserved</i>
15:8	RW-L	00h	<i>Reserved</i>
7:0	RW-L	00h	<i>Reserved</i>



### 1.9.36 DDRMPLL1 - DDR PLL BIOS

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 2C20-2C21h  
 Default Value: 000Ch  
 Access: RW; RO  
 Size: 16 bits

DDR PLL BIOS Registers Settings

Encoding[7:0] Data edge rate in MHz

0Ch: 800 MHz

10h: 1066 MHz

Bit	Access	Default Value	Description
15:12	RO	0000b	Reserved
11	RW	0b	Reserved
10	RW	0b	Reserved
9:8	RW	00b	Reserved
7	RO	0b	Reserved
6:1	RW	000110b	<b>Feedback Divider Ratio (FBRATIO)</b> Program feedback div value.
0	RO	0b	<b>Feedback Divider Ratio[0] (FBRATIO LSB)</b> Program feedback div value.

## 1.10 Device 0 MCHBAR Thermal Management Controls

(Sheet 1 of 2)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Thermal Sensor Control 1	TSC1	1001	1002	0000h	RW-L; RO; RW; AF
Thermal Sensor Status 1	TSS1	1004	1005	0000h	RO
Thermometer Read 1	TR1	1006	1006	FFh	RO
Thermometer Offset 1	TOF1	1007	1007	00h	RW
Relative Thermometer Read 1	RTR1	1008	1008	00h	RO
Thermal Sensor Temperature Trip Point A1	TSTTPA1	1010	1013	00000000h	RW-L; RO
Thermal Sensor Temperature Trip Point B1	TSTTPB1	1014	1017	00000000h	RW-L
Hardware Throttle Control 1	HWTHROTCTRL1	101C	101C	00h	RW-L; RO; RW-O
Thermal Interrupt Status 1	TIS1	101E	101F	0000h	RO; RW1C
Thermal Error Command	TERRCMD	10E4	10E4	00h	RW; RO



(Sheet 2 of 2)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Thermal SMI Command	TSMICMD	10E5	10E5	00h	RO; RW
Thermal SCI Command	TSCICMD	10E6	10E6	00h	RO; RW
Thermal INTR Command	TINTRCMD	10E7	10E7	00h	RO; RW
External Thermal Sensor Control and Status	EXTTSCS	10EC	10ED	0000h	RW-L; RO; RW-O
MCH Thermal Sensor Watch Dog Timer	MCHTSWDT	12D0	12D3	00000000h	RO; RW; RW 1C
Memory TDP Controller Registers	MEMTDPCTW	2D4	2D7	00000000h	RO; RW
Memory TDP Controller Combined RD/WR Thermal Weight Hot Thresholds	MTDPCCRWTWHOT TH	2F0	2F3	00000000h	RW
Memory TDP Controller Combined RD/WR Thermal Weight Hot Thresholds 2	MTDPCCRWTWHOT TH2	2F4	2F7	00000000h	RW
Memory TDP Controller Combined RD/WR Thermal Weight Hot Thresholds 3	MTDPCCRWTWHOT TH3	2F8	2FB	00000000h	RW
Memory TDP Controller Combined RD/WR Thermal Weight Hot Thresholds 4	MTDPCCRWTWHOT TH4	2FC	2FF	00000000h	RW
Memory TDP Controller Hot Throttled Intervals	MTDPCHOTTHINT	300	303	00000000h	RW
Memory TDP Controller Hot Throttled Intervals 2	MTDPCHOTTHINT2	304	307	00000000h	RW
Memory TDP Controller Aux and Throttle-Non Throttle Intervals	MTDPCTLAUXTNTI NT	308	30B	00000000h	RO; RW
Memory TDP Controller Miscellaneous	MTDPCMISC	30C	30F	00000000h	RO; RW
Thermal Sensor Fuses	TSFUSE	1020	1023	00000000h	RO



### 1.10.1 TSC1 - Thermal Sensor Control 1

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 1001-1002h  
 Default Value: 0000h  
 Access: RW-L; RO; RW; RS-WC  
 Size: 16 bits  
 BIOS Optimal Default 0h

This register controls the operation of the internal thermal sensor located in the graphics region of the die.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
15:14	RO	00b	<i>Reserved</i>
13:10	RW	0h	<p><b>Digital Hysteresis Amount (DHA)</b></p> <p>This bit enables the analog hysteresis control to the thermal sensor. When enabled, about 1 degree of hysteresis is applied. This bit should normally be off in thermometer mode since the thermometer mode of the thermal sensor defeats the usefulness of analog hysteresis.</p> <p>0 = hysteresis disabled            1 = analog hysteresis enabled.</p> <p>This setting falls within the same byte as the In Use bit in bit 8. Therefore if this setting is read, software must write a 1 to bit 8 if it does not intend to maintain ownership of the Thermal Sensor resource.</p>
9:9	RO	0h	<i>Reserved</i>
8	RS-WC	0b	<p><b>In Use (IU)</b></p> <p>Software semaphore bit. After a full processor RESET, a read to this bit returns a 0. After the first read, subsequent reads will return a 1.</p> <p>Writing a 1 to this bit will reset the next read value to 0. Writing a 0 to this bit has no effect.</p> <p>Software can poll this bit until it reads a 0, and will then own the usage of the thermal sensor. This bit has no other effect on the hardware, and is only used as a semaphore among various independent software threads that may need to use the thermal sensor.</p> <p>Software that reads this register but does not intend to claim exclusive access of the thermal sensor must write a one to this bit if it reads a 0, in order to allow other software threads to claim it.</p> <p>See also THERM Bit 15, which is an independent additional semaphore bit.</p>
7:1	RO	00h	<i>Reserved</i>



(Sheet 2 of 2)

Bit	Access	Default Value	Description
0	RW-L	0b	<p><b>Thermal Sensor Enable (TSE)</b></p> <p>This bit enables the thermal sensor logic in the core. The thermal sensor circuit EBB is enabled on PWROK. Lockable via TSTTPA1 Bit 30, TIC1 Bit 7.</p> <p>0 = Disabled 1 = Enabled</p>

### 1.10.2 TSS1 - Thermal Sensor Status 1

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 1004-1005h  
 Default Value: 0000h  
 Access: RO  
 Size: 16 bits

This read only register provides trip point and other status of the thermal sensor.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
15:11	RO	00h	<i>Reserved</i>
10	RO	0b	<p><b>Thermometer Mode Output Valid (TMOV)</b></p> <p>0 = The Thermometer mode is off, or the temperature is out of range, or the TR register is being looked at before a temperature conversion has had time to complete.</p> <p>1 = The Thermometer mode is able to converge to a temperature and the TR register is reporting a reasonable estimate of the thermal sensor temperature.</p>
9	RO	0b	<p><b>Direct Catastrophic Comparator Read (DCCR)</b></p> <p>This bit reads the output of the Catastrophic comparator directly, without latching via the Thermometer mode circuit. Used for testing.</p>
8	RO	0b	<i>Reserved</i>
7:6	RO	00b	<i>Reserved</i>
5	RO	0b	<p><b>Catastrophic Trip Indicator (CTI)</b></p> <p>A 1 indicates that the internal thermal sensor temperature is above the catastrophic setting.</p>
4	RO	0b	<p><b>Hot Trip Indicator (HTI)</b></p> <p>A 1 indicates that the internal thermal sensor temperature is above the Hot setting.</p>
3	RO	0b	<p><b>Aux3 Trip Indicator (A3TI)</b></p> <p>A 1 indicates that the internal thermal sensor temperature is above the Aux3 setting.</p>



(Sheet 2 of 2)

Bit	Access	Default Value	Description
2	RO	0b	<b>Aux2 Trip Indicator (A2TI)</b> A 1 indicates that the internal thermal sensor temperature is above the Aux2 setting.
1	RO	0b	<b>Aux1 Trip Indicator (A1TI)</b> A 1 indicates that the internal thermal sensor temperature is above the Aux1 setting.
0	RO	0b	<b>Aux0 Trip Indicator (A0TI)</b> A 1 indicates that the internal thermal sensor temperature is above the Aux0 setting.

### 1.10.3 TR1 - Thermometer Read 1

B/D/F/Type: 0/0/0/MCHBAR  
Address Offset: 1006h  
Default Value: FFh  
Access: RO  
Size: 8 bits

This register generally provides the un-calibrated counter value from the thermometer circuit when the Thermometer mode is enabled. See the temperature tables for the temperature calculations.

Bit	Access	Default Value	Description
7:0	RO	FFh	<b>Thermometer Reading (TR)</b> Provides the current counter value. The current counter value corresponds to thermal sensor temperature if TSS [Thermometer mode Output Valid] = 1. This register has a straight binary encoding that will range from 00h to FFh. Note: when thermometer mode is disabled via TERATE register, TR will read FFh



### 1.10.4 TOF1 - Thermometer Offset 1

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 1007h  
 Default Value: 00h  
 Access: RW  
 Size: 8 bits

This register is used for programming the thermometer offset.

Bit	Access	Default Value	Description
7:0	RW	00h	<b>Thermometer Offset (TOF)</b> This value is used to adjust the current thermometer reading so that the TR value is not relative to a specific trip or calibration point, and is positive going for positive increases in temperature. The initial default value is 00h and software must determine the correct temperature adjustment that corresponds to a zero reading by reading the fuses and referring to the temperature tables, and then programming the computed offset into this register.

### 1.10.5 RTR1 - Relative Thermometer Read 1

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 1008h  
 Default Value: 00h  
 Access: RO  
 Size: 8 bits

This register contains the relative temperature.

Bit	Access	Default Value	Description
7:0	RO	00h	<b>Relative Thermometer Reading (RTR1)</b> In Thermometer mode, this register reports the relative temperature of the thermal sensor. Provides a two's complement value of the thermal sensor relative to TOF. TR and HTPS can both vary between 0 and 255. $RTR1 = TR + TOF$ . See also TSS [Thermometer mode Output Valid] In the Analog mode, the RTR field reports HTPS value.



### 1.10.6 TSTTPA1 - Thermal Sensor Temperature Trip Point A1

B/D/F/Type: 0/0/0/MCHBAR  
Address Offset: 1010-1013h  
Default Value: 00000000h  
Access: RW-L; RO  
Size: 32 bits

This register:

1. Sets the target values for some of the trip points in thermometer mode. See also TST [Direct DAC Connect Test Enable].
2. Reports the relative thermal sensor temperature See also TSTTPB.

Bit	Access	Default Value	Description
31	RW-L	0b	<b>Lock Bit for Aux0, Aux1, Aux2 and Aux3 Trip Points (AUXLOCK)</b> This bit, when written to a 1, locks the Aux x Trip point settings. It is expected that the Aux x Trip point settings can be changed dynamically when this lock is not set.
30	RW-L	0b	<b>Lock Bit for Catastrophic (LBC)</b> This bit, when written to a 1, locks the Catastrophic programming interface, including Bits 7:0 of TSTTPA[15-0], Bits 15 and 9 of TSC, and Bits 10 and 8 of TST1. This bit may only be set to a 0 by a hardware reset. Writing a 0 to this bit has no effect.
29:16	RO	0000h	<i>Reserved</i>
15:8	RW-L	00h	<b>Hot Trip Point Setting (HTPS)</b> Sets the target value for the Hot trip point. Lockable via TSTTPA1 Bit 30.
7:0	RW-L	00h	<b>Catastrophic Trip Point Setting (CTPS)</b> Sets the target for the Catastrophic trip point. See also TST[Direct DAC Connect Test Enable]. Lockable via TSTTPA1 Bit 30.





### 1.10.7 TSTTPB1 - Thermal Sensor Temperature Trip Point B1

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 1014-1017h  
 Default Value: 00000000h  
 Access: RW-L  
 Size: 32 bits

This register sets the target values for some of the trip points in the Thermometer mode. See also TSTTPA1.

Bit	Access	Default Value	Description
31:24	RW-L	00h	<b>Aux 3 Trip Point Setting (A3TPS)</b> Sets the target value for the Aux3 trip point Lockable by TSTTPA1[31].
23:16	RW-L	00h	<b>Aux 2 Trip Point Setting (A2TPS)</b> Sets the target value for the Aux2 trip point Lockable by TSTTPA1[31].
15:8	RW-L	00h	<b>Aux 1 Trip Point Setting (A1TPS)</b> Sets the target value for the Aux1 trip point Lockable by TSTTPA1[31].
7:0	RW-L	00h	<b>Aux 0 Trip Point Setting (A0TPS)</b> Sets the target value for the Aux0 trip point Lockable by TSTTPA1[31].



### 1.10.8 HWTHROTCTRL1 - Hardware Throttle Control 1

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 101Ch  
 Default Value: 00h  
 Access: RW-L; RO; RW-O  
 Size: 8 bits

Bit	Access	Default Value	Description
7	RW-L	0b	<b>Internal Thermal Hardware Throttling Enable (ITHTE)</b> This bit is a master enable for internal thermal sensor-based hardware throttling: 0 = Hardware actions via the internal thermal sensor are disabled. 1 = Hardware actions via the internal thermal sensor are enabled.
6	RO	0b	<i>Reserved</i>
5	RW-L	0b	<b>Use Direct Catastrophic Trip for HOC (UDCTHOC)</b> 0 = Thermometer comparison to catastrophic trip value is used to control THRMTRIPB. 1 = Catastrophic trip output of DTS circuit is used to control THRMTRIPB.
4	RW-L	0b	<b>Throttle Zone Selection (TZS)</b> This bit determines what temperature zones will enable auto-throttling. This register applies to internal thermal sensor throttling. Lockable by bit0 of this register. 0 = Hot, Aux2, and Catastrophic. 1 = Hot and Catastrophic.
3	RW-L	0b	<b>Halt on Catastrophic (HOC)</b> When this bit is set, THRMTRIPB is asserted on catastrophic trip to bring the platform down. A system reboot is required to bring the system out of a halt from the thermal sensor. Once the catastrophic trip point is reached, THRMTRIPB will stay asserted even if the catastrophic trip deasserts before the platform is shut down.
2:1	RO	00b	<i>Reserved</i>
0	RW-O	0b	<b>Hardware Throttling Lock Bit (HTL)</b> This bit locks Bits 7:1 of this register. When this bit is set to a one, the register bits are locked. It may only be set to a 0 by a hardware reset. Writing a 0 to this bit has no effect.



### 1.10.9 TIS1 - Thermal Interrupt Status 1

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 101E-101Fh  
 Default Value: 0000h  
 Access: RO; RWC  
 Size: 16 bits

This register is used to report which specific error condition resulted in the D2F0 or D2F1 ERRSTS [Thermal Sensor event for SMI/SCI/SERR] or memory mapped IIR Thermal Event. SW can examine the current state of the thermal zones by examining the TSS. Software can distinguish internal or external Trip Event by examining TSS.

(Sheet 1 of 3)

Bit	Access	Default Value	Description
15:14	RO	00b	<i>Reserved</i>
13	RWC	0b	<b>Was Catastrophic Thermal Sensor Interrupt Event (WCTSIE)</b> 0 = No trip for this event. 1 = Indicates that a Catastrophic Thermal Sensor trip based on a higher to lower temperature transition thru the trip point. Software must write a 1 to clear this status bit.
12	RWC	0b	<b>Was Hot Thermal Sensor Interrupt Event (WHTSIE)</b> 0 = No trip for this event. 1 = A Hot Thermal Sensor trip occurred based on a higher to lower temperature transition through the trip point. Software must write a 1 to clear this status bit.
11	RWC	0b	<b>Was Aux 3 Thermal Sensor Interrupt Event (WA3TSIE)</b> 0 = No trip for this event. 1 = An Aux3 Thermal Sensor trip occurred based on a higher to lower temperature transition through the trip point. Software must write a 1 to clear this status bit.
10	RWC	0b	<b>Was Aux 2 Thermal Sensor Interrupt Event (WA2TSIE)</b> 0 = No trip for this event. 1 = An Aux2 Thermal Sensor trip occurred based on a higher to lower temperature transition through the trip point. Software must write a 1 to clear this status bit.
9	RWC	0b	<b>Was Aux 1 Thermal Sensor Interrupt Event (WA1TSIE)</b> 0 = No trip for this event. 1 = An Aux1 Thermal Sensor trip occurred based on a higher to lower temperature transition through the trip point. Software must write a 1 to clear this status bit.



(Sheet 2 of 3)

Bit	Access	Default Value	Description
8	RWC	0b	<p><b>Was Aux 0 Thermal Sensor Interrupt Event (WAOTSIE)</b></p> <p>0 = No trip for this event.            1 = An Aux0 Thermal Sensor trip occurred based on a higher to lower temperature transition through the trip point.            Software must write a 1 to clear this status bit.</p>
7:6	RO	00b	<i>Reserved</i>
5	RWC	0b	<p><b>Catastrophic Thermal Sensor Interrupt Event (CTSIE)</b></p> <p>0 = No trip for this event.            1 = A Catastrophic Thermal Sensor trip event occurred based on a lower to higher temperature transition through the trip point.            Software must write a 1 to clear this status bit.</p>
4	RWC	0b	<p><b>Hot Thermal Sensor Interrupt Event (HTSIE)</b></p> <p>0 = No trip for this event.            1 = A Hot Thermal Sensor trip event occurred based on a lower to higher temperature transition through the trip point.            Software must write a 1 to clear this status bit.</p>
3	RWC	0b	<p><b>Aux 3 Thermal Sensor Interrupt Event (A3TSIE)</b></p> <p>0 = No trip for this event.            1 = An Aux3 Thermal Sensor trip event occurred based on a lower to higher temperature transition through the trip point.            Software must write a 1 to clear this status bit.</p>
2	RWC	0b	<p><b>Aux 2 Thermal Sensor Interrupt Event (A2TSIE)</b></p> <p>0 = No trip for this event.            1 = Indicates that an Aux2 Thermal Sensor trip event occurred based on a lower to higher temperature transition thru the trip point.            Software must write a 1 to clear this status bit.</p>
1	RWC	0b	<p><b>Aux 1 Thermal Sensor Interrupt Event (A1TSIE)</b></p> <p>0 = No trip for this event.            1 = Indicates that an Aux1 Thermal Sensor trip event occurred based on a lower to higher temperature transition thru the trip point.            Software must write a 1 to clear this status bit.</p>



(Sheet 3 of 3)

Bit	Access	Default Value	Description
0	RWC	0b	<p><b>Aux 0 Thermal Sensor Interrupt Event (A0TSIE)</b></p> <p>0 = No trip for this event.                      1 = An Aux0 Thermal Sensor trip event occurred based on a lower to higher temperature transition through the trip point.</p> <p>Software must write a 1 to clear this status bit.</p> <p>The following scenario is possible. An interrupt is initiated on a rising temperature trip, the appropriate DMI cycles are generated, and eventually the software services the interrupt and sees a rising temperature trip as the cause in the status bits for the interrupts.</p>

### 1.10.10 TERATE - Thermometer Mode Enable and Rate

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 1070h  
 Default Value: 00h  
 Access: RO; RW  
 Size: 8 bits

This common register helps select between the analog and the thermometer mode and also helps select the DAC settling timer.

Bit	Access	Default Value	Description
7:4	RO	0h	Reserved



Bit	Access	Default Value	Description
3:0	RW	0h	<p><b>Thermometer Mode Enable and Rate (TE)</b></p> <p>If analog thermal sensor mode is not enabled by setting these bits to 0000b, these bits enable the thermometer mode functions and set the Thermometer controller rate. When the Thermometer mode is disabled and TSC1[TSC] =enabled, the analog sensor mode should be fully functional.</p> <p>In the analog sensor mode, the Catastrophic trip is functional. The other trip points are not functional in this mode.</p> <p>When Thermometer mode is enabled, all the trip points (Catastrophic, Hot, Aux0, Aux1, Aux2 will all operate using the programmed trip points and Thermometer mode rate.</p> <p><b>NOTES:</b></p> <ol style="list-style-type: none"> <li>1. When disabling the Thermometer mode while thermometer running, the Thermometer mode controller will finish the current cycle.</li> <li>2. During boot, all other thermometer mode registers (except lock bits) should be programmed appropriately before enabling the Thermometer Mode.</li> </ol> <p>Thermometer rate select (i.e., AST clock select)</p> <p>0000 = Thermometer mode disabled (i.e., analog sensor mode)</p> <p>0001 = Enabled, 2 <math>\mu</math>sec (normal thermometer mode operation, pre-silicon)</p> <p>0010 = Enabled, 4 <math>\mu</math>sec</p> <p>0011 = Enabled, 6 <math>\mu</math>sec</p> <p>0100 = Enabled, 8 <math>\mu</math>sec</p> <p>0101 = Enabled, 10 <math>\mu</math>sec</p> <p>0110 = Enabled, 12 <math>\mu</math>sec</p> <p>0111 = Enabled, 14 <math>\mu</math>sec</p> <p>All others reserved.</p>



### 1.10.11 TERRCMD - Thermal Error Command

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 10E4h  
 Default Value: 00h  
 Access: RO; RW  
 Size: 8 bits

This register select which errors are generate a SERR DMI interface special cycle, as enabled by ERRCMD [SERR Thermal Sensor event]. The SERR and SCI must not be enabled at the same time for the thermal sensor event.

Bit	Access	Default Value	Description
7:6	RO	00b	<i>Reserved</i>
5	RW	0b	<b>SERR on Catastrophic Thermal Sensor Event (CATSERR)</b> 0 = Disable. Reporting of this condition via SERR messaging is disabled. 1 = Does not mask the generation of a SERR DMI cycle on a catastrophic thermal sensor trip.
4	RW	0b	<b>SERR on Hot Thermal Sensor Event (HOTSERR)</b> 0 = Disable. Reporting of this condition via SERR messaging is disabled. 1 = Do not mask the generation of a SERR DMI cycle on a Hot thermal sensor trip.
3	RW	0b	<b>SERR on Aux 3 Thermal Sensor Event (AUX3SERR)</b> 0 = Disable. Reporting of this condition via SERR messaging is disabled. 1 = Do not mask the generation of a SERR DMI cycle on an Aux3 thermal sensor trip.
2	RW	0b	<b>SERR on Aux 2 Thermal Sensor Event (AUX2SERR)</b> 0 = Disable. Reporting of this condition via SERR messaging is disabled. 1 = Do not mask the generation of a SERR DMI cycle on an Aux2 thermal sensor trip.
1	RW	0b	<b>SERR on Aux 1 Thermal Sensor Event (AUX1SERR)</b> 0 = Disable. Reporting of this condition via SERR messaging is disabled. 1 = Do not mask the generation of a SERR DMI cycle on an Aux1 thermal sensor trip.
0	RW	0b	<b>SERR on Aux 0 Thermal Sensor Event (AUX0SERR)</b> 0 = Disable. Reporting of this condition via SERR messaging is disabled. 1 = Do not mask the generation of a SERR DMI cycle on an Aux0 thermal sensor trip.



### 1.10.12 TSMICMD - Thermal SMI Command

B/D/F/Type: 0/0/0/MCHBAR  
Address Offset: 10E5h  
Default Value: 00h  
Access: RO; RW  
Size: 8 bits

This register selects specific errors to generate a SMI DMI cycle, as enabled by the SMI Error Command Register [SMI on Thermal Sensor Trip].

Bit	Access	Default Value	Description
7:6	RO	00b	<i>Reserved</i>
5	RW	0b	<b>SMI on Catastrophic Thermal Sensor Trip (CATSMI)</b> 0 = Disable reporting of this condition via SMI messaging 1 = Does not mask the generation of an SMI DMI cycle on a catastrophic thermal sensor trip.
4	RW	0b	<b>SMI on Hot Thermal Sensor Trip (HOTSMI)</b> 0 = Disable reporting of this condition via SMI messaging 1 = Does not mask the generation of an SMI DMI cycle on a Hot thermal sensor trip.
3	RW	0b	<b>SMI on AUX3 Thermal Sensor Trip (AUX3SMI)</b> 0 = Disable reporting of this condition via SMI messaging 1 = Does not mask the generation of an SMI DMI cycle on an Aux3 thermal sensor trip.
2	RW	0b	<b>SMI on AUX2 Thermal Sensor Trip (AUX2SMI)</b> 0 = Disable reporting of this condition via SMI messaging 1 = Does not mask the generation of an SMI DMI cycle on an Aux2 thermal sensor trip.
1	RW	0b	<b>SMI on AUX1 Thermal Sensor Trip (AUX1SMI)</b> 0 = Disable reporting of this condition via SMI messaging 1 = Does not mask the generation of an SMI DMI cycle on an Aux1 thermal sensor trip.
0	RW	0b	<b>SMI on AUX0 Thermal Sensor Trip (AUX0SMI)</b> 0 = Disable reporting of this condition via SMI messaging 1 = Does not mask the generation of an SMI DMI cycle on an Aux0 thermal sensor trip.





### 1.10.13 TSCICMD - Thermal SCI Command

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 10E6h  
 Default Value: 00h  
 Access: RO; RW  
 Size: 8 bits

This register selects specific errors to generate a SCI DMI cycle, as enabled by the SCI Error Command Register [SCI on Thermal Sensor Trip]. The SCI and SERR must not be enabled at the same time for the thermal sensor event.

Bit	Access	Default Value	Description
7:6	RO	00b	<i>Reserved</i>
5	RW	0b	<b>SCI on Catastrophic Thermal Sensor Trip (CATSCI)</b> 0 = Disable reporting of this condition via SCI messaging. 1 = Does not mask the generation of an SCI DMI cycle on a catastrophic thermal sensor trip.
4	RW	0b	<b>SCI on Hot Thermal Sensor Trip (HOTSCI)</b> 0 = Disable reporting of this condition via SCI messaging. 1 = Does not mask the generation of an SCI DMI cycle on a Hot thermal sensor trip.
3	RW	0b	<b>SCI on AUX 3 Thermal Sensor Trip (AUX3SCI)</b> 0 = Disable reporting of this condition via SCI messaging. 1 = Does not mask the generation of an SCI DMI cycle on an Aux3 thermal sensor trip.
2	RW	0b	<b>SCI on AUX 2 Thermal Sensor Trip (AUX2SCI)</b> 0 = Disable reporting of this condition via SCI messaging. 1 = Does not mask the generation of an SCI DMI cycle on an Aux2 thermal sensor trip.
1	RW	0b	<b>SCI on AUX 1 Thermal Sensor Trip (AUX1SCI)</b> 0 = Disable reporting of this condition via SCI messaging. 1 = Does not mask the generation of an SCI DMI cycle on an Aux1 thermal sensor trip.
0	RW	0b	<b>SCI on AUX 0 Thermal Sensor Trip (AUX0SCI)</b> 0 = Disable reporting of this condition via SCI messaging. 1 = Does not mask the generation of an SCI DMI cycle on an Aux0 thermal sensor trip.



### 1.10.14 TINTRCMD - Thermal INTR Command

B/D/F/Type: 0/0/0/MCHBAR  
Address Offset: 10E7h  
Default Value: 00h  
Access: RO; RW  
Size: 8 bits

This register selects specific errors to generate an INT DMI cycle

Bit	Access	Default Value	Description
7:6	RO	00b	<i>Reserved</i>
5	RW	0b	<b>INTR on Catastrophic Thermal Sensor Trip (CATINTR)</b> 1 = INTR DMI cycle is generated by the memory controller
4	RW	0b	<b>INTR on Hot Thermal Sensor Trip (HOTINTR)</b> 1 = INTR DMI cycle is generated by the memory controller
3	RW	0b	<b>INTR on AUX3 Thermal Sensor Trip (AUX3INTR)</b> 1 = INTR DMI cycle is generated by the memory controller
2	RW	0b	<b>INTR on AUX2 Thermal Sensor Trip (AUX2INTR)</b> 1 = INTR DMI cycle is generated by the memory controller
1	RW	0b	<b>INTR on AUX1 Thermal Sensor Trip (AUX1INTR)</b> 1 = INTR DMI cycle is generated by the memory controller
0	RW	0b	<b>INTR on AUX0 Thermal Sensor Trip (AUX0INTR)</b> 1 = INTR DMI cycle is generated by the memory controller



### 1.10.15 EXTTS - External Thermal Sensor Control and Status

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 10EC-10EDh  
 Default Value: 0000h  
 Access: RO; RW-O; RW-L  
 Size: 16 bits

(Sheet 1 of 2)

Bit	Access	Default Value	Description
15	RW-O	0b	<p><b>External Sensor Enable (ESE)</b></p> <p>Setting this bit to 1 locks the lockable bits in this register. This bit may only be set to a zero by a hardware reset. Once locked, writing a 0 to bit has no effect. EXTTS0 and EXTTS1 input signal pins are dedicated for external thermal sensor use.</p> <p>An asserted External Thermal Sensor Trip signal can also cause a SCI, SMI, SERR or INTR interrupt in the same manner as the Internal Sensor can. A 0 on the pins can be used to trigger throttling.</p> <p>If both internal sensor throttling and external sensor throttling are enabled, either can initiate throttling. The AS0 and AS1 bits of this register allow control of what action is triggered by external sensor trips. The memory controller Throttling select bit controls the type of throttling action that will happen, and the {AS0, AS1} bits control what trip actions will result.</p> <p>0 = External Sensor input is disabled.                      1 = External Sensor input is enabled.</p>
14	RO	0b	<i>Reserved</i>
13	RW-L	0b	<p><b>Select between EXTTS PIN 0 and 1 (EXTPINSEL)</b></p> <p>0 = Use EXTTS Pin 0 for Thermal throttling, based of EXTPMTRIP, EXTTFMX and SD2X.                      1 = Use EXTTS Pin 1 for the above.</p>
12	RW-L	0b	<p><b>EXTTS Based Power Monitor Trip (EXTPMTRIP)</b></p> <p>When this bit is 1, EXTTS Bit 0 can be programmed to look like a power-monitor trip</p> <ul style="list-style-type: none"> <li>• Will be OR'ed with the Global monitor/Gfx monitor so that, when programmed for gfx throttle, when EXTTS# is asserted at the sample point, it will look like a monitor trip and force RP down by the programmed amount</li> <li>• EXTTS# is only sampled on the sampling window for graphics throttling, so even if both the Gfx monitor and global monitor are disabled, the sampling window must be programmed in order to have EXTTS# work as a graphics throttle</li> </ul>
11	RW-L	0b	<p><b>Force DDR on EXTTS bit (EXTTFMX)</b></p> <p>Enables forcing of DDR and PEG to specified MX state in registers EXTTSMXST when the selected EXTTS Bit 0 or 1 (from EXTPINSEL field) is asserted.</p>



(Sheet 2 of 2)

Bit	Access	Default Value	Description
10:8	RW-L	000b	<b>EXTTSS Programmable MX state (EXTTSMXST)</b> MX state to which DDR/PEG to be forced to if EXTTS Bit 0 asserts and Force DDR on EXTTS bit (EXTTFMX) is enabled.
7	RW-L	0b	<i>Reserved</i>
6	RW-L	0b	<b>Throttling Type Select (TTS)</b> Lockable by EXTTS [External Sensor Enable]. If External Thermal Sensor Enable = 1, then: 0 = DRAM throttling based on the settings in the Device 0 MCHBAR Dram Throttling Control register. 1 = Memory controller throttling, based on the settings in the Device 0 MCHBAR Memory controller Throttling Control Register and the Device 2 Graphics Render Throttle Control Register [Catastrophic and Hot Hardware controlled Thermal Throttle Duty Cycle].
5	RW-L	0b	<b>EXTTS1 Action Select (AS1)</b> Lockable by EXTTS [External Sensor Enable]. If External Thermal Sensor Enable = 1, then: 0 = The external sensor trip functions same as a Thermometer mode hot trip. 1 = The external sensor trip functions same as a Thermometer mode aux0 trip.
4	RW-L	0b	<b>EXTTS0 Action Select (AS0)</b> Locatable by EXTTS [External Sensor Enable]. If External Thermal Sensor Enable = 1, then 0 = The external sensor trip functions same as a Thermometer mode catastrophic trip. 1 = The external sensor trip functions same as a Thermometer mode hot trip.
3	RO	0b	<b>EXTTS0 Trip Indicator (SOT1)</b> 0 = The external sensor trip is exceeding the programmed setting of its external thermal sensor. 1 = The external sensor trip is not exceeding the programmed setting of its external thermal sensor.
2	RO	0b	<b>EXTTS1 Trip Indicator (S1T1)</b> 0 = The external sensor trip is exceeding the programmed setting of its external thermal sensor. 1 = The external sensor trip is not exceeding the programmed setting of its external thermal sensor.
1	RO	0b	<i>Reserved</i>
0	RW-L	0b	<b>External Thermal Sensor Signals Routing Control (EXTTSSRC)</b> 0 = Rote all external sensor signals to affect internal thermal sensor registers, as appropriate. 1 = No affect of external sensor signals to internal thermal sensor registers.



### 1.10.16 MCHTSWDT - Memory Controller Thermal Sensor Watch Dog Timer

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 12D0-12D3h  
 Default Value: 00000000h  
 Access: RO; RW; RWC  
 Size: 32 bits

When thermally hot tripped and memory controller throttling is enabled, this register allows the value in the TSWD0[Delta] field to affect the impact of the MCHETT[PEWAT] whenever the MCHTSWDT WDT times out.

Bit	Access	Default Value	Description
31	RW	0b	<b>Enable WDT (ENWDT)</b> 0 = WDT function is not enabled. WDT has no impact to throttling. 1 = WDT function is enabled.
30	RW	0b	Reserved
29	RW	0b	Reserved
28:21	RW	00h	Reserved
20	RW	0b	Reserved
19:16	RW	0000b	Reserved
15:8	RO	00h	Reserved
7:5	RO	000b	Reserved
4	RWC	0b	Reserved
3:0	RW	0h	Reserved

### 1.10.17 MEMTDPCTW - Memory TDP Controller Registers

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 2D4-2D7h  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits

Bit	Access	Default Value	Description
31:24	RW	00h	Reserved
23:16	RW	00h	Reserved
15:8	RW	00h	Reserved
7:3	RO	00h	Reserved
2:0	RW	000b	Reserved



### 1.10.18 MTDPCRWTWHOTTH - Memory TDP Controller Combined RD/WR Thermal Weight Hot Thresholds

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 2F0-2F3h  
 Default Value: 00000000h  
 Access: RW  
 Size: 32 bits

The settings in these registers apply to the combined memory Rd/Wr thermal weight trackers in both ch0 and ch1.

The range of threshold supported allows sub-millisecond differentiation among the thresholds. For example,

let MEMCRWTWHOTTH[38:23]=0007h and MEMCRWTWHOTM1TH[38:23]=0006h;  
 let TW counter operational clock frequency be 333 MHz mb4clk=3 ns period;  
 let TW value[7:0]= FFh;  
 Elapsed time to cross these 2 thresholds =  $(2^{23} / 2^8) * 3 \text{ ns} = 98,304 \text{ ns} = 0.098 \text{ ms}$ .

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31:16	RW	0000h	<p><b>Memory Combined Rd/Wr Thermal Weight Hot Threshold (MemCRWTWHotTh)</b></p> <p>The hot thermal weight threshold used for memory combined Rd/Wr thermal weight tracking. This 16-bit value is compared to the MSB (i.e., [38:23]) of the thermal weight counter value, to determine whether the HOT thermal weight threshold is reached or not.</p> <p>Since there are 8 NearHot-to-Hot levels defined, the minimum that this field should be set to i0008h, to avoid overlapping levels if all 8 NearHot-to-Hot thresholds are to be used. Overlapping levels can certainly be set when less than 8 thresholds are to be used. To use 1 threshold setting only, effectively disabling adaptive throttling, all 8 thresholds can be set to the same value.</p>



(Sheet 2 of 2)

Bit	Access	Default Value	Description
15:0	RW	0000h	<p><b>Memory Combined Rd/Wr Thermal Weight Hot-1 Threshold (MemCRWTWHotM1Th)</b></p> <p>The Hot-1 thermal weight threshold used for memory combined Rd/Wr thermal weight tracking. This 16-bit value is compared to the MSB (i.e., [38:23]) of the thermal weight counter value, to determine whether the HOT-1 thermal weight threshold is reached or not.</p> <p>This threshold level should be set lower than the Hot threshold, but higher than the Hot-2 to Hot-7 threshold levels. Since there are 8 NearHot-to-Hot levels defined, the minimum that this field should be set to is 0007h, to avoid overlapping levels if all 8 NearHot-to-Hot thresholds are to be used. Similarly, the maximum that this field can be set to is FFFEh, if all 8 NearHot-to-Hot thresholds are to be used. Overlapping levels can certainly be set when less than 8 thresholds are to be used. To use 1 threshold setting only, effectively disabling adaptive throttling, all 8 thresholds can be set to the same value.</p>



### 1.10.19 MTDPCRWTWHOTTH2 - Memory TDP Controller Combined RD/Wr Thermal Weight Hot Thresholds 2

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 2F4-2F7h  
 Default Value: 00000000h  
 Access: RW  
 Size: 32 bits

Bit	Access	Default Value	Description
31:16	RW	0000h	<p><b>Memory Combined Rd/Wr Thermal Weight Hot-2 Threshold (MemCRWTWHotM2Th)</b></p> <p>The Hot-2 thermal weight threshold used for memory combined Rd/Wr thermal weight tracking. This 16-bit value is compared to the MSB (i.e., [38:23]) of the thermal weight counter value, to determine whether the HOT-2 thermal weight threshold is reached or not.</p> <p>This threshold level should be set lower than the Hot to Hot-1 threshold, but higher than the Hot-3 to Hot-7 threshold levels. Since there are 8 NearHot-to-Hot levels defined, the minimum that this field should be set to is 0006h, to avoid overlapping levels if all 8 NearHot-to-Hot thresholds are to be used. Similarly, the maximum that this field can be set to is FFFDh, if all 8 NearHot-to-Hot thresholds are to be used. Overlapping levels can certainly be set when less than 8 thresholds are to be used. To use 1 threshold setting only, effectively disabling adaptive throttling, all 8 thresholds can be set to the same value.</p>
15:0	RW	0000h	<p><b>Memory Combined Rd/Wr Thermal Weight Hot-3 Threshold (MemCRWTWHotM3Th)</b></p> <p>The Hot-3 thermal weight threshold used for memory combined Rd/Wr thermal weight tracking. This 16-bit value is compared to the MSB (i.e., [38:23]) of the thermal weight counter value, to determine whether the HOT-3 thermal weight threshold is reached or not.</p> <p>This threshold level should be set lower than the Hot to Hot-2 threshold, but higher than the Hot-4 to Hot-7 threshold levels. Since there are 8 NearHot-to-Hot levels defined, the minimum that this field should be set to is 0005h, to avoid overlapping levels if all 8 NearHot-to-Hot thresholds are to be used. Similarly, the maximum that this field can be set to is FFFCh, if all 8 NearHot-to-Hot thresholds are to be used. Overlapping levels can certainly be set when less than 8 thresholds are to be used. To use 1 threshold setting only, effectively disabling adaptive throttling, all 8 thresholds can be set to the same value.</p>





### 1.10.20 MTDPCRWTWHOTTH3 - Memory TDP Controller Combined RD/WR Thermal Weight Hot Thresholds 3

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 2F8-2FBh  
 Default Value: 00000000h  
 Access: RW  
 Size: 32 bits

Bit	Access	Default Value	Description
31:16	RW	0000h	<p><b>Memory Combined Rd/Wr Thermal Weight Hot-4 Threshold (MemCRWTWHotM4Th)</b></p> <p>The Hot-4 thermal weight threshold used for memory combined Rd/Wr thermal weight tracking. This 16-bit value is compared to the MSB (i.e., [38:23]) of the thermal weight counter value, to determine whether the HOT-4 thermal weight threshold is reached or not.</p> <p>This threshold level should be set lower than the Hot to Hot-3 threshold, but higher than the Hot-5 to Hot-7 threshold levels. Since there are 8 NearHot-to-Hot levels defined, the minimum that this field should be set to is 0004h, to avoid overlapping levels if all 8 NearHot-to-Hot thresholds are to be used. Similarly, the maximum that this field can be set to is FFFBh, if all 8 NearHot-to-Hot thresholds are to be used. Overlapping levels can certainly be set when less than 8 thresholds are to be used. To use 1 threshold setting only, effectively disabling adaptive throttling, all 8 thresholds can be set to the same value.</p>
15:0	RW	0000h	<p><b>Memory Combined Rd/Wr Thermal Weight Hot-5 Threshold (MemCRWTWHotM5Th)</b></p> <p>The Hot-5 thermal weight threshold used for memory combined Rd/wr thermal weight tracking. This 16-bit value is compared to the MSB (i.e., [38:23]) of the thermal weight counter value, to determine whether the HOT-5 thermal weight threshold is reached or not.</p> <p>This threshold level should be set lower than the Hot to Hot-4 threshold, but higher than the Hot-6 to Hot-7 threshold levels. Since there are 8 NearHot-to-Hot levels defined, the minimum that this field should be set to is 0003h, to avoid overlapping levels if all 8 NearHot-to-Hot thresholds are to be used. Similarly, the maximum that this field can be set to is FFFAh, if all 8 NearHot-to-Hot thresholds are to be used. Overlapping levels can certainly be set when less than 8 thresholds are to be used. To use 1 threshold setting only, effectively disabling adaptive throttling, all 8 thresholds can be set to the same value.</p>



### 1.10.21 MTDPCCRWTWHOTTH4 - Memory TDP Controller Combined RD/WR Thermal Weight Hot Thresholds 4

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 2FC-2FFh  
 Default Value: 00000000h  
 Access: RW;  
 Size: 32 bits

Bit	Access	Default Value	Description
31:16	RW	0000h	<p><b>Memory Combined Rd/Wr Thermal Weight Hot-6 Threshold (MemCRWTWHotM6Th)</b></p> <p>The Hot-6 thermal weight threshold used for memory combined Rd/Wr thermal weight tracking. This 16-bit value is compared to the MSB (i.e., [38:23]) of the thermal weight counter value, to determine whether the HOT-6 thermal weight threshold is reached or not.</p> <p>This threshold level should be set lower than the Hot to Hot-5 threshold, but higher than the Hot-7 threshold level. Since there are 8 NearHot-to-Hot levels defined, the minimum that this field should be set to is 0002h, to avoid overlapping levels if all 8 NearHot-to-Hot thresholds are to be used. Similarly, the maximum that this field can be set to is FFF9h, if all 8 NearHot-to-Hot thresholds are to be used. Overlapping levels can certainly be set when less than 8 thresholds are to be used. To use 1 threshold setting only, effectively disabling adaptive throttling, all 8 thresholds can be set to the same value.</p>
15:0	RW	0000h	<p><b>Memory Combined Rd/Wr Thermal Weight Hot-7 Threshold (MemCRWTWHotM7Th)</b></p> <p>The Hot-7 thermal weight threshold used for memory combined Rd/Wr thermal weight tracking. This 16-bit value is compared to the MSB (i.e., [38:23]) of the thermal weight counter value, to determine whether the HOT-7 thermal weight threshold is reached or not.</p> <p>This threshold level should be set lower than the Hot to Hot-6 threshold. Since there are 8 NearHot-to-Hot levels defined, the minimum that this field should be set to is 0001h, to avoid overlapping levels if all 8 NearHot-to-Hot thresholds are to be used. Similarly, the maximum that this field can be set to is FFF8h, if all 8 NearHot-to-Hot thresholds are to be used. Overlapping levels can certainly be set when less than 8 thresholds are to be used. To use 1 threshold setting only, effectively disabling adaptive throttling, all 8 thresholds can be set to the same value.</p>



## 1.10.22 MTDpchotthint - Memory TDP Controller Hot Throttled Intervals

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 300-303h  
 Default Value: 00000000h  
 Access: RW  
 Size: 32 bits

These registers control the duty cycle of throttling. The total throttled + non-throttled interval can be from 32 up to 256 MCCLKS. Since throttling could be done for up to 256 clocks, 8-bit fields are needed to specify the number of clocks being throttled.

Memory will be throttled for the most number of clocks when the Hot threshold is reached, and for progressively fewer number of mb4clks for Hot-1, Hot-2, ..., Hot-7 threshold trips. The register settings should conform to this expected behavior.

Bit	Access	Default Value	Description
31:24	RW	00h	<b>Memory IFC Hot Throttle Interval (MIHOTThrotInt)</b> The number of clocks, memory would be throttled for Rd or Wr operations within the period defined by the MITNTInt register field, when the Hot threshold is tripped.
23:16	RW	00h	<b>Memory IFC Hot Minus 1 Throttle Interval (MIHOTM1ThrotInt)</b> The number of clocks, memory would be throttled for Rd or Wr operations within the period defined by the MITNTInt register field, when the Hot-1 threshold is tripped.
15:8	RW	00h	<b>Memory IFC Hot Minus 2 Throttle Interval (MIHOTM2ThrotIn)</b> The number of clocks, memory would be throttled for Rd or Wr operations within the period defined by the MITNTInt register field, when the Hot-2 threshold is tripped.
7:0	RW	00h	<b>Memory IFC Hot Minus 3 Throttle Interval (MIHOTM3ThrotInt)</b> The number of clocks, memory would be throttled for Rd or Wr operations within the period defined by the MITNTInt register field, when the Hot-3 threshold is tripped.



### 1.10.23 MTDpchotthint2 - Memory TDP Controller Hot Throttled Intervals 2

B/D/F/Type: 0/0/0/MCHBAR  
Address Offset: 304-307h  
Default Value: 00000000h  
Access: RW  
Size: 32 bits

Bit	Access	Default Value	Description
31:24	RW	00h	<b>Memory IFC Hot Minus 4 Throttle Interval (MIHOTM4ThrotInt)</b> The number of clocks, memory would be throttled for Rd or Wr operations within the period defined by the MITNTInt register field, when the Hot-4 threshold is tripped.
23:16	RW	00h	<b>Memory IFC Hot Minus 5 Throttle Interval (MIHOTM5ThrotInt)</b> The number of clocks, memory would be throttled for Rd or Wr operations within the period defined by the MITNTInt register field, when the Hot-5 threshold is tripped.
15:8	RW	00h	<b>Memory IFC Hot Minus 6 Throttle Interval (MIHOTM6ThrotIn)</b> The number of clocks, memory would be throttled for Rd or Wr operations within the period defined by the MITNTInt register field, when the Hot-6 threshold is tripped.
7:0	RW	00h	<b>Memory IFC Hot Minus 7 Throttle Interval (MIHOTM7ThrotInt)</b> The number of clocks, memory would be throttled for Rd or Wr operations within the period defined by the MITNTInt register field, when the Hot-7 threshold is tripped.



### 1.10.24 MTDPCTLAUXTNTINT - Memory TDP Controller Aux and Throttle-NonThrottle Intervals

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 308-30Bh  
 Default Value: 00000000h  
 Access: RW; RO  
 Size: 32 bits

Bit	Access	Default Value	Description
31:24	RO	00h	<i>Reserved</i>
23:20	RW	0h	<i>Reserved</i>
19:16	RW	0h	<i>Reserved</i>
15:10	RO	00h	<i>Reserved</i>
9:8	RW	00b	<b>Memory Interface Throttling Plus Non Throttling Interval (MITNTInt)</b> 00: 32 mb4clks 01: 64 mb4clks 10: 128 mb4clks 11: 256 mb4clks
7:0	RW	00h	<i>Reserved</i>

### 1.10.25 MTDPCMISC - Memory TDP Controller Miscellaneous Control

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 30C-30Fh  
 Default Value: 00000000h  
 Access: RW; RO  
 Size: 32 bits

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31:24	RO	00h	<b>Memory IFC Thermal Weight Counter Read Out (MTWCRO)</b> This field reports the current value of the selected thermal weight counter bits.
23:19	RW	00h	<i>Reserved</i>
18	RW	0b	<i>Reserved</i>
17	RW	0b	<i>Reserved</i>
16	RW	0b	<i>Reserved</i>
15	RO	0b	<i>Reserved</i>
14	RW	0b	<i>Reserved</i>
13	RW	0b	<i>Reserved</i>



(Sheet 2 of 2)

Bit	Access	Default Value	Description
12	RW	0b	Reserved
11	RW	0b	Reserved
10	RW	0b	Reserved
9	RW	0b	Reserved
8	RW	0b	Reserved
7:4	RO	0h	Reserved
3:0	RW	0h	Reserved

### 1.10.26 TSFUSE - Thermal Sensor Fuses

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 1020-1023h  
 Default Value: 00000000h  
 Access: RO  
 Size: 32 bits

This register reads the thermal sensor fuses.

Calibration is a linear model of the form  $y=Mx-B$  where

- y is the calibrated temperature,
- M is the slope correction factor,
- x is the raw temperature from the thermal sensor,
- B is the intercept at the y axis, a constant offset.

The hardware uses either the 8-bit mode by default, or the 10-bit mode when enabled in TS10BITMCTRL.

Bit	Access	Default Value	Description
31:24	RO	00h	Reserved
23:16	RO	00h	<b>OVIDR3F (OVID Range 3 Fuses)</b> Offset to the Voltage Identifier value to use when operating graphics render at LFM. The MSB is the sign bit
15:8	RO	00h	<b>Thermal Sensor M-fuse value (slope) for 8-bit mode (TS8BITSLOPE)</b> Software reads the fuse value directly for temperature calibration correction.
7:0	RO	00h	<b>Thermal Sensor B-fuse value (intercept) for 8-bit mode (TS8BITINTCPT)</b> Software reads the fuse value directly for temperature calibration correction.



## 1.11 MCHBAR Render Thermal Throttling Controls

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Thermal State Control	THERMSTCTL	1174	1177	00000000h	RW; RO
Render Standby State Control	RSTDBYCTL	11B8	11BB	00000000h	RW; RO
VID Control	VIDCTL	11C0	11C3	00000000h	RW;
Watchdog Timer For Thermal Sensor Trip	WDTMTRS	11D4	11D6	000000h	RW
Watchdog Timer Based Px Step Size	WDTSTPSZ	11D7	11D7	00h	RW

### 1.11.1 THERMSTCTL – Render Thermal State Control

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 1174-1177h  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits  
 BIOS Optimal Default 0h

This register bit field shall contain the default value unless otherwise indicated in the BIOS specification.

Bit	Access	Default Value	Description
31	RW	0b	Reserved
30	RW	0b	Reserved
29	RW	0b	Reserved
28	RW	0b	Reserved
27:25	RO	000b	Reserved
24	RW	0b	Reserved
23:20	RO	0h	Reserved
19:16	RW	0h	Reserved
15:12	RO	0h	Reserved
11:8	RW	0h	Reserved
7:4	RW	0000b	Reserved
3:0	RW	0h	Reserved



### 1.11.2 RSTDBYCTL - Render Standby State Control

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 11B8-11BBh  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits  
 RS2 = Render Standby with Context Restore.

Bit	Access	Default Value	Description
31	RW	0b	Reserved
30	RW	0b	<b>RS2 Enable (RS2EN)</b> 0 = RS2 not enabled 1 = RS2 enabled
29	RW	0b	Reserved
28	RW	0b	Reserved
27	RW	0b	Reserved
26	RW	0b	Reserved
25	RW	0b	Reserved
24	RW	0b	Reserved
23	RW	0b	Reserved
22:20	RO	000b	Reserved
19	RW	0b	Reserved
18	RW	0b	Reserved
17	RW	0b	Reserved
16	RW	0b	<b>Allow RS2 when in C0 (RS2INC0)</b>
15:14	RW	00b	Reserved
13:12	RW	00b	Reserved
11	RW	0b	Reserved
10	RW	0b	Reserved
9	RW	0b	Reserved
8	RW	0b	Reserved
7	RW	0b	Reserved
6	RO	0b	Reserved
5:4	RW	00b	Reserved
3	RW	0b	Reserved
2	RW	0b	Reserved
1:0	RO	00b	Reserved





### 1.11.3 VIDCTL - VID Control

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 11C0-11C3h  
 Default Value: 00000000h  
 Access: RW  
 Size: 32 bits

This register bit field shall contain the default value unless otherwise indicated in the BIOS specification.

Bit	Access	Default Value	Description
31:24	RW	00h	Reserved
23:16	RW	00h	Reserved
15:8	RW	00h	Reserved
7:0	RW	00h	Reserved

### 1.11.4 WDTMRTS - Watchdog Timer For Thermal Sensor Trip

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 11D4-11D6h  
 Default Value: 000000h  
 Access: RW;  
 Size: 24 bits

Bit	Access	Default Value	Description
23:0	RW	000000h	Reserved

### 1.11.5 WDTSTPSZ - Watchdog Timer Based Px Step Size

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 11D7h  
 Default Value: 00h  
 Access: RW;  
 Size: 8 bits

Bit	Access	Default Value	Description
7:4	RW	0h	Reserved
3:0	RW	0h	Reserved



## 1.12 Device 0 MCHBAR ACPI Power Management Controls

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
C3/C6 EntryTimers	C3C6ET	1200	1203	00000000h	RW; RO
Self-Refresh Channel Status	SLFRCS	1211	1212	0000h	RW1C-S; RO
PM Memory Subsystem	DSLFRCS	120C	120D	0300h	RO; RW
Power Management Configuration	PMCFG	1210	1210	04h	RW; RW-S

### 1.12.1 C3C6ET - C3/C6 EntryTimers

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 1200-1203h  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits

Bit	Access	Default Value	Description
31:30	RO	00b	<i>Reserved</i>
29:22	RW	00h	<b>C6 Entry Timer (C6ET)</b> Dual purpose timer in 64 core clock granularity. Number of host clocks (133 MHz) to wait between last snoop from PEG, DMI or GFX to allowing (C6) entry from the processor. 00h:0 * 64 = 0 host clocks FFh:255 * 64 = 16320 host clocks MSIs, for the purpose of this register, are handled as snoops.
21:16	RO	00h	<i>Reserved</i>
15:14	RO	00b	<i>Reserved</i>
13:6	RW	00h	<b>C3 Entry Timer (C3ET)</b> Dual purpose timer in 64 core clock granularity. Number of host clocks (133 MHz) to wait between last snoop from PEG, DMI or GFX to allowing (C3) entry from the CPU. 00h:0 * 64 = 0 $\mu$ s FFh:255 * 64 = 16320 host clocks MSIs, for the purpose of this register, are handled as snoops.
5:0	RO	00h	<i>Reserved</i>



### 1.12.2 SLFRCS - Self-Refresh Channel Status

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 1211-1212h  
 Default Value: 0000h  
 Access: RWC-P; RO  
 Size: 16 bits

This register is Reset by PWROK only.

Bit	Access	Default Value	Description
15:2	RO	0000h	Reserved
1	RWC-P	0b	Reserved
0	RWC-P	0b	Reserved

### 1.12.3 DSLFRC - PM Memory Subsystem

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 120C-120Dh  
 Default Value: 0300h  
 Access: RO; RW  
 Size: 16 bits

Bit	Access	Default Value	Description
15:11	RO	00h	Reserved
10	RW	0b	Reserved
9:8	RW	11b	Reserved
7:3	RO	00h	Reserved
2	RW	0b	Reserved
1	RW	0b	Reserved
0	RW	0b	<b>C-state Dependency for Self Refresh (CXSR)</b> 0 = Memory is not allowed to enter self-refresh during C3/C6/ 1 = Memory is allowed to enter self-refresh during C3/C6/



### 1.12.4 PMCFG - Power Management Configuration

B/D/F/Type: 0/0/0/MCHBAR  
 Address Offset: 1210h  
 Default Value: 04h  
 Access: RW; RW-P  
 Size: 8 bits

This Register bit field shall contain the default unless otherwise indicated in the BIOS Specification.

### 1.13 PCI Device1

Device 1 contains the controls associated with the PCI Express x16 root port that is the intended to attach as the point for external graphics. It is typically referred to as PCI EXPRESS-G (PCI Express graphics) port. In addition, it also functions as the virtual PCI-to-PCI bridge.

**Warning:** When reading the PCI Express “conceptual” registers such as this, you may not get a valid value unless the register value is stable.

The PCI Express based specification defines two types of reserved bits.

Reserved and Preserved:

1. Reserved for future RW implementations; software must preserve value read for writes to bits.
2. Reserved and Zero: Reserved for future R/WC/S implementations; software must use 0 for writes to bits.

Unless explicitly documented as Reserved and Zero, all bits marked as reserved are part of the Reserved and Preserved type, which have historically been the typical definition for Reserved.

It is important to note that most (if not all) control bits in this device cannot be modified unless the link is down. Software is required to first Disable the link, then program the registers, and then re-enable the link (which will cause a full-retrain with the new settings).

(Sheet 1 of 3)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Vendor Identification	VID1	0	1	8086h	RO
Device Identification	DID1	2	3	0045h	RO
PCI Command	PCICMD1	4	5	0000h	RO; RW
PCI Status	PCISTS1	6	7	0010h	RO; RWC



(Sheet 2 of 3)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Revision Identification	RID1	8	8	12h	RO
Class Code	CC1	9	B	060400h	RO
Cache Line Size	CL1	C	C	00h	RW
Header Type	HDR1	E	E	01h	RO
Primary Bus Number	PBUSN1	18	18	00h	RO
Secondary Bus Number	SBUSN1	19	19	00h	RW
Subordinate Bus Number	SUBUSN1	1A	1A	00h	RW
I/O Base Address	IOBASE1	1C	1C	F0h	RO; RW
I/O Limit Address	IOLIMIT1	1D	1D	00h	RO; RW
Secondary Status	SSTS1	1E	1F	0000h	RWC; RO
Memory Base Address	MBASE1	20	21	FFF0h	RO; RW
Memory Limit Address	MLIMIT1	22	23	0000h	RO; RW
Prefetchable Memory Base Address	PMBASE1	24	25	FFF1h	RO; RW
Prefetchable Memory Limit Address	PMLIMIT1	26	27	0001h	RO; RW
Prefetchable Memory Base Address Upper	PMBASEU1	28	2B	00000000h	RW
Prefetchable Memory Limit Address Upper	PMLIMITU1	2C	2F	00000000h	RW
Capabilities Pointer	CAPPTR1	34	34	88h	RO
Interrupt Line	INTRLINE1	3C	3C	00h	RW
Interrupt Pin	INTRPIN1	3D	3D	01h	RO
Bridge Control	BCTRL1	3E	3F	0000h	RO; RW
Capabilities List Control	CAPL	7F	7F	02h	RO; RW
Power Management Capabilities	PM_CAPID1	80	83	C8039001h	RO
Power Management Control/Status	PM_CS1	84	87	00000008h	RO; RW-S; RW



(Sheet 3 of 3)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Subsystem ID and Vendor ID Capabilities	SS_CAPID	88	8B	0000800Dh	RO
Subsystem ID and Subsystem Vendor ID	SS	8C	8F	00008086h	RW-O
Message Signaled Interrupts Capability ID	MSI_CAPID	90	91	A005h	RO
Message Control	MC	92	93	0000h	RO; RW
Message Address	MA	94	97	00000000h	RO; RW
Message Data	MD	98	99	0000h	RW
PCI Express-G Capability List	PEG_CAPL	A0	A1	0010h	RO
PCI Express-G Capabilities	PEG_CAP	A2	A3	0142h	RO; RW-O
Device Capabilities	DCAP	A4	A7	00008000h	RO
Device Control	DCTL	A8	A9	0000h	RO; RW
Device Status	DSTS	AA	AB	0000h	RO; RWC
Link Capabilities	LCAP	AC	AF	02214D02h	RO; RW-O
Link Control	LCTL	B0	B1	0000h	RO; RW; RW-SC
Link Status	LSTS	B2	B3	1000h	RWC; RO
Slot Capabilities	SLOTCAP	B4	B7	00040000h	RW-O; RO
Slot Control	SLOTCTL	B8	B9	0000h	RO; RW
Slot Status	SLOTSTS	BA	BB	0000h	RO; RWC
Root Control	RCTL	BC	BD	0000h	RO; RW
Root Status	RSTS	C0	C3	00000000h	RO; RWC
Link Control 2	LCTL2	D0	D1	0002h	RO; RW-S; RW
Link Status 2	LSTS2	D2	D3	0000h	RO
PCI Express-G Legacy Control	PEGLC	EC	EF	00000000h	RO; RW



### 1.13.1 VID1 - Vendor Identification

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 0-1h  
 Default Value: 8086h  
 Access: RO  
 Size: 16 bits

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

Bit	Access	Default Value	Description
15:0	RO	8086h	<b>Vendor Identification (VID1)</b> PCI standard identification for Intel.

### 1.13.2 DID1 - Device Identification

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 2-3h  
 Default Value: 0045h  
 Access: RO  
 Size: 16 bits

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

Bit	Access	Default Value	Description
15:4	RO	004h	<b>Device Identification Number (DID1(UB))</b> Identifier assigned to the processor Device 1 (virtual PCI-to-PCI bridge, PCI Express Graphics port).
3:2	RO	00b	<b>Device Identification Number (DID1(HW))</b> Identifier assigned to the processor Device 1 (virtual PCI-to-PCI bridge, PCI Express Graphics port).
1:0	RO	01b	<b>Device Identification Number (DID1(LB))</b> Identifier assigned to the processor Device 1 (virtual PCI-to-PCI bridge, PCI Express Graphics port).



### 1.13.3 PCICMD1 - PCI Command

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 4-5h  
 Default Value: 0000h  
 Access: RO; RW  
 Size: 16 bits

(Sheet 1 of 2)

Bit	Access	Default Value	Description
15:11	RO	00h	<i>Reserved</i>
10	RW	0b	<p><b>INTA Assertion Disable (INTAAD)</b></p> <p>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.</p> <p>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.</p>
9	RO	0b	<p><b>Fast Back-to-Back Enable (FB2B)</b></p> <p>Not Applicable or Implemented. hard wired to 0.</p>
8	RW	0b	<p><b>SERR# Message Enable (SERRE1)</b></p> <p>Controls Device 1 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.</p> <p>In addition, 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.</p> <p>0 = The SERR message is generated by the processor for Device 1 only under conditions enabled individually through the Device Control Register.            1 = The processor is enabled to generate SERR messages which is sent to the PCH for specific Device 1 error conditions generated/detected on the primary side of the virtual PCI to PCI bridge (not those received by the secondary side). The status of SERRs generated is reported in the PCISTS1 register.</p>
7	RO	0b	<p><i>Reserved</i></p> <p>Not Applicable or Implemented. Hard wired to 0.</p>





(Sheet 2 of 2)

Bit	Access	Default Value	Description
6	RW	0b	<p><b>Parity Error Response Enable (PERRE)</b>                      Controls whether or not the Master Data Parity Error bit in the PCI Status register can be set.                      0 = Master Data Parity Error bit in PCI Status register CANNOT be set.                      1 = Master Data Parity Error bit in PCI Status register CAN be set.</p>
5	RO	0b	<p><b>VGA Palette Snoop (VGAPS)</b>                      Not Applicable or Implemented. Hard wired to 0.</p>
4	RO	0b	<p><b>Memory Write and Invalidate Enable (MWIE)</b>                      Not Applicable or Implemented. Hard wired to 0.</p>
3	RO	0b	<p><b>Special Cycle Enable (SCE)</b>                      Not Applicable or Implemented. hard wired to 0.</p>
2	RW	0b	<p><b>Bus Master Enable (BME)</b>                      Controls the ability of the PEG port to forward Memory and IO Read/Write Requests in the upstream direction.                      0 = This device is prevented from making memory or IO requests to its primary bus. Note that according to the <i>PCI Local Bus Specification</i>, 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, IO writes/reads, peer writes/reads, and MSIs will all be treated as illegal cycles. Writes are forwarded to memory address C0000h with byte enables deasserted. Reads is forwarded to memory address C0000h 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 is issued when the data is available. This bit does not affect forwarding of Completions from the primary interface to the secondary interface.</p>
1	RW	0b	<p><b>Memory Access Enable (MAE)</b>                      0 = All of Device 1's memory space is disabled.                      1 = Enable the Memory and Pre-fetchable memory address ranges defined in the MBASE1, MLIMIT1, PMBASE1, and PMLIMIT1 registers.</p>
0	RW	0b	<p><b>IO Access Enable (IOAE)</b>                      0 = All of Device 1's I/O space is disabled.                      1 = Enable the I/O address range defined in the IOBASE1, and IOLIMIT1 registers.</p>



### 1.13.4 PCISTS1 - PCI Status

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 6-7h  
 Default Value: 0010h  
 Access: RO; RWC  
 Size: 16 bits

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

(Sheet 1 of 2)

Bit	Access	Default Value	Description
15	RO	0b	<b>Detected Parity Error (DPE)</b> Not Applicable or Implemented. Hard wired to 0. Parity (generating poisoned TLPs) is not supported on the primary side of this device (we don't do error forwarding).
14	RWC	0b	<b>Signaled System Error (SSE)</b> 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.
13	RO	0b	<b>Received Master Abort Status (RMAS)</b> Not Applicable or Implemented. Hard wired to 0. The concept of a master abort does not exist on primary side of this device.
12	RO	0b	<b>Received Target Abort Status (RTAS)</b> Not Applicable or Implemented. Hard wired to 0. The concept of a target abort does not exist on primary side of this device.
11	RO	0b	<b>Signaled Target Abort Status (STAS)</b> Not Applicable or Implemented. Hard wired to 0. The concept of a target abort does not exist on primary side of this device.
10:9	RO	00b	<b>DEVSELB Timing (DEVT)</b> This device is not the subtractively decoded device on bus 0. This bit field is therefore hard wired to 00 to indicate that the device uses the fastest possible decode.
8	RO	0b	<b>Master Data Parity Error (PMDPE)</b> Because the primary side of the PCIe graphic's virtual P2P bridge is integrated with the PROCESSOR functionality there is no scenario where this bit will get set. Because hardware will never set this bit, it is impossible for software to have an opportunity to clear this bit or otherwise test that it is implemented. The <i>PCI Local Bus Specification</i> defines it as a R/WC, but for our implementation an RO definition behaves the same way and will meet all Microsoft testing requirements.  This bit can only be set when the Parity Error Enable bit in the PCI Command register is set.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
7	RO	0b	<b>Fast Back-to-Back (FB2B)</b> Not Applicable or Implemented. Hard wired to 0.
6	RO	0b	<i>Reserved</i>
5	RO	0b	<b>66-/60-MHz Capability (CAP66)</b> Not Applicable or Implemented. Hard wired to 0.
4	RO	1b	<b>Capabilities List (CAPL)</b> Indicates that a capabilities list is present. Hard wired to 1.
3	RO	0b	<b>INTA Status (INTAS)</b> 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.
2:0	RO	000b	<i>Reserved</i>



### 1.13.5 RID1 - Revision Identification

B/D/F/Type: 0/1/0/PCI  
Address Offset: 8h  
Default Value: 12h  
Access: RO  
Size: 8 bits

This register contains the revision number of the processor. The Revision ID (RID) is a traditional 8-bit Read Only (RO) register located at offset 08h in the standard PCI header of every PCI/PCI Express compatible device and function.

Following reset, the SRID is returned when the RID is read at offset 08h. The SRID value reflects the actual product stepping. To select the CRID value, BIOS/configuration software writes a key value of 69h to Bus 0, Device 0, Function 0 (DMI device) of the CPU's RID register at offset 08h. This causes the CRID to be returned when the RID is read at offset 08h.

#### Stepping Revision ID (SRID)

This register contains the revision number of the CPU.

The SRID is a 8-bit hardwired value assigned by Intel, based on product's stepping. The SRID is not a directly addressable PCI register. The SRID value is reflected through the RID register when appropriately addressed.

#### Compatible Revision ID (CRID)

The CRID is an 8-bit hardwired value assigned by Intel during manufacturing process. Normally, the value assigned as the CRID will be identical to the SRID value of a previous stepping of the product with which the new product is deemed "compatible".

The CRID is not a directly addressable PCI register. The CRID value is reflected through the RID register when appropriately addressed.

Bit	Access	Default Value	Description
7:0	RO	10h	<b>Revision Identification Number (RID1)</b> This is an 8-bit value that indicates the revision identification number for the processor Device 0. For the C-2 Stepping, these values are: SRID = 12h CRID = 02h



### 1.13.6 CC1 - Class Code

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 9-Bh  
 Default Value: 060400h  
 Access: RO  
 Size: 24 bits

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

Bit	Access	Default Value	Description
23:16	RO	06h	<b>Base Class Code (BCC)</b> Indicates the base class code for this device. This code has the value 06h, indicating a Bridge device.
15:8	RO	04h	<b>Sub-Class Code (SUBCC)</b> Indicates the sub-class code for this device. The code is 04h indicating a PCI to PCI Bridge.
7:0	RO	00h	<b>Programming Interface (PI)</b> 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.

### 1.13.7 CL1 - Cache Line Size

B/D/F/Type: 0/1/0/PCI  
 Address Offset: Ch  
 Default Value: 00h  
 Access: RW  
 Size: 8 bits

Bit	Access	Default Value	Description
7:0	RW	00h	<b>Cache Line Size (Scratch pad)</b> Implemented by PCI Express devices as a read-write field for legacy compatibility purposes but has no impact on any PCI Express device functionality.



### 1.13.8 HDR1 - Header Type

B/D/F/Type: 0/1/0/PCI  
Address Offset: Eh  
Default Value: 01h  
Access: RO  
Size: 8 bits

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

Bit	Access	Default Value	Description
7:0	RO	01h	<b>Header Type Register (HDR)</b> Returns 01 to indicate that this is a single function device with bridge header layout.

### 1.13.9 PBUSN1 - Primary Bus Number

B/D/F/Type: 0/1/0/PCI  
Address Offset: 18h  
Default Value: 00h  
Access: RO  
Size: 8 bits

This register identifies that this “virtual” Host-PCI Express bridge is connected to PCI Bus 0.

Bit	Access	Default Value	Description
7:0	RO	00h	<b>Primary Bus Number (BUSN)</b> Configuration software typically programs this field with the number of the bus on the primary side of the bridge. Since Device 1 is an internal device and its primary bus is always 0, these bits are read only and are hard wired to 0.



### 1.13.10 SBUSN1 - Secondary Bus Number

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 19h  
 Default Value: 00h  
 Access: RW  
 Size: 8 bits

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

Bit	Access	Default Value	Description
7:0	RW	00h	<b>Secondary Bus Number (BUSN)</b> This field is programmed by configuration software with the bus number assigned to PCI Express-G.

### 1.13.11 SUBUSN1 - Subordinate Bus Number

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 1Ah  
 Default Value: 00h  
 Access: RW  
 Size: 8 bits

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

Bit	Access	Default Value	Description
7:0	RW	00h	<b>Subordinate Bus Number (BUSN)</b> This register is programmed by configuration software with the number of the highest subordinate bus that lies behind the Device 1 bridge. When only a single PCI device resides on the PCI Express-G segment, this register will contain the same value as the SBUSN1 register.



### 1.13.12 IOBASE1 - I/O Base Address

B/D/F/Type: 0/1/0/PCI  
Address Offset: 1Ch  
Default Value: F0h  
Access: RO; RW  
Size: 8 bits

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

$$\text{IO\_BASE} = \langle \text{address} \rangle \ll \text{IO\_LIMIT}$$

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 is aligned to a 4-KB boundary.

Bit	Access	Default Value	Description
7:4	RW	Fh	<b>I/O Address Base (IOBASE)</b> Corresponds to A[15:12] of the I/O addresses passed by bridge 1 to PCI Express-G. BIOS must not set this register to 00h otherwise OCF8h/OCFCh accesses is forwarded to the PCI Express hierarchy associated with this device.
3:0	RO	0h	<i>Reserved</i>

### 1.13.13 IOLIMIT1 - I/O Limit Address

B/D/F/Type: 0/1/0/PCI  
Address Offset: 1Dh  
Default Value: 00h  
Access: RO; RW  
Size: 8 bits

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

$$\text{IO\_BASE} = \langle \text{address} \rangle \ll \text{IO\_LIMIT}$$

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 is at the top of a 4-KB aligned address block.

Bit	Access	Default Value	Description
7:4	RW	0h	<b>I/O Address Limit (IOLIMIT)</b> Corresponds to A[15:12] of the I/O address limit of Device 1. Devices between this upper limit and IOBASE1 is passed to the PCI Express hierarchy associated with this device.
3:0	RO	0h	<i>Reserved</i>





### 1.13.14 SSTS1 - Secondary Status

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 1E-1Fh  
 Default Value: 0000h  
 Access: RWC; RO  
 Size: 16 bits

SSTS1 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-to-PCI bridge embedded within processor.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
15	RWC	0b	<b>Detected Parity Error (DPE)</b> 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.
14	RWC	0b	<b>Received System Error (RSE)</b> This bit is set when the Secondary Side for a Type 1 configuration space header device receives an ERR_FATAL or ERR_NONFATAL.
13	RWC	0b	<b>Received Master Abort (RMA)</b> 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.
12	RWC	0b	<b>Received Target Abort (RTA)</b> 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.
11	RO	0b	<b>Signaled Target Abort (STA)</b> Not Applicable or Implemented. Hard wired to 0. The processor does not generate Target Aborts (the processor will never complete a request using the Completer Abort Completion status).
10:9	RO	00b	<b>DEVSELB Timing (DEVT)</b> Not Applicable or Implemented. Hard wired to 0.
8	RWC	0b	<b>Master Data Parity Error (SMDPE)</b> 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.
7	RO	0b	<b>Fast Back-to-Back (FB2B)</b> Not Applicable or Implemented. Hard wired to 0.
6	RO	0b	<i>Reserved</i>



(Sheet 2 of 2)

Bit	Access	Default Value	Description
5	RO	0b	<b>66-/60-MHz Capability (CAP66)</b> Not Applicable or Implemented. Hard wired to 0.
4:0	RO	00h	Reserved

### 1.13.15 MBASE1 - Memory Base Address

B/D/F/Type: 0/1/0/PCI

Address Offset: 20-21h

Default Value: FFF0h

Access: RO; RW

Size: 16 bits

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

$$\text{MEMORY\_BASE} = \langle \text{address} \rangle = \langle \text{MEMORY\_LIMIT} \rangle$$

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 is aligned to a 1-MB boundary.

Bit	Access	Default Value	Description
15:4	RW	FFFh	<b>Memory Address Base (MBASE)</b> Corresponds to A[31:20] of the lower limit of the memory range that is passed to PCI Express-G.
3:0	RO	0h	Reserved

### 1.13.16 MLIMIT1 - Memory Limit Address

B/D/F/Type: 0/1/0/PCI

Address Offset: 22-23h

Default Value: 0000h

Access: RO; RW

Size: 16 bits

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

$$\text{MEMORY\_BASE} = \langle \text{address} \rangle = \langle \text{MEMORY\_LIMIT} \rangle$$

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 FFFFh. Thus, the top of the defined memory address range is at the top of a 1-MB aligned memory block.

**Note:** Memory range covered by MBASE and MLIMIT registers are used to map non-prefetchable PCI Express-G 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 CPU-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.

Bit	Access	Default Value	Description
15:4	RW	000h	<b>Memory Address Limit (MLIMIT)</b> Corresponds to A[31:20] of the upper limit of the address range passed to PCI Express-G.
3:0	RO	0h	<i>Reserved</i>

### 1.13.17 PMBASE1 - Prefetchable Memory Base Address

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 24-25h  
 Default Value: FFF1h  
 Access: RO; RW  
 Size: 16 bits

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

$$\text{PREFETCHABLE\_MEMORY\_BASE} = \langle \text{address} \rangle \langle \text{PREFETCHABLE\_MEMORY\_LIMIT} \rangle$$

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 is aligned to a 1-MB boundary.



Bit	Access	Default Value	Description
15:4	RW	FFFh	<b>Prefetchable Memory Base Address (MBASE)</b> Corresponds to A[31:20] of the lower limit of the memory range that is passed to PCI Express-G.
3:0	RO	1h	<b>64-bit Address Support (64-bit Address Support)</b> 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.

### 1.13.18 PMLIMIT1 - Prefetchable Memory Limit Address

B/D/F/Type: 0/1/0/PCI  
Address Offset: 26-27h  
Default Value: 0001h  
Access: RO; RW  
Size: 16 bits

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

$$\text{PREFETCHABLE\_MEMORY\_BASE} = < \text{address} = < \text{PREFETCHABLE\_MEMORY\_LIMIT}$$

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 FFFFh. Thus, the top of the defined memory address range is at the top of a 1-MB aligned memory block.

**Note:** 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 CPU perspective.

Bit	Access	Default Value	Description
15:4	RW	000h	<b>Prefetchable Memory Address Limit (PMLIMIT)</b> Corresponds to A[31:20] of the upper limit of the address range passed to PCI Express-G.
3:0	RO	1h	<b>64-bit Address Support (RSVD)</b> 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



### 1.13.19 PMBASEU1 - Prefetchable Memory Base Address Upper

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 28-2Bh  
 Default Value: 00000000h  
 Access: RW  
 Size: 32 bits

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 CPU to PCI Express-G prefetchable memory access routing based on the following formula:

$$\text{PREFETCHABLE\_MEMORY\_BASE} = \langle \text{address} \rangle \langle \text{PREFETCHABLE\_MEMORY\_LIMIT} \rangle$$

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 is aligned to a 1-MB boundary.

Bit	Access	Default Value	Description
31:0	RW	00000000h	<b>Prefetchable Memory Base Address (MBASEU)</b> Corresponds to A[63:32] of the lower limit of the prefetchable memory range that is passed to PCI Express-G.

### 1.13.20 PMLIMITU1 - Prefetchable Memory Limit Address Upper

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 2C-2Fh  
 Default Value: 00000000h  
 Access: RW  
 Size: 32 bits

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 CPU to PCI Express-G prefetchable memory access routing based on the following formula:

$$\text{PREFETCHABLE\_MEMORY\_BASE} = \langle \text{address} \rangle \langle \text{PREFETCHABLE\_MEMORY\_LIMIT} \rangle$$



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 FFFFh. Thus, the top of the defined memory address range is at the top of a 1-MB 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 CPU perspective.

Bit	Access	Default Value	Description
31:0	RW	00000000h	<b>Prefetchable Memory Address Limit (MLIMITU)</b> Corresponds to A[63:32] of the upper limit of the prefetchable Memory range that is passed to PCI Express-G.

### 1.13.21 CAPPTR1 - Capabilities Pointer

B/D/F/Type: 0/1/0/PCI  
Address Offset: 34h  
Default Value: 88h  
Access: RO  
Size: ]8 bits

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

Bit	Access	Default Value	Description
7:0	RO	88h	<b>First Capability (CAPPTR1)</b> The first capability in the list is the Subsystem ID and Subsystem Vendor ID Capability.



### 1.13.22 INTRLINE1 - Interrupt Line

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 3Ch  
 Default Value: 00h  
 Access: RW  
 Size: 8 bits

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.

Bit	Access	Default Value	Description
7:0	RW	00h	<b>Interrupt Connection (INTCON)</b> 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.

### 1.13.23 INTRPIN1 - Interrupt Pin

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 3Dh  
 Default Value: 01h  
 Access: RO  
 Size: 8 bits

This register specifies which interrupt pin this device uses.

Bit	Access	Default Value	Description
7:0	RO	01h	<b>Interrupt Pin (INTRPIN)</b> As a single function device, the PCI Express device specifies INTA as its interrupt pin. 01h=INTA.



### 1.13.24 BCTRL1 - Bridge Control

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 3E-3Fh  
 Default Value: 0000h  
 Access: RO; RW  
 Size: 16 bits

This register provides extensions to the PCICMD1 register that are specific to PCI-PCI bridges. The BCTRL provides additional control for the secondary interface (i.e., PCI Express-G) as well as some bits that affect the overall behavior of the “virtual” Host-PCI Express bridge embedded within processor, e.g., VGA compatible address ranges mapping.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
15:12	RO	0h	<i>Reserved</i>
11	RO	0b	<b>Discard Timer SERR# Enable (DTSERRE)</b> Not Applicable or Implemented. Hard wired to 0.
10	RO	0b	<b>Discard Timer Status (DTSTS)</b> Not Applicable or Implemented. Hard wired to 0.
9	RO	0b	<b>Secondary Discard Timer (SDT)</b> Not Applicable or Implemented. Hard wired to 0.
8	RO	0b	<b>Primary Discard Timer (PDT)</b> Not Applicable or Implemented. Hard wired to 0.
7	RO	0b	<b>Fast Back-to-Back Enable (FB2BEN)</b> Not Applicable or Implemented. Hard wired to 0.
6	RW	0b	<b>Secondary Bus Reset (SRESET)</b> 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.
5	RO	0b	<b>Master Abort Mode (MAMODE)</b> Does not apply to PCI Express. Hard wired to 0.
4	RW	0b	<b>VGA 16-bit Decode (VGA16D)</b> 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.





(Sheet 2 of 2)

Bit	Access	Default Value	Description
3	RW	0b	<p><b>VGA Enable (VGAEN)</b></p> <p>Controls the routing of CPU initiated transactions targeting VGA compatible I/O and memory address ranges. See the VGAEN/MDAP table in device 0, offset 97h[0].</p> <p>The VGA enable bit must only be set if there exists a VGA device within the PCI express hierarchy on the secondary side of the bridge. It must not be set if no such device is discovered during PCI enumeration.</p>
2	RW	0b	<p><b>ISA Enable (ISAEN)</b></p> <p>Needed to exclude legacy resource decode to route ISA resources to legacy decode path. Modifies the response by the processor to an I/O access issued by the CPU that target ISA I/O addresses. This applies only to I/O addresses that are enabled by the IOBASE and IOLIMIT registers.</p> <p>0 = All addresses defined by the IOBASE and IOLIMIT for CPU I/O transactions is mapped to PCI Express-G.</p> <p>1 = Processor will not forward to PCI Express-G any I/O transactions addressing the last 768 bytes in each 1-KB block even if the addresses are within the range defined by the IOBASE and IOLIMIT registers.</p>
1	RW	0b	<p><b>SERR Enable (SERREN)</b></p> <p>0 = No forwarding of error messages from secondary side to primary side that could result in an SERR.</p> <p>1 = ERR_COR, ERR_NONFATAL, and ERR_FATAL messages result in SERR message when individually enabled by the Root Control register.</p>
0	RW	0b	<p><b>Parity Error Response Enable (PEREN)</b></p> <p>Controls whether or not the Master Data Parity Error bit in the Secondary Status register is set when the PROCESSOR receives across the link (upstream) a Read Data Completion Poisoned TLP</p> <p>0 = Master Data Parity Error bit in Secondary Status register CANNOT be set.</p> <p>1 = Master Data Parity Error bit in Secondary Status register CAN be set.</p>



### 1.13.25 PM\_CAPID1 - Power Management Capabilities

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 80-83h  
 Default Value: C8039001h  
 Access: RO  
 Size: 32 bits

Bit	Access	Default Value	Description
31:27	RO	19h	<b>PME Support (PMES)</b> 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 latest <i>PCI Power Management Specification</i> for encoding explanation and other power management details.
26	RO	0b	<b>D2 Power State Support (D2PSS)</b> hard wired to 0 to indicate that the D2 power management state is NOT supported.
25	RO	0b	<b>D1 Power State Support (D1PSS)</b> hard wired to 0 to indicate that the D1 power management state is NOT supported.
24:22	RO	000b	<b>Auxiliary Current (AUXC)</b> hard wired to 0 to indicate that there are no 3.3Vaux auxiliary current requirements.
21	RO	0b	<b>Device Specific Initialization (DSI)</b> hard wired to 0 to indicate that special initialization of this device is NOT required before generic class device driver is to use it.
20	RO	0b	<b>Auxiliary Power Source (APS)</b> hard wired to 0.
19	RO	0b	<b>PME Clock (PMECLK)</b> hard wired to 0 to indicate this device does NOT support PMEB generation.
18:16	RO	011b	<b>PCI PM CAP Version (PCIPMCV)</b> Version - A value of 011b indicates that this function complies with the latest revision of the <i>PCI Power Management Interface Specification</i> .
15:8	RO	90h	<b>Pointer to Next Capability (PNC)</b> 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.
7:0	RO	01h	<b>Capability ID (CID)</b> Value of 01h identifies this linked list item (capability structure) as being for PCI Power Management registers.



### 1.13.26 PM\_CS1 - Power Management Control/Status

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 84-87h  
 Default Value: 00000008h  
 Access: RO; RW-S; RW  
 Size: 32 bits

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31:16	RO	0000h	<i>Reserved</i> Not Applicable or Implemented. Hard wired to 0.
15	RO	0b	<b>PME Status (PMESTS)</b> Indicates that this device does not support PMEB generation from D3cold.
14:13	RO	00b	<b>Data Scale (DSCALE)</b> Indicates that this device does not support the power management data register.
12:9	RO	0h	<b>Data Select (DSEL)</b> Indicates that this device does not support the power management data register.
8	RW-S	0b	<b>PME Enable (PMEE)</b> 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]
7:4	RO	0000b	<i>Reserved</i>
3	RO	1b	<b>No Soft Reset (NSR)</b> When set to 1 this bit indicates that the device is transitioning from D3hot to D0 because the power state commands do not perform an internal reset. Config context is preserved. Upon transition no additional 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 D3hot to D0 by a system or bus segment reset will return to the device state D0 un-initialized with only PME context preserved if PME is supported and enabled.
2	RO	0b	<i>Reserved</i>



(Sheet 2 of 2)

Bit	Access	Default Value	Description
1:0	RW	00b	<p><b>Power State (PS)</b></p> <p>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.</p> <p>00: D0 01: D1 (Not supported in this device.) 10: D2 (Not supported in this device.) 11: D3</p> <p>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).</p> <p>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.</p> <p>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).</p> <p>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]</p> <p>There is no additional hardware functionality required to support these Power States.</p>



### 1.13.27 SS\_CAPID - Subsystem ID and Vendor ID Capabilities

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 88-8Bh  
 Default Value: 0000800Dh  
 Access: RO  
 Size: 32 bits

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.

Bit	Access	Default Value	Description
31:16	RO	0000h	<i>Reserved</i>
15:8	RO	80h	<b>Pointer to Next Capability (PNC)</b> This contains a pointer to the next item in the capabilities list which is the PCI Power Management capability.
7:0	RO	0Dh	<b>Capability ID (CID)</b> Value of 0Dh identifies this linked list item (capability structure) as being for SSID/SSVID registers in a PCI-to-PCI Bridge.

### 1.13.28 SS - Subsystem ID and Subsystem Vendor ID

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 8C-8Fh  
 Default Value: 00008086h  
 Access: RW-O  
 Size: 32 bits

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.

Bit	Access	Default Value	Description
31:16	RW-O	0000h	<b>Subsystem ID (SSID)</b> Identifies the particular subsystem and is assigned by the vendor.
15:0	RW-O	8086h	<b>Subsystem Vendor ID (SSVID)</b> Identifies the manufacturer of the subsystem and is the same as the vendor ID which is assigned by the PCI Special Interest Group.



### 1.13.29 MSI\_CAPID - Message Signaled Interrupts Capability ID

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 90-91h  
 Default Value: A005h  
 Access: RO  
 Size: 16 bits

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.

Bit	Access	Default Value	Description
15:8	RO	A0h	<b>Pointer to Next Capability (PNC)</b> This contains a pointer to the next item in the capabilities list which is the PCI Express capability.
7:0	RO	05h	<b>Capability ID (CID)</b> Value of 05h identifies this linked list item (capability structure) as being for MSI registers.

### 1.13.30 MC - Message Control

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 92-93h  
 Default Value: 0000h  
 Access: RO; RW  
 Size: 16 bits

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.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
15:8	RO	00h	<i>Reserved</i>
7	RO	0b	<b>64-bit Address Capable (64AC)</b> hard wired 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 32-b/4-GB limit.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
6:4	RW	000b	<b>Multiple Message Enable (MME)</b> System software programs this field to indicate the actual number of messages allocated to this device. This number is equal to or less than the number actually requested. The encoding is the same as for the MMC field below.
3:1	RO	000b	<b>Multiple Message Capable (MMC)</b> 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
0	RW	0b	<b>MSI Enable (MSIEN)</b> Controls the ability of this device to generate MSIs. 0 = MSI will not be generated. 1 = MSI is generated when we receive PME or HotPlug messages. INTA will not be generated and INTA Status (PCISTS1[3]) will not be set.

### 1.13.31 MA - Message Address

B/D/F/Type: 0/1/0/PCI  
 Address Offset: 94-97h  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits

Bit	Access	Default Value	Description
31:2	RW	00000000h	<b>Message Address (MA)</b> 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.
1:0	RO	00b	<b>Force DWord Align (FDWA)</b> hard wired to 0 so that addresses assigned by system software are always aligned on a dword address boundary.



### 1.13.32 MD - Message Data

B/D/F/Type: 0/1/0/PCI  
Address Offset: 98-99h  
Default Value: 0000h  
Access: RW  
Size: 16 bits

Bit	Access	Default Value	Description
15:0	RW	0000h	<b>Message Data (MD)</b> 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.

### 1.13.33 PEG\_CAPL - PCI Express-G Capability List

B/D/F/Type: 0/1/0/PCI  
Address Offset: A0-A1h  
Default Value: 0010h  
Access: RO  
Size: 16 bits  
Enumerates the PCI Express capability structure.

Bit	Access	Default Value	Description
15:8	RO	00h	<b>Pointer to Next Capability (PNC)</b> 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.
7:0	RO	10h	<b>Capability ID (CID)</b> Identifies this linked list item (capability structure) as being for PCI Express registers.





### 1.13.34 PEG\_CAP - PCI Express-G Capabilities

B/D/F/Type: 0/1/0/PCI  
 Address Offset: A2-A3h  
 Default Value: 0142h  
 Access: RO; RW-O  
 Size: 16 bits

Indicates PCI Express device capabilities.

Bit	Access	Default Value	Description
15	RO	0b	<i>Reserved</i>
14	RO	0b	<i>Reserved</i> Reserved for TCS Routing Supported.
13:9	RO	00h	<b>Interrupt Message Number (IMN)</b> Not Applicable or Implemented. Hard wired to 0.
8	RW-O	1b	<b>Slot Implemented (SI)</b> 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. <b>BIOS Requirement:</b> This field must be initialized appropriately if a slot connection is not implemented.
7:4	RO	4h	<b>Device/Port Type (DPT)</b> hard wired to 4h to indicate root port of PCI Express Root Complex.
3:0	RO	2h	<b>PCI Express Capability Version (PCIECV)</b> hard wired to 2h to indicate compliance to the PCI Express Capabilities Register Expansion ECN.



### 1.13.35 DCAP - Device Capabilities

B/D/F/Type: 0/1/0/PCI  
Address Offset: A4-A7h  
Default Value: 00008000h  
Access: RO  
Size: 32 bits

Indicates PCI Express device capabilities.

Bit	Access	Default Value	Description
31:16	RO	0000h	<i>Reserved</i> Not Applicable or Implemented. Hard wired to 0.
15	RO	1b	<b>Role-Based Error Reporting (RBER)</b> Indicates that this device implements the functionality defined in the Error Reporting ECN as required by the PCI Express Base spec.
14:6	RO	000h	<i>Reserved</i> Not Applicable or Implemented. Hard wired to 0.
5	RO	0b	<b>Extended Tag Field Supported (ETFS)</b> hard wired to indicate support for 5-bit Tags as a Requestor.
4:3	RO	00b	<b>Phantom Functions Supported (PFS)</b> Not Applicable or Implemented. Hard wired to 0.
2:0	RO	000b	<b>Max Payload Size (MPS)</b> hard wired to indicate 128B max supported payload for Transaction Layer Packets (TLP).



### 1.13.36 DCTL - Device Control

B/D/F/Type: 0/1/0/PCI  
 Address Offset: A8-A9h  
 Default Value: 0000h  
 Access: RO; RW  
 Size: 16 bits

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.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
15	RO	0h	<i>Reserved</i>
14:12	RO	000b	<i>Reserved for Max Read Request Size (MRRS)</i>
11	RO	0b	<i>Reserved for Enable No Snoop (RSVD)</i>
10	RO	0b	<i>Reserved</i> Reserved for Auxiliary (AUX) PM Enable ()
9	RO	0b	<i>Reserved</i> Reserved for Phantom Functions Enable ()
8	RO	0b	<i>Reserved</i> Reserved for Extended Tag Field Enable ()
7:5	RW	000b	<b>Max Payload Size (MPS)</b> 000: 128B max supported payload for Transaction Layer Packets (TLP). 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. All other encodings are reserved. Hardware will actually ignore this field. It is writeable only to support compliance testing.
4	RO	0b	<i>Reserved for Enable Relaxed Ordering (RSVD)</i>
3	RW	0b	<b>Unsupported Request Reporting Enable (URRE)</b> 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.
2	RW	0b	<b>Fatal Error Reporting Enable (FERE)</b> 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.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
1	RW	0b	<b>Non-Fatal Error Reporting Enable (NERE)</b> When set, enables signaling of ERR_NONFATAL 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.
0	RW	0b	<b>Correctable Error Reporting Enable (CERE)</b> 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.

### 1.13.37 DSTS - Device Status

B/D/F/Type: 0/1/0/PCI  
 Address Offset: AA-ABh  
 Default Value: 0000h  
 Access: RO; RWC  
 Size: 16 bits

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.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
15:6	RO	000h	<i>Reserved and Zero (RSVD)</i> For future R/WC/S implementations; software must use 0 for writes to bits.
5	RO	0b	<b>Transactions Pending (TP)</b> 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).
4	RO	0b	<i>Reserved</i>
3	RWC	0b	<b>Unsupported Request Detected (URD)</b> 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.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
2	RWC	0b	<b>Fatal Error Detected (FED)</b> 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 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	RWC	0b	<b>Non-Fatal Error Detected (NFED)</b> 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.
0	RWC	0b	<b>Correctable Error Detected (CED)</b> 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.

### 1.13.38 LCAP - Link Capabilities

B/D/F/Type: 0/1/0/PCI  
 Address Offset: AC-AFh  
 Default Value: 02214D02h  
 Access: RO; RW-O  
 Size: 32 bits  
 Indicates PCI Express device specific capabilities.

(Sheet 1 of 4)

Bit	Access	Default Value	Description
31:24	RO	02h	<b>Port Number (PN)</b> Indicates the PCI Express port number for the given PCI Express link. Matches the value in Element Self Description[31:24].
23:22	RO	00b	<i>Reserved</i>



(Sheet 2 of 4)

Bit	Access	Default Value	Description
21	RO	1b	<p><b>Link Bandwidth Notification Capability (LBNC)</b></p> <p>A value of 1b indicates support for the Link Bandwidth Notification status and interrupt mechanisms. This capability is required for all Root Ports and Switch downstream ports supporting Links wider than x1 and/or multiple Link speeds.</p> <p>This field is not applicable and is reserved for Endpoint devices, PCI Express to PCI/PCI-X bridges, and Upstream Ports of Switches.</p> <p>Devices that do not implement the Link Bandwidth Notification capability must hardwire this bit to 0b.</p>
20	RO	0b	<p><b>Data Link Layer Link Active Reporting Capable (DLLARC)</b></p> <p>For a Downstream Port, this bit must be set to 1b if the component supports the optional capability of reporting the DL_Active state of the Data Link Control and Management State Machine. For a hot-plug capable Downstream Port (as indicated by the Hot-Plug Capable field of the Slot Capabilities register), this bit must be set to 1b.</p> <p>For Upstream Ports and components that do not support this optional capability, this bit must be hard wired to 0b.</p>
19	RO	0b	<p><b>Surprise Down Error Reporting Capable (SDERC)</b></p> <p>For a Downstream Port, this bit must be set to 1b if the component supports the optional capability of detecting and reporting a Surprise Down error condition.</p> <p>For Upstream Ports and components that do not support this optional capability, this bit must be hard wired to 0b.</p>
18	RO	0b	<p><b>Clock Power Management (CPM)</b></p> <p>A value of 1b in this bit indicates that the component tolerates the removal of any reference clock(s) when the link is in the L1 and L2/3 Ready link states. A value of 0b indicates the component does not have this capability and that reference clock(s) must not be removed in these link states.</p> <p>This capability is applicable only in form factors that support "clock request" (CLKREQ#) capability.</p> <p>For a multi-function device, each function indicates its capability independently. Power Management configuration software must only permit reference clock removal if all functions of the multifunction device indicate a 1b in this bit.</p>



(Sheet 3 of 4)

Bit	Access	Default Value	Description
17:15	RW-O	010b	<p><b>L1 Exit Latency (L1ELAT)</b></p> <p>Indicates the length of time this Port requires to complete the transition from L1 to L0.</p> <p>000: Less than 1us                      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</p> <p>BIOS Requirement: If this field is required to be any value other than the default, BIOS must initialize it accordingly.</p> <p>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.</p>
14:12	RO	100b	<p><b>L0s Exit Latency (LOSELAT)</b></p> <p>Indicates the length of time this Port requires to complete the transition from L0s to L0.</p> <p>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 μs                      101: 1 μs to less than 2 μs                      110: 2 μs - 4 μs                      111: More than 4 μs</p> <p>The actual value of this field depends on the common Clock Configuration bit (LCTL[6]) register.</p>
11:10	RW-O	11b	<p><b>Active State Link PM Support (ASLPMS)</b></p> <p>ASPM L0s and L1 supported.</p>
9:4	RW-O	10h if x16 device is present 08h if x8 device is present or in bifurcation mode	<p><b>Max Link Width (MLW)</b></p> <p>Indicates the maximum number of lanes supported for this link.</p>



(Sheet 4 of 4)

Bit	Access	Default Value	Description
3:0	RW-O	1h	<b>Max Link Speed (MLS)</b> Supported Link Speed – This field indicates the supported Link speed(s) of the associated Port. Defined encodings are: 0001b2.5-GT/s Link speed supported All other encodings are reserved.





### 1.13.39 LCTL - Link Control

B/D/F/Type: 0/1/0/PCI  
 Address Offset: B0-B1h  
 Default Value: 0000h  
 Access: RO; RW; RW-SC  
 Size: 16 bits  
 Allows control of PCI Express link.

(Sheet 1 of 3)

Bit	Access	Default Value	Description
15:12	RO	0000b	<i>Reserved</i>
11	RW	0b	<p><b>Link Autonomous Bandwidth Interrupt Enable (LABIE)</b>                      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.</p>
10	RW	0b	<p><b>Link Bandwidth Management Interrupt Enable (LBMI E)</b>                      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.</p>
9	RW	0b	<p><b>Hardware Autonomous Width Disable (HAWD)</b>                      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.</p>
8	RO	0b	<p><b>Enable Clock Power Management (ECPM)</b>                      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.</p>



(Sheet 2 of 3)

Bit	Access	Default Value	Description
7	RW	0b	<p><b>Extended Synch (ES)</b></p> <p>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.</p> <p>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.</p>
6	RW	0b	<p><b>Common Clock Configuration (CCC)</b></p> <p>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.</p>
5	RW-SC	0b	<p><b>Retrain Link (RL)</b></p> <p>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.</p> <p>This bit always returns 0 when read. This bit is cleared automatically (no need to write a 0).</p>
4	RW	0b	<p><b>Link Disable (LD)</b></p> <p>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.</p> <p>Writes to this bit are immediately reflected in the value read from the bit, regardless of actual Link state.</p>
3	RO	0b	<p><b>Read Completion Boundary (RCB)</b></p> <p>hard wired to 0 to indicate 64 byte.</p>
2	RO	0b	Reserved (FEDLB)



(Sheet 3 of 3)

Bit	Access	Default Value	Description
1:0	RW	00b	<p><b>Active State PM (ASPM)</b>                      Controls the level of active state power management supported on the given link.                      00: Disabled                      01: L0s Entry Supported                      10: L1 Entry Enabled                      11: L0s and L1 Entry Supported</p> <p><b>Note:</b> "L0s Entry Enabled" indicates the Transmitter entering L0s is supported. The Receiver must be capable of entering L0s even when the field is disabled (00b).                      ASPM L1 must be enabled by software in the Upstream component on a Link prior to enabling ASPM L1 in the Downstream component on that Link. When disabling ASPM L1, software must disable ASPM L1 in the Downstream component on a Link prior to disabling ASPM L1 in the Upstream component on that Link. ASPM L1 must only be enabled on the Downstream component if both components on a Link support ASPM L1.</p>

### 1.13.40 LSTS - Link Status

B/D/F/Type: 0/1/0/PCI  
 Address Offset: B2-B3h  
 Default Value: 1000h  
 Access: RWC; RO  
 Size: 16 bits

Indicates PCI Express link status.

Bit	Access	Default Value	Description
15	RWC	0b	<p><b>Link Autonomous Bandwidth Status (LABWS)</b>                      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.                      This bit must be set when the upstream component receives eight consecutive TS1 or TS2 ordered sets with the Autonomous Change bit set.</p>



Bit	Access	Default Value	Description
14	RWC	0b	<p><b>Link Bandwidth Management Status (LBWMS)</b></p> <p>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.</p> <p><b>Note:</b> 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.</p> <p>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</p> <p>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.</p>
13	RO	0b	<p><b>Data Link Layer Link Active (Optional) (DLLLA)</b></p> <p>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 hard wired to 0b.</p>
12	RO	1b	<p><b>Slot Clock Configuration (SCC)</b></p> <p>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.</p>
11	RO	0b	<p><b>Link Training (LTRN)</b></p> <p>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.</p>
10	RO	0b	<p><b>Undefined (Undefined)</b></p> <p>The value read from this bit is undefined. In previous versions of this specification, this bit was used to indicate a Link Training Error. System software must ignore the value read from this bit. System software is permitted to write any value to this bit.</p>
9:4	RO	00h	<p><b>Negotiated Link Width (NLW)</b></p> <p>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).</p> <p>00h: Reserved            01h: X1            02h: X2            04h: X4            08h: X8            10h: X16            All other encodings are reserved.</p>



Bit	Access	Default Value	Description
3:0	RO	0h	<b>Current Link Speed (CLS)</b> This field indicates the negotiated Link speed of the given PCI Express Link. Defined encodings are: 0001b 2.5 GT/s PCI Express Link All other encodings are reserved. The value in this field is undefined when the Link is not up.

### 1.13.41 SLOTCAP - Slot Capabilities

B/D/F/Type: 0/1/0/PCI  
 Address Offset: B4-B7h  
 Default Value: 00040000h  
 Access: RW-O; RO  
 Size: 32 bits

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

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31:19	RW-O	0000h	<b>Physical Slot Number (PSN)</b> 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.
18	RW-O	1b	<b>No Command Completed Support (NCCS)</b> 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.
17	RO	0b	<b>Reserved for Electromechanical Interlock Present (EIP)</b> When set to 1b, this bit indicates that an Electromechanical Interlock is implemented on the chassis for this slot.
16:15	RW-O	00b	<b>Slot Power Limit Scale (SPLS)</b> 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.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
14:7	RW-O	00h	<p><b>Slot Power Limit Value (SPLV)</b></p> <p>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.</p> <p>If this field is written, the link sends a Set_Slot_Power_Limit message.</p>
6	RO	0b	<p><b>Reserved for Hot-plug Capable (HPC)</b></p> <p>When set to 1b, this bit indicates that this slot is capable of supporting hot-lug operations.</p>
5	RO	0b	<p><b>Reserved for Hot-plug Surprise (HPS)</b></p> <p>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.</p>
4	RO	0b	<p><b>Reserved for Power Indicator Present (PIP)</b></p> <p>When set to 1b, this bit indicates that a Power Indicator is electrically controlled by the chassis for this slot.</p>
3	RO	0b	<p><b>Reserved for Attention Indicator Present (AIP)</b></p> <p>When set to 1b, this bit indicates that an Attention Indicator is electrically controlled by the chassis.</p>
2	RO	0b	<p><b>Reserved for MRL Sensor Present (MSP)</b></p> <p>When set to 1b, this bit indicates that an MRL Sensor is implemented on the chassis for this slot.</p>
1	RO	0b	<p><b>Reserved for Power Controller Present (PCP)</b></p> <p>When set to 1b, this bit indicates that a software programmable Power Controller is implemented for this slot/ adapter (depending on form factor).</p>
0	RO	0b	<p><b>Reserved for Attention Button Present (ABP)</b></p> <p>When set to 1b, this bit indicates that an Attention Button for this slot is electrically controlled by the chassis.</p>



### 1.13.42 SLOTCTL - Slot Control

B/D/F/Type: 0/1/0/PCI  
 Address Offset: B8-B9h  
 Default Value: 0000h  
 Access: RO; RW  
 Size: 16 bits

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

(Sheet 1 of 3)

Bit	Access	Default Value	Description
15:13	RO	000b	<i>Reserved</i>
12	RO	0b	<p><b>Reserved for Data Link Layer State Changed Enable (DLLSCE)</b></p> <p>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.</p> <p>If the Data Link Layer Link Active capability is not implemented, this bit is permitted to be read-only with a value of 0b.</p>
11	RO	0b	<p><b>Reserved for Electromechanical Interlock Control (EIC)</b></p> <p>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.</p>
10	RO	0b	<p><b>Reserved for Power Controller Control (PCC)</b></p> <p>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 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.</p> <p>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.</p> <p>The defined encodings are:                      0b Power On                      1b Power Off</p> <p>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.</p>



(Sheet 2 of 3)

Bit	Access	Default Value	Description
9:8	RO	00b	<p><b>Reserved Power Indicator Control (PIC)</b></p> <p>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.</p> <p>00: Reserved 01: On 10: Blink 11: Off</p> <p>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.</p>
7:6	RO	00b	<p><b>Reserved for Attention Indicator Control (AIC)</b></p> <p>If an Attention Indicator is implemented, writes to this 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.</p> <p>00: Reserved 01: On 10: Blink 11: Off</p> <p>If the Attention Indicator Present bit in the Slot Capabilities register is 0b, this field is permitted to be read only with a value of 00b.</p>
5	RO	0b	<p><b>Reserved for Hot-Plug Interrupt Enable (HPIE)</b></p> <p>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.</p>
4	RO	0b	<p><b>Reserved for Command Completed Interrupt Enable (CCI)</b></p> <p>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.</p> <p>Default value of this field is 0b.</p> <p>If Command Completed notification is not supported, this bit must be hard wired to 0b.</p>





(Sheet 3 of 3)

Bit	Access	Default Value	Description
3	RW	0b	<b>Presence Detect Changed Enable (PDCE)</b> When set to 1b, this bit enables software notification on a presence detect changed event.
2	RO	0b	<b>Reserved for MRL Sensor Changed Enable (MSCE)</b> 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.
1	RO	0b	<b>Reserved for Power Fault Detected Enable (PFDE)</b> 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
0	RO	0b	<b>Reserved for Attention Button Pressed Enable (ABPE)</b> When set to 1b, this bit enables software notification on an attention button pressed event.

### 1.13.43 SLOTSTS - Slot Status

B/D/F/Type: 0/1/0/PCI  
 Address Offset: BA-BBh  
 Default Value: 0000h  
 Access: RO; RWC  
 Size: 16 bits

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

(Sheet 1 of 3)

Bit	Access	Default Value	Description
15:9	RO	0000000b	<b>Reserved and Zero</b> For future R/WC/S implementations; software must use 0 for writes to bits.
8	RO	0b	<b>Reserved for Data Link Layer State Changed (DLLSC)</b> 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.



(Sheet 2 of 3)

Bit	Access	Default Value	Description
7	RO	0b	<p><b>Reserved for Electromechanical Interlock Status (EIS)</b></p> <p>If an Electromechanical Interlock is implemented, this bit indicates the current status of the Electromechanical Interlock.</p> <p>Defined encodings are:</p> <p>0b Electromechanical Interlock Disengaged 1b Electromechanical Interlock Engaged</p>
6	RO	0b	<p><b>Presence Detect State (PDS)</b></p> <p>In band presence detect state:</p> <p>0b: Slot Empty 1b: Card present in slot</p> <p>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.</p> <p>Defined encodings are:</p> <p>0b Slot Empty 1b Card Present in slot</p> <p>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.</p>
5	RO	0b	<p><b>Reserved for MRL Sensor State (MSS)</b></p> <p>This register reports the status of the MRL sensor if it is implemented.</p> <p>Defined encodings are:</p> <p>0b MRL Closed 1b MRL Open</p>
4	RO	0b	<p><b>Reserved for Command Completed (CC)</b></p> <p>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 is complete.</p> <p>If Command Completed notification is not supported, this bit must be hard wired to 0b.</p>



(Sheet 3 of 3)

Bit	Access	Default Value	Description
3	RWC	0b	<b>Presence Detect Changed (PDC)</b> 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.
2	RO	0b	<b>Reserved for MRL Sensor Changed (MSC)</b> 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.
1	RO	0b	<b>Reserved for Power Fault Detected (PFD)</b> 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.
0	RO	0b	<b>Reserved for Attention Button Pressed (ABP)</b> 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.



### 1.13.44 RCTL - Root Control

B/D/F/Type: 0/1/0/PCI  
 Address Offset: BC-BDh  
 Default Value: 0000h  
 Access: RO; RW  
 Size: 16 bits

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.

Bit	Access	Default Value	Description
15:5	RO	000h	<i>Reserved</i>
4	RO	0b	<b>Reserved for CRS Software Visibility Enable (CSVE)</b> 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.
3	RW	0b	<b>PME Interrupt Enable (PMEIE)</b> 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.
2	RW	0b	<b>System Error on Fatal Error Enable (SEFEE)</b> 0 = Controls the Root Complex's response to fatal errors.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.
1	RW	0b	<b>System Error on Non-Fatal Uncorrectable Error Enable (SENFUEE)</b> 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.
0	RW	0b	<b>System Error on Correctable Error Enable (SECEE)</b> 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.



### 1.13.45 RSTS - Root Status

B/D/F/Type: 0/1/0/PCI  
 Address Offset: C0-C3h  
 Default Value: 00000000h  
 Access: RO; RWC  
 Size: 32 bits

Provides information about PCI Express Root Complex specific parameters.

Bit	Access	Default Value	Description
31:18	RO	0000h	<b>Reserved and Zero (RSVD)</b> For future R/WC/S implementations; software must use 0 for writes to bits.
17	RO	0b	<b>PME Pending (PMEP)</b> 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.
16	RWC	0b	<b>PME Status (PMES)</b> 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.
15:0	RO	0000h	<b>PME Requestor ID (PMERID)</b> Indicates the PCI requestor ID of the last PME requestor.



### 1.13.46 LCTL2 - Link Control 2

B/D/F/Type: 0/1/0/PCI  
 Address Offset: D0-D1h  
 Default Value: 0002h  
 Access: RO; RW-S; RW  
 Size: 16 bits

(Sheet 1 of 3)

Bit	Access	Default Value	Description
15:13	RO	000b	<i>Reserved</i>
12	RW-S	0b	<p><b>Compliance De-emphasis (ComplianceDeemphasis)</b></p> <p>This bit sets the de-emphasis level in Polling.Compliance state if the entry occurred due to the Enter Compliance bit being 1b.</p> <p>Encodings:            1b -3.5 dB            0b -6 dB</p> <p>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 0b. 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.</p>
11	RW-S	0b	<p><b>Compliance SOS (composos)</b></p> <p>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. Components that support only the 2.5 GT/s speed are permitted to hardwire this field to 0b.</p>
10	RW-S	0b	<p><b>Enter Modified Compliance (entermodcompliance)</b></p> <p>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.</p>



(Sheet 2 of 3)

Bit	Access	Default Value	Description
9:7	RW-S	000b	<p><b>Transmit Margin (txmargin)</b></p> <p>This field controls the value of the non-de-emphasized 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).</p> <p>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 200-400 mV for full-swing and 100-200 mV for half-swing -                      111: reserved</p> <p>Default value is 000b.</p> <p>Components that support only the 2.5GT/s speed are permitted to hardwire this bit to 0b.</p>
6	RW-S	0b	<p><b>Selectable De-emphasis (selectabledeemphasis)</b></p> <p>Encodings:                      1b) -3.5dB                      0b -6 dB</p> <p>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.</p>
5	RW	0b	<p><b>Hardware Autonomous Speed Disable (HASD)</b></p> <p>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.</p>
4	RW-S	0b	<p><b>Enter Compliance (EC)</b></p> <p>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.</p>



(Sheet 3 of 3)

Bit	Access	Default Value	Description
3:0	RW	2h	<p><b>Target Link Speed (TLS)</b></p> <p>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. Defined encodings are: 0001b 2.5Gb/s Target Link Speed All other encodings are reserved.</p> <p>If a value is written to this field that does not correspond to a speed included in the Supported Link Speeds field, the result is undefined. The default value of this field is the highest link speed supported by the component (as reported in the Supported Link Speeds 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.</p>

### 1.13.47 LSTS2 - Link Status 2

B/D/F/Type: 0/1/0/PCI  
 Address Offset: D2-D3h  
 Default Value: 0000h  
 Access: RO  
 Size: 16 bits

Bit	Access	Default Value	Description
15:1	RO	0000h	<i>Reserved</i>
0	RO	0b	<p><b>Current De-emphasis Level (CURDELVL)</b></p> <p>Current De-emphasis Level – 1b -3.5 dB 0b -6 dB When the Link is operating at 2.5 GT/s speed, this bit is 0b.</p>





### 1.13.48 PEGLC - PCI Express-G Legacy Control

B/D/F/Type: 0/1/0/PCI  
 Address Offset: EC-EFh  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits

Controls functionality that is needed by Legacy (non-PCI Express aware) OS's during run time.

Bit	Access	Default Value	Description
31:3	RO	00000000h	<i>Reserved</i>
2	RW	0b	<b>PME GPE Enable (PMEGPE)</b> 0 = Do not generate GPE PME message when PME is received. 1 = Generate a GPE PME message when PME is received. This enables the processor to support PMEs on the PEG port under legacy OSs.
1	RW	0b	<b>Hot-Plug GPE Enable (HPGPE)</b> 0 = Do not generate GPE Hot-Plug message when Hot-Plug event is received. 1 = Generate a GPE Hot-Plug message when Hot-Plug Event is received. This enables the processor to support Hot-Plug on the PEG port under legacy OSs
0	RW	0b	<b>General Message GPE Enable (GENGPE)</b> 0 = Do not forward received GPE assert/deassert messages. 1 = Forward received GPE assert/deassert messages.



## 1.14 PCI Device1 - Extended Configuration

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Virtual Channel Enhanced Capability Header	VCECH	100	103	00010002h	RW-O; RO
Port VC Capability Register 1	PVCCAP1	104	107	00000000h	RO
Port VC Capability Register 2	PVCCAP2	108	10B	00000000h	RO
Port VC Control	PVCCTL	10C	10D	0000h	RO; RW
VCO Resource Capability	VCORCAP	110	113	00000001h	RO
VCO Resource Control	VCORCTL	114	117	800000FFh	RO; RW
VCO Resource Status	VCORSTS	11A	11B	0002h	RO
PCI Express-G Sequence	PEGSSTS	218	21F	000000000000 0FFFh	RO
PCI Express-G Transmit De-Emphasis Select Register	PEGTXDEM PSEL	DA8	DAB	43E00BF9h	RW; RO

### 1.14.1 VCECH - Virtual Channel Enhanced Capability Header

B/D/F/Type: 0/1/0/MMR  
 Address Offset: 100-103h  
 Default Value: 00010002h  
 Access: RW-O; RO  
 Size: 32 bits

Indicates PCI Express device Virtual Channel capabilities. Extended capability structures for PCI Express devices are located in PCI Express extended configuration space and have different field definitions than standard PCI capability structures.

Bit	Access	Default Value	Description
31:20	RW-O	000h	<b>Pointer to Next Capability (PNC)</b> The Link Declaration Capability is the next in the PCI Express extended capabilities list.
19:16	RO	1h	<b>PCI Express Virtual Channel Capability Version (PCIEVCCV)</b> hard wired 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
15:0	RO	0002h	<b>Extended Capability ID (ECID)</b> Value of 0002 h identifies this linked list item (capability structure) as being for PCI Express Virtual Channel registers.



### 1.14.2 PVCCAP1 - Port VC Capability Register 1

B/D/F/Type: 0/1/0/MMR  
 Address Offset: 104-107h  
 Default Value: 00000000h  
 Access: RO  
 Size: 32 bits

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

Bit	Access	Default Value	Description
31:12	RO	00000h	<i>Reserved</i>
11:10	RO	00b	<i>Reserved</i> Reserved for Port Arbitration Table Entry Size ()
9:8	RO	00b	<i>Reserved</i> Reserved for Reference Clock ()
7	RO	0b	<i>Reserved</i>
6:4	RO	000b	<b>Low Priority Extended VC Count (LPEVCC)</b> 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.
3	RO	0b	<i>Reserved</i>
2:0	RO	000b	<b>Extended VC Count (EVCC):</b> Indicates the number of (extended) Virtual Channels in addition to the default VC supported by the device.



### 1.14.3 PVCCAP2 - Port VC Capability Register 2

B/D/F/Type: 0/1/0/MMR  
Address Offset: 108-10Bh  
Default Value: 00000000h  
Access: RO  
Size: 32 bits

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

Bit	Access	Default Value	Description
31:24	RO	00h	<b>VC Arbitration Table Offset (VCATO)</b> 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).
23:8	RO	0000h	<i>Reserved</i>
7:0	RO	00h	<b>Reserved for VC Arbitration Capability (VCAC)</b>

### 1.14.4 PVCCTL - Port VC Control

B/D/F/Type: 0/1/0/MMR  
Address Offset: 10C-10Dh  
Default Value: 0000h  
Access: RO; RW  
Size: 16 bits

Bit	Access	Default Value	Description
15:4	RO	000h	<i>Reserved</i>
3:1	RW	000b	<b>VC Arbitration Select (VCAS)</b> This field is 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.
0	RO	0b	<b>Reserved for Load VC Arbitration Table</b> 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.



### 1.14.5 VCORCAP - VCO Resource Capability

B/D/F/Type: 0/1/0/MMR  
 Address Offset: 110-113h  
 Default Value: 00000001h  
 Access: RO  
 Size: 32 bits

Bit	Access	Default Value	Description
31:24	RO	00h	<b>Reserved for Port Arbitration Table Offset</b>
23	RO	0b	<i>Reserved</i>
22:16	RO	00h	<b>Reserved for Maximum Time Slots</b>
15	RO	0b	<b>Reject Snoop Transactions (RSNPT)</b> 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 is rejected as an Unsupported Request
14:8	RO	00h	<i>Reserved</i>
7:0	RO	01h	<b>Port Arbitration Capability (PAC)</b> 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-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 MCH default indicates “Non-configurable hardware-fixed arbitration scheme”.



### 1.14.6 VCORCTL - VCO Resource Control

B/D/F/Type: 0/1/0/MMR  
Address Offset: 114-117h  
Default Value: 800000FFh  
Access: RO; RW  
Size: 32 bits

Controls the resources associated with PCI Express Virtual Channel 0.

Bit	Access	Default Value	Description
31	RO	1b	<b>VCO Enable (VCOE)</b> For VCO this is hard wired to 1 and read only as VCO can never be disabled.
30:27	RO	0h	<i>Reserved</i>
26:24	RO	000b	<b>VCO ID (VCOID)</b> Assigns a VC ID to the VC resource. For VCO this is hard wired to 0 and read only.
23:20	RO	0h	<i>Reserved</i>
19:17	RW	000b	<b>Port Arbitration Select (PAS)</b> 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.
16	RO	0b	<i>Reserved</i> Reserved for Load Port Arbitration Table ()
15:8	RO	00h	<i>Reserved</i>
7:1	RW	7Fh	<b>TC/VCO Map (TCVCOM)</b> 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.
0	RO	1b	<b>TC0/VCO Map (TC0VCOM)</b> Traffic Class 0 is always routed to VCO.



### 1.14.7 VCORSTS - VCO Resource Status

B/D/F/Type: 0/1/0/MMR  
 Address Offset: 11A-11Bh  
 Default Value: 0002h  
 Access: RO  
 Size: 16 bits

Reports the Virtual Channel specific status.

Bit	Access	Default Value	Description
15:2	RO	0000h	<b>Reserved and Zero</b>
1	RO	1b	<b>VCO Negotiation Pending (VCONP)</b> 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.
0	RO	0b	<i>Reserved</i> Reserved for Port Arbitration Table Status ()



### 1.14.8 PEGSSTS - PCI Express-G Sequence Status

B/D/F/Type: 0/1/0/MMR  
 Address Offset: 218-21Fh  
 Default Value: 0000000000000FFFh  
 Access: RO  
 Size: 64 bits

PCI Express status reporting that is required by the PCI Express spec.

Bit	Access	Default Value	Description
63:60	RO	0h	<i>Reserved</i>
59:48	RO	000h	<b>Next Transmit Sequence Number (NTSN)</b> Value of the NXT_TRANS_SEQ counter. This counter represents the transmit Sequence number to be applied to the next TLP to be transmitted onto the Link for the first time.
47:44	RO	0h	<i>Reserved</i>
43:32	RO	000h	<b>Next Packet Sequence Number (NPSN)</b> Packet sequence number to be applied to the next TLP to be transmitted or re-transmitted onto the Link.
31:28	RO	0h	<i>Reserved</i>
27:16	RO	000h	<b>Next Receive Sequence Number (NRSN)</b> This is the sequence number associated with the TLP that is expected to be received next.
15:12	RO	0h	<i>Reserved</i>
11:0	RO	FFFh	<b>Last Acknowledged Sequence Number (LASN)</b> This is the sequence number associated with the last acknowledged TLP.

### 1.14.9 PEGTXDEMPSEL - PEG Transmit De-Emphasis Select Register

B/D/F/Type: 0/1/0/MMR  
 Address Offset: DA8-DABh  
 Default Value: 43E00BF9h  
 Access: RW; RO  
 Size: 32 bits

(Sheet 1 of 4)

Bit	Access	Default Value	Description
31	RO	0b	<i>Reserved</i>





(Sheet 2 of 4)

Bit	Access	Default Value	Description																																																																																																			
30:26	RW	10000b	<p><b>Transmit De-emphasis For Gen2 (TXDGEN2SEL)</b>                      Default of TXDGEN2LSEL and TXDGEN2LDRVEN provides -6dB of de-emphasis. The following are the production settings:</p> <table border="1"> <thead> <tr> <th>Input Code</th> <th>De-emphasis</th> <th>Notes</th> </tr> </thead> <tbody> <tr> <td>00000b</td> <td>0dB</td> <td>No de-emphasis</td> </tr> <tr> <td>00001b</td> <td>-0.37dB</td> <td></td> </tr> <tr> <td>00010b</td> <td>-0.76dB</td> <td></td> </tr> <tr> <td>00011b</td> <td>-1.16dB</td> <td></td> </tr> <tr> <td>00100b</td> <td>-1.58dB</td> <td></td> </tr> <tr> <td>00101b</td> <td>-2.03dB</td> <td></td> </tr> <tr> <td>00110b</td> <td>-2.50dB</td> <td></td> </tr> <tr> <td>00111b</td> <td>-3.00dB</td> <td></td> </tr> <tr> <td>01000b</td> <td>-2.03dB</td> <td></td> </tr> <tr> <td>01001b</td> <td>-2.50dB</td> <td></td> </tr> <tr> <td>01010b</td> <td>-3.00dB</td> <td></td> </tr> <tr> <td>01011b</td> <td>-3.52dB</td> <td>Gen1 Default</td> </tr> <tr> <td>01100b</td> <td>-4.08dB</td> <td></td> </tr> <tr> <td>01101b</td> <td>-4.68dB</td> <td></td> </tr> <tr> <td>01110b</td> <td>-5.33dB</td> <td></td> </tr> <tr> <td>01111b</td> <td>-6.02dB</td> <td></td> </tr> <tr> <td>10000b</td> <td>-6.02dB</td> <td>Gen2 Default</td> </tr> <tr> <td>10001b</td> <td>-6.78dB</td> <td></td> </tr> <tr> <td>10010b</td> <td>-7.60dB</td> <td></td> </tr> <tr> <td>10011b</td> <td>-8.52dB</td> <td></td> </tr> <tr> <td>10100b</td> <td>-9.54dB</td> <td></td> </tr> <tr> <td>10101b</td> <td>-10.70dB</td> <td></td> </tr> <tr> <td>10110b</td> <td>-12.04dB</td> <td></td> </tr> <tr> <td>10111b</td> <td>-13.62dB</td> <td></td> </tr> <tr> <td>11000b</td> <td>-10.70dB</td> <td></td> </tr> <tr> <td>11001b</td> <td>-12.04dB</td> <td></td> </tr> <tr> <td>11010b</td> <td>-13.62dB</td> <td></td> </tr> <tr> <td>11011b</td> <td>-15.56dB</td> <td></td> </tr> <tr> <td>11100b</td> <td>-18.06dB</td> <td></td> </tr> <tr> <td>11101b</td> <td>-21.58dB</td> <td></td> </tr> <tr> <td>11110b</td> <td>-27.60dB</td> <td></td> </tr> <tr> <td>11111b</td> <td>N/A</td> <td>Not Valid</td> </tr> </tbody> </table>	Input Code	De-emphasis	Notes	00000b	0dB	No de-emphasis	00001b	-0.37dB		00010b	-0.76dB		00011b	-1.16dB		00100b	-1.58dB		00101b	-2.03dB		00110b	-2.50dB		00111b	-3.00dB		01000b	-2.03dB		01001b	-2.50dB		01010b	-3.00dB		01011b	-3.52dB	Gen1 Default	01100b	-4.08dB		01101b	-4.68dB		01110b	-5.33dB		01111b	-6.02dB		10000b	-6.02dB	Gen2 Default	10001b	-6.78dB		10010b	-7.60dB		10011b	-8.52dB		10100b	-9.54dB		10101b	-10.70dB		10110b	-12.04dB		10111b	-13.62dB		11000b	-10.70dB		11001b	-12.04dB		11010b	-13.62dB		11011b	-15.56dB		11100b	-18.06dB		11101b	-21.58dB		11110b	-27.60dB		11111b	N/A	Not Valid
Input Code	De-emphasis	Notes																																																																																																				
00000b	0dB	No de-emphasis																																																																																																				
00001b	-0.37dB																																																																																																					
00010b	-0.76dB																																																																																																					
00011b	-1.16dB																																																																																																					
00100b	-1.58dB																																																																																																					
00101b	-2.03dB																																																																																																					
00110b	-2.50dB																																																																																																					
00111b	-3.00dB																																																																																																					
01000b	-2.03dB																																																																																																					
01001b	-2.50dB																																																																																																					
01010b	-3.00dB																																																																																																					
01011b	-3.52dB	Gen1 Default																																																																																																				
01100b	-4.08dB																																																																																																					
01101b	-4.68dB																																																																																																					
01110b	-5.33dB																																																																																																					
01111b	-6.02dB																																																																																																					
10000b	-6.02dB	Gen2 Default																																																																																																				
10001b	-6.78dB																																																																																																					
10010b	-7.60dB																																																																																																					
10011b	-8.52dB																																																																																																					
10100b	-9.54dB																																																																																																					
10101b	-10.70dB																																																																																																					
10110b	-12.04dB																																																																																																					
10111b	-13.62dB																																																																																																					
11000b	-10.70dB																																																																																																					
11001b	-12.04dB																																																																																																					
11010b	-13.62dB																																																																																																					
11011b	-15.56dB																																																																																																					
11100b	-18.06dB																																																																																																					
11101b	-21.58dB																																																																																																					
11110b	-27.60dB																																																																																																					
11111b	N/A	Not Valid																																																																																																				



(Sheet 3 of 4)

Bit	Access	Default Value	Description																																																																																																			
25:21	RW	11111b	<b>Transmit Driver Enable For Gen2 (TXDGEN2DRVEN)</b> Default of TXDGEN2SEL and TXDGEN2DRVEN provides -3.5dB of de-emphasis.																																																																																																			
20:17	RO	0000b	<i>Reserved</i>																																																																																																			
16:13	RO	0000b	<i>Reserved</i>																																																																																																			
12:8	RW	01011b	<p><b>Transmit De-emphasis For Gen1 (TXDGEN1SEL)</b>            Default of TXDGEN1SEL and TXDGEN1DRVEN provides -3.5dB of de-emphasis. The following are the production settings:</p> <table border="1"> <thead> <tr> <th>Input Code</th> <th>De-emphasis</th> <th>Notes</th> </tr> </thead> <tbody> <tr> <td>00000b</td> <td>0dB</td> <td>No de-emphasis</td> </tr> <tr> <td>00001b</td> <td>-0.37dB</td> <td></td> </tr> <tr> <td>00010b</td> <td>-0.76dB</td> <td></td> </tr> <tr> <td>00011b</td> <td>-1.16dB</td> <td></td> </tr> <tr> <td>00100b</td> <td>-1.58dB</td> <td></td> </tr> <tr> <td>00101b</td> <td>-2.03dB</td> <td></td> </tr> <tr> <td>00110b</td> <td>-2.50dB</td> <td></td> </tr> <tr> <td>00111b</td> <td>-3.00dB</td> <td></td> </tr> <tr> <td>01000b</td> <td>-2.03dB</td> <td></td> </tr> <tr> <td>01001b</td> <td>-2.50dB</td> <td></td> </tr> <tr> <td>01010b</td> <td>-3.00dB</td> <td></td> </tr> <tr> <td>01011b</td> <td>-3.52dB</td> <td>Gen1 Default</td> </tr> <tr> <td>01100b</td> <td>-4.08dB</td> <td></td> </tr> <tr> <td>01101b</td> <td>-4.68dB</td> <td></td> </tr> <tr> <td>01110b</td> <td>-5.33dB</td> <td></td> </tr> <tr> <td>01111b</td> <td>-6.02dB</td> <td></td> </tr> <tr> <td>10000b</td> <td>-6.02dB</td> <td>Gen2 Default</td> </tr> <tr> <td>10001b</td> <td>-6.78dB</td> <td></td> </tr> <tr> <td>10010b</td> <td>-7.60dB</td> <td></td> </tr> <tr> <td>10011b</td> <td>-8.52dB</td> <td></td> </tr> <tr> <td>10100b</td> <td>-9.54dB</td> <td></td> </tr> <tr> <td>10101b</td> <td>-10.70dB</td> <td></td> </tr> <tr> <td>10110b</td> <td>-12.04dB</td> <td></td> </tr> <tr> <td>10111b</td> <td>-13.62dB</td> <td></td> </tr> <tr> <td>11000b</td> <td>-10.70dB</td> <td></td> </tr> <tr> <td>11001b</td> <td>-12.04dB</td> <td></td> </tr> <tr> <td>11010b</td> <td>-13.62dB</td> <td></td> </tr> <tr> <td>11011b</td> <td>-15.56dB</td> <td></td> </tr> <tr> <td>11100b</td> <td>-18.06dB</td> <td></td> </tr> <tr> <td>11101b</td> <td>-21.58dB</td> <td></td> </tr> <tr> <td>11110b</td> <td>-27.60dB</td> <td></td> </tr> <tr> <td>11111b</td> <td>N/A</td> <td>Not Valid</td> </tr> </tbody> </table>	Input Code	De-emphasis	Notes	00000b	0dB	No de-emphasis	00001b	-0.37dB		00010b	-0.76dB		00011b	-1.16dB		00100b	-1.58dB		00101b	-2.03dB		00110b	-2.50dB		00111b	-3.00dB		01000b	-2.03dB		01001b	-2.50dB		01010b	-3.00dB		01011b	-3.52dB	Gen1 Default	01100b	-4.08dB		01101b	-4.68dB		01110b	-5.33dB		01111b	-6.02dB		10000b	-6.02dB	Gen2 Default	10001b	-6.78dB		10010b	-7.60dB		10011b	-8.52dB		10100b	-9.54dB		10101b	-10.70dB		10110b	-12.04dB		10111b	-13.62dB		11000b	-10.70dB		11001b	-12.04dB		11010b	-13.62dB		11011b	-15.56dB		11100b	-18.06dB		11101b	-21.58dB		11110b	-27.60dB		11111b	N/A	Not Valid
Input Code	De-emphasis	Notes																																																																																																				
00000b	0dB	No de-emphasis																																																																																																				
00001b	-0.37dB																																																																																																					
00010b	-0.76dB																																																																																																					
00011b	-1.16dB																																																																																																					
00100b	-1.58dB																																																																																																					
00101b	-2.03dB																																																																																																					
00110b	-2.50dB																																																																																																					
00111b	-3.00dB																																																																																																					
01000b	-2.03dB																																																																																																					
01001b	-2.50dB																																																																																																					
01010b	-3.00dB																																																																																																					
01011b	-3.52dB	Gen1 Default																																																																																																				
01100b	-4.08dB																																																																																																					
01101b	-4.68dB																																																																																																					
01110b	-5.33dB																																																																																																					
01111b	-6.02dB																																																																																																					
10000b	-6.02dB	Gen2 Default																																																																																																				
10001b	-6.78dB																																																																																																					
10010b	-7.60dB																																																																																																					
10011b	-8.52dB																																																																																																					
10100b	-9.54dB																																																																																																					
10101b	-10.70dB																																																																																																					
10110b	-12.04dB																																																																																																					
10111b	-13.62dB																																																																																																					
11000b	-10.70dB																																																																																																					
11001b	-12.04dB																																																																																																					
11010b	-13.62dB																																																																																																					
11011b	-15.56dB																																																																																																					
11100b	-18.06dB																																																																																																					
11101b	-21.58dB																																																																																																					
11110b	-27.60dB																																																																																																					
11111b	N/A	Not Valid																																																																																																				



(Sheet 4 of 4)

Bit	Access	Default Value	Description
7:3	RW	11111b	<b>Transmit Driver Enable for Gen1 (TXDGEN1DRVEN)</b> Default of TXDGEN1SEL and TXDGEN1DRVEN provides -3.5dB of de-emphasis.
2:1	RO	00b	<i>Reserved</i>
0	RW	1b	<i>Reserved</i>

## 1.15 DMIBAR

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
DMI Virtual Channel Enhanced Capability	DMIVCECH	0	3	00010002h	RW-O; RO
DMI Port VC Capability Register 1	DMIPVCCAP1	4	7	00000000h	RO; RW-O
DMI Port VC Capability Register 2	DMIPVCCAP2	8	B	00000000h	RO
DMI Port VC Control	DMIPVCCTL	C	D	0000h	RO; RW
DMI VC0 Resource Capability	DMIVCORCAP	10	13	00000001h	RO
DMI VC0 Resource Control	DMIVCORCTL0	14	17	800000FFh	RO; RW
DMI VC0 Resource Status	DMIVCORSTS	1A	1B	0002h	RO
DMI VC1 Resource Capability	DMIVC1RCAP	1C	1F	00008001h	RO
DMI VC1 Resource Control	DMIVC1RCTL1	20	23	01000000h	RO; RW
DMI VC1 Resource Status	DMIVC1RSTS	26	27	0002h	RO
DMI Link Capabilities	DMILCAP	84	87	00012C41h	RO; RW-O
DMI Link Control	DMILCTL	88	89	0000h	RO; RW
DMI Link Status	DMILSTS	8A	8B	0001h	RO



### 1.15.1 DMIVCECH - DMI Virtual Channel Enhanced Capability

B/D/F/Type: 0/0/0/DMIBAR  
Address Offset: 0-3h  
Default Value: 00010002h  
Access: RW-O; RO  
Size: 32 bits

Indicates DMI Virtual Channel capabilities.

Bit	Access	Default Value	Description
31:20	RW-O	000h	<b>Pointer to Next Capability (PNC)</b> This field contains the offset to the next PCI Express capability structure in the linked list of capabilities (Link Declaration Capability).
19:16	RO	1h	<b>PCI Express Virtual Channel Capability Version (PCIEVCCV)</b> hard wired to 1 to indicate compliances with the <i>PCI Local Bus Specification</i> . <b>Note:</b> This version does not change for 2.0 compliance.
15:0	RO	0002h	<b>Extended Capability ID (ECID)</b> Value of 0002h identifies this linked list item (capability structure) as being for PCI Express Virtual Channel registers.



### 1.15.2 DMIPVCCAP1 - DMI Port VC Capability Register 1

B/D/F/Type: 0/0/0/DMIBAR  
 Address Offset: 4-7h  
 Default Value: 00000000h  
 Access: RO; RW-O  
 Size: 32 bits

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

Bit	Access	Default Value	Description
31:7	RO	0000000h	<i>Reserved</i>
6:4	RO	000b	<b>Low Priority Extended VC Count (LPEVCC)</b> 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.
3	RO	0b	<i>Reserved</i>
2:0	RW-O	000b	<b>Extended VC Count (EVCC)</b> Indicates the number of (extended) Virtual Channels in addition to the default VC supported by the device. For DMI, only the default Virtual Channel (VC0) is advertised in the Extended VC Capability structure.

### 1.15.3 DMIPVCCAP2 - DMI Port VC Capability Register 2

B/D/F/Type: 0/0/0/DMIBAR  
 Address Offset: 8-Bh  
 Default Value: 00000000h  
 Access: RO  
 Size: 32 bits

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

Bit	Access	Default Value	Description
31:24	RO	00h	<i>Reserved for VC Arbitration Table Offset (RSVD)</i>
23:8	RO	0000h	<i>Reserved</i>
7:0	RO	00h	<i>Reserved for VC Arbitration Capability (VCAC)</i>



### 1.15.4 DMIPVCCTL - DMI Port VC Control

B/D/F/Type: 0/0/0/DMIBAR  
 Address Offset: C-Dh  
 Default Value: 0000h  
 Access: RO; RW  
 Size: 16 bits

Bit	Access	Default Value	Description
15:4	RO	000h	<i>Reserved</i>
3:1	RW	000b	<b>VC Arbitration Select (VCAS)</b> This field is 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 <i>PCI Express Base Specification</i> for more details.
0	RO	0b	<i>Reserved for Load VC Arbitration Table (RSVD)</i>

### 1.15.5 DMIVCORCAP - DMI VCO Resource Capability

B/D/F/Type: 0/0/0/DMIBAR  
 Address Offset: 10-13h  
 Default Value: 00000001h  
 Access: RO  
 Size: 32 bits

Bit	Access	Default Value	Description
31:24	RO	00h	<i>Reserved for Port Arbitration Table Offset (RSVD)</i>
23	RO	0b	<i>Reserved</i>
22:16	RO	00h	<i>Reserved for Maximum Time Slots (RSVD)</i>
15	RO	0b	<b>Reject Snoop Transactions (REJSNPT)</b> 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 is rejected as an Unsupported Request.
14:8	RO	00h	<i>Reserved</i>
7:0	RO	01h	<b>Port Arbitration Capability (PAC)</b> Having only Bit 0 set indicates that the only supported arbitration scheme for this VC is non-configurable hardware-fixed.



### 1.15.6 DMIVCORCTLO - DMI VCO Resource Control

B/D/F/Type: 0/0/0/DMIBAR  
 Address Offset: 14-17h  
 Default Value: 800000FFh  
 Access: RO; RW  
 Size: 32 bits

Controls the resources associated with PCI Express Virtual Channel 0.

Bit	Access	Default Value	Description
31	RO	1b	<b>Virtual Channel 0 Enable (VCOE)</b> For VCO this is hard wired to 1 and read only as VCO can never be disabled.
30:27	RO	0h	<i>Reserved</i>
26:24	RO	000b	<b>Virtual Channel 0 ID (VCOID)</b> Assigns a VC ID to the VC resource. For VCO this is hard wired to 0 and read only.
23:20	RO	0h	<i>Reserved</i>
19:17	RW	000b	<b>Port Arbitration Select (PAS)</b> 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.
16:8	RO	000h	<i>Reserved</i>
7:1	RW	7Fh	<b>Traffic Class/Virtual Channel 0 Map (TCVCOM)</b> 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.
0	RO	1b	<b>Traffic Class 0/Virtual Channel 0 Map (TCOVCOM)</b> Traffic Class 0 is always routed to VCO.



### 1.15.7 DMI VCO RSTS - DMI VCO Resource Status

B/D/F/Type: 0/0/0/DMIBAR  
Address Offset: 1A-1Bh  
Default Value: 0002h  
Access: RO  
Size: 16 bits

Reports the Virtual Channel specific status.

Bit	Access	Default Value	Description
15:2	RO	0000h	<i>Reserved</i> Reserved and Zero for future R/WC/S implementations. Software must use 0 for writes to these bits.
1	RO	1b	<b>Virtual Channel 0 Negotiation Pending (VCONP)</b> 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, software must check whether the VC Negotiation Pending fields for that Virtual Channel are cleared in both Components on a Link.
0	RO	0b	<i>Reserved</i>





### 1.15.8 DMIVC1RCAP - DMI VC1 Resource Capability

B/D/F/Type: 0/0/0/DMIBAR  
 Address Offset: 1C-1Fh  
 Default Value: 00008001h  
 Access: RO  
 Size: 32 bits

Bit	Access	Default Value	Description
31:24	RO	00h	<i>Reserved for Port Arbitration Table Offset (RSVD)</i>
23	RO	0b	<i>Reserved</i>
22:16	RO	00h	<i>Reserved for Maximum Time Slots (RSVD)</i>
15	RO	1b	<b>Reject Snoop Transactions (REJSNPT)</b> 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 is rejected as an Unsupported Request.
14:8	RO	00h	<i>Reserved</i>
7:0	RO	01h	<b>Port Arbitration Capability (PAC)</b> Having only Bit 0 set indicates that the only supported arbitration scheme for this VC is non-configurable hardware-fixed.



### 1.15.9 DMIVC1RCTL1 - DMI VC1 Resource Control

B/D/F/Type: 0/0/0/DMIBAR  
Address Offset: 20-23h  
Default Value: 01000000h  
Access: RO; RW  
Size: 32 bits  
Controls the resources associated with PCI Express Virtual Channel 1.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31	RW	0b	<b>Virtual Channel Enable (VCE)</b> 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.
30:27	RO	0h	<i>Reserved</i>
26:24	RW	001b	<b>Virtual Channel ID (VCID)</b> 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.
23:20	RO	0h	<i>Reserved</i>
19:17	RW	000b	<b>Port Arbitration Select (PAS)</b> 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.
16:8	RO	000h	<i>Reserved</i>



(Sheet 2 of 2)

Bit	Access	Default Value	Description
7:1	RW	00h	<p><b>Traffic Class/Virtual Channel Map (TCVCM)</b>                      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.                      BIOS Requirement: Program this field, including Bit 0, with the value 00100010b (22h), which maps TC1 and TC5 to VC1.</p>
0	RO	0b	<p><b>Traffic Class 0/Virtual Channel 1 Map (TC0VC1M)</b>                      Traffic Class 0 is always routed to VC0.</p>

### 1.15.10 DMIVC1RSTS - DMI VC1 Resource Status

B/D/F/Type: 0/0/0/DMIBAR  
 Address Offset: 26-27h  
 Default Value: 0002h  
 Access: RO  
 Size: 16 bits  
 Reports the Virtual Channel specific status.

Bit	Access	Default Value	Description
15:2	RO	0000h	<i>Reserved</i>
1	RO	1b	<p><b>Virtual Channel 1 Negotiation Pending (VC1NP)</b>                      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.</p>
0	RO	0b	<i>Reserved</i>



### 1.15.11 DMILCAP - DMI Link Capabilities

B/D/F/Type: 0/0/0/DMIBAR  
 Address Offset: 84-87h  
 Default Value: 00012C41h  
 Access: RO; RW-O  
 Size: 32 bits

Indicates DMI specific capabilities.

Bit	Access	Default Value	Description
31:18	RO	0000h	<i>Reserved</i>
17:15	RW-O	010b	<p><b>L1 Exit Latency (L1SELAT)</b>            Indicates the length of time this Port requires to complete the transition from L1 to L0. The value 010b indicates the range of 2 <math>\mu</math>s to less than 4 <math>\mu</math>s.</p> <p>000: Less than 1 <math>\mu</math>s            001: 1 <math>\mu</math>s to less than 2 <math>\mu</math>s            010: 2 <math>\mu</math>s to less than 4 <math>\mu</math>s            011: 4 <math>\mu</math>s to less than 8 <math>\mu</math>s            100: 8 <math>\mu</math>s to less than 16 <math>\mu</math>s            101: 16 <math>\mu</math>s to less than 32 <math>\mu</math>s            110: 32 <math>\mu</math>s-64 <math>\mu</math>s            111: More than 64 <math>\mu</math>s</p> <p>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.</p>
14:12	RW-O	010b	<p><b>LOs Exit Latency (LOSELAT)</b>            Indicates the length of time this Port requires to complete the transition from LOs to L0.</p> <p>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 <math>\mu</math>s            101: 1 <math>\mu</math>s to less than 2 <math>\mu</math>s            110: 2 <math>\mu</math>s-4 <math>\mu</math>s            111: More than 4 <math>\mu</math>s</p>
11:10	RO	11b	<p><b>Active State Link PM Support (ASLPMS)</b>            L0s &amp; L1 entry supported.</p>
9:4	RO	04h	<p><b>Max Link Width (MLW)</b>            Indicates the maximum number of lanes supported for this link.</p>
3:0	RO	1h	<p><b>Max Link Speed (MLS)</b>            hard wired to indicate 2.5 Gb/s.</p>



### 1.15.12 DMILCTL - DMI Link Control

B/D/F/Type: 0/0/0/DMIBAR  
 Address Offset: 88-89h  
 Default Value: 0000h  
 Access: RO; RW  
 Size: 16 bits  
 Allows control of DMI.

Bit	Access	Default Value	Description
15:8	RO	00h	<i>Reserved</i>
7	RW	0b	<b>Extended Synch (EXTSYNC)</b> 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.
6:3	RO	0h	<i>Reserved</i>
2	RO	0b	<i>Reserved</i>
1:0	RW	00b	<b>Active State Power Management Support (ASPMS)</b> Controls the level of active state power management supported on the given link. 00: Disabled 01: L0s Entry Supported 10: L1 Entry Enabled 11: L0s and L1 Entry Supported



### 1.15.13 DMILSTS - DMI Link Status

B/D/F/Type: 0/0/0/DMIBAR  
 Address Offset: 8A-8Bh  
 Default Value: 0001h  
 Access: RO  
 Size: 16 bits  
 Indicates DMI status.

Bit	Access	Default Value	Description
15:10	RO	00h	<i>Reserved</i>
9:4	RO	00h	<b>Negotiated Width (NWID)</b> 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.
3:0	RO	1h	<b>Negotiated Speed (NSPD)</b> Indicates negotiated link speed. 1h: 2.5 Gb/s All other encodings are reserved.

## 1.16 PCI Device 2 Function 0

(Sheet 1 of 3)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Vendor Identification	VID2	0	1	8086h	RO
Device Identification	DID2	2	3	0046h	RO
PCI Command	PCICMD2	4	5	0000h	RO; RW
PCI Status	PCISTS2	6	7	0090h	RO
Revision Identification	RID2	8	8	12h	RO
Class Code	CC	9	B	030000h	RO
Cache Line Size	CLS	C	C	00h	RO
Master Latency Timer	MLT2	D	D	00h	RO



(Sheet 2 of 3)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Header Type	HDR2	E	E	00h	RO
Graphics Translation Table, Memory Mapped Range Address	GTTMMADR	10	17	000000000000 0004h	RO; RW
Graphics Memory Range Address	GMADR	18	1F	000000000000 000Ch	RO; RW-L; RW
I/O Base Address	IOBAR	20	23	00000001h	RO; RW
Subsystem Vendor Identification	SVID2	2C	2D	0000h	RW-O
Subsystem Identification	SID2	2E	2F	0000h	RW-O
Video BIOS ROM Base Address	ROMADR	30	33	00000000h	RO
Capabilities Pointer	CAPPOINT	34	34	90h	RO
Interrupt Line	INTRLINE	3C	3C	00h	RW
Interrupt Pin	INTRPIN	3D	3D	01h	RO
Minimum Grant	MINGNT	3E	3E	00h	RO
Maximum Latency	MAXLAT	3F	3F	00h	RO
Graphics Enhanced Intel® SpeedStep Technology Capability	GGCTL	4C	4F	003F003Fh	RO
Processor Graphics Control Register	MGGC	52	53	0030h	RO
Device Enable	DEVEN	54	57	0000210Bh	RO
Software Scratch Read Write	SSRW	58	5B	00000000h	RW
Base of Stolen Memory	BSM	5C	5F	00000000h	RO
Hardware Scratch Read Write	HSRW	60	61	0000h	RW
Message Control	MC	92	93	0000h	RO; RW
Message Address	MA	94	97	00000000h	RO; RW
Message Data	MD	98	99	0000h	RW



(Sheet 3 of 3)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Power Management Capabilities ID	PMCAPID	D0	D1	A401h	RO
Power Management Capabilities	PMCAP	D2	D3	0022h	RO
Power Management Control/Status	PMCS	D4	D5	0000h	RO; RW
Software SMI	SWSMI	E0	E1	0000h	RW
Graphics System Event Register	GSE	E4	E7	00000000h	RW
ASL Storage	ASLS	FC	FF	00000000h	RW

### 1.16.1 VID2 - Vendor Identification

B/D/F/Type: 0/2/0/PCI

Address Offset: 0-1h

Default Value: 8086h

Access: RO

Size: 16 bits

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

Bit	Access	Default Value	Description
15:0	RO	8086h	<b>Vendor Identification Number (VID)</b> PCI standard identification for Intel.





### 1.16.2 DID2 - Device Identification

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 2-3h  
 Default Value: 0046h  
 Access: RO  
 Size: 16 bits

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

Bit	Access	Default Value	Description
15:0	RO	0046h	<b>Device Identification Number (DID)</b> This is a 16-bit value assigned to the processor Graphics device.

### 1.16.3 PCICMD2 - PCI Command

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 4-5h  
 Default Value: 0000h  
 Access: RO; RW  
 Size: 16 bits

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.

(Sheet 1 of 2)

Bit	Access	Default Value	RST/ PWR	Description
15:11	RO	00h	Core	<i>Reserved</i>
10	RW	0b	FLR, Core	<b>Interrupt Disable (INTDIS)</b> 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.
9	RO	0b	Core	<b>Fast Back-to-Back (FB2B)</b> Not Implemented. Hard wired to 0.
8	RO	0b	Core	<b>SERR Enable (SERRE)</b> Not Implemented. Hard wired to 0.
7	RO	0b	Core	<b>Address/Data Stepping Enable (ADSTEP)</b> Not Implemented. Hard wired to 0.
6	RO	0b	Core	<b>Parity Error Enable (PERRE)</b> Not Implemented. Hard wired 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.



(Sheet 2 of 2)

Bit	Access	Default Value	RST/ PWR	Description
5	RO	0b	Core	<b>Video Palette Snooping (VPS)</b> This bit is hard wired to 0 to disable snooping.
4	RO	0b	Core	<b>Memory Write and Invalidate Enable (MWIE)</b> hard wired to 0. The IGD does not support memory write and invalidate commands.
3	RO	0b	Core	<b>Special Cycle Enable (SCE)</b> This bit is hard wired to 0. The IGD ignores Special cycles.
2	RW	0b	FLR, Core	<b>Bus Master Enable (BME)</b> 0 = Disable IGD bus mastering. 1 = Enable the IGD to function as a PCI compliant master.
1	RW	0b	FLR, Core	<b>Memory Access Enable (MAE)</b> This bit controls the IGD's response to memory space accesses. 0 = Disable. 1 = Enable.
0	RW	0b	FLR, Core	<b>I/O Access Enable (IOAE)</b> This bit controls the IGD's response to I/O space accesses. 0 = Disable. 1 = Enable.

### 1.16.4 PCISTS2 - PCI Status

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 6-7h  
 Default Value: 0090h  
 Access: RO  
 Size: 16 bits

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.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
15	RO	0b	<b>Detected Parity Error (DPE)</b> Since the IGD does not detect parity, this bit is always hard wired to 0.
14	RO	0b	<b>Signaled System Error (SSE)</b> The IGD never asserts SERR#, therefore this bit is hard wired to 0.
13	RO	0b	<b>Received Master Abort Status (RMAS)</b> The IGD never gets a Master Abort, therefore this bit is hard wired to 0.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
12	RO	0b	<b>Received Target Abort Status (RTAS)</b> The IGD never gets a Target Abort, therefore this bit is hard wired to 0.
11	RO	0b	<b>Signaled Target Abort Status (STAS)</b> hard wired to 0. The IGD does not use target abort semantics.
10:9	RO	00b	<b>DEVSEL Timing (DEVT)</b> N/A. These bits are hard wired to "00".
8	RO	0b	<b>Master Data Parity Error Detected (DPD)</b> Since Parity Error Response is hard wired to disabled (and the IGD does not do any parity detection), this bit is hard wired to 0.
7	RO	1b	<b>Fast Back-to-Back (FB2B)</b> hard wired to 1. The IGD accepts fast back-to-back when the transactions are not to the same agent.
6	RO	0b	<b>User Defined Format (UDF)</b> hard wired to 0.
5	RO	0b	<b>66 MHz PCI Capable (66C)</b> N/A - hard wired to 0.
4	RO	1b	<b>Capability List (CLIST)</b> 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.
3	RO	0b	<b>Interrupt Status (INTSTS)</b> 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.
2:0	RO	000b	<i>Reserved</i>



### 1.16.5 RID2 - Revision Identification

B/D/F/Type: 0/2/0/PCI  
Address Offset: 8h  
Default Value: 12h  
Access: RO  
Size: 8 bits

This register contains the revision number of the processor. The Revision ID (RID) is a traditional 8-bit Read Only (RO) register located at offset 08h in the standard PCI header of every PCI/PCI Express compatible device and function.

Following reset, the SRID is returned when the RID is read at offset 08h. The SRID value reflects the actual product stepping. To select the CRID value, BIOS/configuration software writes a key value of 69h to Bus 0, Device 0, Function 0 (DMI device) of the CPU's RID register at offset 08h. This causes the CRID to be returned when the RID is read at offset 08h.

#### Stepping Revision ID (SRID)

This register contains the revision number of the CPU.

The SRID is a 8-bit hardwired value assigned by Intel, based on product's stepping. The SRID is not a directly addressable PCI register. The SRID value is reflected through the RID register when appropriately addressed.

#### Compatible Revision ID (CRID)

The CRID is an 8-bit hardwired value assigned by Intel during manufacturing process. Normally, the value assigned as the CRID will be identical to the SRID value of a previous stepping of the product with which the new product is deemed "compatible".

The CRID is not a directly addressable PCI register. The CRID value is reflected through the RID register when appropriately addressed.

Bit	Access	Default Value	RST/ PWR	Description
7:0	RO	10h	Core	<b>Revision Identification Number (RID)</b> This is an 8-bit value that indicates the revision identification number for the processor Device 0. For the C-2 Stepping, these values are: SRID = 12h CRID = 02h



### 1.16.6 CC - Class Code

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 9-Bh  
 Default Value: 030000h  
 Access: RO  
 Size: 24 bits

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.

Bit	Access	Default Value	Description
23:16	RO	03h	<b>Base Class Code (BCC)</b> This is an 8-bit value that indicates the base class code for the processor. This code has the value 03h, indicating a Display Controller. When MCHBAR Offset 44 Bit 31 is 0 this code has the value 03h, indicating a Display Controller. When MCHBAR Offset 44 Bit 31 is 1 this code has the value 04h, indicating a Multimedia Device.
15:8	RO	00h	<b>Sub-Class Code (SUBCC)</b> When MCHBAR Offset 44 Bit 31 is 0 this value is determined based on Device 0 GGC register, GMS and IVD fields. 00h: VGA compatible 80h: Non VGA (GMS = "0000" or IVD = "1") When MCHBAR offset 44 bit 31 is 1 this value is 80h, indicating other multimedia device.
7:0	RO	00h	<b>Programming Interface (PI)</b> When MCHBAR Offset 44 Bit 31 is 0 this value is 00h, indicating a Display Controller. When MCHBAR Offset 44 Bit 31 is 1 this value is 00h, indicating a NOP.



### 1.16.7 CLS - Cache Line Size

B/D/F/Type: 0/2/0/PCI  
Address Offset: Ch  
Default Value: 00h  
Access: RO  
Size: 8 bits

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

Bit	Access	Default Value	Description
7:0	RO	00h	<b>Cache Line Size (CLS)</b> This field is hard wired to 0's. 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.

### 1.16.8 MLT2 - Master Latency Timer

B/D/F/Type: 0/2/0/PCI  
Address Offset: Dh  
Default Value: 00h  
Access: RO  
Size: 8 bits

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

Bit	Access	Default Value	Description
7:0	RO	00h	<b>Master Latency Timer Count Value (MLTCV)</b> Hard wired to 0's.



### 1.16.9 HDR2 - Header Type

B/D/F/Type: 0/2/0/PCI  
 Address Offset: Eh  
 Default Value: 00h  
 Access: RO  
 Size: 8 bits

This register contains the Header Type of the IGD.

Bit	Access	Default Value	Description
7	RO	0b	<b>Multi Function Status (MFUNC)</b> Indicates if the device is a Multi-Function Device. The value is hard wired to 0 to indicate that this Internal Graphics Device is a single-function device.
6:0	RO	00h	<b>Header Code (H)</b> 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.

### 1.16.10 GTTMMADR - Graphics Translation Table, Memory Mapped Range Address

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 10-17h  
 Default Value: 0000000000000004h  
 Access: RO; RW  
 Size: 64 bits

This register requests allocation for the combined Graphics Translation Table Modification Range and Memory Mapped Range. The range requires 4 MB combined for MMIO and Global GTT aperture, with 512K of that used by MMIO and 2 MB used by GTT. GTTADR will begin at (GTTMMADR + 2 MB) while the MMIO base address is the same as GTTMMADR.

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. There are some exceptions to this - see GTT-TLB in the Programming Interface chapter.

The allocation is for 4 MB and the base address is defined by bits [35:22].



Bit	Access	Default Value	RST/ PWR	Description
63:36	RW	0000000h	FLR, Core	<i>Reserved for Memory Base Address</i> Must be set to 0 since addressing above 64 GB is not supported.
35:22	RW	0000h	FLR, Core	<i>Memory Base Address (MBA)</i> Set by the OS, these bits correspond to address signals [35:22]. 4 MB combined for MMIO and Global GTT table aperture (512 KB for MMIO and 2 MB for GTT).
21:4	RO	00000h	Core	<i>Reserved</i> Hardwired to 0's to indicate at least 4-MB address range.
3	RO	0b	Core	<b>Prefetchable Memory (PREFMEM)</b> Hardwired to 0 to prevent prefetching.
2:1	RO	10b	Core	<b>Memory Type (MEMTYP)</b> 00: To indicate 32-bit base address 01: Reserved 10: To indicate 64 bit base address 11: Reserved
0	RO	0b	Core	<b>Memory/IO Space (MIOS)</b> hard wired to 0 to indicate memory space.

### 1.16.11 GMADR - Graphics Memory Range Address

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 18-1Fh  
 Default Value: 0000000000000000Ch  
 Access: RO; RW-L; RW  
 Size: 64 bits

IGD graphics memory base address is specified in this register.

Software must not change the value in MSAC[2:1] (offset 62h) after writing to the GMADR register.

(Sheet 1 of 2)

Bit	Access	Default Value	RST/ PWR	Description
63:36	RW	0000000h	FLR, Core	<b>Memory Base Address (MBA2)</b> Set by the OS, these bits correspond to address signals [63:36].
35:29	RW	0000000b	FLR, Core	<b>Memory Base Address (MBA)</b> Set by the OS, these bits correspond to address signals [35:29].





(Sheet 2 of 2)

Bit	Access	Default Value	RST/ PWR	Description
28	RW-L	0b	FLR, Core	<b>512-MB Address Mask (512ADMSK)</b> This bit is either part of the Memory Base Address (R/W) or part of the Address Mask (RO), depending on the value of MSAC[2:1]. See MSAC (Device 2, Function 0, Offset 62h) for details.
27	RW-L	0b	FLR, Core	<b>256-MB Address Mask (256ADMSK)</b> This bit is either part of the Memory Base Address (R/W) or part of the Address Mask (RO), depending on the value of MSAC[2:1]. See MSAC (Device 2, Function 0, Offset 62h) for details.
26:4	RO	000000h	Core	<b>Address Mask (ADM)</b> hard wired to 0's to indicate at least 128-MB address range.
3	RO	1b	Core	<b>Prefetchable Memory (PREFMEM)</b> hard wired to 1 to enable prefetching.
2:1	RO	10b	Core	<b>Memory Type (MEMTYP)</b> 00: indicate 32-bit address. 10: Indicate 64-bit address
0	RO	0b	Core	<b>Memory/IO Space (MIOS)</b> hard wired to 0 to indicate memory space.

### 1.16.12 IOBAR - I/O Base Address

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 20-23h  
 Default Value: 00000001h  
 Access: RO; RW  
 Size: 32 bits

This register provides the Base offset of the I/O registers within Device 2. Bits 15:3 are programmable allowing the I/O Base to be located anywhere in 16-bit I/O Address Space. Bits 2:1 are fixed and return zero, Bit 0 is hard wired to a one indicating that 8 bytes of I/O space are decoded. Access to the 8 bytes 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 or if internal graphics is disabled thru the fuse or fuse override mechanisms.

If accesses to this IO bar is allowed then the processor claims all 8-, 16- or 32-bit I/O cycles from the CPU that falls within the 8B claimed.



Bit	Access	Default Value	RST/PWR	Description
31:16	RO	0000h	Core	<i>Reserved</i>
15:3	RW	0000h	FLR, Core	<b>IO Base Address (IOBASE)</b> Set by the OS, these bits correspond to address signals [15:3].
2:1	RO	00b	Core	<b>Memory Type (MEMTYPE)</b> hard wired to 0's to indicate 32-bit address.
0	RO	1b	Core	<b>Memory/IO Space (MIOS)</b> hard wired to 1 to indicate IO space.

### 1.16.13 SVID2 - Subsystem Vendor Identification

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 2C-2Dh  
 Default Value: 0000h  
 Access: RW-O  
 Size: 16 bits

Bit	Access	Default Value	RST/PWR	Description
15:0	RW-O	0000h	Core	<b>Subsystem Vendor ID (SUBVID)</b> 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.

### 1.16.14 SID2 - Subsystem Identification

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 2E-2Fh  
 Default Value: 0000h  
 Access: RW-O  
 Size: 16 bits

Bit	Access	Default Value	RST/PWR	Description
15:0	RW-O	0000h	Core	<b>Subsystem Identification (SUBID)</b> 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.



### 1.16.15 ROMADR - Video BIOS ROM Base Address

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 30-33h  
 Default Value: 00000000h  
 Access: RO  
 Size: 32 bits

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

Bit	Access	Default Value	RST/ PWR	Description
31:18	RO	0000h	Core	<b>ROM Base Address (RBA)</b> hard wired to 0's.
17:11	RO	00h	Core	<b>Address Mask (ADMSK)</b> hard wired to 0's to indicate 256-KB address range.
10:1	RO	000h	Core	<i>Reserved</i> Hardwired to 0's.
0	RO	0b	Core	<b>ROM BIOS Enable (RBE)</b> 0 = ROM not accessible.

### 1.16.16 CAPPOINT - Capabilities Pointer

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 34h  
 Default Value: 90h  
 Access: RO  
 Size: 8 bits

Bit	Access	Default Value	RST/ PWR	Description
7:0	RO	90h	Core	<b>Capabilities Pointer Value (CPV)</b> 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].



### 1.16.17 INTRLINE - Interrupt Line

B/D/F/Type: 0/2/0/PCI  
Address Offset: 3Ch  
Default Value: 00h  
Access: RW  
Size: 8 bits

Bit	Access	Default Value	Description
7:0	RW	00h	<b>Interrupt Connection (INTCON)</b> 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.

### 1.16.18 INTRPIN - Interrupt Pin

B/D/F/Type: 0/2/0/PCI  
Address Offset: 3Dh  
Default Value: 01h  
Access: RO  
Size: 8 bits

Bit	Access	Default Value	Description
7:0	RO	01h	<b>Interrupt Pin (INTRPIN)</b> As a single function device, the IGD specifies INTA# as its interrupt pin. 01h: INTA#.

### 1.16.19 MINGNT - Minimum Grant

B/D/F/Type: 0/2/0/PCI  
Address Offset: 3Eh  
Default Value: 00h  
Access: RO  
Size: 8 bits

Bit	Access	Default Value	Description
7:0	RO	00h	<b>Minimum Grant Value (MGV)</b> The IGD does not burst as a PCI-compliant master.



### 1.16.20 MAXLAT - Maximum Latency

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 3Fh  
 Default Value: 00h  
 Access: RO  
 Size: 8 bits

Bit	Access	Default Value	Description
7:0	RO	00h	<b>Maximum Latency Value (MLV)</b> The IGD has no specific requirements for how often it needs to access the PCI bus.

### 1.16.21 GGCTL - Graphics Enhanced Intel® SpeedStep Technology Capability

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 4C-4Fh  
 Default Value: 003F003Fh  
 Access: RO  
 Size: 32 bits

This register conveys the maximum frequencies and voltages supported by this particular component, as determined by manufacturing testing. This information is intended to be used by software to properly configure the graphics clocking solution. Hardware does not directly use this information or enforce these limits in any way.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31	RO	0b	<i>Reserved</i>
30:24	RO	00h	<b>Max Turbo Graphics Render Voltage ID (MTGRVID)</b> The voltage required to operate at the Maximum Turbo Graphics Render Frequency. The encoding of this field follows the definition of the Intel® MVP 6.5 Voltage Regulator, which is: Voltage = 1.5 V - <field> * 0.0125 V (min of 0 V) Example: 0100000b = 1.1 V
23:22	RO	00b	<i>Reserved</i>
21:16	RO	3Fh	<b>Max Turbo Graphics Render Frequency (MTGRF)</b> The maximum frequency at which the graphics render clock may operate under "Turbo" conditions, which may exceed specified TDP budget of the processor. This field is a binary encoding of the maximum frequency divided by 33.333 MHz. Max Frequency = <field> * 33.333 MHz Example: 011000b = 800 MHz
15	RO	0b	<i>Reserved</i>



(Sheet 2 of 2)

Bit	Access	Default Value	Description
14:8	RO	00h	<b>Max Graphics Render Voltage ID (MGRVID)</b> The voltage required to operate at the Maximum Graphics Render Frequency. The encoding of this field follows the definition of the Intel MVP 6.5 Voltage Regulator, which is: Voltage = 1.5 V - <field> * 0.0125 V (min of 0 V) Example: 0100000b = 1.1 V
7:6	RO	00b	<i>Reserved</i>
5:0	RO	3Fh	<b>Max Graphics Render Frequency (MGRF)</b> The maximum frequency at which the graphics render clock may operate under normal conditions. This field is a binary encoding of the maximum frequency divided by 33.333 MHz. Max Frequency = <field> * 33.333 MHz Example: 011000b = 800 MHz



### 1.16.22 MGGC - Processor Graphics Control Register

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 52-53h  
 Default Value: 0030h  
 Access: RO  
 Size: 16 bits

All the Bits in this register are Intel TXT lockable.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
15:12	RO	0h	<i>Reserved</i>
11:8	RO	0h	<p><b>GTT Graphics Memory Size (GGMS)</b></p> <p>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.</p> <p>GSM is assumed to be a contiguous physical DRAM space with DSM, and BIOS needs to allocate a contiguous memory chunk. Hardware will drive the base of GSM from DSM only using the GSM size programmed in the register.</p> <p>0h: No memory pre-allocated. GTT cycles (Mem and IO) are not claimed.</p> <p>1h: No Intel® VT-d mode, 1 MB of memory pre-allocated for GTT.</p> <p>3h: No Intel VT-d mode, 2 MB of memory pre-allocated for GTT.</p> <p>9h: Intel VT-d mode, 2 MB of memory pre-allocated for 1 MB of Global GTT and 1 MB for Shadow GTT.</p> <p>Ah: Intel VT-d mode, 3 MB of memory pre-allocated for 1.5 MB of Global GTT and 1.5 MB for Shadow GTT.</p> <p>Bh: Intel VT-d mode, 4 MB of memory pre-allocated for 2 MB of Global GTT and 2 MB for Shadow GTT.</p> <p>All unspecified encodings of this register field are reserved, hardware functionality is not guaranteed if used. This register is locked and becomes Read Only when the D_LCK bit in the SMRAM register is set.</p>



(Sheet 2 of 2)

Bit	Access	Default Value	Description
7:4	RO	0011b	<p><b>Graphics Mode Select (GMS)</b></p> <p>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.</p> <p>0000: No memory pre-allocated. 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.</p> <p>Defined values: DVMT (UMA) mode, memory pre-allocated for frame buffer, in quantities as shown in the Encoding table. Values larger than 128 MB include 128 MB for Protected Content Memory.</p> <p>Encoding Description</p> <p>0h No memory pre-allocate</p> <p>5h 32 MB</p> <p>7h 64 MB</p> <p>8h 128 MB</p> <p>Bh 160 MB</p> <p>Ch 224 MB</p> <p>Dh 352 MB</p> <p><b>NOTE:</b> This register is locked and becomes Read Only when the D_LCK bit in the SMRAM register is set.</p> <p><b>BIOS Requirement:</b> BIOS must not set this field to 000 if IVD (Bit 1 of this register) is 0.</p>
3:2	RO	00b	<i>Reserved</i>
1	RO	0b	<p><b>IGD VGA Disable (IVD)</b></p> <p>Enable. Device 2 (IGD) claims VGA memory and IO cycles, the Sub-Class Code within Device 2 Class Code register is 00.</p> <p>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.</p> <p>BIOS Requirement: BIOS must not set this bit to 0 if the GMS field (Bits 6: 4 of this register) pre-allocates no memory. This bit <b>MUST</b> be set to 1 if Device 2 is disabled via register (DEVEN[3] = 0).</p>
0	RO	0b	<i>Reserved</i>





### 1.16.23 DEVEN - Device Enable

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 54-57h  
 Default Value: 0000210Bh  
 Access: RO  
 Size: 32 bits

Allows for enabling/disabling of PCI devices and functions that are within the processor. 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.

Bit	Access	Default Value	Description
31:15	RO	00000h	<i>Reserved</i>
14	RO	0b	<i>Reserved</i>
13	RO	0b	<i>Reserved</i>
12	RO	0b	<i>Reserved</i>
11	RO	0b	<i>Reserved</i>
10	RO	0b	<i>Reserved</i>
9	RO	0b	<i>Reserved</i>
8	RO	1b	<i>Reserved</i>
7:4	RO	0h	<i>Reserved</i>
3	RO	1b	<b>Internal Graphics Engine Function 0 (D2FOEN)</b> 0 = Bus 0 Device 2 Function 0 is disabled and hidden 1 = Bus 0 Device 2 Function 0 is enabled and visible If this processor does not have internal graphics capability then Device 2 Function 0 is disabled and hidden independent of the state of this bit.
2	RO	0b	<i>Reserved</i>
1	RO	1b	<b>PCI Express Port (D1EN)</b> 0 = Bus 0 Device 1 Function 0 is disabled and hidden. 1 = Bus 0 Device 1 Function 0 is enabled and visible. Default value is determined by the device capabilities, SDVO Presence HW strap and the sDVO/PCIe Concurrent HW strap. Device 1 is Disabled on Reset if the SDVO Presence strap was sampled high, and the sDVO/PCIe Concurrent strap was sampled low at the last assertion of PWROK, and is enabled by default otherwise.
0	RO	1b	<b>Host Bridge (D0EN)</b> Bus 0 Device 0 Function 0 may not be disabled and is therefore hard wired to 1.



### 1.16.24 SSRW - Software Scratch Read Write

B/D/F/Type: 0/2/0/PCI  
Address Offset: 58-5Bh  
Default Value: 00000000h  
Access: RW  
Size: 32 bits

Bit	Access	Default Value	RST/ PWR	Description
31:0	RW	00000000h	FLR, Core	<i>Reserved</i>

### 1.16.25 BSM - Base of Stolen Memory

B/D/F/Type: 0/2/0/PCI  
Address Offset: 5C-5Fh  
Default Value: 00000000h  
Access: RO  
Size: 32 bits

Graphics Stolen Memory and TSEG are within DRAM space defined under TOLUD. From the top of low used DRAM, processor claims 1 to 64 MBs of DRAM for internal graphics if enabled.

The base of stolen memory will always be below 4G. This is required to prevent aliasing between stolen range and the reclaim region.

Bit	Access	Default Value	Description
31:20	RO	000h	<b>Base of Stolen Memory (BSM)</b> This register contains Bits 31:20 of the base address of stolen DRAM memory. The host interface determines the base of Graphics Stolen memory by subtracting the graphics stolen memory size from TOLUD. See Device 0 TOLUD for more explanation.
19:0	RO	00000h	<i>Reserved</i>

### 1.16.26 HSRW - Hardware Scratch Read Write

B/D/F/Type: 0/2/0/PCI  
Address Offset: 60-61h  
Default Value: 0000h  
Access: RW  
Size: 16 bits

Bit	Access	Default Value	RST/ PWR	Description
15:0	RW	0000h	FLR, Core	<i>Reserved</i>



### 1.16.27 MSAC - Multi Size Aperture Control

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 62h  
 Default Value: 02h  
 Access: RO; RW; RW-K;  
 Size: 8 bits

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.

The size of the aperture must not be modified by software after its location is written into GMADR (offset 18h).

Bit	Access	Default Value	Description
7:4	RW	0h	<i>Reserved</i> Scratch Bits Only -- Have no physical effect on hardware
3	RO	0b	<b>Reserved</b>
2:1	RW-K	01b	<b>Untrusted Aperture Size (LHSAS)</b> 11: bits [28:27] of GMADR register are made Read only and forced to zero, allowing only 512MB of GMADR 01: bit [28] of GMADR is made R/W and bit [27] of GMADR is forced to zero allowing 256MB of GMADR 00: bits [28:27] of GMADR register are made R/W allowing 128MB of GMADR 10: Illegal programming. These bits are read-only in Intel TXT mode.
0	RO	0h	<i>Reserved</i>

### 1.16.28 MC - Message Control

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 92-93h  
 Default Value: 0000h  
 Access: RO; RW  
 Size: 16 bits

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.



Bit	Access	Default Value	RST/ PWR	Description
15:8	RO	00h	Core	<i>Reserved</i>
7	RO	0b	Core	<p><b>64 Bit Capable (64BCAP)</b>  hard wired 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/4 GB limit.</p>
6:4	RW	000b	FLR, Core	<p><b>Multiple Message Enable (MME)</b>  System software programs this field to indicate the actual number of messages allocated to this device. This number is equal to or less than the number actually requested.  The encoding is the same as for the MMC field below.</p>
3:1	RO	000b	Core	<p><b>Multiple Message Capable (MMC)</b>  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</p>
0	RW	0b	FLR, Core	<p><b>MSI Enable (MSIEN)</b>  Controls the ability of this device to generate MSIs.</p>



### 1.16.29 MA - Message Address

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 94-97h  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits

Bit	Access	Default Value	RST/ PWR	Description
31:2	RW	00000000h	FLR, Core	<b>Message Address (MESSADD)</b> 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.
1:0	RO	00b	Core	<b>Force Dword Align (FDWORD)</b> hard wired to 0 so that addresses assigned by system software are always aligned on a DWORD address boundary.

### 1.16.30 MD - Message Data

B/D/F/Type: 0/2/0/PCI  
 Address Offset: 98-99h  
 Default Value: 0000h  
 Access: RW  
 Size: 16 bits

Bit	Access	Default Value	RST/ PWR	Description
15:0	RW	0000h	FLR, Core	<b>Message Data (MESSDATA)</b> 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.



### 1.16.31 PMCAPID - Power Management Capabilities ID

B/D/F/Type: 0/2/0/PCI  
Address Offset: D0-D1h  
Default Value: A401h  
Access: RO  
Size: 16 bits

Bit	Access	Default Value	Description
15:8	RO	A4h	<b>Next Capability Pointer (NEXT_PTR)</b> This contains a pointer to the next item in the capabilities list.
7:0	RO	01h	<b>Capability Identifier (CAP_ID)</b> SIG defines this ID is 01h for power management.

### 1.16.32 PMCAP - Power Management Capabilities

B/D/F/Type: 0/2/0/PCI  
Address Offset: D2-D3h  
Default Value: 0022h  
Access: RO  
Size: 16 bits

(Sheet 1 of 2)

Bit	Access	Default Value	Description
15:11	RO	00h	<b>PME Support (PMES)</b> This field indicates the power states in which the IGD may assert PME#. Hard wired to 0 to indicate that the IGD does not assert the PME# signal.
10	RO	0b	<b>D2 Support (D2)</b> The D2 power management state is not supported. This bit is hard wired to 0.
9	RO	0b	<b>D1 Support (D1)</b> hard wired to 0 to indicate that the D1 power management state is not supported.
8:6	RO	000b	<i>Reserved</i>
5	RO	1b	<b>Device Specific Initialization (DSI)</b> hard wired to 1 to indicate that special initialization of the IGD is required before generic class device driver is to use it.
4	RO	0b	<i>Reserved</i>
3	RO	0b	<b>PME Clock (PMECLK)</b> hard wired to 0 to indicate IGD does not support PME# generation.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
2:0	RO	010b	<b>Version (VER)</b> hard wired 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.

### 1.16.33 PMCS - Power Management Control/Status

B/D/F/Type: 0/2/0/PCI  
 Address Offset: D4-D5h  
 Default Value: 0000h  
 Access: RO; RW  
 Size: 16 bits

Bit	Access	Default Value	RST/ PWR	Description
15	RO	0b	Core	<b>PME Status (PMESTS)</b> This bit is 0 to indicate that IGD does not support PME# generation from D3 (cold).
14:13	RO	00b	Core	<b>Data Scale (DSCALE)</b> The IGD does not support data register. This bit always returns 00 when read, write operations have no effect.
12:9	RO	0h	Core	<b>Data Select (DSEL)</b> The IGD does not support data register. This bit always returns 0h when read, write operations have no effect.
8	RO	0b	Core	<b>PME Enable (PME_EN)</b> This bit is 0 to indicate that PME# assertion from D3 (cold) is disabled.
7:2	RO	00h	Core	<i>Reserved</i>



Bit	Access	Default Value	RST/ PWR	Description
1:0	RW	00b	FLR, Core	<p><b>Power State (PWRSTAT)</b></p> <p>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. Behavior of the graphics controller in supported states is detailed in the power management section of the Bspec.</p> <p>Bits[1:0] Power state            00: D0 Default            01: D1 Not Supported            10: D2 Not Supported            11: D3</p>

### 1.16.34 SWSMI - Software SMI

B/D/F/Type: 0/2/0/PCI  
 Address Offset: E0-E1h  
 Default Value: 0000h  
 Access: RW  
 Size: 16 bits

As long as there is the potential that DVO port legacy drivers exist which expect this register at this address, Dev2F0address E0h-E1h must be reserved for this register.

Bit	Access	Default Value	Description
15:8	RW	00h	<b>Software Scratch Bits (SWSB)</b>
7:1	RW	00h	<p><b>Software Flag (SWF)</b></p> <p>Used to indicate caller and SMI function desired, as well as return result.</p>
0	RW	0b	<p><b>Processor Software SMI Event (GSSMIE)</b></p> <p>When Set this bit will trigger an SMI. Software must write a 0 to clear this bit.</p>





### 1.16.35 GSE - Graphics System Event Register

B/D/F/Type: 0/2/0/PCI  
 Address Offset: E4-E7h  
 Default Value: 00000000h  
 Access: RW  
 Size: 32 bits

Bit	Access	Default Value	Description
31:24	RW	00h	<b>GSE Scratch Trigger 3 (AST3)</b> When written, this scratch byte triggers an interrupt when IER Bit 0 is enabled and IMR bit 0 is unmasked. If written as part of a 16-bit or 32-bit write, only one interrupt is generated in common.
23:16	RW	00h	<b>GSE Scratch Trigger 2 (AST2)</b> When written, this scratch byte triggers an interrupt when IER Bit 0 is enabled and IMR bit 0 is unmasked. If written as part of a 16-bit or 32-bit write, only one interrupt is generated in common.
15:8	RW	00h	<b>GSE Scratch Trigger 1 (AST1)</b> When written, this scratch byte triggers an interrupt when IER Bit 0 is enabled and IMR bit 0 is unmasked. If written as part of a 16-bit or 32-bit write, only one interrupt is generated in common.
7:0	RW	00h	<b>GSE Scratch Trigger 0 (AST0)</b> When written, this scratch byte triggers an interrupt when IER Bit 0 is enabled and IMR bit 0 is unmasked. If written as part of a 16-bit or 32-bit write, only one interrupt is generated in common.



### 1.16.36 ASLS - ASL Storage

B/D/F/Type: 0/2/0/PCI  
 Address Offset: FC-FFh  
 Default Value: 00000000h  
 Access: RW  
 Size: 32 bits

This software scratch register only needs to be read/write accessible. The exact bit register usage must be worked out in common between System BIOS and driver software, but storage for switching/indicating up to six devices is possible with this amount.

For each device, the ASL control method will require two bits for \_DOD (BIOS detectable yes or no, VGA/NonVGA), one bit for \_DGS (enable/disable requested), and two bits for \_DCS (enabled now/disabled now, connected or not).

Bit	Access	Default Value	Description
31:0	RW	00000000h	<b>Device Switching Storage (DSS)</b> Software controlled usage to support device switching.

### 1.17 Device 2 IO

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
MMIO Address Register	Index	0	3	00000000h	RW
MMIO Data Register	Data	4	7	00000000h	RW



### 1.17.1 Index - MMIO Address Register

B/D/F/Type: 0/2/0/PCI IO  
 Address Offset: 0-3h  
 Default Value: 00000000h  
 Access: RW  
 Size: 32 bits

MMIO\_INDEX: A 32 bit IO write to this port loads the offset of the MMIO register or offset into the GTT that needs to be accessed. An IO Read returns the current value of this register. An 8-/16-bit IO write to this register is completed by the processor but does not update this register.

This mechanism to access internal graphics MMIO registers must not be used to access VGA IO registers which are mapped through the MMIO space. VGA registers must be accessed directly through the dedicated VGA IO ports.

Bit	Access	Default Value	Description
31:2	RW	00000000h	<b>Register/GTT Offset (REGGTO)</b> This field selects any one of the DWORD registers within the MMIO register space of Device 2 if the target is MMIO Registers. This field selects a GTT offset if the target is the GTT.
1:0	RW	00b	<b>Target (TARG)</b> 00: MMIO Registers 01: GTT 1X: Reserved

### 1.17.2 Data - MMIO Data Register

B/D/F/Type: 0/2/0/PCI IO  
 Address Offset: 4-7h  
 Default Value: 00000000h  
 Access: RW  
 Size: 32 bits

MMIO\_DATA: A 32-bit IO write to this port is re-directed to the MMIO register/GTT location pointed to by the MMIO-index register. A 32-bit IO read to this port is re-directed to the MMIO register address pointed to by the MMIO-index register regardless of the target selection in MMIO\_INDEX(1:0). 8- or 16-bit IO writes are completed by the processor and may have un-intended side effects, hence must not be used to access the data port. 8- or 16-bit IO reads are completed normally.

**Note:** If the target field in MMIO Index selects "GTT", reads to MMIO data return is undefined.

Bit	Access	Default Value	Description
31:0	RW	00000000h	<b>MMIO Data Window (DATA)</b>



## 1.18 DMI and PEG VC0/VCp Remap Registers

(Sheet 1 of 3)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Version Register	VER_REG	0	3	00000010h	RO
Capability Register	CAP_REG	8	F	00C9008020630272h	RO
Extended Capability Register	ECAP_REG	10	17	0000000000001000h	RO
Global Command Register	GCMD_REG	18	1B	00000000h	W; WO; RO
Global Status Register	GSTS_REG	1C	1F	00000000h	RO
Root-Entry Table Address Register	RTADDR_REG	20	27	0000000000000000h	RO; RW
Context Command Register	CCMD_REG	28	2F	0000000000000000h	W; RW; RO
Fault Status Register	FSTS_REG	34	37	00000000h	RO; RO-P; RWC-P
Fault Event Control Register	FECTL_REG	38	3B	80000000h	RO; RW
Fault Event Data Register	FEDATA_REG	3C	3F	00000000h	RO; RW
Fault Event Address Register	FEADDR_REG	40	43	00000000h	RO; RW
Fault Event Upper Address Register	FEUADDR_REG	44	47	00000000h	RO
Advanced Fault Log Register	AFLOG_REG	58	5F	0000000000000000h	RO
Protected Memory Enable Register	PMEM_REG	64	67	00000000h	RO; RW
Protected Low-Memory Base Register	PLMBASE_REG	68	6B	00000000h	RO; RW
Protected Low-Memory Limit Register	PLMLIMIT_REG	6C	6F	00000000h	RO; RW
Protected High-Memory Base Register	PHMBASE_REG	70	77	0000000000000000h	RO; RW



(Sheet 2 of 3)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Protected High-Memory Limit Register	PHMLIMIT_REG	78	7F	000000000000 0000h	RO; RW
Invalidation Queue Head Register	IQH_REG	80	87	000000000000 0000h	RO
Invalidation Queue Tail Register	IQT_REG	88	8F	000000000000 0000h	RO
Invalidation Queue Address Register	IOA_REG	90	97	000000000000 0000h	RO; RW
Invalidation Completion Status Register	ICS_REG	9C	9F	00000000h	RO; RWC-P
Invalidation Event Control Register	IECTL_REG	A0	A3	00000000h	RO; RW
Invalidation Event Data Register	IEDATA_REG	A4	A7	00000000h	RW
Invalidation Event Address Register	IEADDR_REG	A8	AB	00000000h	RO; RW
Invalidation Event Upper Address Register	IEUADDR_REG	AC	AF	00000000h	RW
Interrupt Remapping Table Address Register	IRTA_REG	B8	BF	000000000000 0000h	RO; RW
Invalidate Address Register	IVA_REG	100	107	000000000000 0000h	RO; W
IOTLB Invalidate Register	IOTLB_REG	108	10F	000000000000 0000h	RO; RW
Fault Recording Registers	FRCD_REG	200	20F	000000000000 000000000000 00000000h	RWC-P; RO-P; RO
VT-d Completion Resource Dedication	VTCMPLRESR	F00	F03	00060000h	RW-L; RO
VCO/VCp Intel VT-d Fetch Arbiter Control	VTFTCHARBCTL	F04	F07	0000FFFFh	RW-L



(Sheet 3 of 3)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
PEG VT-d Completion Resource Dedication	PEGVTCMPLRESR	F08	F0B	20004000h	RW-L; RO
DMA Remap Engine Policy Control	VTPOLICY	FFC	FFF	00000000h	RW-L

### 1.18.1 VER\_REG - Version Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 0-3h  
 Default Value: 00000010h  
 Access: RO  
 Size: 32 bits

Register to report the architecture version supported. Backward compatibility for the architecture is maintained with new revision numbers, allowing software to load DMA-remapping drivers written for prior architecture versions.

Bit	Access	Default Value	Description
31:8	RO	00000000 00000000 00000000 b	<i>Reserved</i>
7:4	RO	0001b	<b>Major Version number (MAX)</b> Indicates supported architecture version.
3:0	RO	0000b	<b>Minor Version number (MIN)</b> Indicates supported architecture minor version.



## 1.18.2 CAP\_REG - Capability Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 8-Fh  
 Default Value: 00C9008020630272h  
 Access: RO  
 Size: 64 bits

Register to report general DMA remapping hardware capabilities.

(Sheet 1 of 4)

Bit	Access	Default Value	Description
63:56	RO	00h	<i>Reserved</i>
55	RO	1b	<b>DMA Read Draining (DRD)</b> 0 = On IOTLB invalidations, hardware does not support draining of translated DMA read requests queued within the root complex. 1 = On IOTLB invalidations, hardware supports draining of translated DMA read requests queued within the root complex. Indicates supported architecture version.
54	RO	1b	<b>DMA Write Draining (DWD)</b> 0 = On IOTLB invalidations, hardware does not support draining of translated DMA writes queued within the root complex. 1 = On IOTLB invalidations, hardware supports draining of translated DMA writes queued within the root complex.
53:48	RO	001001b	<b>Maximum Address Mask Value (MAMV)</b> The value in this field indicates the maximum supported value for the Address Mask (AM) field in the Invalidation Address (IVA_REG) register.
47:40	RO	0000000 0b	<b>Number of Fault-recording Registers (NFR)</b> This field indicates a value of N-1, where N is the number of fault recording registers supported by hardware. Implementations must support at least one fault recording register (NFR = 0) for each DMA-remapping hardware unit in the platform. The maximum number of fault recording registers per DMA-remapping hardware unit is 256. Bit 40 in the capability register is the least significant bit of the NFR field (47:40).
39	RO	1b	<b>Page Selective Invalidation Support (PSI)</b> 0 = Indicates that the DMAR engine does not support page selective invalidations. 1 = Indicates the DMAR engine does support page-selective IOTLB invalidations. The MAMV field indicates the maximum number of contiguous translations that may be invalidated in a single request.
38	RO	0b	<i>Reserved</i>



(Sheet 2 of 4)

Bit	Access	Default Value	Description
37:34	RO	0000b	<p><b>Super Page Support (SPS)</b></p> <p>This field indicates the super page sizes supported by hardware.</p> <p>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:</p> <ul style="list-style-type: none"> <li>0: 21-bit offset to page frame</li> <li>1: 30-bit offset to page frame</li> <li>2: 39-bit offset to page frame</li> <li>3: 48-bit offset to page frame</li> </ul>
33:24	RO	020h	<p><b>Fault-recording Register Offset (FRO)</b></p> <p>This field specifies the location to the first fault recording register relative to the register base address of this DMA-remapping hardware unit.</p> <p>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 <math>X + (16 * Y)</math>.</p>
23	RO	0b	<p><b>Isochrony (Isoch)</b></p> <ul style="list-style-type: none"> <li>0 = Indicates this DMA-remapping hardware unit has no critical isochronous requesters in its scope.</li> <li>1 = Indicates this DMA-remapping hardware unit has one or more critical isochronous requesters in its scope. To guarantee isochronous performance, software must ensure invalidation operations do not impact active DMA streams. This implies that when DMA is active, software perform page-selective invalidations (instead of coarser invalidations).</li> </ul>
22	RO	1b	<p><b>Zero Length Read (ZLR)</b></p> <ul style="list-style-type: none"> <li>0 = Indicates the remapping hardware unit blocks (and treats as fault) zero length DMA read requests to write-only pages.</li> <li>1 = Indicates the remapping hardware unit supports zero length DMA read requests to write-only pages.</li> </ul>
21:16	RO	100011b	<p><b>Maximum Guest Address Width (MGAW)</b></p> <p>This field indicates the maximum DMA virtual addressability supported by remapping hardware. The Maximum Guest Address Width (MGAW) is computed as <math>(N + 1)</math>, 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.</p> <p>If the value in this field is X, DMA requests to addresses above <math>2^{(X+1)} - 1</math> are always blocked by hardware. 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).</p>
15:13	RO	000b	<i>Reserved</i>





(Sheet 3 of 4)

Bit	Access	Default Value	Description
12:8	RO	00010b	<p><b>Supported Adjusted Guest Address Widths (SAGAW)</b></p> <p>This 5-bit field indicates the supported adjusted guest address widths (which in turn represents the levels of page-table walks) 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:</p> <ul style="list-style-type: none"> <li>0: 30-bit AGAW (2-level page table)</li> <li>1: 39-bit AGAW (3-level page table)</li> <li>2: 48-bit AGAW (4-level page table)</li> <li>3: 57-bit AGAW (5-level page table)</li> <li>4: 64-bit AGAW (6-level page table)</li> </ul> <p>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.</p>
7	RO	0b	<p><b>Caching Mode (CM)</b></p> <p>0 = Hardware does not cache not present and erroneous entries in the context-cache and IOTLB. 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.</p> <p>1 = Hardware may cache not present and erroneous mappings in the context-cache or IOTLB. Any software updates to the DMA-remapping structures (including updates to not-present or erroneous entries) require explicit invalidation.</p> <p>Refer to Section 8.4.9 for more details on caching mode. Hardware implementations are recommended to support operation corresponding to CM=0.</p>
6	RO	1b	<p><b>Protected High-Memory Region (PHMR)</b></p> <p>0 = Indicates protected high-memory region not supported.                      1 = Indicates protected high-memory region is supported.</p> <p>DMA-remapping hardware implementations on Intel VT-d platforms supporting main memory above 4 GB are required to support protected high-memory region.</p>
5	RO	1b	<p><b>Protected Low-Memory Region (PLMR)</b></p> <p>0 = Indicates protected low-memory region not supported.                      1 = Indicates protected low-memory region is supported.</p> <p>DMA-remapping hardware implementations on Intel TXT platforms are required to support protected low-memory region.</p>



(Sheet 4 of 4)

Bit	Access	Default Value	Description
4	RO	1b	<b>Required Write-Buffer Flushing (RWBF)</b> 0 = Indicates no write-buffer flushing needed to ensure changes to memory-resident structures are visible to hardware. 1 = Indicates software must explicitly flush the write buffers (through the Global Command register) to ensure updates made to memory-resident DMA-remapping structures are visible to hardware. Refer Section 9.1 for more details on write buffer flushing requirements.
3	RO	0b	<b>Advanced Fault Logging (AFL)</b> 0 = Indicates advanced fault logging not supported. Only primary fault logging is supported. 1 = Indicates advanced fault logging is supported.
2:0	RO	010b	<b>Number of Domains Supported (ND)</b> 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 4-K domains. 101b: Hardware supports 14-bit domain-ids with support for up to 16-K domains. 110b: Hardware supports 16-bit domain-ids with support for up to 64-K domains. 111b: Reserved.

### 1.18.3 ECAP\_REG - Extended Capability Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 10-17h  
 Default Value: 0000000000001000h  
 Access: RO  
 Size: 64 bits

Register to report DMA-remapping hardware extended capabilities.

(Sheet 1 of 3)

Bit	Access	Default Value	Description
63:32	RO	00000000h	Reserved



(Sheet 2 of 3)

Bit	Access	Default Value	Description
31:24	RO	00h	<p><b>Number of IOTLB Invalidation Units (NIU)</b></p> <p>This field indicates a value of N-1, where N is the number of IOTLB invalidation units supported by hardware. Each IOTLB invalidation unit consists of two registers: A 64-bit IOTLB Invalidation Register (IOTLB_REG), followed by a 64-bit Invalidation Address Register (IVA_REG).</p> <p>Implementations must support at least one IOTLB invalidation unit (NIVU = 0) for each DMA-remapping hardware unit in the platform.</p> <p>The maximum number of IOTLB invalidation register units per DMA-remapping hardware unit is 256.</p>
23:20	RO	0000b	<p><b>Maximum Handle Mask Value (MHMV)</b></p> <p>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).</p> <p>This field is valid only when the IR field is reported as Set.</p>
19:18	RO	00b	<i>Reserved</i>
17:8	RO	010h	<p><b>Invalidation Unit Offset (IVO)</b></p> <p>This field specifies the location to the first IOTLB invalidation unit relative to the register base address of this DMA-remapping hardware unit.</p> <p>If the register base address is X, and the value reported in this field is Y, the address for the first IOTLB invalidation unit is calculated as <math>+(16*Y)</math>.</p> <p>If N is the value reported in NIU field, the address for the last IOTLB invalidation unit is calculated as <math>X+(16*Y)+(16*N)</math>.</p>
7	RO	0b	<p><b>Snoop Control (SC)</b></p> <p>0 = Hardware does not support 1-setting of the SNP field in the page-table entries.</p> <p>1 = Hardware supports the 1-setting of the SNP field in the page-table entries.</p>
6	RO	0b	<p><b>Pass Through (PT)</b></p> <p>0 = Hardware does not support pass-through translation type in context entries.</p> <p>1 = Hardware supports pass-through translation type in context entries.</p>
5	RO	0b	<p><b>Caching Hints (CH)</b></p> <p>0 = Hardware does not support IOTLB caching hints (ALH and EH fields in context-entries are treated as reserved).</p> <p>1 = Hardware supports IOLTb caching hints through the ALH and EH fields in context-entries.</p>



(Sheet 3 of 3)

Bit	Access	Default Value	Description
4	RO	0b	<p><b>Extended Interrupt Mode (EIM)</b></p> <p>0 = Hardware supports only 8-bit APICIDs (Legacy Interrupt Mode) on Intel® 64 and IA-32 platforms and 16-bit APIC-IDs on Itanium® platforms.</p> <p>1 = Hardware supports Extended Interrupt Mode (32-bit APIC-IDs) on Intel® 64 platforms. This field is valid only when the IR field is reported as Set.</p>
3	RO	0b	<p><b>Interrupt Remapping Support (IR)</b></p> <p>0 = Hardware does not support interrupt remapping.</p> <p>1 = Hardware supports interrupt remapping.</p> <p>Implementations reporting this field as Set must also support Queued Invalidation (QI = 1b).</p>
2	RO	0b	<p><b>Device IOTLB Support (DI)</b></p> <p>0 = Hardware does not support device- IOTLBs.</p> <p>1 = Hardware supports Device-IOTLBs.</p> <p>Implementations reporting this field as Set must also support Queued Invalidation (QI = 1b).</p>
1	RO	0b	<p><b>Queued Invalidation Support (QI)</b></p> <p>0 = Hardware does not support queued invalidations.</p> <p>1 = Hardware supports queued invalidations.</p>
0	RO	0b	<p><b>Coherency (C)</b></p> <p>0 = Indicates that hardware accesses to the root, context, and page table structures are non-coherent (non-snoop).</p> <p>1 = Indicates that hardware accesses to the root, context, and page table structures are coherent (snoop).</p> <p>Hardware writes to the advanced fault log is required to be coherent.</p>



### 1.18.4 GCMD\_REG - Global Command Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 18-1Bh  
 Default Value: 00000000h  
 Access: W; WO; RO  
 Size: 32 bits

Register to control DMA-remapping hardware. If multiple control fields in this register need to be modified, software must serialize through multiple writes to this register.

(Sheet 1 of 5)

Bit	Access	Default Value	Description
31	W	0b	<p><b>Translation Enable (TE)</b></p> <p>Software writes to this field to request hardware to enable/disable DMA-remapping hardware.</p> <p>0 = Disable DMA-remapping hardware                      1 = Enable DMA-remapping hardware</p> <p>Hardware reports the status of the translation enable operation through the TES field in the Global Status register. Before enabling (or re-enabling) DMA-remapping hardware through this field, software must:</p> <ul style="list-style-type: none"> <li>• Setup the DMA-remapping structures in memory.</li> <li>• Flush the write buffers (through WBF field), if write buffer flushing is reported as required.</li> <li>• Set the root-entry table pointer in hardware (through S RTP field).</li> <li>• Perform global invalidation of the context-cache and global invalidation of IOTLB</li> <li>• If advanced fault logging supported, setup fault log pointer (through SFL field) and enable advanced fault logging (through EAFL field).</li> </ul> <p>Refer to Section 9 for detailed software requirements. There may be active DMA requests in the platform when software updates this field. Hardware must enable or disable remapping logic only at deterministic transaction boundaries, so that any in-flight transaction is either subject to remapping or not at all.</p> <p>Hardware implementations supporting DMA draining must drain any in-flight translated 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 GSTS_REG.</p> <p>Value returned on read of this field is undefined.</p>



(Sheet 2 of 5)

Bit	Access	Default Value	Description
30	WO	0b	<p><b>Set Root Table Pointer (S RTP)</b></p> <p>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 register. Hardware reports the status of the root table pointer set operation through the RTPS field in the Global Status register. The root table pointer set operation must be performed before enabling or re-enabling (after disabling) DMA remapping hardware.</p> <p>After a root table pointer set operation, software must globally invalidate the context cache followed by global invalidate of IOTLB. This is required to ensure hardware uses only the remapping structures referenced by the new root table pointer, and not any stale cached entries.</p> <p>While DMA-remapping hardware is active, software may update the root table pointer through this field.</p> <p>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.</p> <p>Value returned on read of this field is undefined.</p>
29	W	0b	<p><b>Set Fault Log (SFL)</b></p> <p>This field is valid only for implementations supporting advanced fault logging. If advanced fault logging is not supported, writes to this field are ignored.</p> <p>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.</p> <p>Hardware reports the status of the fault log set 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 hardware is active.</p> <p>Clearing this bit has no effect.</p> <p>Value returned on read of this field is undefined.</p>



(Sheet 3 of 5)

Bit	Access	Default Value	Description
28	W	0b	<p><b>Enable Advanced Fault Logging (EAFL)</b></p> <p>This field is valid only for implementations supporting advanced fault logging. If advanced fault logging is not supported, writes to this field are ignored.</p> <p>Software writes to this field to request hardware to enable or disable advanced fault logging.</p> <p>0 = Disable advanced fault logging. In this case, translation faults are reported through the Fault-recording registers.</p> <p>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 SFL field) before enabling advanced fault logging.</p> <p>Hardware reports the status of the advanced fault logging enable operation through the AFLS field in the Global Status register.</p> <p>Value returned on read of this field is undefined.</p>
27	WO	0b	<p><b>Write Buffer Flush (WBF)</b></p> <p>This bit is valid only for implementations requiring write buffer flushing. If write buffer flushing is not required, writes to this field are ignored.</p> <p>Software sets this field to request hardware to flush the root-complex internal write buffers. This is done to ensure any updates to the memory-resident DMA-remapping structures are not held in any internal write posting buffers. Refer to Section 9.1 for details on write-buffer flushing requirements.</p> <p>Hardware reports the status of the write buffer flushing operation through the WBFS field in the Global Status register.</p> <p>Clearing this bit has no effect.</p> <p>Value returned on read of this field is undefined.</p>
26	W	0b	<p><b>Queued Invalidation Enable (QIE)</b></p> <p>This field is valid only for implementations supporting queued invalidations.</p> <p>Software writes to this field to enable or disable queued invalidations.</p> <p>0 = Disable queued invalidations.</p> <p>1 = Enable use of queued invalidations.</p> <p>Hardware reports the status of queued invalidation enable operation through QIES field in the Global Status register.</p> <p>Refer to Section 6.2.2 for software requirements for enabling/disabling queued invalidations.</p> <p>The value returned on a read of this field is undefined.</p>



(Sheet 4 of 5)

Bit	Access	Default Value	Description
25	W	0b	<p><b>Interrupt Remapping Enable (IRE)</b></p> <p>This field is valid only for implementations supporting interrupt remapping.</p> <p>0 = Disable interrupt-remapping hardware. 1 = Enable interrupt-remapping hardware.</p> <p>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.</p> <p>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.</p> <p>The value returned on a read of this field is undefined.</p>
24	W	0b	<p><b>Set Interrupt Remap Table Pointer (SIRTP)</b></p> <p>This field is valid only for implementations supporting interrupt-remapping.</p> <p>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 register.</p> <p>Hardware reports the status of the interrupt remapping table pointer set operation through the IRTPS field in the Global Status register. The interrupt remap table pointer set operation must be performed before enabling or re-enabling (after disabling) interrupt-remapping hardware through the IRE field.</p> <p>After an interrupt remap table pointer set 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.</p> <p>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.</p> <p>Clearing this bit has no effect. The value returned on a read of this field is undefined.</p>





(Sheet 5 of 5)

Bit	Access	Default Value	Description
23	W	0b	<p><b>Compatibility Format Interrupt (CFI)</b></p> <p>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 Legacy Interrupt Mode is active.</p> <p>0 = Block Compatibility format interrupts.                      1 = Process Compatibility format interrupts as pass-through (bypass interrupt remapping).</p> <p>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.</p> <p>This field is not implemented on Itanium® processor implementations.</p>
22:0	RO	000000h	<i>Reserved</i>

### 1.18.5 GSTS\_REG - Global Status Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 1C-1Fh  
 Default Value: 00000000h  
 Access: RO  
 Size: 32 bits

Register to report general DMA-remapping hardware status.

(Sheet 1 of 3)

Bit	Access	Default Value	Description
31	RO	0b	<p><b>Translation Enable Status (TES)</b></p> <p>This field indicates the status of DMA-remapping hardware.</p> <p>0 = DMA-remapping hardware is not enabled.                      1 = DMA-remapping hardware is enabled.</p>
30	RO	0b	<p><b>Root Table Pointer Status (RTPS)</b></p> <p>This field indicates the status of the root- table pointer in hardware.</p> <p>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.</p>



(Sheet 2 of 3)

Bit	Access	Default Value	Description
29	RO	0b	<p><b>Fault Log Status (FLS)</b></p> <p>This field is valid only for implementations supporting advanced fault logging.</p> <p>This field indicates the status of the fault-log pointer in hardware.</p> <p>This field is cleared by hardware when software sets the SFL field in the Global Command register.</p> <p>This field is set by hardware when hardware completes the set fault-log pointer operation using the value provided in the Advanced Fault Log register.</p>
28	RO	0b	<p><b>Advanced Fault Logging Status (AFLS)</b></p> <p>This field is valid only for implementations supporting advanced fault logging. This field indicates advanced fault logging status.</p> <p>0 = Advanced Fault Logging is not enabled. 1 = Advanced Fault Logging is enabled.</p>
27	RO	0b	<p><b>Write Buffer Flush Status (WBFS)</b></p> <p>This bit is valid only for implementations requiring write buffer flushing.</p> <p>This field indicates the status of the write buffer flush operation.</p> <p>This field is set by hardware when software sets the WBF field in the Global Command register.</p> <p>This field is cleared by hardware when hardware completes the write buffer flushing operation.</p>
26	RO	0b	<p><b>Queued Invalidation Enable Status (QIES)</b></p> <p>This field indicates queued invalidation enable status.</p> <p>0 = Queued invalidation is not enabled 1 = Queued invalidation is enabled</p>
25	RO	0b	<p><b>Interrupt Remapping Enable Status (IRES)</b></p> <p>This field indicates the status of Interrupt-remapping hardware.</p> <p>0 = Interrupt-remapping hardware is not enabled 1 = Interrupt-remapping hardware is enabled</p>
24	RO	0b	<p><b>Interrupt Remapping Table Pointer Status (IRTPS)</b></p> <p>This field indicates the status of the interrupt remapping table pointer in hardware.</p> <p>This field is cleared by hardware when software sets the SIRTP field in the Global Command register.</p> <p>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.</p>



(Sheet 3 of 3)

Bit	Access	Default Value	Description
23	RO	0b	<b>Compatibility Format Interrupt Status (CFIS)</b> 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 Legacy interrupt mode is active. 0 = Compatibility format interrupts are blocked. 1 = Compatibility format interrupts are processed as pass-through (bypassing interrupt remapping).
22:0	RO	000000h	<i>Reserved</i>

### 1.18.6 RTADDR\_REG - Root-Entry Table Address Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 20-27h  
 Default Value: 0000000000000000h  
 Access: RO; RW  
 Size: 64 bits  
 Register providing the base address of root-entry table.

Bit	Access	Default Value	Description
63:12	RW	00000000 000000h	<b>Root Table Address (RTA)</b> This register points to base of page aligned 4-KB-sized root-entry table in system memory. Hardware may ignore and not implement 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.
11:0	RO	000h	<i>Reserved</i>



### 1.18.7 CCMD\_REG - Context Command Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 28-2Fh  
 Default Value: 0000000000000000h  
 Access: W; RW; RO  
 Size: 64 bits

Register to manage context cache. The act of writing the uppermost byte of the CCMD\_REG with ICC field set causes the hardware to perform the context-cache invalidation.

(Sheet 1 of 3)

Bit	Access	Default Value	Description
63	RW	0h	<p><b>Invalidate Context-Cache (ICC)</b></p> <p>Software requests invalidation of context-cache by setting this field. Software must also set the requested invalidation granularity by programming the CIRG field.</p> <p>Software must read back and check the ICC field to be clear to confirm the invalidation is complete. Software must not update this register when this field is set.</p> <p>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 not submit another invalidation request through this register while the ICC field is set.</p> <p>Software must submit a context cache invalidation request through this field only when there are no invalidation requests pending at this DMA-remapping hardware unit. Refer to Section 9 for software programming requirements.</p> <p>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.</p> <p>Hardware implementations reporting write-buffer flushing requirement (RWBF=1 in Capability register) must implicitly perform a write buffer flushing before reporting invalidation complete to software through the ICC field. Refer to Section 9.1 for write buffer flushing requirements.</p>



(Sheet 2 of 3)

Bit	Access	Default Value	Description
62:61	RW	0h	<p><b>Context Invalidation Request Granularity (CIRG)</b></p> <p>Software provides the requested invalidation granularity through this field when setting the ICC field.</p> <p>Following are the encodings for the CIRG field:</p> <p>00: Reserved.</p> <p>01: Global Invalidation request.</p> <p>10: Domain-selective invalidation request. The target domain-id must be specified in the DID field.</p> <p>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.</p> <p>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.</p>
60:59	RO	0h	<p><b>Context Actual Invalidation Granularity (CAIG)</b></p> <p>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).</p> <p>The following are the encodings for the CAIG field:</p> <p>00: Reserved.</p> <p>01: Global Invalidation performed. This could be in response to a global, domain-selective or device-selective invalidation request.</p> <p>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.</p> <p>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.</p>
58:34	RO	00000000h	<i>Reserved</i>



(Sheet 3 of 3)

Bit	Access	Default Value	Description
33:32	W	0h	<b>Function Mask (FM)</b> 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 device(s) specified through the FM and SID fields must correspond to the domain-id specified in the DID field. Value returned on read of this field is undefined.
31:16	W	0000h	<b>Source ID (SID)</b> 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. Value returned on read of this field is undefined.
15:0	RW	0000h	<b>Domain-ID (DID)</b> Indicates the id of the domain whose context-entries needs 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 Bits 15:N where N is the supported domain-id width reported in the capability register.



### 1.18.8 FSTS\_REG - Fault Status Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 34-37h  
 Default Value: 00000000h  
 Access: RO; RO-P; RWC-P  
 Size: 32 bits

Register indicating the primary fault logging status.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31:16	RO	0000h	<i>Reserved</i>
15:8	RO-P	00h	<b>Fault Record Index (FRI)</b> 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. Valid values for this field are from 0 to N, where N is the value reported through NFR field in the Capability register. The value read from this field is undefined when the PPF field is clear.
7	RO	0b	<i>Reserved</i>
6	RWC-P	0b	<b>Invalidation Time-out Error (ITE)</b> 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- IOTLBs implement this bit as RSVD.
5	RWC-P	0b	<b>Invalidation Completion Error (ICE)</b> 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 the Fault Event Control register. Hardware implementations not supporting Device-IOTLBs implement this bit as RSVD
4	RWC-P	0b	<b>Invalidation Queue Error (IQE)</b> 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 RSVD.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
3	RWC-P	0b	<b>Advanced Pending Fault (APF)</b> 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 RSVD.
2	RWC-P	0b	<b>Advanced Fault Overflow (AFO)</b> 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 RSVD.
1	RO-P	0h	<b>Primary Pending Fault (PPF)</b> 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 DMA-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.
0	RWC-P	0h	<b>Primary Fault Overflow (PFO)</b> Hardware sets this field to indicate overflow of fault recording registers. Software writing 1 clears this field.





### 1.18.9 FECTL\_REG - Fault Event Control Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 38-3Bh  
 Default Value: 80000000h  
 Access: RO; RW  
 Size: 32 bits

Register specifying the fault event interrupt message control bits.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31	RW	1h	Interrupt Mask (IM) 0 = No masking of interrupt. When a interrupt condition is detected, hardware issues an interrupt message (using the Fault Event Data & 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.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
30	RO	0h	<p><b>Interrupt Pending (IP)</b></p> <p>Hardware sets the IP field whenever it detects an interrupt condition. Interrupt condition 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. If the PPF field was already set at the time of recording a fault, it is not treated as a new interrupt condition.</p> <p>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 Advanced Fault Log register. If the APF field was already set at the time of detecting/recording a fault, it is not treated as a new interrupt condition.</p> <p>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.</p> <p>The IP field is cleared by hardware as soon as the interrupt message pending condition is serviced. This could be due to either</p> <p>(a) 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.</p> <p>(b) Software servicing the interrupting condition through one of the following ways:</p> <ul style="list-style-type: none"> <li>• 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.</li> <li>• When advanced fault logging is active, software clearing the APF field in Advanced Fault Log register.</li> </ul>
29:0	RO	00000000h	<i>Reserved</i>



### 1.18.10 FEDATA\_REG - Fault Event Data Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 3C-3Fh  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits

Register specifying the interrupt message data.

Bit	Access	Default Value	Description
31:16	RO	0000h	<b>Extended Interrupt Message Data (EIMD)</b> This field is valid only for implementations supporting 32-bit MSI data fields. Hardware implementations supporting only 16-bit MSI data may treat this field as read-only (0).
15:0	RW	0000h	<b>Interrupt Message Data (ID)</b> Data value in the fault-event interrupt message.

### 1.18.11 FEADDR\_REG - Fault Event Address Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 40-43h  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits

Register specifying the interrupt message address.

Bit	Access	Default Value	RST/ PWR	Description
31:2	RW	00000000h	Core	<b>Message Address (MA)</b> When fault events are enabled, the contents of this register specify the DWORD aligned address (Bits 31:2) for the MSI memory write transaction.
1:0	RO	0h	Core	<i>Reserved</i>



### 1.18.12 FEUADDR\_REG - Fault Event Upper Address Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 44-47h  
 Default Value: 00000000h  
 Access: RO  
 Size: 32 bit

Register specifying the interrupt message address. For platforms supporting only interrupt messages in the 32-bit address range, this register is treated as read-only (0).

Bit	Access	Default Value	Description
31:0	RO	00000000h	<b>Message Upper Address (MUA)</b> This register need to be implemented only if hardware supports 64-bit message address. If implemented, the contents of this register specify the upper 32-bits of a 64-bit MSI write transaction. If hardware does not support 64-bit messages, the register is treated as read-only (0).

### 1.18.13 AFLOG\_REG - Advanced Fault Log Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 58-5Fh  
 Default Value: 0000000000000000h  
 Access: RO  
 Size: 64 bits

Register to specify the base address of memory-resident fault-log region. This register is treated as read-only (0) for implementations not supporting advanced translation fault logging (AFL field reported as 0 in the Capability register). This register is sticky and can be cleared only through powergood reset or via software clearing the RW1C fields by writing a 1.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
63:12	RO	00000000 00000h	<b>Fault Log Address (FLA)</b> This field specifies the base of size-aligned fault-log region in system memory. Hardware may ignore and 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 returns value that was last programmed to it.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
11:9	RO	0h	<p><b>Fault Log Size (FLS)</b></p> <p>This field specifies the size of the fault log region pointed by the FLA field. The size of the fault log region is <math>2^X * 4\text{-KB}</math>, where X is the value programmed in this register.</p> <p>When implemented, reads of this field returns value that was last programmed to it.</p>
8:2	RO	00h	<i>Reserved</i>
1	RO	0h	<p><b>Advanced Pending Fault (APF)</b></p> <p>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.</p>
0	RO	0h	<p><b>Advanced Fault Overflow (AFO)</b></p> <p>Hardware sets this field to indicate advanced fault log overflow condition. Software writing 1 to this field clears it.</p>



### 1.18.14 PMEM\_REG - Protected Memory Enable Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 64-67h  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits

Register to enable the DMA protected memory regions setup through the PLMBASE, PLMLIMIT, PHMBASE, PHMLIMIT registers. When LT.CMD.LOCK.PMRC command is invoked, this register is locked (treated RO). When LT.CMD.UNLOCK.PMRC command is invoked, this register is unlocked (treated RW). This register is always treated as RO (0) for implementations not supporting protected memory regions (PLMR and PHMR fields reported as 0 in the Capability register).

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31	RW	0h	<p><b>Enable Protected Memory (EPM)</b></p> <p>This field controls DMA accesses to the protected low-memory and protected high-memory regions.</p> <p>0 = DMA accesses to protected memory regions are handled as follows:</p> <ul style="list-style-type: none"> <li>If DMA-remapping hardware is not enabled, DMA requests (including those to protected regions) are not blocked.</li> <li>If DMA-remapping hardware is enabled, DMA requests are translated per the programming of the DMA-remapping structures. Software may program the DMA-remapping structures to allow or block DMA to the protected memory regions.</li> </ul> <p>1 = DMA accesses to protected memory regions are handled as follows:</p> <ul style="list-style-type: none"> <li>If DMA-remapping hardware is not enabled, DMA to protected memory regions are blocked. These DMA requests are not recorded or reported as DMA-remapping faults.</li> <li>If DMA-remapping hardware is enabled, hardware may or may not block DMA to the protected memory region(s).</li> </ul> <p>Software must not depend on hardware protection of the protected memory regions, and must ensure the DMA-remapping structures are properly programmed to not allow DMA to the protected memory regions. Hardware reports the status of the protected memory enable/disable operation through the PRS field in this register.</p> <p>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.</p>
30:1	RO	00000000h	Reserved



(Sheet 2 of 2)

0	RO	0h	<b>Protected Region Status (PRS)</b> This field indicates the status of protected memory region. 0 = Protected memory region(s) not enabled. 1 = Protected memory region(s) enabled.
---	----	----	---

### 1.18.15 PLMBASE\_REG - Protected Low-Memory Base Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 68-6Bh  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits

Register to setup the base address of DMA protected low-memory region. This register must be setup before enabling protected memory through PMEN\_REG, and must not be updated when protected memory regions are enabled.

When LT.CMD.LOCK.PMRC command is invoked, this register is locked (treated RO). When LT.CMD.UNLOCK.PMRC command is invoked, this register is unlocked (treated RW).

This register is always treated as RO for implementations not supporting protected low memory region (PLMR field reported as 0 in the Capability register). The alignment of the protected low memory region base depends on the number of reserved bits (N) 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 with 0 in the value read back from the register. Bits N:0 of this register is decoded by hardware as all 0's.

Bit	Access	Default Value	Description
31:21	RW	000h	<b>Protected Low-Memory Base (PLMB)</b> This register specifies the base of size aligned, protected low-memory region in system memory. The protected low-memory region has a minimum size of 2 MB and must be size aligned.
20:0	RO	000000h	<i>Reserved</i>



### 1.18.16 PLMLIMIT\_REG - Protected Low-Memory Limit Register

B/D/F/Type: 0/0/0/VCOPREMAP  
Address Offset: 6C-6Fh  
Default Value: 00000000h  
Access: RO; RW  
Size: 32 bits

Register to setup the limit address of DMA protected low-memory region. This register must be setup before enabling protected memory through PMEN\_REG, and must not be updated when protected memory regions are enabled.

When LT.CMD.LOCK.PMRC command is invoked, this register is locked (treated RO).  
When LT.CMD.UNLOCK.PMRC command is invoked, this register is unlocked (treated RW).

This register is always treated as RO for implementations not supporting protected low memory region (PLMR field reported as 0 in the Capability register). The alignment of the protected low memory region limit depends on the number of reserved bits (N) 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 with 0 in the value read back from the register. Bits N:0 of the limit register is decoded by hardware as all 1's. 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.

Bit	Access	Default Value	Description
31:21	RW	000h	<b>Protected Low-Memory Limit (PLML)</b> This register specifies the last host physical address of the DMA protected low-memory region in system memory.
20:0	RO	000000h	<i>Reserved</i>





### 1.18.17 PHMBASE\_REG - Protected High-Memory Base Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 70-77h  
 Default Value: 0000000000000000h  
 Access: RO; RW  
 Size: 64 bits

Register to setup the base address of DMA protected high-memory region. This register must be setup before enabling protected memory through PMEN\_REG, and must not be updated when protected memory regions are enabled.

When LT.CMD.LOCK.PMRC command is invoked, this register is locked (treated RO).  
 When LT.CMD.UNLOCK.PMRC command is invoked, this register is unlocked (treated RW).

This register is always treated as RO for implementations not supporting protected high memory region (PHMR field reported as 0 in the Capability register). The alignment of the protected high memory region base depends on the number of reserved bits (N) 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 0's.

Bit	Access	Default Value	Description
63:21	RW	00000000 000h	<b>Protected High-Memory Base (PHMB)</b> This register specifies the base of size aligned, protected memory region in system memory. Hardware may not utilize Bits 63:HAW, where HAW is the host address width. The protected high-memory region has a minimum size of 2 MB and must be size aligned.
20:0	RO	000000h	<i>Reserved</i>



### 1.18.18 PHMLIMIT\_REG - Protected High-Memory Limit Register

B/D/F/Type: 0/0/0/VCOPREMAP  
Address Offset: 78-7Fh  
Default Value: 0000000000000000h  
Access: RO; RW  
Size: 64 bits

Register to setup the limit address of DMA protected high-memory region. This register must be setup before enabling protected memory through PMEN\_REG, and must not be updated when protected memory regions are enabled.

When LT.CMD.LOCK.PMRC command is invoked, this register is locked (treated RO).  
When LT.CMD.UNLOCK.PMRC command is invoked, this register is unlocked (treated RW).

This register is always treated as RO for implementations not supporting protected high memory region (PHMR field reported as 0 in the Capability register).

The alignment of the protected high memory region limit depends on the number of reserved bits (N) 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 1's.

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.

Bit	Access	Default Value	Description
63:21	RW	00000000 000h	<b>Protected High-Memory Limit (PHML)</b> This register specifies the last host physical address of the DMA protected high-memory region in system memory. Hardware may not utilize Bits 63:HAW, where HAW is the host address width.
20:0	RO	000000h	<i>Reserved</i>



### 1.18.19 IQH\_REG - Invalidation Queue Head Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 80-87h  
 Default Value: 0000000000000000h  
 Access: RO  
 Size: 64 bits

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.

Bit	Access	Default Value	Description
63:19	RO	00000000 0000h	<i>Reserved</i>
18:4	RO	0000h	<b>Queue Head (QH)</b> Specifies the offset (128-bit aligned) to the invalidation queue for the command that is fetched next by hardware. Hardware resets this field to 0 whenever the queued invalidation is disabled (QIES field Clear in the Global Status register).
3:0	RO	0h	<i>Reserved</i>

### 1.18.20 IQT\_REG - Invalidation Queue Tail Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 88-8Fh  
 Default Value: 0000000000000000h  
 Access: RO  
 Size: 64 bits

Register indicating the invalidation queue tail. This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

Bit	Access	Default Value	Description
63:19	RO	00000000 0000h	<i>Reserved</i>
18:4	RO	0000h	<b>Queue Tail (QT)</b> Specifies the offset (128-bit aligned) to the invalidation queue for the command that is written next by software.
3:0	RO	0h	<i>Reserved</i>



### 1.18.21 IQA\_REG - Invalidation Queue Address Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 90-97h  
 Default Value: 0000000000000000h  
 Access: RO; RW  
 Size: 64 bits

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.

When supported, writing to this register causes the Invalidation Queue Head and Invalidation Queue Tail registers to be reset to 0h.

Bit	Access	Default Value	Description
63:12	RW	00000000 00000h	<b>Invalidation Queue Base Address (IQA)</b> This field points to the base of 4-KB aligned invalidation request queue. Hardware ignores and 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.
11:3	RO	000h	<i>Reserved</i>
2:0	RW	0h	<b>Queue Size (QS)</b> This field specifies the size of the invalidation request queue. A value of X in this field indicates an invalidation request queue of (X+1) 4-KB pages. The number of entries in the invalidation queue is $2^{(X + 8)}$ .

### 1.18.22 ICS\_REG - Invalidation Completion Status Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 9C-9Fh  
 Default Value: 00000000h  
 Access: RO; RWC-P  
 Size: 32 bits

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.

Bit	Access	Default Value	Description
31:1	RO	00000000h	<i>Reserved</i>
0	RWC-P	0b	<b>Invalidation Wait Descriptor Complete (IWC)</b> Indicates completion of Invalidation Wait Descriptor with Interrupt Flag (IF) field Set. Hardware implementations not supporting queued invalidations implement this field as RsvdZ.



### 1.18.23 IECTL\_REG - Invalidation Event Control Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: A0-A3h  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits

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.

Bit	Access	Default Value	Description
31	RW	0b	<p><b>Interrupt Mask (IM)</b></p> <p>0 = No masking of interrupt. When a invalidation event condition is detected, hardware issues an interrupt message (using the Invalidation Event Data &amp; Invalidation Event Address register values).</p> <p>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.</p>
30	RO	0b	<p><b>Interrupt Pending (IP)</b></p> <p>Hardware sets the IP field whenever it detects an interrupt condition. Interrupt condition is defined as:</p> <ul style="list-style-type: none"> <li>An Invalidation Wait Descriptor with Interrupt Flag (IF) field Set completed, setting the IWC field in the Invalidation Completion Status register.</li> <li>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.</li> </ul> <p>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:</p> <ul style="list-style-type: none"> <li>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.</li> <li>Software servicing the IWC field in the Invalidation Completion Status register.</li> </ul>
29:0	RO	00000000h	Reserved



### 1.18.24 IEDATA\_REG - Invalidation Event Data Register

B/D/F/Type: 0/0/0/VCOPREMAP  
Address Offset: A4-A7h  
Default Value: 00000000h  
Access: RW  
Size: 32 bits

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.

Bit	Access	Default Value	Description
31:16	RW	0000h	<b>Extended Interrupt Message Data (EIMD)</b> 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.
15:0	RW	0000h	<b>Interrupt Message Data (IMD)</b> Data value in the interrupt request.

### 1.18.25 IEADDR\_REG - Invalidation Event Address Register

B/D/F/Type: 0/0/0/VCOPREMAP  
Address Offset: A8-ABh  
Default Value: 00000000h  
Access: RO; RW  
Size: 32 bits

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.

Bit	Access	Default Value	Description
31:2	RW	00000000h	<b>Message Address (MA)</b> When fault events are enabled, the contents of this register specify the DWORD-aligned address (Bits 31:2) for the interrupt request.
1:0	RO	00b	<i>Reserved</i>



### 1.18.26 IEUADDR\_REG - Invalidation Event Upper Address Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: AC-AFh  
 Default Value: 00000000h  
 Access: RW  
 Size: 32 bits

Register specifying the Invalidation Event interrupt message upper address. This register is treated as RsvdZ by implementations reporting both Queued Invalidation (QI) and Extended Interrupt Mode (EIM) as not supported in the Extended Capability register.

Bit	Access	Default Value	Description
31:0	RW	00000000h	<b>Message Upper Address (MUA)</b> Hardware implementations supporting Queued Invalidation and Extended Interrupt Mode are required to implement this register. Hardware implementations not supporting Queued Invalidation and Extended Interrupt Mode may treat this field as RSVD.

### 1.18.27 IRTA\_REG - Interrupt Remapping Table Address Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: B8-BFh  
 Default Value: 0000000000000000h  
 Access: RO; RW  
 Size: 64 bits

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.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
63:12	RW	000000000000h	<b>Interrupt Remapping Table Address (IRTA)</b> This field points to the base interrupt remapping table. Hardware ignores and does not implement 63:HAW, where HAW is the host address width. Reads of this field returns last programmed to it.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
11	RW	0b	<b>Extended Interrupt Mode Enable (EIMI)</b> 0 = Legacy interrupt 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 is treated as reserved. On Itanium® processor platforms hardware interprets low 16- bits of Destination-ID field in the IRTEs and treats the high 16-bits as reserved. 1 = Intel®64 platform is operating in Extended Interrupt Mode. Hardware interprets all 32-bits of the Destination- ID field in the IRTEs. Hardware reporting Extended Interrupt Mode (EIM) as Clear in the Capability register treats this field as RsvdZ.
10:4	RO	00h	<i>Reserved</i>
3:0	RW	0h	<b>Size (S)</b> 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.

### 1.18.28 IVA\_REG - Invalidate Address Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 100-107h  
 Default Value: 0000000000000000h  
 Access: RO; W  
 Size: 64 bits

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. Value returned on reads of this register is undefined. There is an IVA\_REG for each IOTLB Invalidation unit supported by hardware.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
63:12	W	00000000 00000h	<b>Address (Addr)</b> Software provides the DMA address that needs to be page-selectively invalidated. To request a page-selective invalidation request to hardware, software must first write the appropriate fields in this register, and then issue appropriate page-selective invalidate command through the IOTLB_REG. Hardware ignores Bits 63:N, where N is the maximum guest address width (MGAW) supported. Value returned on read of this field is undefined.
11:7	RO	00h	<i>Reserved</i>





(Sheet 2 of 2)

Bit	Access	Default Value	Description																														
6	W	0h	<p><b>Invalidation Hint (IH)</b></p> <p>The field provides hint to hardware to preserve or flush the non-leaf (page-directory) entries that may be cached in hardware.</p> <p>0 = Software may have modified both leaf and non-leaf page-table entries corresponding to mappings specified in the ADDR and AM fields.</p> <p>On a pages elective invalidation request, hardware must flush both the cached leaf and non-leaf page-table Value returned on read of this field is undefined. Entries corresponding to mappings specified by ADDR and AM fields.</p> <p>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.</p>																														
5:0	W	00h	<p><b>Address Mask (AM)</b></p> <p>The value in this field specifies the number of low order bits of the ADDR field that must be masked for the invalidation operation. Mask field enables software to request invalidation of contiguous mappings for size-aligned regions.</p> <p>For example: Mask Value ADDR bits masked Pages invalidated:</p> <table border="1" data-bbox="776 1163 1237 1591"> <thead> <tr> <th>Mask Value</th> <th>ADDR Bits</th> <th>Pages Invalidated</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>None</td> <td>1</td> </tr> <tr> <td>1</td> <td>12</td> <td>2</td> </tr> <tr> <td>2</td> <td>13:12</td> <td>8</td> </tr> <tr> <td>3</td> <td>14:12</td> <td>16</td> </tr> <tr> <td>4</td> <td>15:12</td> <td>32</td> </tr> <tr> <td>5</td> <td>16:12</td> <td>64</td> </tr> <tr> <td>6</td> <td>17:12</td> <td>128</td> </tr> <tr> <td>7</td> <td>18:12</td> <td>256</td> </tr> <tr> <td>8</td> <td>19:12</td> <td>512</td> </tr> </tbody> </table> <p>Hardware implementations report the maximum supported mask value through the Capability register. Value returned on read of this field is undefined.</p>	Mask Value	ADDR Bits	Pages Invalidated	0	None	1	1	12	2	2	13:12	8	3	14:12	16	4	15:12	32	5	16:12	64	6	17:12	128	7	18:12	256	8	19:12	512
Mask Value	ADDR Bits	Pages Invalidated																															
0	None	1																															
1	12	2																															
2	13:12	8																															
3	14:12	16																															
4	15:12	32																															
5	16:12	64																															
6	17:12	128																															
7	18:12	256																															
8	19:12	512																															



### 1.18.29 IOTLB\_REG - IOTLB Invalidate Register

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 108-10Fh  
 Default Value: 0000000000000000h  
 Access: RO; RW  
 Size: 64 bits

Register to control page-table entry caching. The act of writing the upper byte of the IOTLB\_REG with IVT field set causes the hardware to perform the IOTLB invalidation. There is an IOTLB\_REG for each IOTLB Invalidation unit supported by hardware.

(Sheet 1 of 3)

Bit	Access	Default Value	Description
63	RW	0h	<p><b>Invalidate IOTLB (IVT)</b></p> <p>Software requests IOTLB invalidation by setting this field. Software must also set the requested invalidation granularity by programming the IIRG field.</p> <p>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.</p> <p>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 through any of the IOTLB invalidation units when there is a context-cache invalidation request pending at this DMA-remapping hardware unit.</p> <p>When more than one IOTLB invalidation units are supported by a DMA-remapping hardware unit, software may submit IOTLB invalidation request through any of the currently free units while there are pending requests on other units.</p> <p>Hardware implementations reporting write-buffer flushing requirement (RWBF=1 in Capability register) must implicitly perform a write buffer flushing before reporting invalidation complete to software through the IVT field.</p>



(Sheet 2 of 3)

Bit	Access	Default Value	Description
62:60	RW	0h	<p><b>IOTLB Invalidation Request Granularity (IIRG)</b></p> <p>When requesting hardware to invalidate the IOTLB (by setting the IVT field), software writes the requested invalidation granularity through this IIRG field. Following are the encodings for the IIRG field. 000: Reserved.</p> <p>001: Global invalidation request.</p> <p>010: Domain-selective invalidation request. The target domain-id must be specified in the DID field.</p> <p>011: Domain-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.</p> <p>100 - 111: Reserved.</p> <p>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.</p>
59:57	RO	0h	<p><b>IOTLB Actual Invalidation Granularity (IAIG)</b></p> <p>Hardware reports the granularity at which an invalidation request was processed through this field at the time of reporting invalidation completion (by clearing the IVT field).</p> <p>The following are the encodings for the IAIG field:</p> <p>000: 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.</p> <p>001: Global Invalidation performed. This could be in response to a global, domain-selective, domain-page-selective, or device-page-selective invalidation request.</p> <p>010: Domain-selective invalidation performed using the domain-id specified by software in the DID field. This could be in response to a domain-selective, domain-page-selective, or device-page-selective invalidation request.</p> <p>011: 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 domain-page-selective or device-page-selective invalidation request.</p> <p>100 - 111: Reserved.</p>
56:50	RO	00h	<i>Reserved</i>



(Sheet 3 of 3)

Bit	Access	Default Value	Description
49	RW	000000h	<p><b>Drain Reads (DR)</b></p> <p>This field is ignored by hardware if the DRD field is reported as clear in the Capability register. When DRD field is reported as set in the Capability register, the following encodings are supported for this field:</p> <p>0 = Hardware may complete the IOTLB invalidation without draining any translated DMA reads that are queued in the root-complex and yet to be processed.</p> <p>1 = Hardware must drain all/relevant translated DMA reads that are queued in the root-complex before indicating IOTLB invalidation completion to software. A DMA read request to system memory is defined as drained when root-complex has finished fetching all of its read response data from memory.</p>
48	RW	00h	<p><b>Drain Writes (DW)</b></p> <p>This field is ignored by hardware if the DWD field is reported as clear in the Capability register. When DWD field is reported as set in the Capability register, the following encodings are supported for this field:</p> <p>0 = Hardware may complete the IOTLB invalidation without draining any translated DMA writes that are queued in the root-complex for processing.</p> <p>1 = Hardware must drain all/relevant translated DMA writes that are queued in the root-complex before indicating IOTLB invalidation completion to software. A DMA write request to system memory is defined as drained when the effects of the write is visible to processor accesses to all addresses targeted by the DMA write.</p>
47:32	RW	0000h	<p><b>Domain-ID (DID)</b></p> <p>Indicates the id of the domain whose IOTLB entries needs to be selectively invalidated. This field must be programmed by software for domain-selective, domain page-selective and device-page-selective invalidation requests.</p> <p>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.</p> <p>Hardware may ignore and not implement Bits 47: (32+N) where N is the supported domain-id width reported in the capability register.</p>
31:0	RO	00000000h	<i>Reserved</i>



### 1.18.30 FRCD\_REG - Fault Recording Registers

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: 200-20Fh  
 Default Value: 00000000000000000000000000000000h  
 Access: RWC-P; RO-P; RO  
 Size: 128 bits

Registers to record DMA-remapping 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. These registers are sticky and can be cleared only through powergood reset or via software clearing the RW1C fields by writing a 1.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
127	RWC-P	0h	<b>Fault (F)</b> 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 the PADDR, SID, FR and T fields. When this field is set, hardware may collapse additional faults from the same requestor (SID). Software writes the value read from this field to clear it.
126	RO-P	0h	<b>Type (T)</b> Type of the faulted DMA request. 0 = DMA write 1 = DMA read request. This field is relevant only when the F field is set.
125:124	RO-P	00b	<b>Address Type (AT)</b> 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.
123:104	RO	0000000000 0000000000 b	<i>Reserved</i>
103:96	RO-P	00h	<b>Fault Reason (FR)</b> Reason for the fault. Appendix B enumerates the various translation fault reason encodings. This field is relevant only when the F field is set.
95:80	RO	0000h	<i>Reserved</i>
79:64	RO-P	0000h	<b>Source Identifier (SID)</b> Requester-id of the faulted DMA request. This field is relevant only when the F field is set.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
63:12	RO-P	0000000000 000h	<b>Page Address (PADDR)</b> This field contains the address (page-granular) in the faulted DMA request. Hardware may treat Bits 63:N as reserved (0), where N is the maximum guest address width (MGAW) supported. This field is relevant only when the F field is set.
11:0	RO	000h	<i>Reserved</i>

### 1.18.31 VTCMPLRESR - VT-d Completion Resource Dedication

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: F00-F03h  
 Default Value: 00060000h  
 Access: RW-L; RO  
 Size: 32 bits

This register provides a programmable interface to dedicate the DMI Completion Tracking Queue resources to DMI VCO Read, DMI VCO Write, DMI VC1 and DMI VCp Intel VT-d fetch and PEG Completion Tracking Queue resources to PEG VCO read and PEG VCO write Intel VT-d fetch.

(Sheet 1 of 3)

Bit	Access	Default Value	Description
31:20	RW-L	000h	<i>Reserved</i>
19:16	RO	6h	<b>DMI VT-d Completion Tracking Queue Resource Available (DMIVTCTRA)</b> Number of entries available in DMI VT-d Completion Tracking Queue. 1-based. The values programmed in the fields below must not be greater than the value advertised in this field.



(Sheet 2 of 3)

Bit	Access	Default Value	Description
15:12	RW-L	0h	<p><b>DMI VC1 VT-d Completion Tracking Queue Resource Threshold (DMIVC1CTQRT)</b></p> <p>One based minimum threshold value used to throttle DMI VC1 VT-d fetch. When the number of free DMI VT-d Completion Tracking Queue entries equals or falls below the value programmed in this field, DMI VC1 VT-d fetch is throttled until the number of free DMI Completion Tracking Queue entries rise above this threshold.</p> <p>For example:</p> <p>0000 – Throttle DMI VC1 VT-d Fetch when there is no entry left.</p> <p>0001 – Throttle DMI VC1 VT-d Fetch when there is 1 or less entry left.</p> <p>0010 - Throttle DMI VC1 VT-d Fetch when there are 2 or less entry left.</p> <p>0011 – Throttle DMI VC1 VT-d Fetch when there are 3 or less entry left.</p> <p>0100 - Throttle DMI VC1 VT-d Fetch when there are 4 or less entry left.</p> <p>0101 - 1111 - Reserved.</p>
11:8	RW-L	0h	<p><b>DMI VCp VT-d Completion Tracking Queue Resource Threshold (DMIVCPCTQRT)</b></p> <p>One based minimum threshold value used to throttle DMI VCp VT-d fetch. When the number of free DMI VT-d Completion Tracking Queue entries equals or falls below the value programmed in this field, DMI VCp VT-d fetch is throttled until the number of free DMI Completion Tracking Queue entries rise above this threshold.</p> <p>For example:</p> <p>0000 – Throttle DMI VCp VT-d Fetch when there is no entry left.</p> <p>0001 – Throttle DMI VCp VT-d Fetch when there is 1 or less entry left.</p> <p>0010 - Throttle DMI VCp VT-d Fetch when there are 2 or less entry left.</p> <p>0011 – Throttle DMI VCp VT-d Fetch when there are 3 or less entry left.</p> <p>0100 - Throttle DMI VCp VT-d Fetch when there are 4 or less entry left.</p> <p>0101 - 1111 - Reserved.</p>



(Sheet 3 of 3)

Bit	Access	Default Value	Description
7:4	RW-L	0h	<p><b>DMI VCO Write VT-d Completion Tracking Queue Resource Threshold (DMIVCOWRCTQRT)</b></p> <p>One based minimum threshold value used to throttle DMI VCO Write VT-d fetch. When the number of free DMI VT-d Completion Tracking Queue entries equals or falls below the value programmed in this field, DMI VCO Write VT-d fetch is throttled until the number of free DMI Completion Tracking Queue entries rise above this threshold.</p> <p>For example:</p> <p>0000 – Throttle DMI VCO Write VT-d Fetch when there is no entry left.</p> <p>0001 – Throttle DMI VCO Write VT-d Fetch when there is 1 or less entry left.</p> <p>0010 - Throttle DMI VCO Write VT-d Fetch when there are 2 or less entry left.</p> <p>0011 – Throttle DMI VCO Write VT-d Fetch when there are 3 or less entry left.</p> <p>0100 - Throttle DMI VCO Write VT-d Fetch when there are 4 or less entry left.</p> <p>0101 - 1111 - Reserved.</p>
3:0	RW-L	0h	<p><b>DMI VCO Read VT-d Completion Tracking Queue Resource Threshold (DMIVCORDCTQRT)</b></p> <p>1-based minimum threshold value used to throttle DMI VCO Read VT-d fetch. When the number of free DMI VT-d Completion Tracking Queue entries equals or falls below the value programmed in this field, DMI VCO Read VT-d fetch is throttled until the number of free DMI Completion Tracking Queue entries rise above this threshold.</p> <p>For example:</p> <p>0000 – Throttle DMI VCO Read VT-d Fetch when there is no entry left.</p> <p>0001 – Throttle DMI VCO Read VT-d Fetch when there is 1 or less entry left.</p> <p>0010 - Throttle DMI VCO Read VT-d Fetch when there are 2 or less entry left.</p> <p>0011 – Throttle DMI VCO Read VT-d Fetch when there are 3 or less entry left.</p> <p>0100 - Throttle DMI VCO Read VT-d Fetch when there are 4 or less entry left.</p> <p>0101 - 1111 - Reserved.</p>





### 1.18.32 VTFTCHARBCTL - VC0/VCp Intel VT-d Fetch Arbiter Control

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: F04-F07h  
 Default Value: 0000FFFFh  
 Access: RW-L  
 Size: 32 bits

This register controls the relative grant count given to each of the DMI VC0, DMI VC1 and PEG VC0 VT-d fetch requests.

Bit	Access	Default Value	Description
31:16	RW-L	0000h	<i>Intel Reserved</i>
15:12	RW-L	Fh	<i>Intel Reserved</i>
11:8	RW-L	Fh	<b>PEG VC0 VT-d Fetch Grant Count (PEGVCOGNTCNT)</b> The arbiter will continue to grant PEG VC0 VT-d fetch as long as the grant count value in this field is greater than zero.
7:4	RW-L	Fh	<b>DMI VCp VT-d Fetch Grant Count (DMIVCPGNTCNT)</b> The arbiter will continue to grant DMI VCp VT-d fetch as long as the grant count value in this field is greater than zero and there is no higher priority VT-d fetch request. Arbitration will switch to PEG VC0 VT-d fetch request if the grant count corresponding to PEG VC0 VT-d fetch is greater than zero and the VT-d fetch request corresponding to PEG VC0 stream is available.
3:0	RW-L	Fh	<b>DMI VC0 VT-d Fetch Grant Count (DMIVCOGNTCNT)</b> The arbiter will continue to grant DMI VC0 VT-d fetch as long as the grant count value in this field is greater than zero and there is no higher priority VT-d fetch requests. Arbitration will switch to DMI VCp or PEG VC0 VT-d fetch requests if the grant count corresponding to those VT-d fetch is greater than zero and the VT-d fetch requests corresponding to those streams are available.



### 1.18.33 PEGVTCMPLRESR - PEG VT-d Completion Resource Dedication

B/D/F/Type: 0/0/0/VCOPREMAP  
 Address Offset: F08-F0Bh  
 Default Value: 20004000h  
 Access: RW-L; RO  
 Size: 32 bits

This register provides a programmable interface to dedicate the PEG0 and Completion Tracking Queue resources to PEG0 VCO read, PEG0 VCO writes VT-d fetch.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31:30	RW-L	00b	<b>PEG Completion Tracking Queue Resource Sharing Mode (PCTQRSM)</b> 11: Intel Reserved 01: Intel Reserved 00: PEG0 is assigned all of the resources in the Completion Tracking Queue.
29:25	RO	10000b	<i>Intel Reserved</i>
24:20	RW-L	00000b	<i>Intel Reserved</i>
19:15	RW-L	00000b	<i>Intel Reserved</i>
14:10	RO	10000b	<b>PEG0 VT-d Completion Tracking Queue Resource Available (PEGOVTCTRA)</b> Number of entries available in PEG0 VT-d Completion Tracking Queue. 1-based. The values programmed in the fields below must not be greater than the value advertised in this field.
9:5	RW-L	00000b	<b>PEG0 VCO Write VT-d Completion Tracking Queue Resource Threshold (PEGOVCOWRCTQRT)</b> 1-based minimum threshold value used to throttle PEG0 VCO Write VT-d fetch. When the number of free PEG0 VT-d Completion Tracking Queue entries equals or falls below the value programmed in this field, PEG0 VCO Write VT-d fetch is throttled until the number of free PEG0 Completion Tracking Queue entries rise above this threshold. For example: 00000 – Throttle PEG0 VCO Write VT-d Fetch when there is no entry left. 00001 – Throttle PEG0 VCO Write VT-d Fetch when there is 1 or less entry left. 00010 - Throttle PEG0 VCO Write VT-d Fetch when there are 2 or less entry left. 00011 – Throttle PEG0 VCO Write VT-d Fetch when there are 3 or less entry left. 00100 - Throttle PEG0 VCO Write VT-d Fetch when there are 4 or less entry left.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
4:0	RW-L	00000b	<p><b>PEGO VCO Read VT-d Completion Tracking Queue Resource Threshold (PEGOVCORDCTQRTCT)</b></p> <p>1-based minimum threshold value used to throttle PEGO VCO Read VT-d fetch. When the number of free PEGO VT-d Completion Tracking Queue entries equals or falls below the value programmed in this field, PEGO VCO Read VT-d fetch is throttled until the number of free PEGO Completion Tracking Queue entries rise above this threshold.</p> <p>For example:</p> <p>00000 – Throttle PEGO VCO Read VT-d Fetch when there is no entry left.</p> <p>00001 – Throttle PEGO VCO Read VT-d Fetch when there is 1 or less entry left.</p> <p>00010 - Throttle PEGO VCO Read VT-d Fetch when there are 2 or less entry left.</p> <p>00011 – Throttle PEGO VCO Read VT-d Fetch when there are 3 or less entry left.</p> <p>00100 - Throttle PEGO VCO Read VT-d Fetch when there are 4 or less entry left.</p>



## 1.19 DMI VC1 REMAP Registers

(Sheet 1 of 2)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Version Register	VER_REG	0	3	00000010h	RO
Capability Register	CAP_REG	8	F	00C9008020E3 0272h	RO
Extended Capability Register	ECAP_REG	10	17	000000000000 1000h	RO
Global Command Register	GCMD_REG	18	1B	00000000h	RO; W
Global Status Register	GSTS_REG	1C	1F	00000000h	RO
Root-Entry Table Address Register	RTADDR_REG	20	27	000000000000 0000h	RO; RW
Context Command Register	CCMD_REG	28	2F	000000000000 0000h	RW-SC; RW; RO; W
Fault Status Register	FSTS_REG	34	37	00000000h	RWC-P; RO-P; RO
Fault Event Control Register	FECTL_REG	38	3B	80000000h	RO; RW
Fault Event Data Register	FEDATA_REG	3C	3F	00000000h	RO; RW
Fault Event Address Register	FEADDR_REG	40	43	00000000h	RO; RW
Fault Event Upper Address Register	FEUADDR_REG	44	47	00000000h	RO
Advanced Fault Log Register	AFLOG_REG	58	5F	000000000000 0000h	RO
Protected Memory Enable Register	PMEN_REG	64	67	00000000h	RO; RW
Protected Low-Memory Base Register	PLMBASE_REG	68	6B	00000000h	RO; RW
Protected Low-Memory Limit Register	PLMLIMIT_REG	6C	6F	00000000h	RO; RW
Protected High-Memory Base Register	PHMBASE_REG	70	77	000000000000 0000h	RO; RW



(Sheet 2 of 2)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Protected High-Memory Limit Register	PHMLIMIT_REGISTER	78	7F	0000000000000000h	RO; RW
Invalidation Queue Head Register	IQH_REG	80	87	0000000000000000h	RO
Invalidation Queue Tail Register	IQT_REG	88	8F	0000000000000000h	RO
Invalidation Queue Address Register	IOA_REG	90	97	0000000000000000h	RO
Invalidation Completion Status Register	ICS_REG	9C	9F	00000000h	RO
Invalidation Event Control Register	IECTL_REG	A0	A3	00000000h	RO
Invalidation Event Data Register	IEDATA_REG	A4	A7	00000000h	RO
Invalidation Event Address Register	IEADDR_REG	A8	AB	00000000h	RO
Invalidation Event Upper Address Register	IEUADDR_REG	AC	AF	00000000h	RO
Interrupt Remapping Table Address Register	IRTA_REG	B8	BF	0000000000000000h	RO
Invalidate Address Register	IVA_REG	100	107	0000000000000000h	RO; W
IOTLB Invalidate Register	IOTLB_REG	108	10F	0000000000000000h	RO; RW; RW-SC
Fault Recording Registers	FRCD_REG	200	20F	00000000000000000000000000000000h	RWC-P; RO-P; RO
DMA Remap Engine Policy Control	VTPOLICY	FFC	FFF	00000000h	RO; RW-L-K; RW-L



### 1.19.1 VER\_REG - Version Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: 0-3h  
 Default Value: 00000010h  
 Access: RO  
 Size: 32 bits

Register to report the architecture version supported. Backward compatibility for the architecture is maintained with new revision numbers, allowing software to load DMA-remapping drivers written for prior architecture versions.

Bit	Access	Default Value	Description
31:8	RO	000000h	<i>Reserved</i>
7:4	RO	1h	<b>Major Version number (MAX)</b> Indicates supported architecture version.
3:0	RO	0h	<b>Minor Version number (MIN)</b> Indicates supported architecture minor version.

### 1.19.2 CAP\_REG - Capability Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: 8-Fh  
 Default Value: 00C9008020E30272h  
 Access: RO  
 Size: 64 bits

Register to report general DMA remapping hardware capabilities.

(Sheet 1 of 5)

Bit	Access	Default Value	Description
63:56	RO	00h	<i>Reserved</i>
55	RO	1b	<b>DMA Read Draining (DRD)</b> 0 = On IOTLB invalidations, hardware does not support draining of DMA read requests. 1 = On IOTLB invalidations, hardware supports draining of DMA read requests. Refer to the Intel VT-d specification.
54	RO	1b	<b>DMA Write Draining (DWD)</b> 0 = On IOTLB invalidations, hardware does not support draining of DMA writes. 1 = On IOTLB invalidations, hardware supports draining of DMA writes. Refer to the Intel VT-d specification Section 6.3 for description of DMA draining.



(Sheet 2 of 5)

Bit	Access	Default Value	Description
53:48	RO	09h	<p><b>Maximum Address Mask Value (MAMV)</b></p> <p>The value in this field indicates the maximum supported value for the Address Mask (AM) field in the Invalidation Address (IVA_REG) register.</p> <p>This field is valid only when the PSI field is reported as Set.</p>
47:40	RO	00h	<p><b>Number of Fault-recording Registers (NFR)</b></p> <p>This field indicates a value of N-1, where N is the number of fault recording registers supported by hardware.</p> <p>Implementations must support at least one fault recording register (NFR = 0) for each DMA remapping hardware unit in the platform.</p> <p>The maximum number of fault recording registers per DMA-remapping hardware unit is 256.</p>
39	RO	1b	<p><b>Page Selective Invalidation Support (PSI)</b></p> <p>0 = Hardware supports only domain and global invalidates for IOTLB.</p> <p>1 = Hardware supports page selective, domain, and global invalidates for IOTLB and hardware must support a minimum MAMV value of 9.</p>
38	RO	0b	<i>Reserved</i>
37:34	RO	0h	<p><b>Super Page Support (SPS)</b></p> <p>This field indicates the super page sizes supported by hardware.</p> <p>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:</p> <ul style="list-style-type: none"> <li>0: 21-bit offset to page frame</li> <li>1: 30-bit offset to page frame</li> <li>2: 39-bit offset to page frame</li> <li>3: 48-bit offset to page frame</li> </ul> <p>Hardware implementations supporting a specific super-page size must support all smaller super-page sizes. i.e., the only valid values for this field are 0001b, 0011b, 0111b, 1111b.</p>
33:24	RO	020h	<p><b>Fault-recording Register Offset (FRO)</b></p> <p>This field specifies the location to the first fault recording register relative to the register base address of this DMA-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 <math>X + (16 * Y)</math>.</p>



(Sheet 3 of 5)

Bit	Access	Default Value	Description
23	RO	1b	<p><b>Isochrony (Isoch)</b></p> <p>0 = Indicates this DMA-remapping hardware unit has no critical isochronous requesters in its scope.            1 = Indicates this DMA-remapping hardware unit has one or more critical isochronous requesters in its scope.</p> <p>To guarantee isochronous performance, software must ensure invalidation operations do not impact active DMA streams from such requesters. This implies that when DMA is active, software perform page-selective invalidations (instead of coarser invalidations).</p>
22	RO	1b	<p><b>Zero Length Read (ZLR)</b></p> <p>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.</p>
21:16	RO	23h	<p><b>Maximum Guest Address Width (MGAW)</b></p> <p>This field indicates the maximum DMA virtual addressability supported by remapping hardware.</p> <p>The Maximum Guest Address Width (MGAW) is computed as <math>(N+1)</math>, where N is the value reported in this field. For example, a hardware implementation supporting 48-bit MGAW reports a value of 47 (<math>101111b = 2Fh</math>) in this field.</p> <p>If the value in this field is X, DMA requests to addresses above <math>2^{(X+1)}-1</math> are always blocked by hardware.</p> <p>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).</p> <p>Implementations are recommended to support MGAW at least equal to the physical addressability (host address width) of the platform.</p>
15:13	RO	000b	<i>Reserved</i>





(Sheet 4 of 5)

Bit	Access	Default Value	Description
12:8	RO	02h	<p><b>Supported Adjusted Guest Address Widths (SAGAW)</b></p> <p>This 5-bit field indicates the supported adjusted guest address widths (which in turn represents the levels of page-table walks for the 4-KB page size) supported by the hardware implementation.</p> <p>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:</p> <ul style="list-style-type: none"> <li>0: 30-bit AGAW (2-level page table)</li> <li>1: 39-bit AGAW (3-level page table)</li> <li>2: 48-bit AGAW (4-level page table)</li> <li>3: 57-bit AGAW (5-level page table)</li> <li>4: 64-bit AGAW (6-level page table)</li> </ul> <p>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.</p>
7	RO	0b	<p><b>Caching Mode (CM)</b></p> <p>0 = Hardware does not cache not present and erroneous entries in any of the remapping caches. Invalidation 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.</p> <p>1 = Hardware may cache not present and erroneous mappings in the remapping caches. Any software updates to the DMA-remapping structures (including updates to not-present or erroneous entries) require explicit invalidation.</p> <p>Refer to VT-d specification Section 6.1 for more details on caching mode.</p> <p>Hardware implementations are required to support operation corresponding to CM=0.</p>
6	RO	1b	<p><b>Protected High-Memory Region (PHMR)</b></p> <p>0 = Indicates protected high-memory region not supported. 1 = Indicates protected high-memory region is supported.</p>
5	RO	1b	<p><b>Protected Low-Memory Region (PLMR)</b></p> <p>0 = Indicates protected low-memory region not supported. 1 = Indicates protected low-memory region is supported.</p>
4	RO	1b	<p><b>Required Write-Buffer Flushing (RWBF)</b></p> <p>0 = Indicates no write-buffer flushing needed to ensure changes to memory-resident structures are visible to hardware.</p> <p>1 = Indicates software must explicitly flush the write buffers to ensure updates made to memory-resident DMA-remapping structures are visible to hardware. Refer VT-d specification Section 11.1 for more details on write buffer flushing requirements.</p>



(Sheet 5 of 5)

Bit	Access	Default Value	Description
3	RO	0b	<b>Advanced Fault Logging (AFL)</b> 0 = Indicates advanced fault logging not supported. Only primary fault logging is supported. 1 = Indicates advanced fault logging is supported.
2:0	RO	010b	<b>Number of Domains Supported (ND)</b> 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.

### 1.19.3 ECAP\_REG - Extended Capability Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: 10-17h  
 Default Value: 0000000000001000h  
 Access: RO  
 Size: 64 bits

Register to report DMA-remapping hardware extended capabilities.

(Sheet 1 of 3)

Bit	Access	Default Value	Description
63:24	RO	00000000 00h	<i>Reserved</i>
23:20	RO	0h	<b>Maximum Handle Mask Value (MHMV)</b> 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 is reported as Set.
19:18	RO	00b	<i>Reserved</i>



(Sheet 2 of 3)

Bit	Access	Default Value	Description
17:8	RO	010h	<p><b>Invalidation Unit Offset (IVO)</b></p> <p>This field specifies the location to the first IOTLB registers relative to the register base address of this DMA-remapping hardware unit.</p> <p>If the register base address is X, and the value reported in this field is Y, the address for the first IOTLB register is calculated as <math>X + (16 * Y)</math>.</p>
7	RO	0b	<p><b>Snoop Control (SC)</b></p> <p>0 = Hardware does not support setting the SNP field to 1 in the page-table entries.                      1 = Hardware supports setting the SNP field to 1 in the page-table entries.</p>
6	RO	0b	<p><b>Pass Through (PT)</b></p> <p>0 = Hardware does not support pass-through translation type in context entries.                      1 = Hardware supports pass-through translation type in context entries.</p>
5	RO	0b	<p><b>Caching Hints (CH)</b></p> <p>0 = Hardware does not support IOTLB caching hints (ALH and EH fields in context-entries are treated as reserved).                      1 = Hardware supports IOLTb caching hints through the ALH and EH fields in context-entries.</p>
4	RO	0b	<p><b>Extended Interrupt Mode (EIM)</b></p> <p>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).</p> <p>Itanium® processor platforms support 16-bit APICIDs and always report this field as 0.</p> <p>This field is valid only when the IR field is reported as Set.</p>
3	RO	0b	<p><b>Interrupt Remapping Support (IR)</b></p> <p>0 = Hardware does not support interrupt remapping.                      1 = Hardware supports interrupt remapping.</p> <p>Implementations reporting this field as Set must also support Queued Invalidation (QI = 1b).</p>
2	RO	0b	<p><b>Device IOTLB Support (DI)</b></p> <p>1 = Hardware does not support device-IOTLBs.                      0 = Hardware supports Device-IOTLBs.</p> <p>Implementations reporting this field as Set must also support Queued Invalidation (QI = 1b).</p>
1	RO	0b	<p><b>Queued Invalidation Support (QI)</b></p> <p>0 = Hardware does not support queued invalidations.                      1 = Hardware supports queued invalidations.</p>



(Sheet 3 of 3)

Bit	Access	Default Value	Description
0	RO	0b	<p><b>Coherency (C)</b></p> <p>This field indicates if hardware access to the root, context, page-table and interrupt remap structures are coherent (snooped) or not.</p> <p>0 = Indicates hardware accesses to remapping structures are incoherent.</p> <p>1 = Indicates hardware accesses to remapping structures are coherent.</p> <p>Hardware access to advanced fault log and invalidation queue are always coherent.</p>

### 1.19.4 GCMD\_REG - Global Command Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: 18-1Bh  
 Default Value: 00000000h  
 Access: RO; W  
 Size: 32 bits

Register to control DMA-remapping hardware. If multiple control fields in this register need to be modified, software must serialize the modifications through multiple writes to this register.

(Sheet 1 of 5)

Bit	Access	Default Value	Description
31	W	0b	<p><b>Translation Enable (TE)</b></p> <p>Software writes to this field to request hardware to enable/disable DMA-remapping hardware.</p> <p>0 = Disable DMA-remapping          1 = Enable DMA-remapping</p> <p>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 deterministic transaction boundaries, so that any in-flight transaction is either subject to remapping or not at all.</p> <p>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 GSTS_REG.</p> <p>Value returned on read of this field is undefined.</p>



(Sheet 2 of 5)

Bit	Access	Default Value	Description
30	W	0b	<p><b>Set Root Table Pointer (SRTP)</b></p> <p>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 register.</p> <p>Hardware reports the status of the root table pointer set operation through the RTPS field in the Global Status register.</p> <p>The root table pointer set operation must be performed before enabling or re-enabling (after disabling) DMA-remapping through the TE field.</p> <p>After a root table pointer set operation, software must globally invalidate the context cache followed by global invalidate of IOTLB. This is required to ensure hardware uses only the remapping structures referenced by the new root table pointer, and not any stale cached entries.</p> <p>While DMA-remapping hardware is active, software may update the root table pointer through this field.</p> <p>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.</p> <p>Clearing this bit has no effect.</p> <p>Value returned on read of this field is undefined.</p>
29	W	0b	<p><b>Set Fault Log (SFL)</b></p> <p>This field is valid only for implementations supporting advanced fault logging.</p> <p>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.</p> <p>Hardware reports the status of the fault log set operation through the FLS field in the Global Status register.</p> <p>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.</p> <p>Clearing this bit has no effect.</p> <p>Value returned on read of this field is undefined.</p>



(Sheet 3 of 5)

Bit	Access	Default Value	Description
28	W	0b	<p><b>Enable Advanced Fault Logging (EAFL)</b></p> <p>This field is valid only for implementations supporting advanced fault logging.</p> <p>Software writes to this field to request hardware to enable or disable advanced fault logging.</p> <p>0 = Disable advanced fault logging. In this case, translation faults are reported through the Fault Recording registers.</p> <p>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 SFL field) before enabling advanced fault logging.</p> <p>Hardware reports the status of the advanced fault logging enable operation through the AFLS field in the Global Status register.</p> <p>Value returned on read of this field is undefined.</p>
27	W	0b	<p><b>Write Buffer Flush (WBF)</b></p> <p>This bit is valid only for implementations requiring write buffer flushing.</p> <p>Software sets this field to request hardware to flush the root-complex internal write buffers. This is done to ensure any updates to the memory-resident DMA-remapping structures are not held in any internal write posting buffers. Refer to Intel VT-d specification Section 11.1 for details on write-buffer flushing requirements.</p> <p>Hardware reports the status of the write buffer flushing operation through the WBFS field in the Global Status register.</p> <p>Clearing this bit has no effect.</p> <p>Value returned on read of this field is undefined.</p>
26	RO	0b	<p><b>Queued Invalidation Enable (QIE)</b></p> <p>This field is valid only for implementations supporting queued invalidations.</p> <p>Software writes to this field to enable or disable queued invalidations.</p> <p>0 = Disable queued invalidations.</p> <p>1 = Enable use of queued invalidations.</p> <p>Hardware reports the status of queued invalidation enable operation through QIES field in the Global Status register.</p> <p>Refer to Section 6.2.2 for software requirements for enabling/disabling queued invalidations.</p> <p>The value returned on a read of this field is undefined.</p>



(Sheet 4 of 5)

Bit	Access	Default Value	Description
25	RO	0b	<p><b>Interrupt Remapping Enable (IRE)</b></p> <p>This field is valid only for implementations supporting interrupt remapping.</p> <p>0 = Disable interrupt-remapping hardware 1 = Enable interrupt-remapping hardware</p> <p>Hardware reports the status of the interrupt remapping enable operation through the IRES field in the Global Status register.</p> <p>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.</p> <p>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.</p> <p>The value returned on a read of this field is undefined.</p>
24	RO	0b	<p><b>Set Interrupt Remap Table Pointer (SIRTP)</b></p> <p>This field is valid only for implementations supporting interrupt-remapping.</p> <p>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 register.</p> <p>Hardware reports the status of the interrupt remapping table pointer set operation through the IRTPS field in the Global Status register. The interrupt remap table pointer set operation must be performed before enabling or re-enabling (after disabling) interrupt-remapping hardware through the IRE field.</p> <p>After an interrupt remap table pointer set 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.</p> <p>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.</p> <p>Clearing this bit has no effect. The value returned on a read of this field is undefined.</p>



(Sheet 5 of 5)

Bit	Access	Default Value	Description
23	RO	0b	<p><b>Compatibility Format Interrupt (CFI)</b></p> <p>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 disabled.</p> <p>0 = Block Compatibility format interrupts. 1 = Process Compatibility format interrupts as pass-through (bypass interrupt remapping).</p> <p>Hardware reports the status of updating this field through the CFIS field in the Global Status register.</p> <p>Refer to Section 5.4.1 for details on Compatibility Format interrupt requests.</p> <p>The value returned on a read of this field is undefined.</p> <p>This field is not implemented on Itanium® processor implementations.</p>
22:0	RO	000000h	Reserved

### 1.19.5 GSTS\_REG - Global Status Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: 1C-1Fh  
 Default Value: 00000000h  
 Access: RO  
 Size: 32 bits  
 Register to report general DMA-remapping hardware status.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31	RO	0b	<p><b>Translation Enable Status (TES)</b></p> <p>This field indicates the status of DMA-remapping hardware.</p> <p>0 = DMA-remapping hardware is not enabled 1 = DMA-remapping hardware is enabled</p>
30	RO	0b	<p><b>Root Table Pointer Status (RTPS)</b></p> <p>This field indicates the status of the root-table pointer in hardware.</p> <p>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.</p>





(Sheet 2 of 2)

Bit	Access	Default Value	Description
29	RO	0b	<b>Fault Log Status (FLS)</b> This field is cleared by hardware when software sets the SFL field in the Global Command register. This field is set by hardware when hardware completes the set fault-log pointer operation using the value provided in the Advanced Fault Log register.
28	RO	0b	<b>Advanced Fault Logging Status (AFLS)</b> This field is valid only for implementations supporting advanced fault logging. This field indicates advanced fault logging status. 0 = Advanced Fault Logging is not enabled 1 = Advanced Fault Logging is enabled
27	RO	0b	<b>Write Buffer Flush Status (WBFS)</b> This bit is valid only for implementations requiring write buffer flushing. This field indicates the status of the write buffer flush operation. This field is set by hardware when software sets the WBF field in the Global Command register. This field is cleared by hardware when hardware completes the write buffer flushing operation.
26	RO	0b	<b>Queued Invalidation Enable Status (QIES)</b> This field indicates queued invalidation enable status. 0 = queued invalidation is not enabled 1 = queued invalidation is enabled
25	RO	0b	<b>Interrupt Remapping Enable Status (IRES)</b> This field indicates the status of Interrupt-remapping hardware. 0 = Interrupt-remapping hardware is not enabled 1 = Interrupt-remapping hardware is enabled
24	RO	0b	<b>Interrupt Remapping Table Pointer Status (IRTPS)</b> 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.
23	RO	0b	<b>Compatibility Format Interrupt Status (CFIS)</b> 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 disabled. 0 = Compatibility format interrupts are blocked. 1 = Compatibility format interrupts are processed as pass-through (bypassing interrupt remapping).
22:0	RO	000000h	<i>Reserved</i>



### 1.19.6 RTADDR\_REG - Root-Entry Table Address Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
Address Offset: 20-27h  
Default Value: 0000000000000000h  
Access: RO; RW  
Size: 64 bits

Register providing the base address of the root-entry table.

Bit	Access	Default Value	Description
63:12	RW	00000000 00000h	<b>Root Table Address (RTA)</b> This register points to base of page aligned, 4-KB-sized root-entry table in system memory. Hardware may ignore and not implement 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.
11:0	RO	000h	<i>Reserved</i>



### 1.19.7 CCMD\_REG - Context Command Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: 28-2Fh  
 Default Value: 0000000000000000h  
 Access: RW-SC; RW; RO; W  
 Size: 64 bits

Register to manage context cache. The act of writing the uppermost byte of the CCMD\_REG with ICC field set causes the hardware to perform the context-cache invalidation.

(Sheet 1 of 3)

Bit	Access	Default Value	Description
63	RW-SC	0b	<p><b>Invalidate Context-Cache (ICC)</b></p> <p>Software requests invalidation of context-cache by setting this field. Software must also set the requested invalidation granularity by programming the CIRG field.</p> <p>Software must read back and check the ICC field to be clear to confirm the invalidation is complete. Software must not update this register when this field is set.</p> <p>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 not submit another invalidation request through this register while the ICC field is set.</p> <p>Software must submit a context cache invalidation request through this field only when there are no invalidation requests pending at this DMA-remapping hardware unit. Refer to Intel VT-d specification Section 11 for software programming requirements.</p> <p>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.</p> <p>Hardware implementations reporting write-buffer flushing requirement (RWBF=1 in Capability register) must implicitly perform a write buffer flush before invalidating the context-cache.</p> <p>Refer to Intel VT-d specification Section 11.1 for write buffer flushing requirements.</p>



(Sheet 2 of 3)

Bit	Access	Default Value	Description
62:61	RW	00b	<p><b>Context Invalidation Request Granularity (CIRG)</b></p> <p>Software provides the requested invalidation granularity through this field when setting the ICC field. Following are the encodings for the CIRG field:</p> <ul style="list-style-type: none"> <li>00: Reserved.</li> <li>01: Global Invalidation request.</li> <li>10: Domain-selective invalidation request. The target domain-id must be specified in the DID field.</li> <li>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.</li> </ul> <p>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.</p>
60:59	RO	00b	<p><b>Context Actual Invalidation Granularity (CAIG)</b></p> <p>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 the CAIG field:</p> <ul style="list-style-type: none"> <li>00: Reserved.</li> <li>01: Global Invalidation performed. This could be in response to a global, domain-selective or device-selective invalidation request.</li> <li>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.</li> <li>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.</li> </ul>
58:34	RO	0000000h	<i>Reserved</i>



(Sheet 3 of 3)

Bit	Access	Default Value	Description
33:32	W	00b	<p><b>Function Mask (FM)</b></p> <p>Software may use the Function Mask to perform device-selective invalidations on behalf of devices supporting PCI Express Phantom Functions.</p> <p>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.</p> <p>The following encodings are defined for this field:</p> <ul style="list-style-type: none"> <li>00: No bits in the SID field masked.</li> <li>01: Mask most significant bit of function number in the SID field.</li> <li>10: Mask two most significant bit of function number in the SID field.</li> <li>11: Mask all three bits of function number in the SID field.</li> </ul> <p>The context-entries corresponding to all the source-ids specified through the FM and SID fields must have the domain-id specified in the DID field.</p> <p>Value returned on read of this field is undefined.</p>
31:16	W	0000h	<p><b>Source ID (SID)</b></p> <p>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.</p> <p>Value returned on read of this field is undefined.</p>
15:0	RW	0000h	<p><b>Domain-ID (DID)</b></p> <p>Indicates the ID of the domain whose context-entries needs to be selectively invalidated. This field must be programmed by software for both domain-selective and device-selective invalidation requests.</p> <p>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.</p> <p>Hardware may ignore and not implement bits 15:N where N is the supported domain-id width reported in the capability register.</p>



### 1.19.8 FSTS\_REG - Fault Status Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: 34-37h  
 Default Value: 00000000h  
 Access: RWC-P; RO-P; RO  
 Size: 32 bits  
 Register indicating the various error status.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31:16	RO	0000h	<i>Reserved</i>
15:8	RO-P	00h	<b>Fault Record Index (FRI)</b> 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.
7	RO	0b	<i>Reserved</i>
6	RWC-P	0b	<b>Invalidation Time-out Error (ITE)</b> 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-IOTLBs implement this bit as RSVD.
5	RWC-P	0b	<b>Invalidation Completion Error (ICE)</b> 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 the Fault Event Control register. Hardware implementations not supporting Device-IOTLBs implement this bit as RSVD.
4	RO	0b	<b>Invalidation Queue Error (IQE)</b> 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 RSVD.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
3	RWC-P	0b	<p><b>Advanced Pending Fault (APF)</b></p> <p>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.</p> <p>Software writing 1 to this field clears it. Hardware implementations not supporting advanced fault logging implement this bit as RSVD.</p>
2	RWC-P	0b	<p><b>Advanced Fault Overflow (AFO)</b></p> <p>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.</p> <p>Software writing 1 to this field clears it. Hardware implementations not supporting advanced fault logging implement this bit as RSVD.</p>
1	RO-P	0b	<p><b>Primary Pending Fault (PPF)</b></p> <p>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 DMA-remapping hardware unit.</p> <p>0 = No pending faults in any of the fault recording registers. 1 = One or more fault recording registers has pending faults.</p> <p>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.</p>
0	RWC-P	0b	<p><b>Primary Fault Overflow (PFO)</b></p> <p>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.</p>



### 1.19.9 FECTL\_REG - Fault Event Control Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
Address Offset: 38-3Bh  
Default Value: 80000000h  
Access: RO; RW  
Size: 32 bits

Register specifying the fault event interrupt message control bits. Intel VT-d specification Section 7.3 describes hardware handling of fault events.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31	RW	1b	<b>Interrupt Mask (IM)</b> 0 = No masking of interrupt. When a interrupt condition is detected, hardware issues an interrupt message (using the Fault Event Data & 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.





(Sheet 2 of 2)

Bit	Access	Default Value	Description
30	RO	0b	<p><b>Interrupt Pending (IP)</b></p> <p>Hardware sets the IP field whenever it detects an interrupt condition, which is defined as:</p> <ul style="list-style-type: none"> <li>• 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 the Fault Status register.</li> <li>• 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.</li> <li>• Hardware detected error associated with the Invalidation Queue, setting the IQE field in the Fault Status register.</li> <li>• Hardware detected invalid Device-IOTLB invalidation completion, setting the ICE field in the Fault Status register.</li> <li>• Hardware detected Device-IOTLB invalidation completion time-out, setting the ITE field in the Fault Status register.</li> </ul> <p>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.</p> <p>The IP field is kept Set by hardware while the interrupt message is held pending. The interrupt message could be held pending due to the interrupt mask (IM field) being Set or other transient hardware conditions.</p> <p>The IP field is cleared by hardware as soon as the interrupt message pending condition is serviced.</p> <p>This could be due to either:</p> <ul style="list-style-type: none"> <li>• Hardware issuing the interrupt message due to either a change in the transient hardware condition that caused the interrupt message to be held pending, or due to software clearing the IM field.</li> <li>• Software servicing all the pending interrupt status fields in the Fault Status register as follows.</li> <li>• 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 the Fault Status register to be evaluated as Clear.</li> <li>• Software clearing other status fields in the Fault Status register by writing back the value read from the respective fields.</li> </ul>
29:0	RO	00000000h	<i>Reserved</i>



### 1.19.10 FEDATA\_REG - Fault Event Data Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
Address Offset: 3C-3Fh  
Default Value: 00000000h  
Access: RO; RW  
Size: 32 bits

Register specifying the interrupt message data.

Bit	Access	Default Value	Description
31:16	RO	0000h	<b>Extended Interrupt Message Data (EIMD)</b> This field is valid only for implementations supporting 32-bit MSI data fields. Hardware implementations supporting only 16-bit MSI data may treat this field as read-only (0).
15:0	RW	0000h	<b>Interrupt message Data (ID)</b> Data value in the interrupt request. Software requirements for programming this register are described in Intel VT-d specification Section 5.7.

### 1.19.11 FEADDR\_REG - Fault Event Address Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
Address Offset: 40-43h  
Default Value: 00000000h  
Access: RO; RW  
Size: 32 bits

Register specifying the interrupt message address.

Bit	Access	Default Value	Description
31:2	RW	00000000h	<b>Message Address (MA)</b> When fault events are enabled, the contents of this register specify the DWORD aligned address (Bits 31:2) for the interrupt request. Software requirements for programming this register are described in Intel VT-d specification Section 5.7.
1:0	RO	00b	<i>Reserved</i>



### 1.19.12 FEUADDR\_REG - Fault Event Upper Address Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: 44-47h  
 Default Value: 00000000h  
 Access: RO  
 Size: 32 bits

Register specifying the interrupt message upper address. This register is treated as RsvdZ by implementations reporting Extended Interrupt Mode (EIM) as not supported in the Extended Capability register.

Bit	Access	Default Value	Description
31:0	RO	00000000h	<b>Message Upper Address (MUA)</b> Hardware implementations supporting Extended Interrupt Mode are required to implement this register. Software requirements for programming this register are described in Intel VT-d specification Section 5.7. Hardware implementations not supporting Extended Interrupt Mode may treat this field as RsvdZ.

### 1.19.13 AFLOG\_REG - Advanced Fault Log Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: 58-5Fh  
 Default Value: 0000000000000000h  
 Access: RO  
 Size: 64 bits

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).

(Sheet 1 of 2)

Bit	Access	Default Value	Description
63:12	RO	0000000000000h	<b>Fault Log Address (FLA)</b> This field specifies the base of 4-KB aligned fault-log region in system memory. Hardware may ignore and 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 returns value that was last programmed to it.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
11:9	RO	000b	<b>Fault Log Size (FLS)</b> This field specifies the size of the fault log region pointed by the FLA field. The size of the fault log region is $(2^X) * 4$ KB, where X is the value programmed in this register. 000: 4 KB 001: 8 KB 010: 16 KB 011: 32 KB 100: 64 KB 101: 128 KB 110: 256 KB 111: 512 KB When implemented, reads of this field returns value that was last programmed to it.
8:0	RO	000h	<i>Reserved</i>

#### 1.19.14 PMEN\_REG - Protected Memory Enable Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
Address Offset: 64-67h  
Default Value: 00000000h  
Access: RO; RW  
Size: 32 bits

Register to enable the DMA-protected memory regions set up through the PLMBASE, PLMLIMIT, PHMBASE, PHMLIMIT registers. This register is 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.



Bit	Access	Default Value	Description
31	RW	0b	<p><b>Enable Protected Memory (EPM)</b></p> <p>This field controls DMA accesses to the protected low-memory and protected high-memory regions.</p> <p>0 = Protected memory regions are disabled. 1 = Protected memory regions are enabled.</p> <p>DMA requests accessing protected memory regions are handled as follows:</p> <ul style="list-style-type: none"> <li>• When DMA remapping is not enabled, all DMA requests accessing protected memory regions are blocked.</li> <li>• When DMA remapping is enabled:                             <ul style="list-style-type: none"> <li>• DMA requests processed as pass-through (Translation Type value of 10b in Context-Entry) and accessing the protected memory regions are blocked.</li> <li>• DMA requests with translated address (AT=10b) and accessing the protected memory regions are blocked.</li> <li>• 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.</li> </ul> </li> </ul> <p>Remapping hardware access to the remapping structures are not subject to protected memory region checks.</p> <p>DMA requests blocked due to protected memory region violation are not recorded or reported as remapping faults.</p> <p>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.</p>
30:1	RO	00000000h	<i>Reserved</i>
0	RO	0b	<p><b>Protected Region Status (PRS)</b></p> <p>This field indicates the status of protected memory region.</p> <p>0 = Protected memory region(s) not enabled. 1 = Protected memory region(s) enabled.</p>



### 1.19.15 PLMBASE\_REG - Protected Low-Memory Base Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
Address Offset: 68-6Bh  
Default Value: 00000000h  
Access: RO; RW  
Size: 32 bits

Register to set up the base address of DMA-protected low-memory region below 4 GB. 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 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 bit position with 0 in the value read back from the register. Bits N:0 of this register are decoded by hardware as all 0's.

Software must setup the protected low memory region below 4 GB. Intel VT-d specification Section 10.4.18 describes the Protected Low-Memory Limit register and hardware decoding of these registers.

Software must not modify this register when protected memory regions are enabled. (PRS field Set in PMEN\_REG).

Bit	Access	Default Value	Description
31:21	RW	000h	<b>Protected Low-Memory Base (PLMB)</b> This register specifies the base of protected low-memory region in system memory.
20:0	RO	000000h	<i>Reserved</i>



### 1.19.16 PLMLIMIT\_REG - Protected Low-Memory Limit Register

B/D/F/Type:	0/0/0/DMIVC1REMAP
Address Offset:	6C-6Fh
Default Value:	00000000h
Access:	RO; RW
Size:	32 bits

Register to setup the limit address of DMA protected low-memory region below 4 GB. This register must be setup 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 0 in the Capability register).

The alignment of the protected low memory region limit depends on the number of reserved bits (N) 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 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 & 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).

Bit	Access	Default Value	Description
31:21	RW	000h	<b>Protected Low-Memory Limit (PLML)</b> This register specifies the last host physical address of the DMA protected low-memory region in system memory.
20:0	RO	000000h	<i>Reserved</i>



### 1.19.17 PHMBASE\_REG - Protected High-Memory Base Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
Address Offset: 70-77h  
Default Value: 0000000000000000h  
Access: RO; RW  
Size: 64 bits

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 4 GB.

Intel VT-d specification Section 10.4.20 describes the Protected High-Memory Limit register and hardware decoding of these registers.

Software must not modify this register when protected memory regions are enabled. (PRS field Set in PMEN\_REG).

Bit	Access	Default Value	Description
63:21	RW	00000000 000h	<b>Protected High-Memory Base (PHMB)</b> 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.
20:0	RO	000000h	<i>Reserved</i>





### 1.19.18 PHMLIMIT\_REG - Protected High-Memory Limit Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: 78-7Fh  
 Default Value: 0000000000000000h  
 Access: RO; RW  
 Size: 64 bits

Register to setup the limit address of DMA protected high-memory region. This register must be setup 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 0 in the Capability register).

The alignment of the protected high memory region limit depends on the number of reserved bits (N) 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).

Bit	Access	Default Value	Description
63:21	RW	00000000 000h	<b>Protected High-Memory Limit (PHML)</b> This register specifies the last host physical address of the DMA protected high-memory region in system memory. Hardware may not utilize Bits 63:HAW, where HAW is the host address width.
20:0	RO	000000h	<i>Reserved</i>



### 1.19.19 IQH\_REG - Invalidation Queue Head Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
Address Offset: 80-87h  
Default Value: 0000000000000000h  
Access: RO  
Size: 64 bits

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.

Bit	Access	Default Value	Description
63:19	RO	00000000 00000h	<i>Reserved</i>
18:4	RO	0000h	<b>Queue Head (QH)</b> Specifies the offset (128-bit aligned) to the invalidation queue for the command that is fetched next by hardware. Hardware resets this field to 0 whenever the queued invalidation is disabled (QIES field Clear in the Global Status register).
3:0	RO	0h	<i>Reserved</i>

### 1.19.20 IQT\_REG - Invalidation Queue Tail Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
Address Offset: 88-8Fh  
Default Value: 0000000000000000h  
Access: RO  
Size: 64 bits

Register indicating the invalidation queue tail. This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

Bit	Access	Default Value	Description
63:19	RO	00000000 00000h	<i>Reserved</i>
18:4	RO	0000h	<b>Queue Tail (QT)</b> Specifies the offset (128-bit aligned) to the invalidation queue for the command that is written next by software.
3:0	RO	0h	<i>Reserved</i>



### 1.19.21 IQA\_REG - Invalidation Queue Address Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: 90-97h  
 Default Value: 0000000000000000h  
 Access: RO  
 Size: 64 bits

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.

When supported, writing to this register causes the Invalidation Queue Head and Invalidation Queue Tail registers to be reset to 0h.

Bit	Access	Default Value	Description
63:12	RO	00000000 00000h	<b>Invalidation Queue Base Address (IQA)</b> This field points to the base of 4-KB 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.
11:3	RO	000h	<i>Reserved</i>
2:0	RO	000b	<b>Queue Size (QS)</b> This field specifies the size of the invalidation request queue. A value of X in this field indicates an invalidation request queue of $(2^X)$ 4-KB pages. The number of entries in the invalidation queue is $2^{(X + 8)}$ .



### 1.19.22 ICS\_REG - Invalidation Completion Status Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
Address Offset: 9C-9Fh  
Default Value: 00000000h  
Access: RO  
Size: 32 bits

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.

Bit	Access	Default Value	Description
31:1	RO	00000000h	<i>Reserved</i>
0	RO	0b	<b>Invalidation Wait Descriptor Complete (IWC)</b> Indicates completion of Invalidation Wait Descriptor with Interrupt Flag (IF) field Set. Hardware implementations not supporting queued invalidations implement this field as Rsvd.

### 1.19.23 IECTL\_REG - Invalidation Event Control Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
Address Offset: A0-A3h  
Default Value: 00000000h  
Access: RO  
Size: 32 bits

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.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31	RO	0b	<b>Interrupt Mask (IM)</b> 0 = No masking of interrupt. When a invalidation event condition is detected, hardware issues an interrupt message (using the Invalidation Event Data & Invalidation Event Address register values). 1 = When implemented, 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.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
30	RO	0b	<p><b>Interrupt Pending (IP)</b></p> <p>Hardware sets the IP field whenever it detects an interrupt condition. Interrupt condition is defined as:</p> <ul style="list-style-type: none"> <li>An Invalidation Wait Descriptor with Interrupt Flag (IF) field Set completed, setting the IWC field in the Invalidation Completion Status register.</li> <li>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.</li> </ul> <p>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:</p> <ul style="list-style-type: none"> <li>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.</li> <li>Software servicing the IWC field in the Invalidation Completion Status register.</li> </ul>
29:0	RO	00000000h	Reserved

### 1.19.24 IEDATA\_REG - Invalidation Event Data Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: A4-A7h  
 Default Value: 00000000h  
 Access: RO  
 Size: 32 bits

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.

Bit	Access	Default Value	Description
31:16	RO	0000h	<p><b>Extended Interrupt Message Data (EIMD)</b></p> <p>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.</p>
15:0	RO	0000h	<p><b>Interrupt Message Data (IMD)</b></p> <p>Data value in the interrupt request. Software requirements for programming this register are described in Intel VT-d specification Section 5.7.</p>



### 1.19.25 IEADDR\_REG - Invalidation Event Address Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
Address Offset: A8-ABh  
Default Value: 00000000h  
Access: RO  
Size: 32 bits

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.

Bit	Access	Default Value	Description
31:2	RO	00000000h	<b>Message Address (MA)</b> When fault events are enabled, the contents of this register specify the DWORD-aligned address (Bits 31:2) for the interrupt request. Software requirements for programming this register are described in Intel VT-d specification Section 5.7.
1:0	RO	00b	<i>Reserved</i>

### 1.19.26 IEUADDR\_REG - Invalidation Event Upper Address Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
Address Offset: AC-AFh  
Default Value: 00000000h  
Access: RO  
Size: 32 bits

Register specifying the Invalidation Event interrupt message upper address. This register is treated as RsvdZ by implementations reporting both Queued Invalidation (QI) and Extended Interrupt Mode (EIM) as not supported in the Extended Capability register.

Bit	Access	Default Value	Description
31:0	RO	00000000h	<b>Message Upper Address (MUA)</b> Hardware implementations supporting Queued Invalidation and Extended Interrupt Mode are required to implement this register. Software requirements for programming this register are described in Section 5.7. Hardware implementations not supporting Queued Invalidation and Extended Interrupt Mode may treat this field as RSVD.



### 1.19.27 IRTA\_REG - Interrupt Remapping Table Address Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: B8-BFh  
 Default Value: 0000000000000000h  
 Access: RO  
 Size: 64 bits

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.

Bit	Access	Default Value	Description
63:12	RO	00000000 000000h	<b>Interrupt Remapping Table Address (IRTA)</b> This field points to the base of the 4-KB aligned interrupt remapping table. Hardware ignores and not implement Bits 63:HAW, where HAW is the width. Reads of this field returns last programmed to it.
11	RO	0b	<b>Extended Interrupt Mode Enable (EIMI)</b> 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. On Itanium platforms hardware interprets low 16-bits of Destination-ID field in the IRTEs and treats the high 16-bits as reserved. 1 = x2APIC mode is active. Hardware interprets all 32-bits of the Destination-ID field in the IRTEs. Hardware reporting Extended Interrupt Mode (EIM) as Clear in the Capability register treats this field as RsvdZ.
10:4	RO	00h	<i>Reserved</i>
3:0	RO	0h	<b>Size (S)</b> 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.



### 1.19.28 IVA\_REG - Invalidate Address Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: 100-107h  
 Default Value: 0000000000000000h  
 Access: W;RO  
 Size: 64 bits

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. Value returned on reads of this register is undefined.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
63:12	W	00000000 00000h	<p><b>Address (ADDR)</b></p> <p>Software provides the DMA address that needs to be page-selectively invalidated. To request a page-selective invalidation request to hardware, software must first write the appropriate fields in this register, and then issue appropriate page-selective invalidate command through the IOTLB_REG.</p> <p>Hardware ignores Bits 63:N, where N is the maximum guest address width (MGAW) supported.</p> <p>Value returned on read of this field is undefined.</p>
11:7	RO	00h	<i>Reserved</i>
6	W	0b	<p><b>Invalidation Hint (IH)</b></p> <p>The field provides hint to hardware to preserve or flush the non-leaf (page-directory) entries that may be cached in hardware.</p> <p>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 to mappings specified by ADDR and AM fields.</p> <p>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.</p> <p>Value returned on read of this field is undefined.</p>





(Sheet 2 of 2)

Bit	Access	Default Value	Description																														
5:0	W	00h	<p><b>Address Mask (AM)</b></p> <p>The value in this field specifies the number of low order bits of the ADDR field that must be masked for the invalidation operation. Mask field enables software to request invalidation of contiguous mappings for size-aligned regions. For example:</p> <p>Mask Value ADDR bits masked Pages invalidated:</p> <table border="1"> <thead> <tr> <th>Mask Value</th> <th>ADDR Bits</th> <th>Pages Invalidated</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>None</td> <td>1</td> </tr> <tr> <td>1</td> <td>12</td> <td>2</td> </tr> <tr> <td>2</td> <td>13:12</td> <td>8</td> </tr> <tr> <td>3</td> <td>14:12</td> <td>16</td> </tr> <tr> <td>4</td> <td>15:12</td> <td>32</td> </tr> <tr> <td>5</td> <td>16:12</td> <td>64</td> </tr> <tr> <td>6</td> <td>17:12</td> <td>128</td> </tr> <tr> <td>7</td> <td>18:12</td> <td>256</td> </tr> <tr> <td>8</td> <td>19:12</td> <td>512</td> </tr> </tbody> </table> <p>Hardware implementations report the maximum supported mask value through the Capability register. Value returned on read of this field is undefined.</p>	Mask Value	ADDR Bits	Pages Invalidated	0	None	1	1	12	2	2	13:12	8	3	14:12	16	4	15:12	32	5	16:12	64	6	17:12	128	7	18:12	256	8	19:12	512
Mask Value	ADDR Bits	Pages Invalidated																															
0	None	1																															
1	12	2																															
2	13:12	8																															
3	14:12	16																															
4	15:12	32																															
5	16:12	64																															
6	17:12	128																															
7	18:12	256																															
8	19:12	512																															



### 1.19.29 IOTLB\_REG - IOTLB Invalidate Register

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: 108-10Fh  
 Default Value: 0000000000000000h  
 Access: RO; RW; RW-SC  
 Size: 64 bits

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.

(Sheet 1 of 3)

Bit	Access	Default Value	Description
63	RW-SC	0b	<p><b>Invalidate IOTLB (IVT)</b></p> <p>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.</p> <p>Software must not submit IOTLB invalidation requests when there is a context-cache invalidation request pending at this DMA-remapping hardware unit.</p> <p>Refer to Section 11 for software programming requirements.</p> <p>Hardware implementations reporting write-buffer flushing requirement (RWBF=1 in Capability register) must implicitly perform a write buffer flush before invalidating the IOTLB.</p> <p>Refer to Section 11.1 for write buffer flushing requirements.</p>



(Sheet 2 of 3)

Bit	Access	Default Value	Description
62:60	RW	000b	<p><b>IOTLB Invalidation Request Granularity (IIRG)</b></p> <p>When requesting hardware to invalidate the IOTLB (by setting the IVT field), software writes the requested invalidation granularity through this IIRG field. Following are the encodings for the IIRG field.</p> <p>000: Reserved.</p> <p>001: Global invalidation request.</p> <p>010: Domain-selective invalidation request. The target domain-id must be specified in the DID field.</p> <p>011: Domain-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.</p> <p>100 - 111: Reserved.</p> <p>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.</p>
59:57	RO	000b	<p><b>IOTLB Actual Invalidation Granularity (IAIG)</b></p> <p>Hardware reports the granularity at which an invalidation request was processed through this field at the time of reporting invalidation completion (by clearing the IVT field). The following are the encodings for the IAIG field.</p> <p>000: 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.</p> <p>001: Global Invalidation performed. This could be in response to a global, domain-selective, or page-selective invalidation request.</p> <p>010: 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 page-selective invalidation request.</p> <p>011: 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 domain-page-selective invalidation request.</p> <p>100 - 111: Reserved.</p>
56:50	RO	00h	<i>Reserved</i>



(Sheet 3 of 3)

Bit	Access	Default Value	Description
49	RW	0b	<p><b>Drain Reads (DR)</b></p> <p>This field is ignored by hardware if the DRD field is reported as clear in the Capability register.</p> <p>When DRD field is reported as set in the Capability register, the following encodings are supported for this field:</p> <p>0 = Hardware may complete the IOTLB invalidation without draining any translated DMA reads that are queued in the root-complex and yet to be processed.</p> <p>1 = Hardware must drain all/relevant translated DMA reads that are queued in the root-complex before indicating IOTLB invalidation completion to software.</p> <p>Refer to Intel VT-d specification section 6.3 for description of DMA draining.</p>
48	RW	0b	<p><b>Drain Writes (DW)</b></p> <p>This field is ignored by hardware if the DWD field is reported as clear in the Capability register.</p> <p>When DWD field is reported as set in the Capability register, the following encodings are supported for this field:</p> <p>0 = Hardware may complete the IOTLB invalidation without draining any translated DMA writes that are queued in the root-complex for processing.</p> <p>1 = Hardware must drain all/relevant translated DMA writes that are queued in the root-complex before indicating IOTLB invalidation completion to software.</p> <p>Refer to Intel VT-d specification section 6.3 for description of DMA draining.</p>
47:32	RW	0000h	<p><b>Domain-ID (DID)</b></p> <p>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 domain-page-selective invalidation requests.</p> <p>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.</p> <p>Hardware may ignore and not implement bits 47:(32+N) where N is the supported domain-id width reported in the capability register.</p>
31:0	RO	00000000h	<i>Reserved</i>



### 1.19.30 FRCD\_REG - Fault Recording Registers

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: 200-20Fh  
 Default Value: 00000000000000000000000000000000h  
 Access: RWC-P; RO-P; RO  
 Size: 128 bits

Registers 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.

These registers are sticky and can be cleared only through powergood reset or via software clearing the RWC fields by writing a 1.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
127	RWC-P	0b	<b>Fault (F)</b> 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. Refer to Intel VT-d specification Section 7.2.1 for hardware details of primary fault logging.
126	RO-P	0b	<b>Type (T)</b> Type of the faulted request: 0 = Write request 1 = Read 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.
125:124	RO-P	00b	<b>Address Type (AT)</b> 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.
123:104	RO	00000h	<i>Reserved</i>
103:96	RO-P	00h	<b>Fault Reason (FR)</b> Reason for the fault. Intel VT-d specification Appendix A enumerates the various translation fault reason encodings. This field is relevant only when the F field is set.
95:80	RO	0000h	<i>Reserved</i>



(Sheet 2 of 2)

Bit	Access	Default Value	Description
79:64	RO-P	0000h	<b>Source Identifier (SID)</b> Requester-id associated with the fault condition. This field is relevant only when the F field is Set.
63:12	RO-P	00000000 00000h	<b>Fault Info (FI)</b> When the Fault Reason (FR) field indicates one of the DMA-remapping fault conditions, Bits 63:12 of this field contains the page address in the faulted DMA request. Hardware treat 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.
11:0	RO	000h	<i>Reserved</i>

### 1.19.31 VTPOLICY - DMA Remap Engine Policy Control

B/D/F/Type: 0/0/0/DMIVC1REMAP  
 Address Offset: FFC-FFFh  
 Default Value: 00000000h  
 Access: RO; RW-L-K; RW-L  
 Size: 32 bits

This registers contains all the policy bits related to the DMA remap engine.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31	RW-L-K	0b	<b>DMA Remap Engine Policy Lock-Down (DMAR_LCKDN)</b> 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 is read-only. This bit can only be clear through platform reset.
30:5	RO	0000000h	<i>Reserved</i>
4	RW-L	0b	<b>TLB Lookup Policy TLB Invalidation (LKUPPTLBINV)</b> DMI Intel High Definition Audio Remap Engine TLB Lookup Policy On TLB Invalidation: 0 = Continue to perform TLB lookup to DMI Intel® High Definition Audio remap engine during TLB Invalidation Window. 1 = Mask all TLB Lookup to DMI Intel High Definition Audio remap engine during TLB Invalidation Window. TLB Invalidation Window refers to the period from when the TLB Invalidation is initiated until all the outstanding DMA read and write cycles at the point of TLB Invalidation are initiated are Globally Ordered.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
3	RW-L	0b	<b>DMI VC1 Hit Queue Throttling (DMIVC1HTQT)</b> 0 = No throttling at the outlet of the DMI VC1 Hit Queue. 1 = Throttle the outlet DMI VC1 Hit Queue to fill up the queue.
2	RW-L	0b	<b>DMIVC1 TLB Disable (DMIVC1TLBDIS)</b> 0 = Normal mode, DMIVC1 TLBs are enabled and normal hit/miss flows are followed. 1 = DMIVC1 TLBs are disabled and each GPA request will result in a miss and a root walk is requested from Intel VT-d Dispatcher.
1	RW-L	0b	<b>Global IOTLB Invalidation Promotion (GLBIOTLBINV)</b> This bit controls the IOTLB Invalidation behavior of the DMA remap engine. 0 = Normal operation. 1 = Any type of IOTLB Invalidation (valid or invalid) is promoted to Global IOTLB Invalidation.
0	RW-L	0b	<b>Global Context Invalidation Promotion (GLBCTXINV)</b> This bit controls the Context Invalidation behavior of the DMA remap engine. 0 = Normal operation. 1 = Any type of Context Invalidation (valid or invalid) is promoted to Global Context Invalidation.



## 1.20 GFXVTBAR

(Sheet 1 of 2)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Version Register	VER_REG	0	3	00000010h	RO
Capability Register	CAP_REG	8	F	00C0000020230272h	RO
Extended Capability Register	ECAP_REG	10	17	0000000000001000h	RO
Global Command Register	GCMD_REG	18	1B	00000000h	W; RO; RW
Global Status Register	GSTS_REG	1C	1F	00000000h	RO
Root-Entry Table Address Register	RTADDR_REG	20	27	0000000000000000h	RO; RW
Context Command Register	CCMD_REG	28	2F	0800000000000000h	RO; RW
Fault Status Register	FSTS_REG	34	37	00000000h	RO; RWC-P; RO-P
Fault Event Control Register	FECTL_REG	38	3B	80000000h	RO; RW
Fault Event Data Register	FEDATA_REG	3C	3F	00000000h	RO; RW
Fault Event Address Register	FEADDR_REG	40	43	00000000h	RO; RW
Fault Event Upper Address Register	FEUADDR_REG	44	47	00000000h	RO
Advanced Fault Log Register	AFLOG_REG	58	5F	0000000000000000h	RO
Protected Memory Enable Register	PMEN_REG	64	67	00000000h	RO; RW
Protected Low Memory Base Register	PLMBASE_REG	68	6B	00000000h	RO; RW
Protected Low Memory Limit Register	PLMLIMIT_REG	6C	6F	00000000h	RO; RW
Protected High Memory Base Register	PHMBASE_REG	70	77	0000000000000000h	RO; RW
Protected High Memory Limit Register	PHMLIMIT_REG	78	7F	0000000000000000h	RO; RW
Invalidation Queue Head	IQH_REG	80	87	0000000000000000h	RO
Invalidation Queue Tail	IQT_REG	88	8F	0000000000000000h	RO





(Sheet 2 of 2)

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
Invalidation Queue Address	IQA_REG	90	97	0000000000000000h	RO
Invalidation Completion Status	ICS_REG	9C	9F	00000000h	RO
Invalidation Completion Event Control	IECTL_REG	A0	A3	80000000h	RO
Invalidation Completion Event Data	IEDATA_REG	A4	A7	00000000h	RO
Invalidation Completion Event Address	IEADDR_REG	A8	AB	00000000h	RO
Invalidation Completion Event Upper Address	IEUADDR_REG	AC	AF	00000000h	RO
Interrupt Remapping Table Address	IRTA_REG	B8	BF	0000000000000000h	RO
Invalidate Address Register	IVA_REG	100	107	0000000000000000h	RO
IOTLB Invalidate Register	IOTLB_REG	108	10F	0200000000000000h	RO; RW
Fault Recording Registers	FRCD_REG	200	20F	0000000000000000 0000000000000000h	RO-P; RO; RWC-P
VT-d Policy	VTPOLICY	FFC	FFF	40000000h	RW-L; RW-O; RO



### 1.20.1 VER\_REG - Version Register

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 0-3h  
 Default Value: 00000010h  
 Access: RO  
 Size: 32 bits

Register to report the architecture version supported. Backward compatibility for the architecture is maintained with new revision numbers, allowing software to load DMA-remapping drivers written for prior architecture versions.

Bit	Access	Default Value	Description
31:8	RO	000000h	<i>Reserved</i>
7:4	RO	1h	<b>Major Version Number (MAX)</b> Indicates supported architecture version.
3:0	RO	0h	<b>Minor Version Number (MIN)</b> Indicates supported architecture minor version.

### 1.20.2 CAP\_REG - Capability Register

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 8-Fh  
 Default Value: 00C0000020230272h  
 Access: RO  
 Size: 64 bits

Register to report general DMA remapping hardware capabilities.

(Sheet 1 of 5)

Bit	Access	Default Value	Description
63:56	RO	00h	<i>Reserved</i>
55	RO	1b	<b>DMA Read Draining (DRD)</b> 0 = On IOTLB invalidations, hardware does not support draining of translated DMA read requests queued within the root complex. 1 = On IOTLB invalidations, hardware supports draining of translated DMA read requests queued within the root complex. Indicates supported architecture version.
54	RO	1b	<b>DMA Write Draining (DWD)</b> 0 = On IOTLB invalidations, hardware does not support draining of translated DMA writes queued within the root complex. 1 = On IOTLB invalidations, hardware supports draining of translated DMA writes queued within the root complex.



(Sheet 2 of 5)

Bit	Access	Default Value	Description
53:48	RO	00h	<p><b>Maximum Address Mask Value (MAMV)</b></p> <p>The value in this field indicates the maximum supported value for the Address Mask (AM) field in the Invalidation Address (IVA_REG) register.</p>
47:40	RO	00h	<p><b>Number of Fault-recording Registers (NFR)</b></p> <p>Number of fault recording registers is computed as <math>N+1</math>, where <math>N</math> is the value reported in this field.</p> <p>Implementations must support at least one fault recording register (<math>NFR = 0</math>) for each DMA remapping hardware unit in the platform.</p> <p>The maximum number of fault recording registers per DMA-remapping hardware unit is 256.</p>
39	RO	0b	<p><b>Page-Selective Invalidation Support (PSI)</b></p> <p>0 = Hardware supports only domain and global invalidates for IOTLB.</p> <p>1 = Hardware supports page selective, domain, and global invalidates for IOTLB and hardware must support a minimum MAMV value of 9.</p>
38	RO	0b	<i>Reserved</i>
37:34	RO	0h	<p><b>Super-Page Support (SPS)</b></p> <p>This field indicates the super page sizes supported by hardware.</p> <p>A value of 1 in any of these bits indicates the corresponding super-page size is supported.</p> <p>The super-page sizes corresponding to various bit positions within this field are:</p> <p>0: 21-bit offset to page frame (2 MB)</p> <p>1: 30-bit offset to page frame (1 GB)</p> <p>2: 39-bit offset to page frame (512 GB)</p> <p>3: 48-bit offset to page frame (1 TB)</p>
33:24	RO	020h	<p><b>Fault-recording Register Offset (FRO)</b></p> <p>This field specifies the location to the first fault recording register relative to the register base address of this DMA-remapping hardware unit.</p> <p>If the register base address is <math>X</math>, and the value reported in this field is <math>Y</math>, the address for the first fault recording register is calculated as <math>X+(16*Y)</math>.</p>
23	RO	0b	<p><b>Isochrony (ISOCH)</b></p> <p>0 = Indicates this DMA-remapping hardware unit has no critical isochronous requesters in its scope.</p> <p>1 = Indicates this DMA-remapping hardware unit has one or more critical isochronous requesters in its scope. To guarantee isochronous performance, software must ensure invalidation operations do not impact active DMA streams from such requesters. This implies that when DMA is active, software perform page-selective invalidations (instead of coarser invalidations).</p>



(Sheet 3 of 5)

Bit	Access	Default Value	Description
22	RO	0b	<p><b>Zero Length Read (ZLR)</b></p> <p>0 = Indicates the remapping hardware unit blocks (and treats as fault) zero length DMA read requests to write-only pages.</p> <p>1 = Indicates the remapping hardware unit supports zero length DMA read requests to write-only pages.</p>
21:16	RO	23h	<p><b>Maximum Guest Address Width (MGAW)</b></p> <p>This field indicates the maximum DMA virtual addressability supported by remapping hardware.</p> <p>The Maximum Guest Address Width (MGAW) is computed as <math>(N+1)</math>, 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.</p> <p>If the value in this field is X, not translated and translated DMA requests to addresses above <math>2^{(X+1)} - 1</math> are always blocked by hardware. Translation requests to address above <math>2^{(X+1)} - 1</math> from allowed devices return a null Translation Completion Data Entry with R=W=0.</p> <p>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).</p>
15:13	RO	000b	<i>Reserved</i>
12:8	RO	02h	<p><b>Supported Adjusted Guest Address Width (SAGAW)</b></p> <p>This 5-bit field indicates the supported adjusted guest address widths (which in turn represents the levels of page-table walks for the 4-KB base page size) supported by the hardware implementation.</p> <p>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:</p> <ul style="list-style-type: none"> <li>0: 30-bit AGAW (2-level page-table)</li> <li>1: 39-bit AGAW (3-level page-table)</li> <li>3: 57-bit AGAW (5-level page-table)</li> <li>4: 64-bit AGAW (6-level page-table)</li> </ul> <p>Software must ensure that the adjusted guest address width used to set up the page tables is one of the supported guest address widths reported in this field.</p>



(Sheet 4 of 5)

Bit	Access	Default Value	Description
7	RO	0b	<p><b>Caching Mode (CM)</b></p> <p>0 = Not-present and erroneous entries are not cached in any of the remapping caches. Invalidation is 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.</p> <p>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.</p> <p>Hardware implementations of this architecture must support a value of 0 in this field. Refer to Section 6.1 for more details on Caching Mode.</p>
6	RO	1b	<p><b>Protected High-Memory Region (PHMR)</b></p> <p>0 = Indicates protected high-memory region is not supported.</p> <p>1 = Indicates protected high-memory region is supported.</p> <p>DMA-remapping hardware implementations on Intel VT-d platforms supporting main memory above 4 GB are required to support protected high-memory region.</p>
5	RO	1b	<p><b>Protected Low-Memory Region (PLMR)</b></p> <p>0 = Indicates protected low-memory region is not supported.</p> <p>1 = Indicates protected low-memory region is supported.</p> <p>DMA-remapping hardware implementations on Intel TXT platforms are required to support protected low-memory region.</p>
4	RO	1b	<p><b>Required Write-Buffer Flushing (RWBF)</b></p> <p>0 = Indicates no write-buffer flushing is needed to ensure changes to memory-resident structures are visible to hardware.</p> <p>1 = Indicates software must explicitly flush the write buffers to ensure updates made to memory-resident remapping structures are visible to hardware. Refer to Section 11.1 for more details on write buffer flushing requirements.</p>
3	RO	0b	<p><b>Advanced Fault Logging (AFL)</b></p> <p>0 = Indicates advanced fault logging is not supported. Only primary fault logging is supported.</p> <p>1 = Indicates advanced fault logging is supported.</p>



(Sheet 5 of 5)

Bit	Access	Default Value	Description																		
2:0	RO	010b	<p><b>Number of Domains Supported (ND)</b></p> <p>000b: Hardware supports 4-bit domain-IDs with support for up to 16 domains.</p> <p>001b: Hardware supports 6-bit domain-IDs with support for up to 64 domains.</p> <p>010b: Hardware supports 8-bit domain-IDs with support for up to 256 domains.</p> <p>011b: Hardware supports 10-bit domain-IDs with support for up to 1024 domains.</p> <p>100b: Hardware supports 12-bit domain-IDs with support for up to 4K domains.</p> <p>100b: Hardware supports 14-bit domain-IDs with support for up to 16K domains.</p> <p>110b: Hardware supports 16-bit domain-IDs with support for up to 64K domains.</p> <p>111b: Reserved.</p> <table border="1" data-bbox="737 924 1235 1293"> <thead> <tr> <th>Encoding</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000</td> <td>4-bit IDs</td> </tr> <tr> <td>010</td> <td>8-bit IDs</td> </tr> <tr> <td>011</td> <td>10-bit IDs</td> </tr> <tr> <td>100</td> <td>12-bit IDs</td> </tr> <tr> <td>101</td> <td>14-bit IDs</td> </tr> <tr> <td>110</td> <td>16-bit IDs</td> </tr> <tr> <td>111</td> <td>Reserved</td> </tr> <tr> <td>001</td> <td>6-bit IDs</td> </tr> </tbody> </table>	Encoding	Description	000	4-bit IDs	010	8-bit IDs	011	10-bit IDs	100	12-bit IDs	101	14-bit IDs	110	16-bit IDs	111	Reserved	001	6-bit IDs
Encoding	Description																				
000	4-bit IDs																				
010	8-bit IDs																				
011	10-bit IDs																				
100	12-bit IDs																				
101	14-bit IDs																				
110	16-bit IDs																				
111	Reserved																				
001	6-bit IDs																				



### 1.20.3 ECAP\_REG - Extended Capability Register

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 10-17h  
 Default Value: 0000000000001000h  
 Access: RO  
 Size: 64 bits

Register to report DMA-remapping hardware extended capabilities.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
63:24	RO	00000000 00h	<i>Reserved</i>
23:20	RO	0h	<b>Maximum Handle Mask Value (MHMV)</b> 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 is reported as Set.
19:18	RO	00b	<i>Reserved</i>
17:8	RO	010h	<b>Invalidation Unit Offset (IVO)</b> This field specifies the offset to the IOTLB invalidation 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 IOTLB invalidation register is calculated as X+(16*Y).
7	RO	0b	<b>Snoop Control (SC)</b> 0 = Hardware does not support setting the SNP field to 1 in the page-table entries. 1 = Hardware supports setting the SNP field to 1 in the page-table entries.
6	RO	0b	<b>Pass Through (PT)</b> 0 = Hardware does not support pass through translation type in context entries. 1 = Hardware supports pass-through translation type in context entries.
5	RO	0b	<b>Caching Hints (CH)</b> 0 = Hardware does not support IOTLB caching hints (ALH and EH fields in context-entries are treated as reserved). 1 = Hardware supports IOLTb caching hints through the ALH and EH fields in context-entries.
4	RO	0b	<b>Extended Interrupt Mode (EIM)</b> 0 = Hardware supports only 8-bit APICIDs (Legacy Interrupt Mode) on Intel®64 and IA-32 platforms and 16-bit APIC-IDs on Itanium platforms. 1 = Hardware supports Extended Interrupt Mode (32-bit APIC-IDs) on Intel®64 platforms. This field is valid only when the IR field is reported as Set.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
3	RO	0b	<b>Interrupt Remapping (IR)</b> 0 = Hardware does not support interrupt remapping. 1 = Hardware supports interrupt remapping. Implementations reporting this field as Set must also support Queued Invalidation (QI = 1b).
2	RO	0b	<b>Device IOTLB Support (DI)</b> 0 = Hardware does not support device-IOTLBs. 1 = Hardware supports Device-IOTLBs. Implementations reporting this field as Set must also support Queued Invalidation (QI = 1b).
1	RO	0b	<b>Queued Invalidation (QI)</b> 0 = Hardware does not support queued invalidations. 1 = Hardware supports queued invalidations.
0	RO	0b	<b>Coherency (C)</b> 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 incoherent. 1 = Indicates hardware accesses to remapping structures are coherent. Hardware access to advanced fault log and invalidation queue is always coherent.





### 1.20.4 GCMD\_REG - Global Command Register

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 18-1Bh  
 Default Value: 00000000h  
 Access: W; RO; RW  
 Size: 32 bits

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.

(Sheet 1 of 5)

Bit	Access	Default Value	Description
31	RW	0b	<p><b>Translation Enable (TE)</b></p> <p>Software writes to this field to request hardware to enable/disable DMA-remapping hardware.</p> <p>0 = Disable DMA-remapping hardware                      1 = Enable DMA-remapping hardware</p> <p>Hardware reports the status of the translation enable operation through the TES field in the Global Status register. Before enabling (or re-enabling) DMA-remapping hardware through this field, software must:</p> <ul style="list-style-type: none"> <li>• Setup the DMA-remapping structures in memory</li> <li>• Flush the write buffers (through WBF field), if write buffer flushing is reported as required.</li> <li>• Set the root-entry table pointer in hardware (through S RTP field).</li> <li>• Perform global invalidation of the context-cache and global invalidation of IOTLB.</li> <li>• If advanced fault logging supported, setup fault log pointer (through SFL field) and enable advanced fault logging (through EAFL field).</li> </ul> <p>Refer to Section 9 for detailed software requirements.</p> <p>There may be active DMA requests in the platform when software updates this field. Hardware must enable or disable remapping logic only at deterministic transaction boundaries, so that any in-flight transaction is either subject to remapping or not at all.</p> <p>Hardware implementations supporting DMA draining must drain any in-flight translated 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 GSTS_REG.</p> <p>Value returned on read of this field is undefined.</p>



(Sheet 2 of 5)

Bit	Access	Default Value	Description
30	W	0b	<p><b>Set Root Table Pointer (SRTP)</b></p> <p>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 register. Hardware reports the status of the “root table pointer set” operation through the RTPS field in the Global Status register.</p> <p>The root table pointer set operation must be performed before enabling or re-enabling (after disabling) DMA remapping through the TE field.</p> <p>After a “root table pointer set” operation, software must globally invalidate the context cache and then globally invalidate the IOTLB. This is required to ensure hardware uses only the remapping structures referenced by the new root table pointer, and not any stale cached entries.</p> <p>While DMA remapping 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.</p> <p>Clearing this bit has no effect. The value returned on a read of this field is undefined.</p>
29	RO	0b	<p><b>Set Fault Log (SFL)</b></p> <p>This field is valid only for implementations supporting advanced fault logging.</p> <p>Software sets this field to request hardware to set/update the fault-log pointer used by hardware.</p> <p>The fault-log pointer is specified through Advanced Fault Log register.</p> <p>Hardware reports the status of the fault log set operation through the FLS field in the Global Status register.</p> <p>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.</p> <p>Clearing this bit has no effect.</p> <p>The value returned on read of this field is undefined.</p>



(Sheet 3 of 5)

Bit	Access	Default Value	Description
28	RO	0b	<p><b>Enable Advanced Fault Logging (EAFL)</b></p> <p>This field is valid only for implementations supporting advanced fault logging.</p> <p>Software writes to this field to request hardware to enable or disable advanced fault logging.</p> <p>0 = Disable advanced fault logging. In this case, translation faults are reported through the Fault Recording registers.</p> <p>1 = Enable use of memory-resident fault log.</p> <p>When enabled, translation faults are recorded in the memory-resident log. The fault log pointer must be set in hardware (through SFL field) before enabling advanced fault logging.</p> <p>Hardware reports the status of the advanced fault logging enable operation through the AFLS field in the Global Status register.</p> <p>Value returned on read of this field is undefined.</p>
27	W	0b	<p><b>Write Buffer Flush (WBF)</b></p> <p>This bit is valid only for implementations requiring write buffer flushing.</p> <p>Software sets this field to request hardware to 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.</p> <p>Refer to Section 11.1 for details on write-buffer flushing requirements.</p> <p>Hardware reports the status of the write buffer flushing operation through the WBFS field in the Global Status register.</p> <p>Clearing this bit has no effect.</p> <p>Value returned on read of this field is undefined.</p>
26	RO	0b	<p><b>Queued Invalidation Enable (QIE)</b></p> <p>This field is valid only for implementations supporting queued invalidations.</p> <p>Software writes to this field to enable or disable queued invalidations.</p> <p>0 = Disable queued invalidations.</p> <p>1 = Enable use of queued invalidations.</p> <p>Hardware reports the status of queued invalidation enable operation through QIES field in the Global Status register.</p> <p>Refer to Section 6.2.2 for software requirements for enabling/disabling queued invalidations.</p> <p>The value returned on a read of this field is undefined.</p>



(Sheet 4 of 5)

Bit	Access	Default Value	Description
25	RO	0b	<p><b>Interrupt Remapping Enable (IRE)</b></p> <p>This field is valid only for implementations supporting interrupt remapping.</p> <p>0 = Disable interrupt-remapping hardware 1 = Enable interrupt-remapping hardware</p> <p>Hardware reports the status of the interrupt remapping enable operation through the IRES field in the Global Status register.</p> <p>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.</p> <p>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.</p> <p>The value returned on a read of this field is undefined.</p>
24	RO	0b	<p><b>Set Interrupt Remap Table Pointer (SIRTP)</b></p> <p>This field is valid only for implementations supporting interrupt-remapping.</p> <p>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 register.</p> <p>Hardware reports the status of the interrupt remapping table pointer set operation through the IRTPS field in the Global Status register.</p> <p>The interrupt remap table pointer set operation must be performed before enabling or re-enabling (after disabling) interrupt-remapping hardware through the IRE field.</p> <p>After an interrupt remap table pointer set 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.</p> <p>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.</p> <p>Clearing this bit has no effect. The value returned on a read of this field is undefined.</p>



(Sheet 5 of 5)

Bit	Access	Default Value	Description
23	RO	0b	<p><b>Compatibility Format Interrupt (CFI)</b></p> <p>This field is valid only for Intel®64 implementations supporting interrupt-remapping.</p> <p>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 Legacy Interrupt Mode is active.</p> <p>0 = Block Compatibility format interrupts. 1 = Process Compatibility format interrupts as pass-through (bypass interrupt remapping).</p> <p>Hardware reports the status of updating this field through the CFIS field in the Global Status register.</p> <p>The value returned on a read of this field is undefined.</p> <p>This field is not implemented on Itanium implementations.</p>
22:0	RO	000000h	<i>Reserved</i>

### 1.20.5 GSTS\_REG - Global Status Register

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 1C-1Fh  
 Default Value: 00000000h  
 Access: RO  
 Size: 32 bits  
 Register to report general remapping hardware status.

(Sheet 1 of 3)

Bit	Access	Default Value	Description
31	RO	0b	<p><b>Translation Enable Status (TES)</b></p> <p>This field indicates the status of DMA-remapping hardware.</p> <p>0 = DMA-remapping hardware is not enabled 1 = DMA-remapping hardware is enabled</p>
30	RO	0b	<p><b>Root Table Pointer Status (RTPS)</b></p> <p>This field indicates the status of the root-table pointer in hardware.</p> <p>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.</p>



(Sheet 2 of 3)

Bit	Access	Default Value	Description
29	RO	0b	<p><b>Fault Log Status (FLS)</b></p> <p>This field:</p> <p>0 = Is cleared by hardware when software Sets the SFL field in the Global Command register.</p> <p>1 = Is Set by hardware when hardware completes the set fault-log pointer operation using the value provided in the Advanced Fault Log register.</p>
28	RO	0b	<p><b>Advanced Fault Logging Status (AFLS)</b></p> <p>This field is valid only for implementations supporting advanced fault logging. It indicates the advanced fault logging status:</p> <p>0 = Advanced Fault Logging is not enabled</p> <p>1 = Advanced Fault Logging is enabled</p>
27	RO	0b	<p><b>Write Buffer Flush Status (WBFS)</b></p> <p>This field is valid only for implementations requiring write buffer flushing. This field indicates the status of the write buffer flush command. It is:</p> <p>0 = Cleared by hardware when hardware completes the write buffer flushing operation.</p> <p>1 = Set by hardware when software sets the WBF field in the Global Command register.</p>
26	RO	0b	<p><b>Queued Invalidation Enable Status (QIES)</b></p> <p>This field indicates queued invalidation enable status.</p> <p>0 = Queued invalidation is not enabled</p> <p>1 = Queued invalidation is enabled</p>
25	RO	0b	<p><b>Interrupt Remapping Enable Status (IRES)</b></p> <p>This field indicates the status of Interrupt-remapping hardware.</p> <p>0 = Interrupt-remapping hardware is not enabled</p> <p>1 = Interrupt-remapping hardware is enabled</p>
24	RO	0b	<p><b>Interrupt Remapping Table Pointer Status (IRTPS)</b></p> <p>This field indicates the status of the interrupt remapping table pointer in hardware.</p> <p>This field is cleared by hardware when software sets the SIRTTP field in the Global Command register.</p> <p>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.</p>
23	RO	0b	<p><b>Compatibility Format Interrupt Status (CIFS)</b></p> <p>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 Legacy interrupt mode is active.</p> <p>0 = Compatibility format interrupts are blocked.</p> <p>1 = Compatibility format interrupts are processed as pass-through (bypassing interrupt remapping).</p>



(Sheet 3 of 3)

Bit	Access	Default Value	Description
22:0	RO	000000h	Reserved

### 1.20.6 RTADDR\_REG - Root-Entry Table Address Register

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 20-27h  
 Default Value: 0000000000000000h  
 Access: RO; RW  
 Size: 64 bits  
 Register providing the base address of root-entry table.

Bit	Access	Default Value	Description
63:36	RO	0000000h	Reserved
35:12	RW	000000h	<p><b>Root Table Address (RTA)</b></p> <p>This register points to base of page aligned, 4-KB-sized root-entry table in system memory. Hardware may ignore and not implement Bits 63:HAW, where HAW is the host address width.</p> <p>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.</p> <p>Reads of this register return the value that was last programmed to it.</p>
11:0	RO	000h	Reserved



### 1.20.7 CCMD\_REG - Context Command Register

B/D/F/Type: 0/2/0/GFXVTBAR  
Address Offset: 28-2Fh  
Default Value: 0800000000000000h  
Access: RO; RW  
Size: 64 bits

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.

(Sheet 1 of 3)

Bit	Access	Default Value	Description
63	RW	0b	<b>Invalidate Context Cache (ICC)</b> 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 a write-buffer flushing requirement (RWBF=1 in the Capability register) must implicitly perform a write buffer flush before invalidating the context-cache..





(Sheet 2 of 3)

Bit	Access	Default Value	Description
62:61	RW	00b	<p><b>Context Invalidation Request Granularity (CIRG)</b></p> <p>Software provides the requested invalidation granularity through this field when setting the ICC field:</p> <p>00: Reserved.</p> <p>01: Global Invalidation request.</p> <p>10: Domain-selective invalidation request. The target domain-id must be specified in the DID field.</p> <p>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.</p> <p>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.</p>
60:59	RO	01b	<p><b>Context Actual Invalidation Granularity (CAIG)</b></p> <p>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:</p> <p>00: Reserved.</p> <p>01: Global Invalidation performed. This could be in response to a global, domain-selective, or device-selective invalidation request.</p> <p>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.</p> <p>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.</p>
58:34	RO	0000000h	<i>Reserved</i>



(Sheet 3 of 3)

Bit	Access	Default Value	Description
33:32	RO	00b	<b>Function Mask (FM)</b> 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 device(s) specified through the FM and SID fields must correspond to the domain-ID specified in the DID field. Value returned on read of this field is undefined.
31:16	RO	0000h	<b>Source-ID (SID)</b> 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. Value returned on read of this field is undefined.
15:0	RW	0000h	<b>Domain-ID (DID)</b> 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 ignores (and may not implement) bits 15:N where N is the supported domain-id width reported in the capability register.



### 1.20.8 FSTS\_REG - Fault Status Register

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 34-37h  
 Default Value: 00000000h  
 Access: RO; RWC-P; RO-P  
 Size: 32 bits

Register indicating the various error statuses.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31:16	RO	0000h	<i>Reserved</i>
15:8	RO-P	00h	<p><b>Fault Record Index (FRI)</b></p> <p>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.</p>
7	RO	0b	<i>Reserved</i>
6	RO	0b	<p><b>Invalidation Time-out Error (ITE)</b></p> <p>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.</p> <p>Hardware implementations not supporting Device-IOTLBs implement this bit as RsvdZ.</p>
5	RO	0b	<p><b>Invalidation Completion Error (ICE)</b></p> <p>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 the Fault Event Control register.</p> <p>Hardware implementations not supporting Device-IOTLBs implement this bit as RsvdZ.</p>
4	RO	0b	<p><b>Invalidation Queue Error (IQE)</b></p> <p>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.</p> <p>Hardware implementations not supporting queued invalidations implement this bit as RsvdZ.</p>



(Sheet 2 of 2)

Bit	Access	Default Value	Description
3	RO	0b	<p><b>Advanced Pending Fault (APF)</b></p> <p>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.</p> <p>Software writing 1 to this field clears it. Hardware implementations not supporting advanced fault logging implement this bit as RsvdZ.</p>
2	RO	0b	<p><b>Advanced Fault Overflow (AFO)</b></p> <p>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.</p> <p>Software writing 1 to this field clears it. Hardware implementations not supporting advanced fault logging implement this bit as RsvdZ.</p>
1	RO-P	0b	<p><b>Primary Pending Fault (PPF)</b></p> <p>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.</p> <p>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.</p>
0	RWC-P	0b	<p><b>Primary Fault Overflow (PFO)</b></p> <p>Hardware sets this field to indicate overflow of the 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.</p>



### 1.20.9 FECTL\_REG - Fault Event Control Register

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 38-3Bh  
 Default Value: 80000000h  
 Access: RO; RW  
 Size: 32 bits

Register specifying the fault event interrupt message control bits. Section 7.3 describes hardware handling of fault events.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31	RW	1b	<b>Interrupt Mask (IM)</b> 0 = No masking of interrupts. When an interrupt condition is detected, hardware issues an interrupt message (using the Fault Event Data & 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.



(Sheet 2 of 2)

Bit	Access	Default Value	Description
30	RO	0b	<p><b>Interrupt Pending (IP)</b></p> <p>Hardware sets the IP field whenever it detects an interrupt condition, which is defined as:</p> <ul style="list-style-type: none"> <li>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 the Fault Status register.</li> <li>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.</li> <li>Hardware detected error associated with the Invalidation Queue, setting the IQE field in the Fault Status register.</li> <li>Hardware detected invalid Device-IOTLB invalidation completion, setting the ICE field in the Fault Status register.</li> <li>Hardware detected Device-IOTLB invalidation completion time-out, setting the ITE field in the Fault Status register.</li> </ul> <p>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.</p> <p>The IP field is kept Set by hardware while the interrupt message is held pending. The interrupt message could be held pending due to the interrupt mask (IM field) being Set or other transient hardware conditions.</p> <p>The IP field is cleared by hardware as soon as the interrupt message pending condition is serviced.</p> <p>This could be due to either:</p> <ul style="list-style-type: none"> <li>Hardware issuing the interrupt message due to either a change in the transient hardware condition that caused the interrupt message to be held pending, or due to software clearing the IM field.</li> <li>Software servicing all the pending interrupt status fields in the Fault Status register as follows.</li> <li>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 the Fault Status register to be evaluated as Clear.</li> <li>Software clearing other status fields in the Fault Status register by writing back the value read from the respective fields.</li> </ul>
29:0	RO	00000000 h	<i>Reserved</i>



### 1.20.10 FEDATA\_REG - Fault Event Data Register

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 3C-3Fh  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits

Register specifying the interrupt message data.

Bit	Access	Default Value	Description
31:16	RO	0000h	Extended Interrupt Message Data (EID) 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 RsvdZ.
15:0	RW	0000h	Interrupt message data (ID) Data value in the interrupt request.

### 1.20.11 FEADDR\_REG - Fault Event Address Register

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 40-43h  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits

Register specifying the interrupt message address.

Bit	Access	Default Value	Description
31:2	RW	00000000h	<b>Message Address (MA)</b> When fault events are enabled, the contents of this register specify the DWORD-aligned address (bits 31:2) for the interrupt request.
1:0	RO	00b	<i>Reserved</i>



### 1.20.12 FEUADDR\_REG - Fault Event Upper Address Register

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 44-47h  
 Default Value: 00000000h  
 Access: RO  
 Size: 32 bits

Register specifying the interrupt message upper address. This register is treated as RsvdZ by implementations reporting Extended Interrupt Mode (EIM) as not supported in the Extended Capability register.

Bit	Access	Default Value	Description
31:0	RO	00000000h	<b>Message Upper Address (MUA)</b> 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.

### 1.20.13 AFLOG\_REG - Advanced Fault Log Register

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 58-5Fh  
 Default Value: 0000000000000000h  
 Access: RO  
 Size: 64 bits

Register to specify the base address of memory-resident fault-log region.

This register is treated as read-only (0) for implementations not supporting advanced translation fault logging (AFL field reported as 0 in the Capability register).

Bit	Access	Default Value	Description
63:12	RO	00000000 000000h	<b>Fault Log Address (FLA)</b> This field specifies the base of 4-KB aligned fault-log region in system memory. Hardware ignores and 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.
11:9	RO	000b	<b>Fault Log Size (FLS)</b> This field specifies the size of the fault log region pointed to by the FLA field. The size of the fault log region is $(2^{X}) * 4$ -KB, where X is the value programmed in this register. When implemented, reads of this field return the value that was last programmed to it.
8:0	RO	000h	<i>Reserved</i>





### 1.20.14 PMEN\_REG - Protected Memory Enable Register

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 64-67h  
 Default Value: 00000000h  
 Access: RO; RW  
 Size: 32 bits

Register to enable the DMA-protected memory regions set up through the PLMBASE, PLMLIMIT, PHMBASE, PHMLIMIT registers. This register is always treated as RO (0) for implementations not supporting protected memory regions (PLMR and PHMR fields reported as 0 in the Capability register).

Protected memory regions may be used by software to securely initialize remapping structures in memory.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31	RW	0b	<p><b>Enable Protected Memory Region (EPM)</b></p> <p>This field controls DMA accesses to the protected low-memory and protected high memory regions.</p> <p>0 = DMA accesses to protected memory regions are handled as follows:</p> <ul style="list-style-type: none"> <li>• If DMA remapping is not enabled, DMA requests (including those to protected regions) are not blocked.</li> <li>• If DMA remapping is enabled, DMA requests are translated per the programming of the DMA remapping structures. Software may program the DMA-remapping structures to allow or block DMA to the protected memory regions.</li> </ul> <p>1 = DMA accesses to protected memory regions are handled as follows:</p> <p>If DMA remapping is not enabled, DMA requests to protected memory regions are blocked. These DMA requests are not recorded or reported as DMA-remapping faults.</p> <p>If DMA remapping is enabled, hardware may or may not block DMA to the protected memory region(s). Software must not depend on hardware protection of the protected memory regions, and must ensure the DMA-remapping structures are properly programmed to not allow DMA to the protected memory regions.</p> <p>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.</p>
30:1	RO	00000000h	<i>Reserved</i>



(Sheet 2 of 2)

Bit	Access	Default Value	Description
0	RO	0b	<b>Protected Region Status (PRS)</b> This field indicates the status of protected memory region(s) 0 = Protected memory region(s) disabled. 1 = Protected memory region(s) enabled.

### 1.20.15 PLMBASE\_REG - Protected Low Memory Base Register

B/D/F/Type: 0/2/0/GFXVTBAR

Address Offset: 68-6Bh

Default Value: 00000000h

Access: RO; RW

Size: 32 bits

Register to set up the base address of DMA-protected low-memory region below 4 GB. This register must be set up before enabling protected memory through PMEN\_REG, and must not be updated when protected memory regions are enabled.

When the LT CMD.LOCK.PMRC command is invoked, this register is locked (treated as RO). When the LT CMD.UNLOCK.PMRC command is invoked, this register is unlocked (treated as RW).

This register is always treated as RO for implementations not supporting protected low memory region (PLMR field reported as 0 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 bit position with 0 in the value read back from the register. Bits N:0 of this register are decoded by hardware as all 0s.

Software must setup the protected low memory region below 4 GB.

Bit	Access	Default Value	Description
31:21	RW	000h	<b>Protected Low-Memory Base (PLMB)</b> This register specifies the base of protected low-memory region in system memory.
20:0	RO	000000h	<i>Reserved</i>



### 1.20.16 PLMLIMIT\_REG - Protected Low Memory Limit Register

B/D/F/Type:	0/2/0/GFXVTBAR
Address Offset:	6C-6Fh
Default Value:	00000000h
Access:	RO; RW
Size:	32 bits

Register to set up the limit address of DMA-protected low-memory region below 4 GB. This register must be set up before enabling protected memory through PMEN\_REG, and must not be updated when protected memory regions are enabled.

When the LT CMD.LOCK.PMRC command is invoked, this register is locked (treated as RO). When the LT CMD.UNLOCK.PMRC command is invoked, this register is unlocked (treated as RW).

This register is always treated as RO for implementations not supporting protected low memory region (PLMR field reported as 0 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 are decoded by hardware as all 1s.

The Protected low-memory base and limit registers function as follows:

- Programming the protected low-memory base and limit registers 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.

Bit	Access	Default Value	Description
31:21	RW	000h	<b>Protected Low-Memory Limit (PLML)</b> This register specifies the last host physical address of the DMA-protected low-memory region in system memory.
20:0	RO	000000h	<i>Reserved</i>



### 1.20.17 PHMBASE\_REG - Protected High Memory Base Register

B/D/F/Type: 0/2/0/GFXVTBAR  
Address Offset: 70-77h  
Default Value: 0000000000000000h  
Access: RO; RW  
Size: 64 bits

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.

When the LT CMD.LOCK.PMRC command is invoked, this register is locked (treated as RO). When the LT CMD.UNLOCK.PMRC command is invoked, this register is unlocked (treated as RW).

This register is always treated as RO for implementations not supporting protected high memory region (PHMR field reported as 0 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 0's.

Software may setup the protected high memory region either above or below 4 GB.

Bit	Access	Default Value	Description
63:36	RO	0000000h	Reserved
35:21	RW	0000h	<b>Protected High-Memory Base (PHMB)</b> 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.
20:0	RO	000000h	Reserved



### 1.20.18 PHMLIMIT\_REG - Protected High Memory Limit Register

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 78-7Fh  
 Default Value: 0000000000000000h  
 Access: RO; RW  
 Size: 64 bits

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.

When the LT CMD.LOCK.PMRC command is invoked, this register is locked (treated as RO). When the LT CMD.UNLOCK.PMRC command is invoked, this register is unlocked (treated as RW).

This register is always treated as RO for implementations not supporting protected high memory region (PHMR field reported as 0 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 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 are decoded by hardware as all 1s.

The protected high-memory base and limit registers function 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.

Bit	Access	Default Value	Description
63:36	RO	0000000h	<i>Reserved</i>
35:21	RW	0000h	<b>Protected High-Memory Limit (PHML)</b> 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.
20:0	RO	000000h	<i>Reserved</i>



### 1.20.19 IQH\_REG - Invalidation Queue Head

B/D/F/Type: 0/2/0/GFXVTBAR  
Address Offset: 80-87h  
Default Value: 0000000000000000h  
Access: RO  
Size: 64 bits

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.

Bit	Access	Default Value	Description
63:19	RO	00000000 00000h	<i>Reserved</i>
18:4	RO	0000h	<b>Queue Head (QH)</b> Specifies the offset (128-bit aligned) to the invalidation queue for the command that is fetched next by hardware. Hardware resets this field to 0 whenever the queued invalidation is disabled (QIES field Clear in the Global Status register).
3:0	RO	0h	<i>Reserved</i>

### 1.20.20 IQT\_REG - Invalidation Queue Tail

B/D/F/Type: 0/2/0/GFXVTBAR  
Address Offset: 88-8Fh  
Default Value: 0000000000000000h  
Access: RO  
Size: 64 bits

Register indicating the invalidation queue tail. This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

Bit	Access	Default Value	Description
63:19	RO	00000000 00000h	<i>Reserved</i>
18:4	RO	0000h	<b>Queue Tail (QT)</b> Specifies the offset (128-bit aligned) to the invalidation queue for the command that is written next by software.
3:0	RO	0h	<i>Reserved</i>



### 1.20.21 IQA\_REG - Invalidation Queue Address

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 90-97h  
 Default Value: 0000000000000000h  
 Access: RO  
 Size: 64 bits

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.

When supported, writing to this register causes the Invalidation Queue Head and Invalidation Queue Tail registers to be reset to 0h.

Bit	Access	Default Value	Description
63:12	RO	00000000 00000h	<b>Invalidation Queue Address (IQA)</b> This field points to the base of 4-KB aligned invalidation request queue. Hardware ignores and 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.
11:3	RO	000h	<i>Reserved</i>
2:0	RO	000b	<b>Queue Size (QS)</b> This field specifies the size of the invalidation request queue. A value of X in this field indicates an invalidation request queue of (X+1) 4-KB pages. The number of entries in the invalidation queue is $2^{(X + 8)}$ .



### 1.20.22 ICS\_REG - Invalidation Completion Status

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 9C-9Fh  
 Default Value: 00000000h  
 Access: RO  
 Size: 32 bits

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.

Bit	Access	Default Value	Description
31:1	RO	00000000h	<i>Reserved</i>
0	RO	0b	<b>Invalidation Wait Descriptor Complete (IWC)</b> Indicates completion of Invalidation Wait Descriptor with Interrupt Flag (IF) field Set. Hardware implementations not supporting queued invalidations implement this field as RsvdZ.

### 1.20.23 IECTL\_REG - Invalidation Completion Event Control

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: A0-A3h  
 Default Value: 80000000h  
 Access: RO  
 Size: 32 bits

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.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
31	RO	1b	<b>Interrupt Mask (IM)</b> 0 = No masking of interrupt. When a invalidation event condition is detected, hardware issues an interrupt message (using the Invalidation Event Data & Invalidation 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.





(Sheet 2 of 2)

Bit	Access	Default Value	Description
30	RO	0b	<p><b>Interrupt Pending (IP)</b></p> <p>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.</p>
29:0	RO	00000000h	<i>Reserved</i>

### 1.20.24 IEDATA\_REG - Invalidation Completion Event Data

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: A4-A7h  
 Default Value: 00000000h  
 Access: RO  
 Size: 32 bits

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.

Bit	Access	Default Value	Description
31:16	RO	0000h	<p><b>Extended Interrupt Message Data (EIMD)</b></p> <p>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 RsvdZ.</p>
15:0	RO	0000h	<p><b>Interrupt Message Data (IMD)</b></p> <p>Data value in the interrupt request.</p>



### 1.20.25 IEADDR\_REG - Invalidation Completion Event Address

B/D/F/Type: 0/2/0/GFXVTBAR  
Address Offset: A8-ABh  
Default Value: 00000000h  
Access: RO  
Size: 32 bits

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.

### 1.20.26 IEUADDR\_REG - Invalidation Completion Event Upper Address

B/D/F/Type: 0/2/0/GFXVTBAR  
Address Offset: AC-AFh  
Default Value: 00000000h  
Access: RO  
Size: 32 bits

Register specifying the Invalidation Event interrupt message upper address. This register is treated as RsvdZ by implementations reporting both Queued Invalidation (QI) and Extended Interrupt Mode (EIM) as not supported in the Extended Capability register.

Bit	Access	Default Value	Description
31:0	RO	00000000h	<b>Message Upper Address (MUA)</b> Hardware implementations supporting Queued Invalidations and Extended Interrupt Mode are required to implement this register. Hardware implementations not supporting Queued Invalidations and Extended Interrupt Mode may treat this field as RsvdZ.



### 1.20.27 IRTA\_REG - Interrupt Remapping Table Address

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: B8-BFh  
 Default Value: 0000000000000000h  
 Access: RO  
 Size: 64 bits

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.

Bit	Access	Default Value	Description
63:12	RO	00000000 000000h	<b>Interrupt Remapping Table Address (IRTA)</b> This field points to the base of 4-KB aligned interrupt remapping table. Hardware ignores and not implement Bits 63:HAW, where HAW is the host address width. Reads of this field returns value that was last programmed to it.
11	RO	0b	<b>Extended Interrupt Mode Enable (EIME)</b> 0 = Legacy interrupt 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 is treated as reserved. On Itanium™ platforms hardware interprets the low 16 bits of the Destination-ID field in the IRTEs and treats the high 16 bits as reserved. 1 = Intel® 64 platform is operating in Extended Interrupt Mode. Hardware interprets all 32 bits of the Destination-ID field in the IRTEs. Hardware reporting Extended Interrupt Mode (EIM) as Clear in the Capability register treats this field as RsvdZ.
10:4	RO	00h	<i>Reserved</i>
3:0	RO	0h	<b>Size (S)</b> 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.



### 1.20.28 IVA\_REG - Invalidate Address Register

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 100-107h  
 Default Value: 0000000000000000h  
 Access: RO  
 Size: 64 bits

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. A value returned on a read of this register is undefined.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
63:36	RO	0000000h	<i>Reserved</i>
35:12	RO	000000h	<b>Address (ADDR)</b> 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 appropriate page-selective invalidate command through the IOTLB_REG. Hardware ignores Bits 63:N, where N is the maximum guest address width (MGAW) supported. Value returned on read of this field is undefined.
11:7	RO	00h	<i>Reserved</i>
6	RO	0b	<b>Invalidation Hint (IH)</b> The field provides hints 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 to the 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 the mappings specified by the ADDR and AM fields. A value returned on a read of this field is undefined.



(Sheet 2 of 2)

Bit	Access	Default Value	Description																														
5:0	RO	00h	<p><b>Address Mask (AM)</b></p> <p>The value in this field specifies the number of low order bits of the ADDR field that must be masked for the invalidation operation. Mask field enables software to request invalidation of contiguous mappings for size-aligned regions. For example:</p> <table border="1"> <thead> <tr> <th>Mask Value</th> <th>ADDR Bits</th> <th>Pages Invalidated</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>None</td> <td>1</td> </tr> <tr> <td>1</td> <td>12</td> <td>2</td> </tr> <tr> <td>2</td> <td>13:12</td> <td>8</td> </tr> <tr> <td>3</td> <td>14:12</td> <td>16</td> </tr> <tr> <td>4</td> <td>15:12</td> <td>32</td> </tr> <tr> <td>5</td> <td>16:12</td> <td>64</td> </tr> <tr> <td>6</td> <td>17:12</td> <td>128</td> </tr> <tr> <td>7</td> <td>18:12</td> <td>256</td> </tr> <tr> <td>8</td> <td>19:12</td> <td>512</td> </tr> </tbody> </table> <p>Hardware implementations report the maximum supported mask value through the Capability register. Value returned on read of this field is undefined.</p>	Mask Value	ADDR Bits	Pages Invalidated	0	None	1	1	12	2	2	13:12	8	3	14:12	16	4	15:12	32	5	16:12	64	6	17:12	128	7	18:12	256	8	19:12	512
Mask Value	ADDR Bits	Pages Invalidated																															
0	None	1																															
1	12	2																															
2	13:12	8																															
3	14:12	16																															
4	15:12	32																															
5	16:12	64																															
6	17:12	128																															
7	18:12	256																															
8	19:12	512																															



### 1.20.29 IOTLB\_REG - IOTLB Invalidate Register

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 108-10Fh  
 Default Value: 0200000000000000h  
 Access: RO; RW  
 Size: 64 bits

Register to invalidate IOTLB. The act of writing the upper byte of the IOTLB\_REG with the IVT field Set causes the hardware to perform the IOTLB invalidation.

(Sheet 1 of 3)

Bit	Access	Default Value	Description
63	RW	0b	<p><b>Invalidate IOTLB (IVT)</b></p> <p>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 a write-buffer flushing requirement (RWBF=1 in Capability register) must implicitly perform a write buffer flushing before invalidating the IOTLB.</p>
62:60	RW	000b	<p><b>IOTLB Invalidation Request Granularity (IIRG)</b></p> <p>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.</p> <p>000: Reserved.</p> <p>001: Global invalidation request.</p> <p>010: Domain-selective invalidation request. The target domain-id must be specified in the DID field.</p> <p>011: Page-selective invalidation request.</p> <p>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.</p> <p>100 - 111: Reserved.</p> <p>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 that time, the granularity at which actual invalidation was performed is reported through the IAIG field.</p>



(Sheet 2 of 3)

Bit	Access	Default Value	Description
59:57	RO	001b	<p><b>IOTLB Actual Invalidation Granularity (IAIG)</b></p> <p>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.</p> <p>000: 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.</p> <p>001: Global Invalidation performed. This could be in response to a global, domain-selective, or page-selective invalidation request.</p> <p>010: 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 page-selective invalidation request.</p> <p>011: 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.</p> <p>100 - 111: Reserved.</p>
56:50	RO	00h	<i>Reserved</i>
49	RW	0b	<p><b>Drain Reads (DR)</b></p> <p>This field is ignored by hardware if the DRD field is reported as clear in the Capability register.</p> <p>When DRD field is reported as set in the Capability register, the following encodings are supported for this field:</p> <p>0 = Hardware may complete the IOTLB invalidation without draining DMA read requests.</p> <p>1 = Hardware must drain DMA read requests.</p>
48	RW	0b	<p><b>Drain Writes (DW)</b></p> <p>This field is ignored by hardware if the DWD field is reported as clear in the Capability register.</p> <p>When DWD field is reported as set in the Capability register, the following encodings are supported for this field:</p> <p>0 = Hardware may complete the IOTLB invalidation without draining DMA write requests.</p> <p>1 = Hardware must drain relevant translated DMA write requests.</p>



(Sheet 3 of 3)

Bit	Access	Default Value	Description
47:32	RW	0000h	<b>Domain-ID (DID)</b> 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 implement Bits 47:(32+N), where N is the supported domain-id width reported in the Capability register.
31:0	RO	00000000h	<i>Reserved</i>

### 1.20.30 FRCD\_REG - Fault Recording Registers

B/D/F/Type: 0/2/0/GFXVTBAR  
 Address Offset: 200-20Fh  
 Default Value: 00000000000000000000000000000000h  
 Access: RO-P; RO; RWC-P  
 Size: 128 bits

Registers 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.

These registers are sticky and can be cleared only through power good reset or by software clearing the RW1C fields by writing a 1.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
127	RWC-P	0b	<b>Fault (F)</b> 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.
126	RO-P	0b	<b>Type (T)</b> Type of the faulted request: 0 = Write request 1 = Read 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.





(Sheet 2 of 2)

Bit	Access	Default Value	Description
125:124	RO	00b	<p><b>Address Type (AT)</b></p> <p>This field captures the AT field from the faulted DMA request.</p> <p>Hardware implementations not supporting Device-IOTLBs (DI field Clear in Extended Capability register) treat this field as RsvdZ.</p> <p>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.</p>
123:104	RO	00000h	<i>Reserved</i>
103:96	RO-P	00h	<p><b>Fault Reason (FR)</b></p> <p>Reason for the fault. Appendix A enumerates the various translation fault reason encodings.</p> <p>This field is relevant only when the F field is Set.</p>
95:80	RO	0000h	<i>Reserved</i>
79:64	RO-P	0000h	<p><b>Source Identifier (SID)</b></p> <p>Requester-id associated with the fault condition. This field is relevant only when the F field is Set.</p>
63:36	RO	0000000h	<p><b>Output Enable for TXPlus-side Transmit Buffer (TXPOE)</b></p> <p>If TOE is set, this bit controls the TXPOEB enable directly, otherwise if this bit is cleared it this bit directs the state machine to keep this bit cleared. Enables P-side of final TX driver.</p> <p>0 = Enable 1 = Disable</p>
35:12	RO-P	000000h	<p><b>Page Address (PADDR)</b></p> <p>When the Fault Reason (FR) field indicates one of the DMA-remapping fault conditions, Bits 63: 12 of this field contains the page address in the faulted DMA request. Hardware treat Bits 63: N as reserved (0), where N is the maximum guest address width (MGAW) supported.</p> <p>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.</p> <p>This field is relevant only when the F field is Set.</p>
11:0	RO	000h	<i>Reserved</i>



## 1.21 Intel® Trusted Execution Technology (Intel® TXT) Specific Registers

The public space registers are mapped to the address range starting at FED30000H and are available before, during and after a measured environment launch.

Register Name	Register Symbol	Register Start	Register End	Default Value	Access
TXT Device ID Register	TXT.DID	110	117	0000000FA0008086h	RO
TXT DMA Protected Range	TXT.DPR	330	337	0000000000000000h	RW-L; RW-L-K; RO
TXT Processor Public Key Hash Lower Half	TXT.PUBLIC.KEY.LOWER	400	40F	73A13C69E7DCF24C384C652BA19DA250h	RO
TXT Processor Public Key Hash Upper Half	TXT.PUBLIC.KEY.UPPER	410	41F	D884C70067DFC104BFDF8368D7254DBBh	RO

### 1.21.1 TXT.DID - TXT Device ID Register

B/D/F/Type: 0/0/0/TXT Specific  
 Address Offset: 110-117h  
 Default Value: 0000000FA0008086h  
 Access: RO  
 Size: 64 bits  
 Contains the TXT ID for the Processor.

(Sheet 1 of 2)

Bit	Access	Default Value	Description
63:48	RO	0000h	Reserved



(Sheet 2 of 2)

Bit	Access	Default Value	Description
47:32	RO	000Fh	<b>Revision ID (TXT.RID)</b> For the initial stepping of the component, the value is 0001h. The value is a bit-mask for compatibility with prior steppings. For the B-0 stepping, this value is 0003h. For the C-0 stepping, this value is 0007h. For the C-2 stepping this value is 000Fh
31:16	RO	A000h	<b>Device ID (TXT.DID)</b> 0xA000
15:0	RO	8086h	<b>Vendor ID (TXT.VID)</b> This register field contains the PCI standard identification for Intel, 8086h.

### 1.21.2 TXT.DPR - DMA Protected Range

B/D/F/Type: 0/0/0/TXT Specific  
 Address Offset: 330-337h  
 Default Value: 0000000000000000h  
 Access: RO; RW-L; RW-L-K  
 Size: 64 bits  
 DMA protected range register.

Bit	Access	Default Value	Description
63:32	RO	0000000h	<i>Reserved</i>
31:20	RO	000h	<b>Top of DMA Protected Range (TopOfDPR)</b> Top address + 1 of DPR. This is the base of TSEG. Bits 19:0 of the BASE reported here are 0x0_0000.
19:12	RO	00h	<i>Reserved</i>
11:4	RW-L	00h	<b>DMA Protected Memory Size (DPR.SIZE)</b> This is the size of memory, in MB, that will be protected from DMA accesses. A value of 0x00 in this field means no additional memory is protected. The maximum amount of memory that will be protected is 255 MB.
3:1	RO	000b	<i>Reserved</i>
0	RW-L-K	0b	<b>Lock (LOCK)</b> Bits 19:0 are locked down in this register when this bit is set. This bit is a write-once bit. If BIOS writes a '0' to the bit, then it can not be written to a '1' on subsequent writes. BIOS must write the entire register with the correct values and set this bit with that write.



### 1.21.3 TXT.PUBLIC.KEY.LOWER - TXT Processor Public Key Hash Lower Half

B/D/F/Type: 0/0/0/TXT Specific  
 Address Offset: 400-40Fh  
 Default Value: 73A13C69E7DCF24C384C652BA19DA250h  
 Access: RO  
 Size: 128 bits

These registers hold the hash of the Processor's public key. It's 256 bits (32 Bytes).

Bit	Access	Default Value	Description
127:0	RO	73A13C69E7DCF24C384C652BA19DA250h	<b>Public Key Hash Lower half (TXT.PUBLIC.KEYHASH)</b> This is a 256 bit (32 byte) field that contains the hash of the Processor's public key. The value of the Processor's public key differs between Production mode and Debug mode. Debug Mode Public Key lower half: 73A13C69E7DCF24C384C652BA19DA250h Production Mode Public Key lower half: C8012D55129B7568DF3979FC2B8BDE54h

### 1.21.4 TXT.PUBLIC.KEY.UPPER - TXT Processor Public Key Hash Upper Half

B/D/F/Type: 0/0/0/TXT Specific  
 Address Offset: 410-41Fh  
 Default Value: D884C70067DFC104BFDF8368D7254DBBh  
 Access: RO  
 Size: 128 bits

These registers hold the hash of the Processor's public key. It's 256 bits (32 Bytes).

Bit	Access	Default Value	Description
127:0	RO	D884C70067DFC104BFDF8368D7254DBBh	<b>Public Key Hash Upper half (TXT.PUBLIC.KEYHASH)</b> This is a 256 bit (32 byte) field that contains the hash of the Processor's public key. The value of the Processor's public key differs between Production mode and Debug mode. Debug Mode Public Key upper half: D884C70067DFC104BFDF8368D7254DBBh Production Mode Public Key upper half: 1337927100B39E8EC9166899A0E12BE0h