

8th and 9th Generation Intel[®] Core[™] Processor Families and Intel[®] Xeon[®] E Processor Family


Datasheet, Volume 2 of 2

Supporting 8th Generation Intel[®] Core[™] Processor Families, Intel[®] Pentium[®] Processors, Intel[®] Celeron[®] Processors for U/H/S Platforms, formerly known as Coffee Lake

Supporting 9th Generation Intel[®] Core[™] Processor Families H/S Platforms, formerly known as Coffee Lake Refresh

May 2020

Revision 003



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

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

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. **No computer system can be absolutely secure.** Check with your system manufacturer or retailer or learn more at intel.com.

Intel technologies may require enabled hardware, specific software, or services activation. Check with your system manufacturer or retailer.

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

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

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

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or visit www.intel.com/design/literature.htm. No computer system can be absolutely secure.

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

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

Copyright © 2018-2020, Intel Corporation. All rights reserved.



Contents


1	Introduction	13
2	Processor Configuration Register Definitions and Address Ranges	14
2.1	Register Terminology	14
2.2	PCI Devices and Functions	15
2.3	System Address Map	17
2.4	Legacy Address Range	21
2.4.1	DOS Range (0h – 9_FFFFh)	22
2.4.2	Legacy Video Area / Compatible SMRAM Area (A_0h – B_FFFFh)	22
2.4.3	Legacy Video Area	22
2.4.4	Monochrome Adapter (MDA) Range	22
2.4.5	Compatible SMRAM Address Range	22
2.4.6	Programmable Attribute Map (PAM) (C_0h – F_FFFFh)	22
2.5	Main Memory Address Range (1 MB – TOLUD)	24
2.5.1	ISA Hole (15 MB –16 MB)	25
2.5.2	1 MB to TSEGMB	25
2.5.3	TSEG	25
2.5.4	Protected Memory Range (PMR) - (programmable)	25
2.5.5	DRAM Protected Range (DPR)	26
2.5.6	Pre-allocated Memory	26
2.6	PCI Memory Address Range (TOLUD – 4 GB)	27
2.6.1	APIC Configuration Space (FEC0_0h – FECF_FFFFh)	29
2.6.2	HSEG (FEDA_0h – FEDB_FFFFh)	29
2.6.3	MSI Interrupt Memory Space (FEE0_0h – FEEF_FFFFh)	29
2.6.4	High BIOS Area	29
2.7	Main Memory Address Space (4 GB to TOUUD)	30
2.7.1	Top of Memory (TOM)	30
2.7.2	Top of Upper Usable DRAM (TOUUD)	30
2.7.3	Top of Low Usable DRAM (TOLUD)	30
2.7.4	TSEG_BASE	30
2.7.5	Memory Re-claim Background	31
2.7.6	Indirect Accesses to MCHBAR Registers	31
2.7.7	Memory Remapping	32
2.7.8	Hardware Remap Algorithm	32
2.8	PCI Express* Configuration Address Space	32
2.9	Graphics Memory Address Ranges	32
2.9.1	IOBAR Mapped Access to Device 2 MMIO Space	33
2.9.2	Trusted Graphics Ranges	33
2.10	System Management Mode (SMM)	33
2.11	SMM and VGA Access Through GTT TLB	34
2.12	Intel® Management Engine (Intel® ME) Stolen Memory Accesses	34
2.13	I/O Address Space	34
2.13.1	PCI Express* I/O Address Mapping	35
2.14	Direct Media Interface (DMI) Interface Decode Rules	36
2.14.1	DMI Accesses to the Processor that Cross Device Boundaries	36
2.14.2	Traffic Class (TC) / Virtual Channel (VC) Mapping Details	37
2.15	PCI Express* Interface Decode Rules	39
2.15.1	TC/VC Mapping Details	39
2.16	Legacy VGA and I/O Range Decode Rules	40
2.17	I/O Mapped Registers	44
3	Host Bridge/DRAM Registers	45



3.1	Vendor Identification (VID)—Offset 0h	46
3.2	Device Identification (DID)—Offset 2h	47
3.3	PCI Command (PCICMD)—Offset 4h	47
3.4	PCI Status (PCISTS)—Offset 6h	48
3.5	Revision Identification (RID)—Offset 8h	50
3.6	Class Code (CC)—Offset 9h	50
3.7	Header Type (HDR)—Offset Eh	51
3.8	Subsystem Vendor Identification (SVID)—Offset 2Ch	52
3.9	Subsystem Identification (SID)—Offset 2Eh	52
3.10	Capabilities Pointer (CAPPTR)—Offset 34h	53
3.11	PCI Express* Egress Port Base Address (PXPEPBAR)—Offset 40h	53
3.12	Host Memory Mapped Register Range Base (MCHBAR)—Offset 48h	54
3.13	GMCH Graphics Control Register (GGC)—Offset 50h	55
3.14	Device Enable (DEVEN)—Offset 54h	56
3.15	Protected Audio Video Path Control (PAVPC)—Offset 58h	58
3.16	DMA Protected Range (DPR)—Offset 5Ch	59
3.17	PCI Express Register Range Base Address (PCIEXBAR)—Offset 60h	60
3.18	Root Complex Register Range Base Address (DMIBAR)—Offset 68h	62
3.19	Manageability Engine Base Address Register (MESEG)—Offset 70h	63
3.20	Manageability Engine Limit Address Register (MESEG)—Offset 78h	64
3.21	Programmable Attribute Map 0 (PAM0)—Offset 80h	65
3.22	Programmable Attribute Map 1 (PAM1)—Offset 81h	66
3.23	Programmable Attribute Map 2 (PAM2)—Offset 82h	67
3.24	Programmable Attribute Map 3 (PAM3)—Offset 83h	68
3.25	Programmable Attribute Map 4 (PAM4)—Offset 84h	69
3.26	Programmable Attribute Map 5 (PAM5)—Offset 85h	70
3.27	Programmable Attribute Map 6 (PAM6)—Offset 86h	71
3.28	Legacy Access Control (LAC)—Offset 87h	72
3.29	System Management RAM Control (SMRAMC)—Offset 88h	75
3.30	Remap Base Address Register (REMAPBASE)—Offset 90h	76
3.31	Remap Limit Address Register (REMAPLIMIT)—Offset 98h	77
3.32	Top of Memory (TOM)—Offset A0h	78
3.33	Top of Upper Usable DRAM (TOUUD)—Offset A8h	78
3.34	Base Data of Stolen Memory (BDSM)—Offset B0h	79
3.35	Base of GTT stolen Memory (BGSM)—Offset B4h	80
3.36	TSEG Memory Base (TSEGMB)—Offset B8h	81
3.37	Top of Low Usable DRAM (TOLUD)—Offset BCh	81
3.38	Error Status (ERRSTS)—Offset C8h	83
3.39	Error Command (ERRCMD)—Offset CAh	84
3.40	SMI Command (SMICMD)—Offset CCh	84
3.41	SCI Command (SCICMD)—Offset CEh	85
3.42	Scratchpad Data (SKPD)—Offset DCh	86
3.43	Capabilities A (CAPID0)—Offset E4h	87
3.44	Capabilities B (CAPID0)—Offset E8h	88
3.45	Capabilities C (CAPID0)—Offset Ech	89
4	Processor Graphics Registers	91
4.1	Vendor Identification (VID2)—Offset 0h	92
4.2	Device Identification (DID2)—Offset 2h	92
4.3	PCI Command (PCICMD)—Offset 4h	93
4.4	PCI Status (PCISTS2)—Offset 6h	94
4.5	Revision Identification (RID2)—Offset 8h	95
4.6	Class Code (CC)—Offset 9h	96
4.7	Cache Line Size (CLS)—Offset Ch	96
4.8	Master Latency Timer (MLT2)—Offset Dh	97



4.9	Header Type (HDR2)—Offset Eh	97
4.10	Graphics Translation Table, Memory Mapped Range Address (GTTMMADR)—Offset 10h	98
4.11	Graphics Memory Range Address (GMADR)—Offset 18h.....	99
4.12	I/O Base Address (IOBAR)—Offset 20h	100
4.13	Subsystem Vendor Identification (SVID2)—Offset 2Ch.....	101
4.14	Subsystem Identification (SID2)—Offset 2Eh	102
4.15	Video BIOS ROM Base Address (ROMADR)—Offset 30h.....	102
4.16	Capabilities Pointer (CAPPOINT)—Offset 34h	103
4.17	Interrupt Line (INTRLINE)—Offset 3Ch	103
4.18	Interrupt Pin (INTRPIN)—Offset 3Dh	104
4.19	Minimum Grant (MINGNT)—Offset 3Eh.....	105
4.20	Maximum Latency (MAXLAT)—Offset 3Fh	105
4.21	Capabilities A (CAPID0)—Offset 44h	106
4.22	Capabilities B (CAPID0)—Offset 48h	107
4.23	Device Enable (DEVEN0)—Offset 54h.....	108
4.24	Base Data of Stolen Memory (BDSM)—Offset 5Ch.....	110
4.25	Multi Size Aperture Control (MSAC)—Offset 62h	110
4.26	PCI Express Capability Header (PCIECAPHDR)—Offset 70h.....	112
4.27	Message Signaled Interrupts Capability ID (MSI)—Offset ACh.....	113
4.28	Message Control (MC)—Offset AEh	113
4.29	Message Address (MA)—Offset B0h	114
4.30	Message Data (MD)—Offset B4h.....	115
4.31	Power Management Capabilities ID (PMCAPID)—Offset D0h	115
4.32	Power Management Capabilities (PMCAP)—Offset D2h.....	116
4.33	Power Management Control/Status (PMCS)—Offset D4h	117
5	Dynamic Power Performance Management (DPPM) Registers.....	119
5.1	Device Enable (DEVEN)—Offset 54h.....	119
5.2	Capabilities A (CAPID0)—Offset E4h	121
5.3	Capabilities B (CAPID0)—Offset E8h	122
6	DMIBAR Registers	124
6.1	DMI Virtual Channel Enhanced Capability (DMIVCECH)—Offset 0h.....	125
6.2	DMI Port VC Capability Register 1 (DMIPVCCAP1)—Offset 4h.....	125
6.3	DMI Port VC Capability Register 2 (DMIPVCCAP2)—Offset 8h.....	126
6.4	DMI Port VC Control (DMIPVCCTL)—Offset Ch	127
6.5	DMI VC0 Resource Capability (DMIVC0RCAP)—Offset 10h.....	128
6.6	DMI VC0 Resource Control (DMIVC0RCTL)—Offset 14h.....	129
6.7	DMI VC0 Resource Status (DMIVC0RSTS)—Offset 1Ah	130
6.8	DMI VC1 Resource Capability (DMIVC1RCAP)—Offset 1Ch.....	131
6.9	DMI VC1 Resource Control (DMIVC1RCTL)—Offset 20h.....	132
6.10	DMI VC1 Resource Status (DMIVC1RSTS)—Offset 26h	133
6.11	DMI VCm Resource Capability (DMIVCMRCAP)—Offset 34h	134
6.12	DMI VCm Resource Control (DMIVCMRCTL)—Offset 38h	135
6.13	DMI VCm Resource Status (DMIVCMRSTS)—Offset 3Eh	136
6.14	DMI Root Complex Link Declaration (DMIRCLDECH)—Offset 40h.....	137
6.15	DMI Element Self Description (DMIESD)—Offset 44h.....	138
6.16	DMI Link Entry 1 Description (DMILE1D)—Offset 50h	139
6.17	DMI Link Entry 1 Address (DMILE1A)—Offset 58h.....	140
6.18	DMI Link Upper Entry 1 Address (DMILUE1A)—Offset 5Ch	140
6.19	DMI Link Entry 2 Description (DMILE2D)—Offset 60h	141
6.20	DMI Link Entry 2 Address (DMILE2A)—Offset 68h.....	142
6.21	Link Capabilities (LCAP)—Offset 84h	142
6.22	Link Control (LCTL)—Offset 88h	144
6.23	DMI Link Status (LSTS)—Offset 8Ah	145



6.24	Link Control 2 (LCTL2)—Offset 98h	146
6.25	Link Status 2 (LSTS2)—Offset 9Ah	148
6.26	DMI Uncorrectable Error Status (DMIUESTS)—Offset 1C4h	149
6.27	DMI Uncorrectable Error Mask (DMIUEMSK)—Offset 1C8h	150
6.28	DMI Uncorrectable Error Severity (DMIUESEV)—Offset 1CCh	151
6.29	DMI Correctable Error Status (DMICESTS)—Offset 1D0h	152
6.30	DMI Correctable Error Mask (DMICEMSK)—Offset 1D4h	153
7	MCHBAR Registers	155
7.1	MCHBAR_CH0_CR_TC_PRE_0_0_0_MCHBAR—Offset 4000h	158
7.2	MCHBAR_CH0_CR_SC_GS_CFG_0_0_0_MCHBAR—Offset 401Ch	159
7.3	ECCERRLOG0—Offset 4048h	160
7.4	ECC Error Log 1 (ECCERRLOG1)—Offset 404Ch	161
7.5	MCHBAR_CH0_CR_TC_ODT_0_0_0_MCHBAR—Offset 4070h	162
7.6	Refresh Parameters (TC)—Offset 4238h	163
7.7	Refresh Timing Parameters (TC)—Offset 423Ch	164
7.8	Power Management DIMM Idle Energy (PM)—Offset 4260h	165
7.9	Power Management DIMM Power Down Energy (PM)—Offset 4264h	166
7.10	Power Management DIMM Activate Energy (PM)—Offset 4268h	167
7.11	Power Management DIMM RdCas Energy (PM)—Offset 426Ch	168
7.12	Power Management DIMM WrCas Energy (PM)—Offset 4270h	169
7.13	MCHBAR_CH1_CR_TC_PRE_0_0_0_MCHBAR—Offset 4400h	170
7.14	MCHBAR_CH0_CR_SC_GS_CFG_0_0_0_MCHBAR—Offset 441Ch	171
7.15	ECCERRLOG0—Offset 4448h	172
7.16	ECC Error Log 1 (ECCERRLOG1)—Offset 444Ch	173
7.17	MCHBAR_CH0_CR_TC_ODT_0_0_0_MCHBAR—Offset 4470h	174
7.18	Refresh Parameters (TC)—Offset 4638h	175
7.19	Refresh Timing Parameters (TC)—Offset 463Ch	176
7.20	Power Management DIMM Idle Energy (PM)—Offset 4660h	177
7.21	Power Management DIMM Power Down Energy (PM)—Offset 4664h	178
7.22	Power Management DIMM Activate Energy (PM)—Offset 4668h	179
7.23	Power Management DIMM RdCas Energy (PM)—Offset 466Ch	180
7.24	Power Management DIMM WrCas Energy (PM)—Offset 4670h	181
7.25	MCSCHEDES_CR_SC_GS_CFG_0_0_0_MCHBAR—Offset 4C1Ch	182
7.26	PM—Offset 4C40h	183
7.27	ECCERRLOG0—Offset 4C48h	184
7.28	ECC Error Log 1 (ECCERRLOG1)—Offset 4C4Ch	185
7.29	MCSCHEDES_CR_TC_ODT_0_0_0_MCHBAR—Offset 4C70h	186
7.30	Refresh Parameters (TC)—Offset 4E38h	187
7.31	Refresh Timing Parameters (TC)—Offset 4E3Ch	188
7.32	Power Management DIMM Idle Energy (PM)—Offset 4E60h	188
7.33	Power Management DIMM Power Down Energy (PM)—Offset 4E64h	189
7.34	Power Management DIMM Activate Energy (PM)—Offset 4E68h	190
7.35	Power Management DIMM RdCas Energy (PM)—Offset 4E6Ch	191
7.36	Power Management DIMM WrCas Energy (PM)—Offset 4E70h	192
7.37	Address Decoder Inter Channel Configuration Register (MAD)—Offset 5000h	193
7.38	Address Decoder Intra Channel Configuration Register (MAD)—Offset 5004h	194
7.39	Address Decoder Intra Channel Configuration Register (MAD)—Offset 5008h	195
7.40	Address Decode DIMM Parameters. (MAD)—Offset 500Ch	196
7.41	Address Decode DIMM Parameters (MAD)—Offset 5010h	198
7.42	MCDECS_CR_MRC_REVISION_0_0_0_MCHBAR_MCMAIN—Offset 5034h	199
7.43	Request Count from GT (DRAM)—Offset 5040h	199
7.44	Request Count from IA (DRAM)—Offset 5044h	200
7.45	Request Count from IO (DRAM)—Offset 5048h	201
7.46	RD Data Count (DRAM)—Offset 5050h	201




7.47	WR Data Count (DRAM)—Offset 5054h.....	202
7.48	Self Refresh Configuration Register (PM)—Offset 5060h.....	203
7.49	Address Compare for ECC Error Injection (ECC)—Offset 5090h.....	203
7.50	Address Mask for ECC Error Injection (ECC)—Offset 5094h.....	204
7.51	NCDECS_CR_GFXVTBAR_0_0_0_MCHBAR_NCU—Offset 5400h.....	205
7.52	NCDECS_CR_VTDPVC0BAR_0_0_0_MCHBAR_NCU—Offset 5410h.....	205
7.53	PACKAGE—Offset 5820h.....	206
7.54	PKG—Offset 5828h.....	208
7.55	PKG—Offset 5830h.....	208
7.56	PKG—Offset 5838h.....	209
7.57	PKG—Offset 5840h.....	209
7.58	PKG—Offset 5848h.....	210
7.59	PKG—Offset 5858h.....	210
7.60	DDR—Offset 5880h.....	211
7.61	DRAM—Offset 5884h.....	213
7.62	DRAM—Offset 5888h.....	213
7.63	DDR—Offset 588Ch.....	214
7.64	DDR—Offset 5890h.....	215
7.65	DDR—Offset 5894h.....	215
7.66	DDR—Offset 5898h.....	216
7.67	DDR—Offset 589Ch.....	217
7.68	DDR—Offset 58A0h.....	217
7.69	PACKAGE—Offset 58A8h.....	219
7.70	DDR—Offset 58B0h.....	220
7.71	DDR—Offset 58B4h.....	221
7.72	DDR—Offset 58C0h.....	221
7.73	DDR—Offset 58C8h.....	222
7.74	DDR—Offset 58D0h.....	223
7.75	DDR—Offset 58D4h.....	223
7.76	DDR—Offset 58D8h.....	224
7.77	DDR—Offset 58DCh.....	225
7.78	PACKAGE—Offset 58F0h.....	225
7.79	IA—Offset 58FCh.....	226
7.80	GT—Offset 5900h.....	228
7.81	SA—Offset 5918h.....	230
7.82	GT—Offset 5948h.....	231
7.83	EDRAM—Offset 594Ch.....	232
7.84	Package—Offset 5978h.....	232
7.85	PP0—Offset 597Ch.....	233
7.86	PP1—Offset 5980h.....	233
7.87	RP—Offset 5994h.....	234
7.88	RP—Offset 5998h.....	235
7.89	SSKPD—Offset 5D10h.....	236
7.90	BIOS_RESET_CPL—Offset 5DA8h.....	236
7.91	PCU_CR_MC_BIOS_REQ_0_0_0_MCHBAR_PCU—Offset 5E00h.....	237
7.92	CONFIG—Offset 5F3Ch.....	238
7.93	CONFIG—Offset 5F40h.....	239
7.94	CONFIG—Offset 5F48h.....	240
7.95	CONFIG—Offset 5F50h.....	241
7.96	TURBO—Offset 5F54h.....	242
7.97	Package Thermal DPPM Status (PKG)—Offset 6200h.....	243
7.98	Memory Thermal DPPM Status (DDR)—Offset 6204h.....	244
8	GFXVTBAR Registers.....	246
8.1	Version Register (VER)—Offset 0h.....	247



8.2	Capability Register (CAP)—Offset 8h	247
8.3	Extended Capability Register (ECAP)—Offset 10h	250
8.4	Global Command Register (GCMD)—Offset 18h.....	252
8.5	Global Status Register (GSTS)—Offset 1Ch	255
8.6	Root-Entry Table Address Register (RTADDR)—Offset 20h.....	256
8.7	Context Command Register (CCMD)—Offset 28h.....	257
8.8	Fault Status Register (FSTS)—Offset 34h	259
8.9	Fault Event Control Register (FECTL)—Offset 38h.....	260
8.10	Fault Event Data Register (FEDATA)—Offset 3Ch.....	262
8.11	Fault Event Address Register (FEADDR)—Offset 40h.....	262
8.12	Fault Event Upper Address Register (FEUADDR)—Offset 44h.....	263
8.13	Advanced Fault Log Register (AFLOG)—Offset 58h	263
8.14	Protected Memory Enable Register (PMEN)—Offset 64h.....	264
8.15	Protected Low-Memory Base Register (PLMBASE)—Offset 68h	265
8.16	Protected Low-Memory Limit Register (PLMLIMIT)—Offset 6Ch.....	266
8.17	Protected High-Memory Base Register (PHMBASE)—Offset 70h	267
8.18	Protected High-Memory Limit Register (PHMLIMIT)—Offset 78h	268
8.19	Invalidation Queue Head Register (IQH)—Offset 80h	269
8.20	Invalidation Queue Tail Register (IQT)—Offset 88h.....	270
8.21	Invalidation Queue Address Register (IQA)—Offset 90h	271
8.22	Invalidation Completion Status Register (ICS)—Offset 9Ch.....	272
8.23	Invalidation Event Control Register (IECTL)—Offset A0h	273
8.24	Invalidation Event Data Register (IEDATA)—Offset A4h.....	274
8.25	Invalidation Event Address Register (IEADDR)—Offset A8h.....	274
8.26	Invalidation Event Upper Address Register (IEUADDR)—Offset ACh	275
8.27	Interrupt Remapping Table Address Register (IRTA)—Offset B8h.....	276
8.28	Fault Recording Low Register (FRCDL)—Offset 400h.....	277
8.29	Fault Recording High Register (FRCDH)—Offset 408h.....	277
8.30	Invalidate Address Register (IVA)—Offset 500h	279
8.31	IOTLB Invalidate Register (IOTLB)—Offset 508h	280
8.32	DMA Remap Engine Policy Control (ARCHDIS)—Offset FF0h.....	282
8.33	DMA Remap Engine Policy Control (UARCHDIS)—Offset FF4h.....	284
9	PXPEPBAR Registers	286
9.1	EP VC 0 Resource Control (EPVC0RCTL)—Offset 14h.....	286
10	VCOPREMAP Registers.....	288
10.1	Version Register (VER)—Offset 0h.....	289
10.2	Capability Register (CAP)—Offset 8h	289
10.3	Extended Capability Register (ECAP)—Offset 10h	292
10.4	Global Command Register (GCMD)—Offset 18h.....	294
10.5	Global Status Register (GSTS)—Offset 1Ch	297
10.6	Root-Entry Table Address Register (RTADDR)—Offset 20h.....	298
10.7	Context Command Register (CCMD)—Offset 28h.....	299
10.8	Fault Status Register (FSTS)—Offset 34h	301
10.9	Fault Event Control Register (FECTL)—Offset 38h.....	302
10.10	Fault Event Data Register (FEDATA)—Offset 3Ch.....	304
10.11	Fault Event Address Register (FEADDR)—Offset 40h.....	304
10.12	Fault Event Upper Address Register (FEUADDR)—Offset 44h.....	305
10.13	Advanced Fault Log Register (AFLOG)—Offset 58h	305
10.14	Protected Memory Enable Register (PMEN)—Offset 64h.....	306
10.15	Protected Low-Memory Base Register (PLMBASE)—Offset 68h	307
10.16	Protected Low-Memory Limit Register (PLMLIMIT)—Offset 6Ch.....	308
10.17	Protected High-Memory Base Register (PHMBASE)—Offset 70h	309
10.18	Protected High-Memory Limit Register (PHMLIMIT)—Offset 78h	310
10.19	Invalidation Queue Head Register (IQH)—Offset 80h	311



10.20	Invalidation Queue Tail Register (IQT)—Offset 88h	312
10.21	Invalidation Queue Address Register (IQA)—Offset 90h	313
10.22	Invalidation Completion Status Register (ICS)—Offset 9Ch	313
10.23	Invalidation Event Control Register (IECTL)—Offset A0h	314
10.24	Invalidation Event Data Register (IEDATA)—Offset A4h	315
10.25	Invalidation Event Address Register (IEADDR)—Offset A8h	316
10.26	Invalidation Event Upper Address Register (IEUADDR)—Offset ACh	317
10.27	Interrupt Remapping Table Address Register (IRTA)—Offset B8h	317
10.28	Fault Recording Low Register (FRCDL)—Offset 400h	318
10.29	Fault Recording High Register (FRCDH)—Offset 408h	319
10.30	Invalidate Address Register (IVA)—Offset 500h	320
10.31	IOTLB Invalidate Register (IOTLB)—Offset 508h	321
11	PCI Express* Controllers Registers (D1:F0-2)	324
11.1	Vendor Identification (VID)—Offset 0h	326
11.2	Device Identification (DID)—Offset 2h	326
11.3	PCI Command (PCICMD)—Offset 4h	327
11.4	PCI Status (PCISTS)—Offset 6h	328
11.5	Revision Identification (RID)—Offset 8h	330
11.6	Class Code (CC)—Offset 9h	331
11.7	Cache Line Size (CL)—Offset Ch	331
11.8	Header Type (HDR)—Offset Eh	332
11.9	Primary Bus Number (PBUSN)—Offset 18h	332
11.10	Secondary Bus Number (SBUSN)—Offset 19h	333
11.11	Subordinate Bus Number (SUBUSN)—Offset 1Ah	333
11.12	I/O Base Address (IOBASE)—Offset 1Ch	334
11.13	I/O Limit Address (IOLIMIT)—Offset 1Dh	334
11.14	Secondary Status (SSTS)—Offset 1Eh	335
11.15	Memory Base Address (MBASE)—Offset 20h	336
11.16	Memory Limit Address (MLIMIT)—Offset 22h	337
11.17	Prefetchable Memory Base Address (PMBASE)—Offset 24h	338
11.18	Prefetchable Memory Limit Address (PMLIMIT)—Offset 26h	339
11.19	Prefetchable Memory Base Address Upper (PMBASEU)—Offset 28h	339
11.20	Prefetchable Memory Limit Address Upper (PMLIMITU)—Offset 2Ch	340
11.21	Capabilities Pointer (CAPPTR)—Offset 34h	341
11.22	Interrupt Line (INTRLINE)—Offset 3Ch	342
11.23	Interrupt Pin (INTRPIN)—Offset 3Dh	342
11.24	Bridge Control (BCTRL)—Offset 3Eh	343
11.25	Power Management Capabilities (PM)—Offset 80h	344
11.26	Power Management Control/Status (PM)—Offset 84h	345
11.27	Subsystem ID and Vendor ID Capabilities (SS)—Offset 88h	346
11.28	Subsystem ID and Subsystem Vendor ID (SS)—Offset 8Ch	347
11.29	Message Signaled Interrupts Capability ID (MSI)—Offset 90h	348
11.30	Message Control (MC)—Offset 92h	348
11.31	Message Address (MA)—Offset 94h	349
11.32	Message Data (MD)—Offset 98h	350
11.33	PCI Express-G Capability List (PEG)—Offset A0h	350
11.34	PCI Express-G Capabilities (PEG)—Offset A2h	351
11.35	Device Capabilities (DCAP)—Offset A4h	352
11.36	Device Control (DCTL)—Offset A8h	352
11.37	Device Status (DSTS)—Offset AAh	354
11.38	Link Capability (LCAP)—Offset ACh	355
11.39	Link Control (LCTL)—Offset B0h	356
11.40	Link Status (LSTS)—Offset B2h	358
11.41	Slot Capabilities (SLOTCAP)—Offset B4h	359



11.42	Slot Control (SLOTCTL)—Offset B8h	360
11.43	Slot Status (SLOTSTS)—Offset BAh	362
11.44	Root Control (RCTL)—Offset BCh	364
11.45	Root Status (RSTS)—Offset C0h	365
11.46	Device Capabilities 2 (DCAP2)—Offset C4h	366
11.47	Device Control 2 (DCTL2)—Offset C8h	368
11.48	Link Control 2 (LCTL2)—Offset D0h	369
11.49	Link Status 2 (LSTS2)—Offset D2h	371
11.50	Port VC Capability Register 1 (PVCCAP1)—Offset 104h	372
11.51	Port VC Capability Register 2 (PVCCAP2)—Offset 108h	373
11.52	Port VC Control (PVCCTL)—Offset 10Ch	373
11.53	VC0 Resource Capability (VC0RCAP)—Offset 110h	374
11.54	VC0 Resource Control (VC0RCTL)—Offset 114h	375
11.55	VC0 Resource Status (VC0RSTS)—Offset 11Ah	376
11.56	Advanced Error Extended Reporting Capability Header (AECH)—Offset 1C0h	377
11.57	PEG Uncorrectable Error Status—Offset 1C4h	377
11.58	PEG Uncorrectable Error Mask—Offset 1C8h	378
11.59	PEG Uncorrectable Error Severity—Offset 1CCh	379
11.60	PEG Correctable Error Status—Offset 1D0h	379
11.61	PEG Correctable Error Mask—Offset 1D4h	380
11.62	PEG Advanced Error Capabilities and Control—Offset 1D8h	381
11.63	PEG Header Log—Offset 1DCh, 1E0h, 1E4h, 1E8h	381
11.64	PEG Root Error Command—Offset 1ECh	382
11.65	PEG Root Error Status—Offset 1F0h	382
11.66	PEG Error Source Identification—Offset 1F4h	383
12	GTTMMADR Registers	384
12.1	Top of Low Usable DRAM (MTOLUD)—Offset 108000h	384
12.2	Top of Upper Usable DRAM (MTOUUD)—Offset 108080h	385
12.3	Base Data of Stolen Memory (MBDSM)—Offset 1080C0h	386
12.4	Base of GTT stolen Memory (MBGSM)—Offset 108100h	387
12.5	Protected Memory Enable Register (MPMEN)—Offset 108180h	388
12.6	Protected Low-Memory Base Register (MPLMBASE)—Offset 1081C0h	389
12.7	Protected Low-Memory Limit Register (MPLMLIMIT)—Offset 108200h	390
12.8	Protected High-Memory Base Register (MPHMBASE)—Offset 108240h	391
12.9	Protected High-Memory Limit Register (MPHMLIMIT)—Offset 108280h	392
12.10	Protected Audio Video Path Control (MPAVPC)—Offset 1082C0h	393
12.11	Global Command Register (MGCMD)—Offset 108300h	395
13	EDRAMBAR Registers	398
13.1	(3250h) GT_PMON_RD_EDRAM_HIT	398
13.2	(3254h) GT_PMON_RD_EDRAM_MISS	398
13.3	(3258h) GT_PMON_PTLWR_EDRAM_HIT	398
13.4	(325Ch) GT_PMON_PTLWR_EDRAM_MISS	399
13.5	(3260h) GT_PMON_WR_EDRAM_HIT	399
13.6	(3264h) GT_PMON_WR_EDRAM_MISS	399
13.7	(3268h) GT_PMON_CNT_MISC_CTRL	399



Figures

2-1	Conceptual Platform PCI Configuration Diagram	17
2-2	System Address Range Example.....	20
2-3	DOS Legacy Address Range	21
2-4	PAM Region Space.....	23
2-5	Main Memory Address Range	24
2-6	PCI Memory Address Range.....	28
2-7	Example: DMI Upstream VC0 Memory Map	38
2-8	PEG Upstream VC0 Memory Map	40


Tables

2-1	Register Attributes and Terminology	14
2-2	Register Attribute Modifiers	15
2-3	PCI Devices and Functions.....	16
2-4	PCI Device Enumeration	16
2-5	SMM Regions.....	33
2-6	Processor Graphics Frame Buffer Accesses.....	41
2-7	Processor Graphics VGA I/O Mapping	42
2-8	VGA and MDA IO Transaction Mapping	43
2-9	MDA Resources	43
3-1	Summary of Bus: 0, Device: 0, Function: 0 (CFG)	45
4-1	Summary of Bus: 0, Device: 2, Function: 0 (CFG)	91
5-1	Summary of Bus: 0, Device: 4, Function: 0 (CFG)	119
6-1	Summary of Bus: 0, Device: 0, Function: 0 (MEM).....	124
7-1	Summary of Bus: 0, Device: 0, Function: 0 (MEM).....	155
8-1	Summary of Bus: 0, Device: 0, Function: 0 (MEM).....	246
9-1	Summary of Bus: 0, Device: 0, Function: 0 (MEM).....	286
10-1	Summary of Bus: 0, Device: 0, Function: 0 (MEM).....	288
11-1	Summary of Bus: 0, Device: 1, Function: 0 (CFG)	324
12-1	Summary of Bus: 0, Device: 2, Function: 0 (MEM).....	384
13-1	Summary of EDRAMBAR Registers	398



Revision History

Revision Number	Description	Release Date
001	<ul style="list-style-type: none">Initial Release	April 2018
002	<ul style="list-style-type: none">Revision Update only	October 2019
003	<ul style="list-style-type: none">Added PCIe controller register in Chapter 11, "Advanced Error Extended Reporting Capability Header (AECH)—Offset 1C0h"Merged PCIe controller to a single chapterUpdated register descriptions	May 2020



1 Introduction

This is Volume 2 of the Datasheet and provides register information for the processor.

Refer to document #337344 for the 8th and 9th Generation Intel® Core™ Processor Families Datasheet- Volume 1 of 2

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

§ §



2 Processor Configuration Register Definitions and Address Ranges

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

Note: PCI Express* (PCIe) Graphics and DMI do not apply to U Processors.

2.1 Register Terminology

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

Table 2-1. Register Attributes and Terminology

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



Table 2-2. Register Attribute Modifiers

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

2.2 PCI Devices and Functions

The processor contains five (H/S) and four (U/Y) PCI devices within a single component. The configuration registers for the devices are mapped as devices residing on PCI Bus 0.

- Device 0: Host Bridge / DRAM Controller / LLC Controller 0 – Logically this device appears as a PCI device residing on PCI bus 0. Device 0 contains the standard PCI header registers, PCI Express base address register, DRAM control (including thermal/throttling control), configuration for the DMI, and other processor specific registers.
- Device 1: Host-PCI Express* Bridge – Logically this device appears as a "virtual" PCI-to-PCI bridge residing on PCI bus 0, and is compliant with the PCI-to-PCI Bridge Architecture Specification, Revision 1.2. Device 1 is a multi-function device consisting of three functions (0, 1, and 2). Device 1 contains the standard PCI-to-PCI bridge registers and the standard PCI Express/PCI configuration registers.
- Device 2: Processor Graphics – Logically, this device appears as a PCI device residing on PCI Bus 0. Physically, Device 2 contains the configuration registers for 3D, 2D, and display functions. In addition, Device 2 is located in two separate physical locations – GT and Display Engine.
- Device 5: Imaging Unit (IMGU) – Logically, this device appears as a PCI device residing on PCI Bus 0. Physically, Device 5 contains the configuration registers for the Imaging Unit.
- Device 8: Gaussian Mixture Model Device (GMM) – Logically, this device appears as a PCI device residing on PCI Bus 0. Physically, Device 8 contains the configuration registers for the Gaussian Mixture Model Device.

Note: PCI Express* (PCIe) Graphics and DMI do not apply to U Processors.



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

The PCI Express controllers (PEG10, PEG11, and PEG12) appear to system software to be real PCI buses behind PCI-to-PCI bridges that are devices resident on PCI Bus 0. This is shown in the following figure.

Table 2-3. PCI Devices and Functions

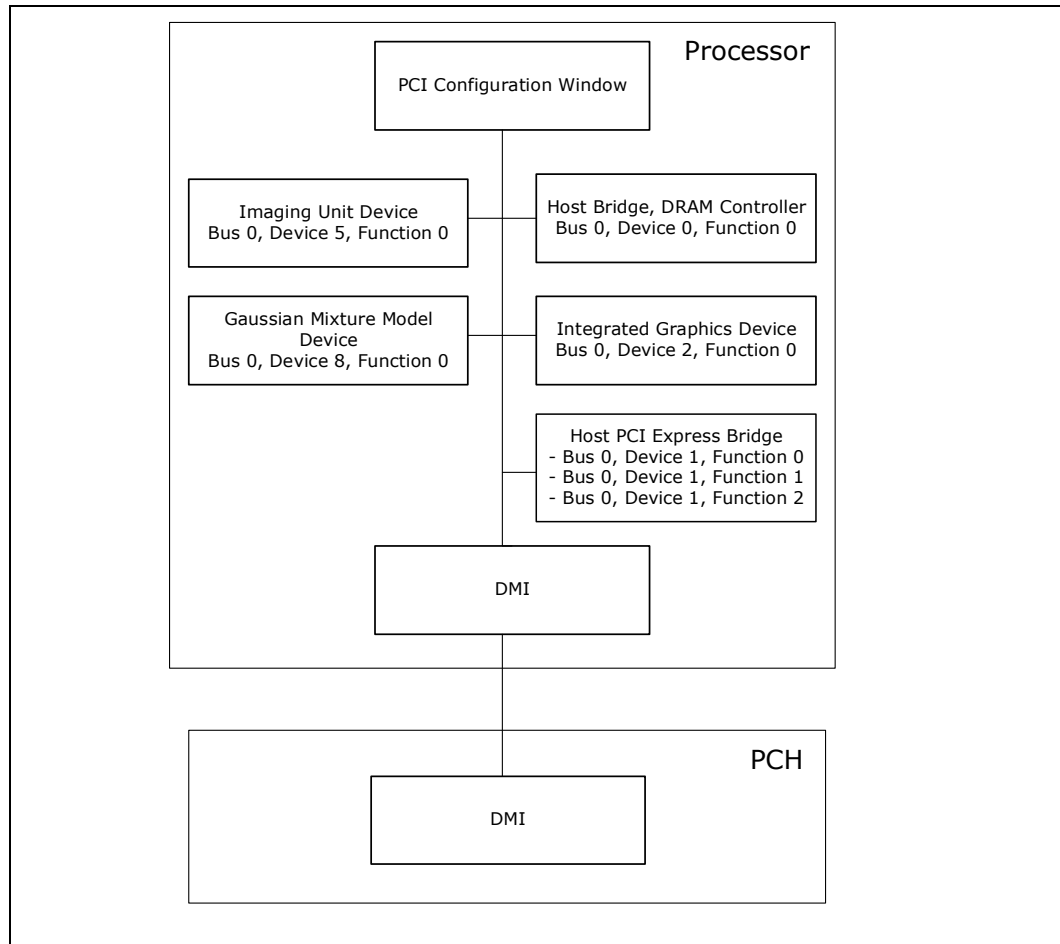
Description	DID					Device/ Function
	U-Processor Line	H-Processor Line	S-Processor Line			
Package	BGA 1528	BGA1440	LGA1151			
Segment	Mobile	Halo	Desktop	Work Station	Server	
HOST and DRAM Controller	2 Cores-3ECCh	4 Cores-3E10h	2 Cores-3E0Fh	4 Cores-3E18h	4 Cores-3E33h	0/0
			4 Cores-3E1Fh			
	4 Cores-3ED0h	6 Cores-3EC4h	6 Cores-3EC2h	6 Cores-3EC6h	6 Cores-3ECAh	
		8 Cores-3E20h	8 Cores-3E30h	8 Cores-3E31h	8 Cores-3E32h	
PCI Express* Controller (x16 PCIe)	N/A	1901h	1901h			1/0
PCI Express* Controller (x8 PCIe)	N/A	1905h	1905h			1/1
PCI Express* Controller (x4 PCIe)	N/A	1909h	1909h			1/2
Processor Graphics	Refer Note below					2/0
Gaussian Mixture Model	1911h	1911h	1911h			8/0
Note: For more information on Graphics Device ID refer to Intel Product Specifications https://ark.intel.com/						

Table 2-4. PCI Device Enumeration

Bus ID [7:0]	Device ID [4:0]	Function ID [2:0]	Endpoint	PCI Device ID U/ H/S - Processor Lines
0x00	00000b (0)	000b (0)	Host Bridge	Refer the above "PCI Devices and Functions" table
0x00	00001b (1)	000b (0)	PEG Root Port 10 - x16 controller	
0x00	00001b (1)	001b (1)	PEG Root Port 11 - x8 controller	
0x00	00001b (1)	010b (2)	PEG Root Port 12 - x4 controller	
0x00	00010b (2)	000b (0)	Processor Graphics	
0x00	00101b (5)	000b (0)	Imaging Unit	
0x00	01000b (8)	000b (0)	Gaussian Mixture Model	



Figure 2-1. Conceptual Platform PCI Configuration Diagram



Note: PCI Express* (PCIe) and DMI do not apply to U Processors.

2.3 System Address Map

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

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

The processor supports PEG port upper prefetchable base/limit registers. This allows the PEG unit to claim I/O accesses above 32 bit. Addressing of greater than 4 GB is allowed on either the DMI Interface or PCI Express interface. The processor supports a maximum of 32 GB of DRAM. No DRAM memory will be accessible above 32 GB. DRAM capacity is limited by the number of address pins available. There is no hardware lock to prevent more memory from being inserted than is addressable.



When running in Processor Graphics mode, processor initiated 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 using fence registers) are not supported from the 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 Processor Graphics device (Processor Graphics). 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 Processor Graphics are related to the PCI Express bus or the Processor Graphics device respectively. The processor does not remap APIC or any other memory spaces above TOLUD (Top of Low Usable DRAM). The TOLUD register is set to the appropriate value by BIOS. The remapbase/remaplimit registers remap logical accesses bound for addresses above 4 GB onto physical addresses that fall within DRAM.

The Address Map includes a number of programmable ranges:

- Device 0:
 - PXPEPBAR – PxP egress port registers. (4 KB window)
 - MCHBAR – Memory mapped range for internal MCH registers. (32 KB window)
 - DMIBAR – This window is used to access registers associated with the processor/PCH Serial Interconnect (DMI) register memory range. (4 KB window)
 - GGC.GMS – Graphics Mode Select. Used to select the amount of main memory that is pre-allocated to support the Processor Graphics device in VGA (non-linear) and Native (linear) modes. (0 – 512 MB options).
 - GGC.GGMS – GTT Graphics Memory Size. Used to select the amount of main memory that is pre-allocated to support the Processor Graphics Translation Table. (0 – 2 MB options).
- For each of the following device functions
- Device 1, Function 0: (PCIe x16 Controller)
- Device 1, Function 1: (PCIe x8 Controller)
- Device 1, Function 2: (PCIe x4 Controller)
- Device 2, Function 0: (Processor Graphics (Processor Graphics))
 - IOBAR – I/O access window for Processor Graphics. Through this window address/data register pair, using I/O semantics, the Processor Graphics and Processor Graphics instruction port registers can be accessed. This allows accessing the same registers as GTTMMADR. The IOBAR can be used to issue writes to the GTTMMADR or the GTT Table.
 - GMADR – Processor Graphics translation window (128 MB, 256 MB, 512 MB window).
 - GTTMMADR – This register requests a 4 MB allocation for combined Graphics Translation Table Modification Range and Memory Mapped Range. GTTADR will be at GTTMMADR + 2 MB while the MMIO base address will be the same as GTTMMADR

The rules for the above programmable ranges are:

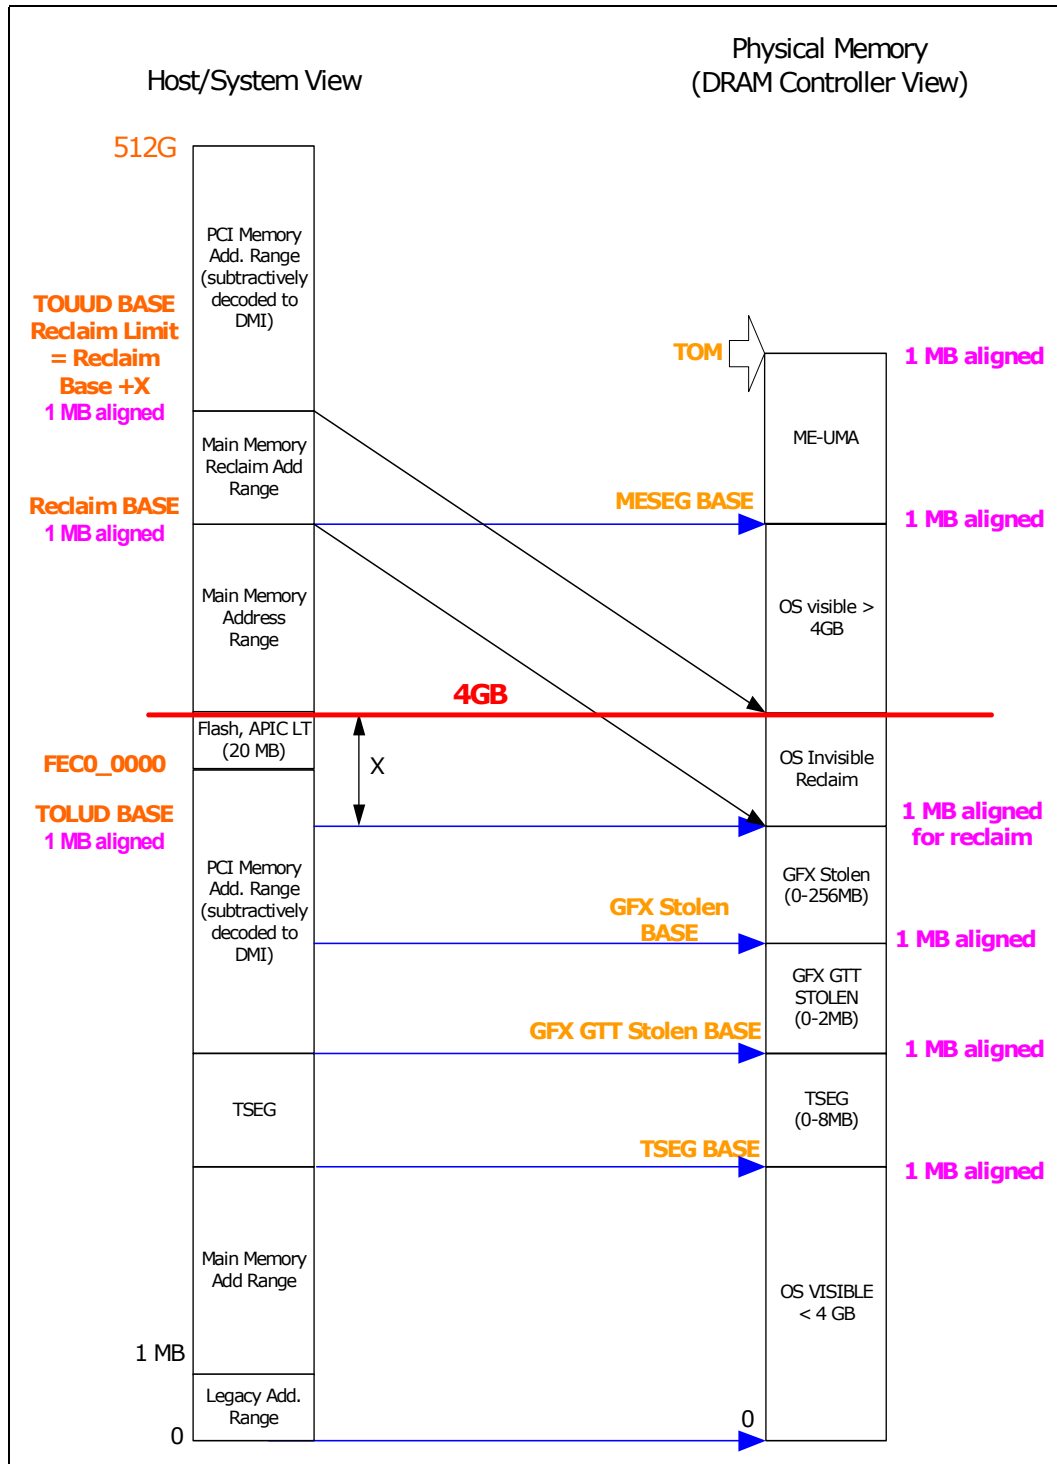
1. For security reasons, the processor will now positively decode (FFE0_0h to FFFF_FFFFh) to DMI. This ensures the boot vector and BIOS execute off the PCH.



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



Figure 2-2. System Address Range Example





2.4 Legacy Address Range


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

- 0 – 640 KB - DOS Area
- 640 – 768 KB - Legacy Video Buffer Area
- 768 – 896 KB in 16 KB sections (total of 8 sections) – Expansion Area
- 896 – 960 KB in 16 KB sections (total of 4 sections) – Extended System BIOS Area
- 960 KB – 1 MB Memory, System BIOS Area

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

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

Figure 2-3. DOS Legacy Address Range





2.4.1 DOS Range (0h – 9_FFFFh)

The DOS area is 640 KB (0000_0h – 0009_FFFFh) in size and is always mapped to the main memory.

2.4.2 Legacy Video Area / Compatible SMRAM Area (A_0h – B_FFFFh)

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

- Legacy Video Area: The legacy 128 KB VGA memory range, frame buffer, at 000A_0h – 000B_FFFFh, can be mapped to Processor Graphics (Device 2), to PCI Express (Device 1), and/or to the DMI Interface.
- Monochrome Adapter (MDA) Range: Legacy support requires the ability to have a second graphics controller (monochrome) in the system. The monochrome adapter may be mapped to Processor Graphics, PCI Express or DMI. Like the Legacy Video Area, decode priority is given first to Processor Graphics, then to PCI Express, and finally to DMI.
- Compatible SMRAM Address Range.

2.4.3 Legacy Video Area

The legacy 128 KB VGA memory range, frame buffer at 000A_0h – 000B_FFFFh, can be mapped to Processor Graphics (Device 2), to PCI Express (Device 1), and/or to the DMI Interface.

2.4.4 Monochrome Adapter (MDA) Range

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

2.4.5 Compatible SMRAM Address Range

When compatible SMM space is enabled, SMM-mode CBO accesses to this range route to physical system DRAM at 00_000A_0h – 00_000B_FFFFh.

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

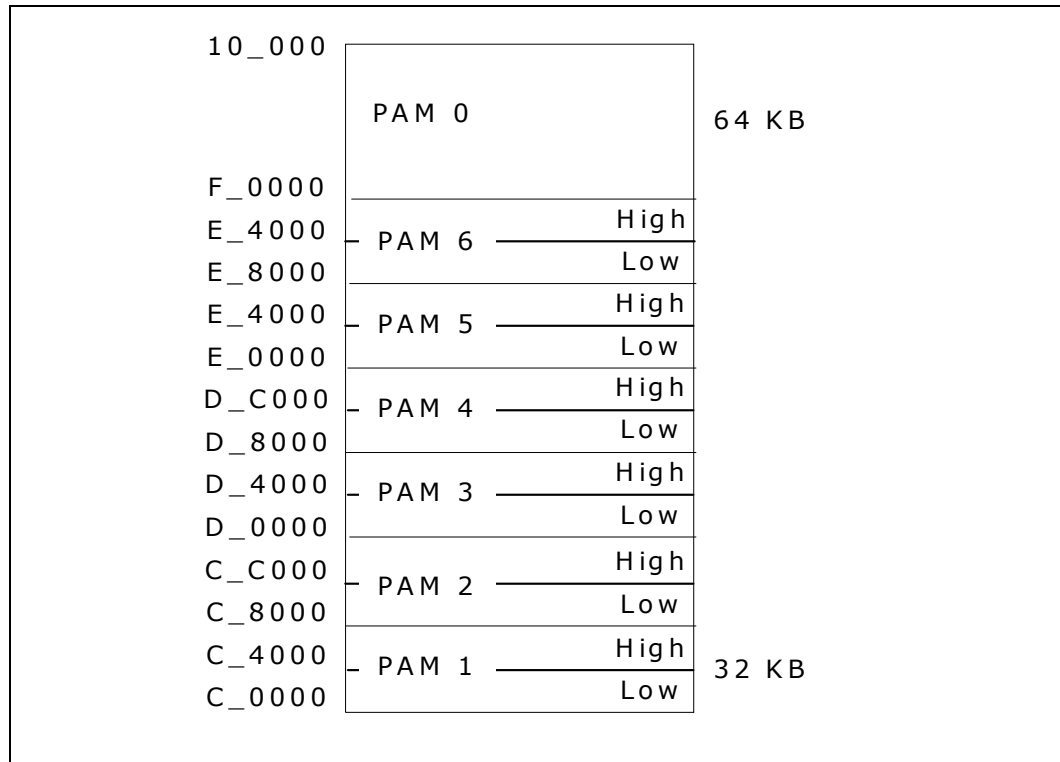
The processor always positively decodes internally mapped devices, namely the Processor Graphics and PCI Express. Subsequent decoding of regions mapped to PCI Express or the DMI Interface depends on the Legacy VGA configuration bits (VGA Enable and MDAP). This region is also the default for SMM space.

2.4.6 Programmable Attribute Map (PAM) (C_0h – F_FFFFh)

PAM is a legacy BIOS ROM area in MMIO. It is overlaid with DRAM and used as a faster ROM storage area. It has a fixed base address (000C_0h) and fix size of 256 KB. The 13 sections from 768 KB to 1 MB comprise what is also known as the PAM Memory Area. Each section has Read enable and Write enable attributes.



Figure 2-4. PAM Region Space



The PAM registers are mapped in Device 0 configuration space.

- ISA Expansion Area (C_0h – D_FFFFh)
- Extended System BIOS Area (E_0h – E_FFFFh)
- System BIOS Area (F_0h – F_FFFFh)

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

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

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

Graphics translated requests to this region are not allowed. If such a mapping error occurs, the request will be routed to C_0h. Writes will have the byte enables de-asserted.



2.5 Main Memory Address Range (1 MB – TOLUD)

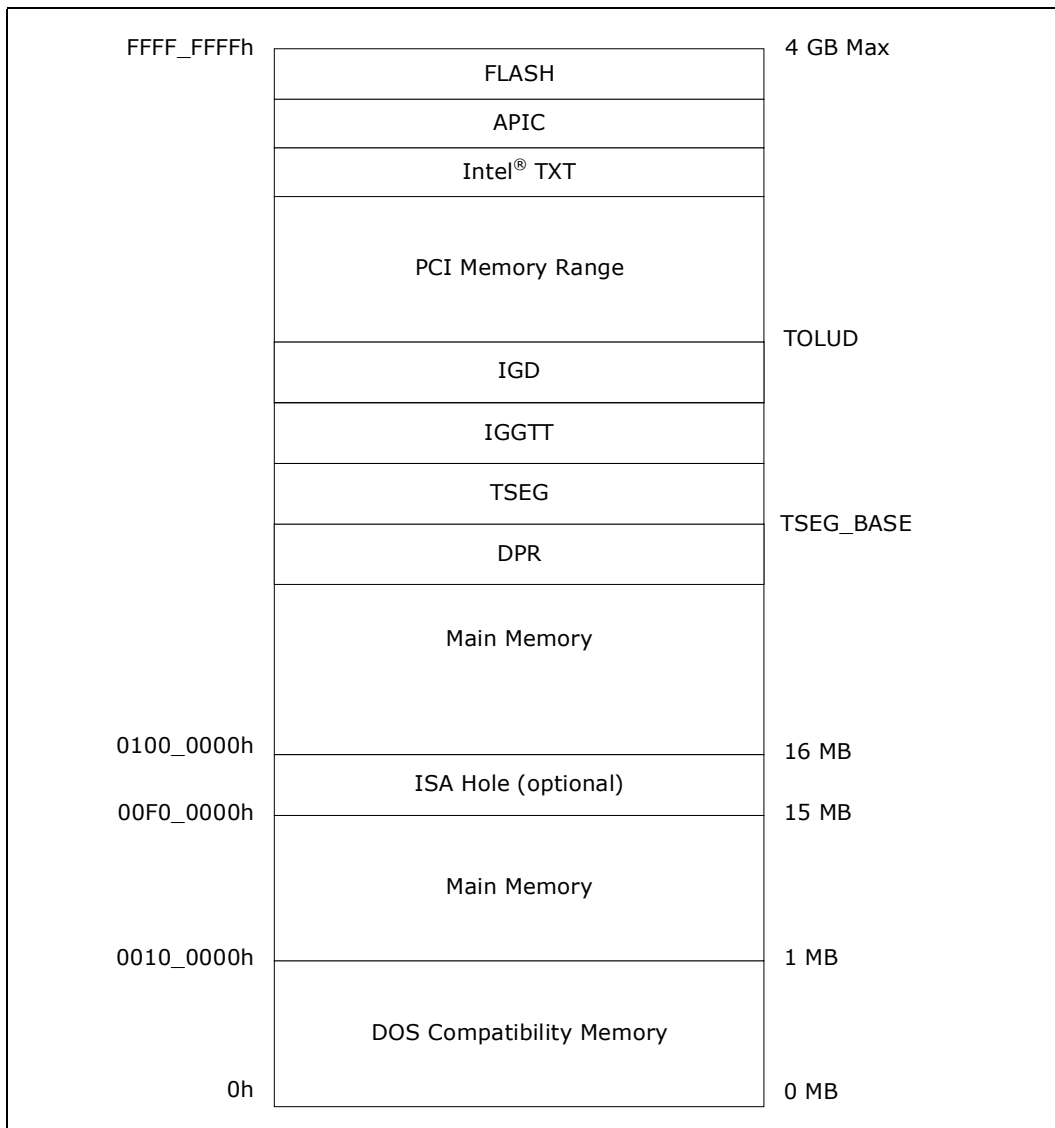
This address range extends from 1 MB to the top of Low Usable physical memory that is permitted to be accessible by the processor (as programmed in the TOLUD register). The processor will route all addresses within this range to the DRAM unless it falls into the optional TSEG, optional ISA Hole, or optional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

This address range is divided into two sub-ranges:

- 1 MB to TSEGMB
- TSEGMB to TOLUD

TSEGMB indicates the TSEG Memory Base address.

Figure 2-5. Main Memory Address Range





2.5.1 ISA Hole (15 MB –16 MB)

The ISA Hole (starting at address F0_0h) is enabled in the Legacy Access Control Register in Device 0 configuration space. If no hole is created, the processor will route the request to DRAM. If a hole is created, the processor will route the request to DMI, since the request does not target DRAM. These downstream requests will be sent to DMI (subtractive decoding).

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

2.5.2 1 MB to TSEGMB

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

2.5.3 TSEG

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

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

$$\text{TSEGMB} = \text{TOLUD} - \text{DSM SIZE} - \text{GSM SIZE} - \text{TSEG SIZE}$$

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

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

Non-processor originated accesses are not allowed to SMM space. PCI-Express, DMI, and Processor Graphics originated cycles to enabled SMM space are handled as invalid cycle type with reads and writes to location C_0h and byte enables turned off for writes.

2.5.4 Protected Memory Range (PMR) - (programmable)

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

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



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

- **Protected Low-memory Region:** This is defined as the protected memory region below 4 GB to hold the MVMM code/private data, and the initial DMA-remapping structures that control DMA to host physical addresses below 4 GB. DMA-remapping hardware implementations on platforms supporting Intel TXT are required to support protected low-memory region 5.
- **Protected High-memory Region:** This is defined as a variable sized protected memory region above 4 GB, enough to hold the initial DMA-remapping structures for managing DMA accesses to addresses above 4 GB. DMA-remapping hardware implementations on platforms supporting Intel TXT are required to support protected high-memory region 6, if the platform supports main memory above 4 GB.

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

2.5.5 DRAM Protected Range (DPR)

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

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

The system will set up:

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

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

If there were cycles from a rogue device to the new region, then those cycles could use the previous decode until the new decode can ensure PV. No flushing of cycles is required.

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

2.5.6 Pre-allocated Memory

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



2.6 PCI Memory Address Range (TOLUD – 4 GB)

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

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

Device 0 exceptions are:

1. Addresses decoded to the egress port registers (PXPEPBAR).
2. Addresses decoded to the memory mapped range for internal MCH registers (MCHBAR).
3. Addresses decoded to the registers associated with the MCH/PCH Serial Interconnect (DMI) register memory range. (DMIBAR).

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

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

In Processor Graphics configurations, there are exceptions to this rule:

6. Addresses decode to the Processor Graphics translation window (GMADR).
7. Addresses decode to the Processor Graphics translation table or Processor Graphics registers. (GTTMMADR).

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

8. Addresses decoded to the memory mapped window to Graphics Intel VT remap engine registers (GFXVTBAR).
9. Addresses decoded to the memory mapped window to DMI VC1 Intel VT remap engine registers (DMIVC1BAR).
10. Addresses decoded to the memory mapped window to PEG/DMI VC0 Intel VT remap engine registers (VTDPVC0BAR).
11. Tcm accesses (to Intel ME stolen memory) from PCH do not go through Intel VT remap engines.

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

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 Processor Graphics and the PCI Express ports **should NOT** overlap with these ranges.


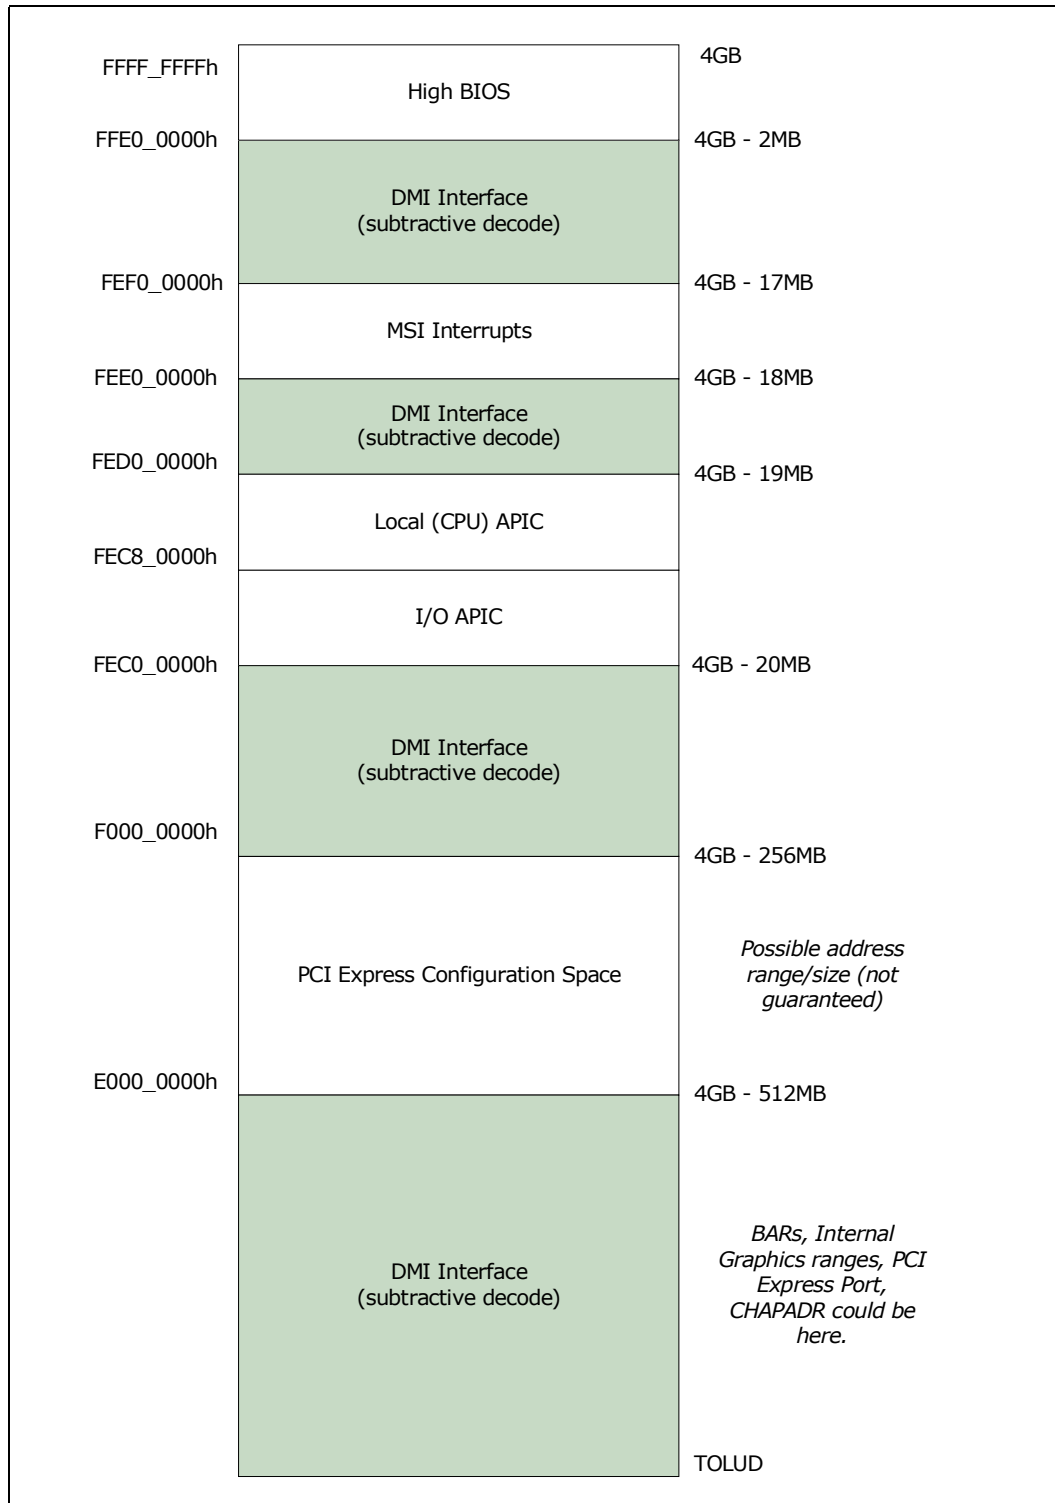


Figure 2-6. PCI Memory Address Range





2.6.1 APIC Configuration Space (FEC0_0h – FECF_FFFFh)

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

The IOAPIC spaces are used to communicate with IOAPIC interrupt controllers that may be populated in the system. Since it is difficult to relocate an interrupt controller using plug-and-play software, fixed address decode regions have been allocated for them. Processor accesses to the default IOAPIC region (FEC0_0h to FEC7_FFFFh) are always forwarded to DMI.

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

Memory requests to this range would then be forwarded to the PCI Express* port. This mode is intended for the entry Workstation/Server SKU of the PCH, and would be disabled in typical Desktop systems. When disabled, any access within the entire APIC Configuration space (FEC0_0h to FECF_FFFFh) is forwarded to DMI.

2.6.2 HSEG (FEDA_0h – FEDB_FFFFh)

This decode range is not supported on this processor platform.

2.6.3 MSI Interrupt Memory Space (FEE0_0h – FEEF_FFFFh)


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

2.6.4 High BIOS Area

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

The top 2 MB (FFE0_0h – FFFF_FFFFh) of the PCI Memory Address Range is reserved for System BIOS (High BIOS), extended BIOS for PCI devices, and the A20 alias of the system BIOS.

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



2.7 Main Memory Address Space (4 GB to TOUUD)

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

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

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

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

2.7.1 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).

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

2.7.2 Top of Upper Usable DRAM (TOUUD)

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

2.7.3 Top of Low Usable DRAM (TOLUD)

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

2.7.4 TSEG_BASE

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



2.7.5 Memory Re-claim Background

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

- High BIOS
- TSEG
- GFX stolen
- GTT stolen
- XAPIC
- Local APIC
- MSI Interrupts
- Mbase/Mlimit
- Pmbase/Pmlimit
- Memory Mapped IO space that supports only 32B addressing

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

2.7.6 Indirect Accesses to MCHBAR Registers

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

- Direct MCHBAR access decode:
 - Cycle to memory from processor
 - Hits MCHBAR base, AND
 - MCHBAR is enabled, AND
 - Within MMIO space (above and below 4 GB)
- GTTMMADR (10h – 13FFFh) range -> MCHBAR decode:
 - Cycle to memory from processor, AND
 - Device 2 (Processor Graphics) is enabled, AND
 - Memory accesses for device 2 is enabled, AND
 - Targets GFX MMIO Function 0, AND
 - MCHBAR is enabled or cycle is a read. If MCHBAR is disabled, only read access is allowed
- MCHTMBAR -> MCHBAR (Thermal Monitor)
 - Cycle to memory from processor, AND
 - Targets MCHTMBAR base
- IOBAR -> GTTMMADR -> MCHBAR
 - Follows IOBAR rules. Refer GTTMMADR information above as well



2.7.7 Memory Remapping

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

2.7.8 Hardware Remap Algorithm

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

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

2.8 PCI Express* Configuration Address Space

PCIEXBAR is located in Device 0 configuration space. The processor detects memory accesses targeting PCIEXBAR. BIOS should assign this address range such that it will not conflict with any other address ranges.

2.9 Graphics Memory Address Ranges

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

- The Graphics Memory Aperture Base Register (GMADR) is used to access graphics memory allocated using the graphics translation table.
- The Graphics Translation Table Base Register (GTTADR) is used to access the translation table and graphics control registers. This is part of the GTTMADR register.

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

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

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



2.9.1 IOBAR Mapped Access to Device 2 MMIO Space

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

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

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

The result of accesses through IOBAR can be:

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

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

This mechanism to access Processor Graphics MMIO registers should NOT be used to access VGA I/O registers that are mapped through the MMIO space. VGA registers should be accessed directly through the dedicated VGA I/O ports.

2.9.2 Trusted Graphics Ranges

Trusted graphics ranges are NOT supported.

2.10 System Management Mode (SMM)

The Core handles all SMM mode transaction routing. The platform does not support HSEG, and the processor will does not allow I/O devices access to CSEG/TSEG/HSEG ranges.

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

Table 2-5. SMM Regions

SMM Space Enabled	Transaction Address Space	DRAM Space (DRAM)
Compatible (C)	000A_0h to 000B_FFFFh	000A_0h to 000B_FFFFh
TSEG (T)	(TOLUD - STOLEN - TSEG) to TOLUD - STOLEN	(TOLUD - STOLEN - TSEG) to TOLUD - STOLEN



2.11 SMM and VGA Access Through GTT TLB

Accesses through GTT TLB address translation SMM DRAM space are not allowed. Writes will be routed to memory address 000C_0h with byte enables de-asserted and reads will be routed to Memory address 000C_0h. If a GTT TLB translated address hits SMM DRAM space, an error is recorded.

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.

PCI Express and DMI Interface write accesses through the GMADR range will not be snooped. Only PCI Express and DMI accesses to GMADR linear range (defined using fence registers) are supported. PCI Express and DMI Interface tileY and tileX writes to GMADR are not supported. If, when translated, the resulting physical address is to enable SMM DRAM space, the request will be remapped to address 000C_0h with de-asserted byte enables.

PCI Express and DMI Interface read accesses to the GMADR range are not supported. Therefore, there are no address translation concerns. PCI Express and DMI Interface reads to GMADR will be remapped to address 000C_0h. The read will complete with UR (unsupported request) completion status.

GTT fetches are always decoded (at fetch time) to ensure fetch is not in SMM (actually, anything above base of TSEG or 640 KB - 1 MB). Thus, the fetches will be invalid and go to address 000C_0h. This is not specific to PCI Express or DMI; it also applies to processor or Processor Graphics engines.

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

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

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

2.13 I/O Address Space

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

The processor allows 64K+3 bytes to be addressed within the I/O space. The upper 3 locations can be accessed only during I/O address wrap-around.



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

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

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

I/O reads that lie within 8-byte boundaries but cross 4-byte boundaries are issued from the processor as one transaction. The reads will be split into two separate transactions. I/O writes that lie within 8-byte boundaries but cross 4-byte boundaries will be split into two transactions by the processor.

2.13.1 PCI Express* I/O Address Mapping

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

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

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

$$\text{I/O_Base_Address} \leq \text{processor I/O Cycle Address} \leq \text{I/O_Limit_Address}$$

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

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



2.14 Direct Media Interface (DMI) Interface Decode Rules

Note: DMI does not apply to U Processors.

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

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

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

- All snoop memory read and write accesses to Main DRAM including PAM region (except stolen memory ranges, TSEG, A0h – BFFFFh space).
- Write accesses to enabled VGA range, MBASE/MLIMIT, and PMBASE/PMLIMIT will be routed as peer cycles to the PCI Express interface.
- Write accesses above the top of usable DRAM and below 4 GB (not decoding to PCI Express or GMADR space) will be treated as master aborts.
- Read accesses above the top of usable DRAM and below 4 GB (not decoding to PCI Express) will be treated as unsupported requests.
- Reads and accesses above the TOUUD will be treated as unsupported requests on VC0.

DMI Interface memory read accesses that fall between TOLUD and 4 GB are considered invalid and will master abort. These invalid read accesses will be reassigned to address 000C_0h and dispatch to DRAM. Reads will return unsupported request completion. Writes targeting PCI Express space will be treated as peer-to-peer cycles.

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

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

2.14.1 DMI Accesses to the Processor that Cross Device Boundaries

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

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


If the starting transaction address hits an invalid address, the entire transaction will be remapped to address 000C_0h and dispatched to DRAM. A single unsupported request completion will result.



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

- VC0 (enabled by default)
 - Snoop port and Non-snoop Asynchronous transactions are supported.
 - Internal Graphics GMADR writes can occur. These writes will NOT be snooped regardless of the snoop not required (SNR) bit.
 - Processor Graphics GMADR reads (unsupported).
 - Peer writes can occur. The SNR bit is ignored.
 - MSI can occur. These will route and be sent to the cores as Intlogical/IntPhysical interrupts regardless of the SNR bit.
 - VLW messages can occur. These will route and be sent to the cores as VLW messages regardless of the SNR bit.
 - MCTP messages can occur. These are routed in a peer fashion.
- VC1 (Optionally enabled)
 - Supports non-snoop transactions only. (Used for isochronous traffic). The PCI Express* Egress port (PXPEPBAR) should also be programmed appropriately.
 - The snoop not required (SNR) bit should be set. Any transaction with the SNR bit not set will be treated as an unsupported request.
 - MSI and peer transactions are treated as unsupported requests.
 - No "pacer" arbitration or TWRR arbitration will occur. Never remaps to different port. (PCH takes care of Egress port remapping). The PCH meters TCm Intel ME accesses and Intel® High Definition Audio (Intel® HD Audio) TC1 access bandwidth.
 - Processor Graphics GMADR writes and GMADR reads are not supported.
- VCm accesses
 - VCm access only map to Intel ME stolen DRAM. These transactions carry the direct physical DRAM address (no redirection or remapping of any kind will occur). This is how the PCH Intel ME accesses its dedicated DRAM stolen space.
 - DMI block will decode these transactions to ensure only Intel ME stolen memory is targeted, and abort otherwise.
 - VCm transactions will only route non-snoop.
 - VCm transactions will not go through VTd remap tables.
 - The remapbase/remaplimit registers do not apply to VCm transactions.

Figure 2-7. Example: DMI Upstream VC0 Memory Map





2.15 PCI Express* Interface Decode Rules

Note: PCI Express* (PCIe) does not apply to U Processors.

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


- Reads: Sent to DRAM address 000C_0h (non-snooped) and will return "unsuccessful completion".
- Writes: Sent to DRAM address 000C_0h (non-snooped) with byte enables all disabled Peer writes from PEG to DMI are not supported.

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

2.15.1 TC/VC Mapping Details

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

Figure 2-8. PEG Upstream VC0 Memory Map



2.16 Legacy VGA and I/O Range Decode Rules

The legacy 128 KB VGA memory range 000A_0h – 000B_FFFFh can be mapped to Processor Graphics (Device 2), PCI Express (Device 1 Functions), and/or to the DMI interface depending on the programming of the VGA steering bits. Priority for VGA mapping is constant in that the processor always decodes internally mapped devices first. Internal to the processor, decode precedence is always given to Processor Graphics. The processor always positively decodes internally mapped devices, namely



the Processor Graphics. Subsequent decoding of regions mapped to either PCI Express port or the DMI Interface depends on the Legacy VGA configurations bits (VGA Enable and MDAP).

For the remainder of this section, PCI Express can refer to either the device 1 port functions.

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

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

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

Table 2-6. Processor Graphics Frame Buffer Accesses

Memory Access GR06(3:2)	A0h - AFFFFh	B0h - B7FFFh MDA	B8000h - BFFFFh
00	Processor Graphics	Processor Graphics	Processor Graphics
01	Processor Graphics	PCI Express bridge or DMI interface	PCI Express bridge or DMI interface
10	PCI Express bridge or DMI interface	Processor Graphics	PCI Express bridge or DMI interface
11	PCI Express bridge or DMI interface	PCI Express bridge or DMI interface	Processor Graphics

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

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

- Processor Graphics controller in Device 2 is enabled through register DEVEN.D2EN bit 4.
- Processor Graphics VGA in Device 0 Function 0 is enabled through register GGC bit 1.
- Processor Graphics's I/O accesses (PCICMD2 04 – 05h, IOAE bit 0) in Device 2 are enabled.
- VGA I/O decodes for Processor Graphics uses 16 address bits (15:0) there is no aliasing. This is different when compared to a bridge device (Device 1) that used only 10 address bits (A 9:0) for VGA I/O decode.



- VGA I/O input/output address select (VGA Miscellaneous Output register - MSR Register, bit 0) is used to select mapping of I/O access as defined in the following table.

Table 2-7. Processor Graphics VGA I/O Mapping

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

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

For regions mapped outside of the Processor Graphics (or if Processor Graphics is disabled), the legacy VGA memory range A0h – BFFFFh are mapped to the DMI Interface or PCI Express depending on the programming of the VGA Enable bit in the BCTRL configuration register in the PEG configuration space, and the MDAPxx bits in the Legacy Access Control (LAC) register in Device 0 configuration space. The same register controls mapping VGA I/O address ranges. The VGA I/O range is defined as addresses where A[9:0] are in the ranges 3B0h to 3BBh and 3C0h to 3DFh (inclusive of ISA address aliases – A[15:10] are not decoded). The function and interaction of these two bits is described below:

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

- Memory accesses in the range 0A0h to 0BFFFFh
- I/O addresses where A[9:0] are in the ranges 3B0h to 3BBh and 3C0h to 3DFh (including ISA address aliases – A[15:10] are not decoded)

When this bit is set to a "1":

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

When this bit is set to a "0":

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

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



Table 2-8. VGA and MDA IO Transaction Mapping

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

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

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

Table 2-9. MDA Resources

Range Type	Address
Memory	0B0h – 0B7FFFh
I/O	3B4h, 3B5h, 3B8h, 3B9h, 3BAh, 3BFh (Including ISA address aliases, A[15:10] are not used in decode)

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

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

Summary of decode priority:

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

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



- Else, if MDA Present (if VGA on PEG is enabled), DMI gets:
 - x3B4,5,8,9,A,F (any access with any of these bytes enabled, regardless of the other BEs)
- Else, if VGA on PEG is enabled, PEG gets:
 - x3B0h – x3BBh
 - x3C0h – x3CFh
 - x3D0h – x3DFh
- Else, if ISA Enable=1, DMI gets:
 - upper 768 bytes of each 1K block
- Else, IOBASE/IOLIMIT apply.

2.17 I/O Mapped Registers

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

§ §



3 Host Bridge/DRAM Registers

Table 3-1. Summary of Bus: 0, Device: 0, Function: 0 (CFG) (Sheet 1 of 2)

Offset	Size (Bytes)	Register Name (Register Symbol)	Default Value
0–1h	2	Vendor Identification (VID)—Offset 0h	8086h
2–3h	2	Device Identification (DID)—Offset 2h	3EXXh
4–5h	2	PCI Command (PCICMD)—Offset 4h	6h
6–7h	2	PCI Status (PCISTS)—Offset 6h	90h
8–8h	1	Revision Identification (RID)—Offset 8h	0h
9–Bh	3	Class Code (CC)—Offset 9h	60h
E–Eh	1	Header Type (HDR)—Offset Eh	0h
2C–2Dh	2	Subsystem Vendor Identification (SVID)—Offset 2Ch	0h
2E–2Fh	2	Subsystem Identification (SID)—Offset 2Eh	0h
34–34h	1	Capabilities Pointer (CAPPTR)—Offset 34h	E0h
40–47h	8	PCI Express* Egress Port Base Address (PXPEPBAR)—Offset 40h	0h
48–4Fh	8	Host Memory Mapped Register Range Base (MCHBAR)—Offset 48h	0h
50–51h	2	GMCH Graphics Control Register (GGC)—Offset 50h	500h
54–57h	4	Device Enable (DEVEN)—Offset 54h	84BFh
58–5Bh	4	Protected Audio Video Path Control (PAVPC)—Offset 58h	0h
5C–5Fh	4	DMA Protected Range (DPR)—Offset 5Ch	0h
60–67h	8	PCI Express Register Range Base Address (PCIEXBAR)—Offset 60h	0h
68–6Fh	8	Root Complex Register Range Base Address (DMIBAR)—Offset 68h	0h
70–77h	8	Manageability Engine Base Address Register (MESEG)—Offset 70h	7FFF00h
78–7Fh	8	Manageability Engine Limit Address Register (MESEG)—Offset 78h	0h
80–80h	1	Programmable Attribute Map 0 (PAM0)—Offset 80h	0h
81–81h	1	Programmable Attribute Map 1 (PAM1)—Offset 81h	0h
82–82h	1	Programmable Attribute Map 2 (PAM2)—Offset 82h	0h
83–83h	1	Programmable Attribute Map 3 (PAM3)—Offset 83h	0h
84–84h	1	Programmable Attribute Map 4 (PAM4)—Offset 84h	0h
85–85h	1	Programmable Attribute Map 5 (PAM5)—Offset 85h	0h
86–86h	1	Programmable Attribute Map 6 (PAM6)—Offset 86h	0h
87–87h	1	Legacy Access Control (LAC)—Offset 87h	0h
88–88h	1	System Management RAM Control (SMRAMC)—Offset 88h	2h
90–97h	8	Remap Base Address Register (REMAPBASE)—Offset 90h	7FFF00h
98–9Fh	8	Remap Limit Address Register (REMAPLIMIT)—Offset 98h	0h
A0–A7h	8	Top of Memory (TOM)—Offset A0h	7FFF00h
A8–AFh	8	Top of Upper Usable DRAM (TOUUD)—Offset A8h	0h
B0–B3h	4	Base Data of Stolen Memory (BDSM)—Offset B0h	0h
B4–B7h	4	Base of GTT stolen Memory (BGSM)—Offset B4h	100h



Table 3-1. Summary of Bus: 0, Device: 0, Function: 0 (CFG) (Continued) (Sheet 2 of 2)

Offset	Size (Bytes)	Register Name (Register Symbol)	Default Value
B8-BBh	4	TSEG Memory Base (TSEGMB)—Offset B8h	0h
BC-BFh	4	Top of Low Usable DRAM (TOLUD)—Offset BCh	100h
C8-C9h	2	Error Status (ERRSTS)—Offset C8h	0h
CA-CBh	2	Error Command (ERRCMD)—Offset CAh	0h
CC-CDh	2	SMI Command (SMICMD)—Offset CCh	0h
CE-CFh	2	SCI Command (SCICMD)—Offset CEh	0h
DC-DFh	4	Scratchpad Data (SKPD)—Offset DCh	0h
E4-E7h	4	Capabilities A (CAPID0)—Offset E4h	0h
E8-EBh	4	Capabilities B (CAPID0)—Offset E8h	0h
EC-EFh	4	Capabilities C (CAPID0)—Offset ECh	0h

3.1 Vendor Identification (VID)—Offset 0h


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

Access Method


Type: CFG
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + 0h

Default: 8086h



Bit Range	Default and Access	Field Name (ID): Description
15:0	8086h RO	VID: Vendor Identification Number: PCI standard identification for Intel.



3.2 Device Identification (DID)—Offset 2h

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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + 2h

Default: 3EXXh

15	12	8	4	0
0	0	1	1	1
1	1	1	0	X
X	X	X	X	X
DID_MSB				DID_SKU

Bit Range	Default and Access	Field Name (ID): Description
15:8	3Eh RO	DID_MSB: Device Identification Number MSB: This is the upper part of device identification assigned to the processor.
7:0	XXh ROV	DID_SKU: Device Identification Number SKU: This is the lower part of device identification assigned to the processor.

3.3 PCI Command (PCICMD)—Offset 4h

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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + 4h

Default: 6h

15	12	8	4	0
0	0	0	0	0
0	0	0	0	0
0	0	0	0	0
RSVD		FB2B	SERRE	ADSTEP
		PERRE	VGASNOOP	MWIE
		SCE	BME	MAE
				IOAE



Bit Range	Default and Access	Field Name (ID): Description
15:10	0h RO	Reserved (RSVD): Reserved.
9	0h RO	FB2B: Fast Back-to-Back Enable: This bit controls whether or not the master can do fast back-to-back write. Since device 0 is strictly a target this bit is not implemented and is hardwired to 0. Writes to this bit position have no effect.
8	0h RW	SERRE: SERR Enable: This bit is a global enable bit for Device 0 SERR messaging. The Processor communicates the SERR condition by sending an SERR message over DMI to the PCH. 1: The Processor is enabled to generate SERR messages over DMI for specific Device 0 error conditions that are individually enabled in the ERRCMD and DMIUEMSK registers. The error status is reported in the ERRSTS, PCISTS, and DMIUEST registers. 0: The SERR message is not generated by the Host for Device 0. This bit only controls SERR messaging for Device 0. Other integrated devices have their own SERRE bits to control error reporting for error conditions occurring in each device. The control bits are used in a logical OR manner to enable the SERR DMI message mechanism. OPI N/A
7	0h RO	ADSTEP: Address/Data Stepping Enable: Address/data stepping is not implemented in the CPU, and this bit is hardwired to 0. Writes to this bit position have no effect.
6	0h RW	PERRE: OPI - N/A Parity Error Enable: Controls whether or not the Master Data Parity Error bit in the PCI Status register can bet set. 0: Master Data Parity Error bit in PCI Status register can NOT be set. 1: Master Data Parity Error bit in PCI Status register CAN be set.
5	0h RO	VGASNOOP: VGA Palette Snoop Enable: The Processor does not implement this bit and it is hardwired to a 0. Writes to this bit position have no effect.
4	0h RO	MWIE: Memory Write and Invalidate Enable: The Processor will never issue memory write and invalidate commands. This bit is therefore hardwired to 0. Writes to this bit position will have no effect.
3	0h RO	SCE: Special Cycle Enable: The Processor does not implement this bit and it is hardwired to a 0. Writes to this bit position have no effect.
2	1h RO	BME: Bus Master Enable: The Processor is always enabled as a master on the backbone. This bit is hardwired to a '1'. Writes to this bit position have no effect.
1	1h RO	MAE: Memory Access Enable: The Processor always allows access to main memory, except when such access would violate security principles. Such exceptions are outside the scope of PCI control. This bit is not implemented and is hardwired to 1. Writes to this bit position have no effect.
0	0h RO	IOAE: I/O Access Enable: This bit is not implemented in the Processor and is hardwired to a 0. Writes to this bit position have no effect.

3.4 PCI Status (PCISTS)—Offset 6h

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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + 6h

Default: 90h



15			12				8			4		0
0	0	0	0	0	0	0	0	1	0	0	1	0
DPE	SSE	RMAS	RTAS	STAS	DEVT	DPD	FB2B	RSVD	MC66	CLIST	RSVD	

Bit Range	Default and Access	Field Name (ID): Description
15	0h RW1C	DPE: Detected Parity Error: This bit is set when this Device receives a Poisoned TLP.
14	0h RW1C	SSE: Signaled System Error: This bit is set to 1 when Device 0 generates an SERR message over DMI for any enabled Device 0 error condition. Device 0 error conditions are enabled in the PCICMD, ERRCMD, and DMIUEMSK registers. Device 0 error flags are read/reset from the PCISTS, ERRSTS, or DMIUEST registers. Software clears this bit by writing a 1 to it.
13	0h RW1C	RMAS: Received Master Abort Status: This bit is set when the Processor generates a DMI request that receives an Unsupported Request completion packet. Software clears this bit by writing a 1 to it.
12	0h RW1C	RTAS: Received Target Abort Status: This bit is set when the Processor generates a DMI request that receives a Completer Abort completion packet. Software clears this bit by writing a 1 to it.
11	0h RO	STAS: Signaled Target Abort Status: The Processor will not generate a Target Abort DMI completion packet or Special Cycle. This bit is not implemented and is hardwired to a 0. Writes to this bit position have no effect.
10:9	0h RO	DEVT: DEVSEL Timing: These bits are hardwired to "00". Writes to these bit positions have no effect. Device 0 does not physically connect to PCI_A. These bits are set to "00" (fast decode) so that optimum DEVSEL timing for PCI_A is not limited by the Host.
8	0h RW1C	DPD: Master Data Parity Error Detected: This bit is set when DMI received a Poisoned completion from PCH. This bit can only be set when the Parity Error Enable bit in the PCI Command register is set.
7	1h RO	FB2B: Fast Back-to-Back: This bit is hardwired to 1. Writes to these bit positions have no effect. Device 0 does not physically connect to PCI_A. This bit is set to 1 (indicating fast back-to-back capability) so that the optimum setting for PCI_A is not limited by the Host.
6	0h RO	Reserved (RSVD): Reserved.
5	0h RO	MC66: 66 MHz Capable: Does not apply to PCI Express. should be hardwired to 0.
4	1h RO	CLIST: Capability List: This bit is hardwired to 1 to indicate to the configuration software that this device/function implements a list of new capabilities. A list of new capabilities is accessed via register CAPPTR at configuration address offset 34h. Register CAPPTR contains an offset pointing to the start address within configuration space of this device where the Capability Identification register resides.
3:0	0h RO	Reserved (RSVD): Reserved.




3.5 Revision Identification (RID)—Offset 8h

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

Access Method

Type: CFG
 (Size: 8 bits) **Offset:** [B:0, D:0, F:0] + 8h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
7:4	0h RO	RID_MSB: Revision Identification Number MSB: Four MSB of RID
3:0	0h RO	RID: Revision Identification Number: Four LSB of RID


3.6 Class Code (CC)—Offset 9h


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

Access Method

Type: CFG
 (Size: 24 bits) **Offset:** [B:0, D:0, F:0] + 9h

Default: 60h





Bit Range	Default and Access	Field Name (ID): Description
23:16	6h RO	BCC: Base Class Code: This is an 8-bit value that indicates the base class code for the Host Bridge device. This code has the value 06h, indicating a Bridge device.
15:8	0h RO	SUBCC: Sub-Class Code: This is an 8-bit value that indicates the category of Bridge into which the Host Bridge device falls. The code is 00h indicating a Host Bridge.
7:0	0h RO	PI: Programming Interface: This is an 8-bit value that indicates the programming interface of this device. This value does not specify a particular register set layout and provides no practical use for this device.

3.7 Header Type (HDR)—Offset Eh


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

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:0, F:0] + Eh

Default: 0h



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



3.8 Subsystem Vendor Identification (SVID)—Offset 2Ch


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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + 2Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
15:0	0h RW_O	SUBVID: Subsystem Vendor ID: This field should be programmed during boot-up to indicate the vendor of the system board. After it has been written once, it becomes read only.

3.9 Subsystem Identification (SID)—Offset 2Eh

This value is used to identify a particular subsystem.

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + 2Eh

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
15:0	0h RW_O	SUBID: Subsystem ID: This field should be programmed during BIOS initialization. After it has been written once, it becomes read only.



3.10 Capabilities Pointer (CAPPTR)—Offset 34h


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

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:0, F:0] + 34h

Default: E0h



Bit Range	Default and Access	Field Name (ID): Description
7:0	E0h RO	CAPPTR: Capabilities Pointer: Pointer to the offset of the first capability ID register block. In this case the first capability is the product-specific Capability Identifier (CAPID0).

3.11 PCI Express* Egress Port Base Address (PXPEPBAR)—Offset 40h

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

Access Method

Type: CFG
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 40h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:12	0h RW	PXPEPBAR: This field corresponds to bits 38 to 12 of the base address PCI Express Egress Port MMIO configuration space. BIOS will program this register resulting in a base address for a 4KB block of contiguous memory address space. This register ensures that a naturally aligned 4KB space is allocated within the first 512GB of addressable memory space. System Software uses this base address to program the PCI Express Egress Port MMIO register set. All the bits in this register are locked in Intel TXT mode.
11:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW	PXPEPBAREN: 0: PXPEPBAR is disabled and does not claim any memory 1: PXPEPBAR memory mapped accesses are claimed and decoded appropriately This register is locked by Intel TXT.

3.12 Host Memory Mapped Register Range Base (MCHBAR)—Offset 48h

This is the base address for the Host Memory Mapped Configuration space. There is no physical memory within this 32KB window that can be addressed. The 32KB reserved by this register does not alias to any PCI 2.3 compliant memory mapped space. On reset, the Host MMIO Memory Mapped Configuration space is disabled and should be enabled by writing a 1 to MCHBAREN [Dev 0, offset48h, bit 0].

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


The register space contains memory control, initialization, timing, and buffer strength registers; clocking registers; and power and thermal management registers.

Access Method

Type: CFG
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 48h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:15	0h RW	MCHBAR: This field corresponds to bits 38 to 15 of the base address Host Memory Mapped configuration space. BIOS will program this register resulting in a base address for a 32KB block of contiguous memory address space. This register ensures that a naturally aligned 32KB space is allocated within the first 512GB of addressable memory space. System Software uses this base address to program the Host Memory Mapped register set. All the bits in this register are locked in Intel TXT mode.
14:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW	MCHBAREN: 0: MCHBAR is disabled and does not claim any memory 1: MCHBAR memory mapped accesses are claimed and decoded appropriately This register is locked in Intel TXT mode.

3.13 GMCH Graphics Control Register (GGC)—Offset 50h


All the bits in this register are Intel TXT lockable.

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + 50h

Default: 500h





Bit Range	Default and Access	Field Name (ID): Description
15:8	5h RW_L	<p>GMS: This field is used to select the amount of Main Memory that is pre-allocated to support the Processor Graphics device in VGA (non-linear) and Native (linear) modes. The BIOS ensures that memory is pre-allocated only when Processor Graphics is enabled.</p> <p>This register is also Intel TXT lockable.</p> <p>Hardware does not clear or set any of these bits automatically based on Processor Graphics being disabled/enabled.</p> <p>BIOS Requirement: BIOS should not set this field to 0h if IVD (bit 1 of this register) is 0.</p>
7:6	0h RW_L	<p>GGMS: This field is used to select the amount of Main Memory that is pre-allocated to support the Processor Graphics Translation Table. The BIOS ensures that memory is pre-allocated only when Processor 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 derive the base of GSM from DSM only using the GSM size programmed in the register.</p> <p>Hardware functionality in case of programming this value to Reserved is not guaranteed.</p>
5:3	0h RO	Reserved (RSVD): Reserved.
2	0h RW_L	<p>VAMEN: Enables the use of the iGFX engines for Versatile Acceleration.</p> <p>1: iGFX engines are in Versatile Acceleration Mode. Device 2 Class Code is 048000h.</p> <p>0:- iGFX engines are in iGFX Mode. Device 2 Class Code is 030h.</p>
1	0h RW_L	<p>IVD:</p> <p>0: Enable. Device 2 (Processor Graphics) claims VGA memory and IO cycles, the Sub-Class Code within Device 2 Class Code register is 00.</p> <p>1: Disable. Device 2 (Processor Graphics) 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 should not set this bit to 0 if the GMS field (bits 7:3 of this register) pre-allocates no memory.</p> <p>This register is locked by Intel TXT lock.</p>
0	0h RW_KL	GGCLCK: When set to 1b, this bit will lock all bits in this register.

3.14 Device Enable (DEVEN)—Offset 54h

Allows for enabling/disabling of PCI devices and functions that are within the Processor package. The table below the bit definitions describes the behavior of all combinations of transactions to devices controlled by this register.


All the bits in this register are Intel TXT Lockable.

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 54h

Default: 84BFh



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15	1h RW_L	D8EN: 0: Bus 0 Device 8 is disabled and not visible. 1: Bus 0 Device 8 is enabled and visible. This bit will be set to 0b and remain 0b if Device 8 capability is disabled.
14	0h RW	D7EN: 0: Bus 0 Device 7 is disabled and not visible. 1: Bus 0 Device 7 is enabled and visible. Non-production BIOS code should provide a setup option to enable Bus 0 Device 7. When enabled, Bus 0 Device 7 should be initialized in accordance to standard PCI device initialization procedures.
13	0h RW	D6EN: Reserved (RSVD)
12:11	0h RO	Reserved (RSVD): Reserved.
10	1h RW_L	D5EN: 0: Bus 0 Device 5 is disabled and not visible. 1: Bus 0 Device 5 is enabled and visible. This bit will be set to 0b and remain 0b if Device 5 capability is disabled.
9:8	0h RO	Reserved (RSVD): Reserved.
7	1h RW_L	D4EN: 0: Bus 0 Device 4 is disabled and not visible. 1: Bus 0 Device 4 is enabled and visible. This bit will be set to 0b and remain 0b if Device 4 capability is disabled.
6	0h RO	Reserved (RSVD): Reserved.
5	1h RW_L	D3EN: 0: Bus 0 Device 3 is disabled and hidden 1: Bus 0 Device 3 is enabled and visible This bit will be set to 0b and remain 0b if Device 3 capability is disabled.
4	1h RW_L	D2EN: 0: Bus 0 Device 2 is disabled and hidden 1: Bus 0 Device 2 is enabled and visible This bit will be set to 0b and remain 0b if Device 2 capability is disabled.
3	1h RW_L	D1F0EN: 0: Bus 0 Device 1 Function 0 is disabled and hidden. 1: Bus 0 Device 1 Function 0 is enabled and visible. This bit will be set to 0b and remain 0b if PEG10 capability is disabled.

Bit Range	Default and Access	Field Name (ID): Description
2	1h RW_L	D1F1EN: 0: Bus 0 Device 1 Function 1 is disabled and hidden. 1: Bus 0 Device 1 Function 1 is enabled and visible.
1	1h RW_L	D1F2EN: 0: Bus 0 Device 1 Function 2 is disabled and hidden. 1: Bus 0 Device 1 Function 2 is enabled and visible.
0	1h RO	DOEN: Bus 0 Device 0 Function 0 may not be disabled and is therefore hardwired to 1.

3.15 Protected Audio Video Path Control (PAVPC)—Offset 58h

All the bits in this register are locked by Intel TXT. When locked the R/W bits are RO.

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 58h

Default: 0h

3	2	2	2	1	1	8	4	0					
1	8	4	0	6	2								
0	0	0	0	0	0	0	0	0					
0	0	0	0	0	0	0	0	0					
PCMBASE				RSVD2			ASMFEN	RSVD1	OVTATTACK	HVYMODESEL	PAVPLCK	PAVPE	PCME

Bit Range	Default and Access	Field Name (ID): Description
31:20	0h RW_L	PCMBASE: Sizes supported in the processor: 1M, 2M, 4M and 8M. Base value programmed (from Top of Stolen Memory) itself defines the size of the WOPCM. Separate WOPCM size programming is redundant information and not required. Default 1M size programming. 4M recommended for the processor. This register is locked (becomes read-only) when PAVPE = 1b.
19:7	0h RW_L	RSVD2: These bits are reserved for future use.
6	0h RW_L	ASMFEN: ASMF method enabled 0b Disabled (default). 1b Enabled. This register is locked when PAVPLCK is set.
5	0h RW_L	RSVD1: These bits are reserved for future use.



Bit Range	Default and Access	Field Name (ID): Description
4	0h RW_L	OVTATTACK: Override of Unsolicited Connection State Attack and Terminate. 0: Disable Override. Attack Terminate allowed. 1: Enable Override. Attack Terminate disallowed. This register bit is locked when PAVPE is set.
3	0h RW_L	HVYMODESEL: This bit is applicable only for PAVP2 operation mode. This bit is also applicable for PAVP3 mode only if the per-App memory config is disabled due to the clearing of bit 9 in the Crypto Function Control_1 register (address 0x320F0). 0: Lite Mode (Non-Serpent mode) 1: Serpent Mode For enabled PAVP3 mode, this one type boot time programming has been replaced by per-App programming (through the Media Crypto Copy command). Note that PAVP2 or PAVP3 mode selection is done by programming bit 8 of the MFX_MODE - Video Mode register.
2	0h RW_KL	PAVPLCK: This bit locks all writeable contents in this register when set (including itself). Only a hardware reset can unlock the register again. This lock bit needs to be set only if PAVP is enabled (bit 1 of this register is asserted).
1	0h RW_L	PAVPE: 0: PAVP functionality is disabled. 1: PAVP functionality is enabled. This register is locked when PAVPLCK is set.
0	0h RW_L	PCME: This field enables Protected Content Memory within Graphics Stolen Memory. This memory is the same as the WOPCM area, whose size is defined by bit 5 of this register. This register is locked when PAVPLOCK is set. A value of 0 in this field indicates that Protected Content Memory is disabled, and cannot be programmed in this manner when PAVP is enabled. A value of 1 in this field indicates that Protected Content Memory is enabled, and is the only programming option available when PAVP is enabled. (Note that the processor legacy Lite mode programming of PCME bit = 0 is not supported. For non-PAVP3 Mode, even for Lite mode configuration, this bit should be programmed to 1 and HVYMODESEL = 0). This bit should always be programmed to 1 if bits 1 and 2 (PAVPE and PAVP lock bits) are both set. With per-App Memory configuration support, the range check for the WOPCM memory area should always happen when this bit is set, regardless of Lite or Serpent mode, or PAVP2 or PAVP3 mode programming.

3.16 DMA Protected Range (DPR)—Offset 5Ch

DMA protected range register.

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5Ch

Default: 0h

3	2	2	2	1	1	8	4	0			
1	8	4	0	6	2						
0	0	0	0	0	0	0	0	0			
0	0	0	0	0	0	0	0	0			
0	0	0	0	0	0	0	0	0			
0	0	0	0	0	0	0	0	0			
TopOfDPR				RSVD		DPRSIZE		RSVD	EPM	PRS	LOCK



Bit Range	Default and Access	Field Name (ID): Description
31:20	0h ROV	TopOfDPR: Top address + 1 of DPR. This is the base of TSEG. Bits 19:0 of the BASE reported here are 0x0_0000.
19:12	0h RO	Reserved (RSVD): Reserved.
11:4	0h RW_L	<p>DPRSIZE: 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. The amount of memory reported in this field will be protected from all DMA accesses, including translated Processor accesses and graphics. The top of the protected range is the BASE of TSEG -1.</p> <p>Note: If TSEG is not enabled, then the top of this range becomes the base of stolen graphics, or ME stolen space or TOLUD, whichever would have been the location of TSEG, assuming it had been enabled.</p> <p>The DPR range works independently of any other range, including the NoDMA.TABLE protection or the PMRC checks in VTd, and is done post any VTd translation or Intel TXT NoDMA lookup. Therefore incoming cycles are checked against this range after the VTd translation and faulted if they hit this protected range, even if they passed the VTd translation or were clean in the NoDMA lookup.</p> <p>All the memory checks are OR'ed with respect to NOT being allowed to go to memory. So if either PMRC, DPR, NoDMA table lookup, NoDMA.TABLE.PROTECT OR a VTd translation disallows the cycle, then the cycle is not allowed to go to memory. Or in other words, all the above checks should pass before a cycle is allowed to DRAM.</p>
3	0h RO	Reserved (RSVD): Reserved.
2	0h RW_L	<p>EPM: This field controls DMA accesses to the DMA Protected Range (DPR) region.</p> <p>0: DPR is disabled 1: DPR is enabled. All DMA requests accessing DPR region are blocked. HW reports the status of DPR enable/disable through the PRS field in this register.</p>
1	0h ROV	<p>PRS: This field indicates the status of DPR.</p> <p>0: DPR protection disabled 1: DPR protection enabled</p>
0	0h RW_KL	LOCK: All bits which may be updated by SW in this register are locked down when this bit is set.

3.17 PCI Express Register Range Base Address (PCIEXBAR)—Offset 60h

This is the base address for the PCI Express configuration space. This window of addresses contains the 4KB of configuration space for each PCI Express device that can potentially be part of the PCI Express Hierarchy associated with the Uncore. There is no actual physical memory within this window of up to 256MB that can be addressed. The actual size of this range is determined by a field in this register.

Each PCI Express Hierarchy requires a PCI Express BASE register. The Uncore supports one PCI Express Hierarchy. The region reserved by this register does not alias to any PCI2.3 compliant memory mapped space. For example, the range reserved for MCHBAR is outside of PCIEXBAR space.

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



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

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

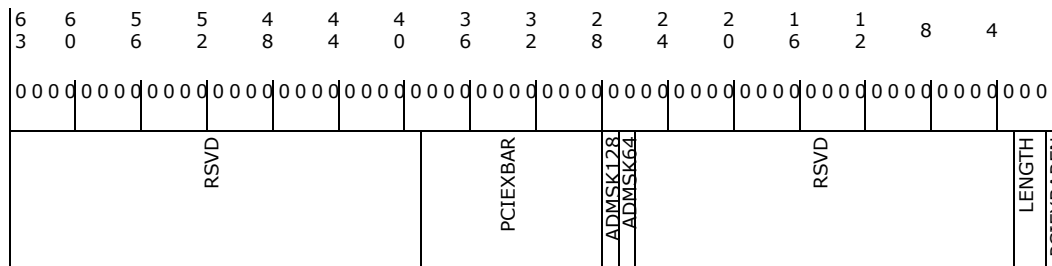
All the bits in this register are locked in Intel TXT mode.

Access Method

Type: CFG
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 60h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:28	0h RW	PCIEXBAR: This field corresponds to bits 38 to 28 of the base address for PCI Express enhanced configuration space. BIOS will program this register resulting in a base address for a contiguous memory address space. The size of the range is defined by bits [2:1] of this register. This Base address shall be assigned on a boundary consistent with the number of buses (defined by the Length field in this register) above TOLUD and still within the 39-bit addressable memory space. The address bits decoded depend on the length of the region defined by this register. This register is locked by Intel TXT. The address used to access the PCI Express configuration space for a specific device can be determined as follows: PCI Express Base Address + Bus Number * 1MB + Device Number * 32KB + Function Number * 4KB This address is the beginning of the 4KB space that contains both the PCI compatible configuration space and the PCI Express extended configuration space.
27	0h RW_V	ADMSK128: This bit is either part of the PCI Express Base Address (R/W) or part of the Address Mask (RO, read 0b), depending on the value of bits [2:1] in this register.
26	0h RW_V	ADMSK64: This bit is either part of the PCI Express Base Address (R/W) or part of the Address Mask (RO, read 0b), depending on the value of bits [2:1] in this register.



Bit Range	Default and Access	Field Name (ID): Description
25:3	0h RO	Reserved (RSVD): Reserved.
2:1	0h RW	<p>LENGTH: This field describes the length of this region.</p> <p>00: 256MB (buses 0-255). Bits 38:28 are decoded in the PCI Express Base Address Field.</p> <p>01: 128MB (buses 0-127). Bits 38:27 are decoded in the PCI Express Base Address Field.</p> <p>10: 64MB (buses 0-63). Bits 38:26 are decoded in the PCI Express Base Address Field.</p> <p>11: Reserved.</p> <p>This register is locked by Intel TXT.</p>
0	0h RW	<p>PCIEXBAREN:</p> <p>0: The PCIEXBAR register is disabled. Memory read and write transactions proceed as if there were no PCIEXBAR register. PCIEXBAR bits 38:26 are R/W with no functionality behind them.</p> <p>1: The PCIEXBAR register is enabled. Memory read and write transactions whose address bits 38:26 match PCIEXBAR will be translated to configuration reads and writes within the Uncore. These Translated cycles are routed as shown in the above table.</p>

3.18 Root Complex Register Range Base Address (DMIBAR)—Offset 68h


This is the base address for the Root Complex configuration space. This window of addresses contains the Root Complex Register set for the PCI Express Hierarchy associated with the Host Bridge. There is no physical memory within this 4KB window that can be addressed. The 4KB reserved by this register does not alias to any PCI 2.3 compliant memory mapped space. On reset, the Root Complex configuration space is disabled and should be enabled by writing a 1 to DMIBAREN [Dev 0, offset 68h, bit 0] All the bits in this register are locked in Intel TXT mode.

Access Method

Type: CFG
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 68h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.



Bit Range	Default and Access	Field Name (ID): Description
38:12	0h RW	DMIBAR: This field corresponds to bits 38 to 12 of the base address DMI configuration space. BIOS will program this register resulting in a base address for a 4KB block of contiguous memory address space. This register ensures that a naturally aligned 4KB space is allocated within the first 512GB of addressable memory space. System Software uses this base address to program the DMI register set. All the Bits in this register are locked in Intel TXT mode.
11:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW	DMIBAREN: 0: DMIBAR is disabled and does not claim any memory 1: DMIBAR memory mapped accesses are claimed and decoded appropriately This register is locked by Intel TXT.

3.19 Manageability Engine Base Address Register (MESEG)—Offset 70h

This register determines the Base Address register of the memory range that is pre-allocated to the Manageability Engine. Together with the MESEG_MASK register it controls the amount of memory allocated to the ME.

This register should be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be 0. Thus, the bottom of the defined memory address range will be aligned to a 1MB boundary.

This register is locked by Intel TXT.


Note: BIOS should program MESEG_BASE and MESEG_MASK so that ME Stolen Memory is carved out from TOM.

Access Method

Type: CFG
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 70h

Default: 7FFFF00h





Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:20	7FFFFh RW_L	MEBASE: Corresponds to A[38:20] of the base address memory range that is allocated to the ME.
19:0	0h RO	Reserved (RSVD): Reserved.

3.20 Manageability Engine Limit Address Register (MESEG)—Offset 78h

This register determines the Mask Address register of the memory range that is pre-allocated to the Manageability Engine. Together with the MESEG_BASE register it controls the amount of memory allocated to the ME.

This register is locked by Intel TXT.

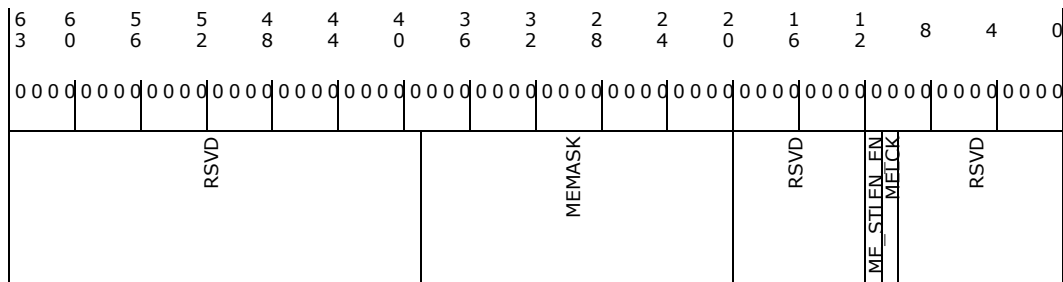
Note: BIOS should program MESEG_BASE and MESEG_MASK so that ME Stolen Memory is carved out from TOM.

Access Method

Type: CFG
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 78h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:20	0h RW_L	MEMASK: This field indicates the bits that should match MEBASE in order to qualify as an ME Memory Range access. For example, if the field is set to 7FFFFh, then ME Memory is 1MB in size. Another example is that if the field is set to 7FFFEh, then ME Memory is 2MB in size. Mask value should be such that once a bit is set to 1 all the more significant bit should be 1. It is not legal to set up mask with 0 and 1's interspersed. In other words, the size of ME Memory Range is limited to power of 2 times 1MB. MEBASE should be naturally aligned to the size of ME region.



Bit Range	Default and Access	Field Name (ID): Description
19:12	0h RO	Reserved (RSVD): Reserved.
11	0h RW_L	ME_STLEN_EN: Indicates whether the ME stolen Memory range is enabled or not.
10	0h RW_KL	MELCK: This field indicates whether all bits in the MESEG_BASE and MESEG_MASK registers are locked. When locked, updates to any field for these registers should be dropped.
9:0	0h RO	Reserved (RSVD): Reserved.

3.21 Programmable Attribute Map 0 (PAM0)—Offset 80h

This register controls the read, write and shadowing attributes of the BIOS range from F_0h to F_FFFFh. The Uncore allows programmable memory attributes on 13 legacy memory segments of various sizes in the 768KB to 1MB address range. Seven Programmable Attribute Map (PAM) registers are used to support these features. Cacheability of these areas is controlled via the MTRR register in the core.

Two bits are used to specify memory attributes for each memory segment. These bits apply to host accesses to the PAM areas. These attributes are:

RE - Read Enable. When RE=1, the host read accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when RE=0, the host read accesses are directed to DMI.

WE - Write Enable. When WE=1, the host write accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when WE=0, the host read accesses are directed to DMI.

The RE and WE attributes permit a memory segment to be Read Only, Write Only, Read/Write or Disabled. For example, if a memory segment has RE=1 and WE=0, the segment is Read Only.

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:0, F:0] + 80h

Default: 0h

7	4	0
0	0	0
RSVD	HTENABLE	Lock



Bit Range	Default and Access	Field Name (ID): Description
7:6	0h RO	Reserved (RSVD): Reserved.
5:4	0h RW_L	HIENABLE: This field controls the steering of read and write cycles that address the BIOS area from 0F_0h to 0F_FFFFh. 00: DRAM Disabled. All accesses are directed to DMI. 01: Read Only. All reads are sent to DRAM, all writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM, all reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.
3:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW_KL	Lock: If this bit is set, all of the PAM* registers are locked (cannot be written)

3.22 Programmable Attribute Map 1 (PAM1)—Offset 81h

This register controls the read, write and shadowing attributes of the BIOS range from C_0h to C_7FFFh. The Uncore allows programmable memory attributes on 13 legacy memory segments of various sizes in the 768KB to 1MB address range. Seven Programmable Attribute Map (PAM) registers are used to support these features. Cacheability of these areas is controlled via the MTRR register in the core.

Two bits are used to specify memory attributes for each memory segment. These bits apply to host accesses to the PAM areas. These attributes are:

RE - Read Enable. When RE=1, the host read accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when RE=0, the host read accesses are directed to DMI.

WE - Write Enable. When WE=1, the host write accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when WE=0, the host read accesses are directed to DMI.

The RE and WE attributes permit a memory segment to be Read Only, Write Only, Read/Write or Disabled. For example, if a memory segment has RE=1 and WE=0, the segment is Read Only.

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:0, F:0] + 81h

Default: 0h

7		4		0
0	0	0	0	0
RSVD		HIENABLE		LOENABLE



Bit Range	Default and Access	Field Name (ID): Description
7:6	0h RO	Reserved (RSVD): Reserved.
5:4	0h RW_L	HIENABLE: This field controls the steering of read and write cycles that address the BIOS area from 0C_4000h to 0C_7FFFh. 00: DRAM Disabled. All accesses are directed to DMI. 01: Read Only. All reads are sent to DRAM, all writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM, all reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.
3:2	0h RO	Reserved (RSVD): Reserved.
1:0	0h RW_L	LOENABLE: This field controls the steering of read and write cycles that address the BIOS area from 0C0h to 0C3FFFh. 00: DRAM Disabled. All reads are sent to DRAM. All writes are forwarded to DMI. 01: Read Only. All reads are sent to DRAM. All writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM. All reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.

3.23 Programmable Attribute Map 2 (PAM2)—Offset 82h

This register controls the read, write and shadowing attributes of the BIOS range from C_8000h to C_FFFFh. The Uncore allows programmable memory attributes on 13 legacy memory segments of various sizes in the 768KB to 1MB address range. Seven Programmable Attribute Map (PAM) registers are used to support these features. Cacheability of these areas is controlled via the MTRR register in the core.

Two bits are used to specify memory attributes for each memory segment. These bits apply to host accesses to the PAM areas. These attributes are:

RE - Read Enable. When RE=1, the host read accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when RE=0, the host read accesses are directed to DMI.

WE - Write Enable. When WE=1, the host write accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when WE=0, the host read accesses are directed to DMI.

The RE and WE attributes permit a memory segment to be Read Only, Write Only, Read/Write or Disabled. For example, if a memory segment has RE=1 and WE=0, the segment is Read Only.

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:0, F:0] + 82h

Default: 0h



7	4	0
0	0	0
RSVD	HIENABLE	LOENABLE

Bit Range	Default and Access	Field Name (ID): Description
7:6	0h RO	Reserved (RSVD): Reserved.
5:4	0h RW_L	HIENABLE: This field controls the steering of read and write cycles that address the BIOS area from 0CC000h to 0CBFFFh. 00: DRAM Disabled. All accesses are directed to DMI. 01: Read Only. All reads are sent to DRAM, all writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM, all reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.
3:2	0h RO	Reserved (RSVD): Reserved.
1:0	0h RW_L	LOENABLE: This field controls the steering of read and write cycles that address the BIOS area from 0C8000h to 0CBFFFh. 00: DRAM Disabled. All reads are sent to DRAM. All writes are forwarded to DMI. 01: Read Only. All reads are sent to DRAM. All writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM. All reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.

3.24 Programmable Attribute Map 3 (PAM3)—Offset 83h

This register controls the read, write and shadowing attributes of the BIOS range from D0h to D7FFFh. The Uncore allows programmable memory attributes on 13 legacy memory segments of various sizes in the 768KB to 1MB address range. Seven Programmable Attribute Map (PAM) registers are used to support these features. Cacheability of these areas is controlled via the MTRR register in the core.

Two bits are used to specify memory attributes for each memory segment. These bits apply to host accesses to the PAM areas. These attributes are:

RE - Read Enable. When RE=1, the host read accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when RE=0, the host read accesses are directed to DMI.

WE - Write Enable. When WE=1, the host write accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when WE=0, the host read accesses are directed to DMI.

The RE and WE attributes permit a memory segment to be Read Only, Write Only, Read/Write or Disabled. For example, if a memory segment has RE=1 and WE=0, the segment is Read Only.

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:0, F:0] + 83h



Default: 0h

7	0	0	0	4	0	0	0	0
0	0	0	0	0	0	0	0	0
RSVD				HIENABLE		RSVD		LOENABLE

Bit Range	Default and Access	Field Name (ID): Description
7:6	0h RO	Reserved (RSVD): Reserved.
5:4	0h RW_L	HIENABLE: This field controls the steering of read and write cycles that address the BIOS area from 0D4000h to 0D7FFFh. 00: DRAM Disabled. All accesses are directed to DMI. 01: Read Only. All reads are sent to DRAM, all writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM, all reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.
3:2	0h RO	Reserved (RSVD): Reserved.
1:0	0h RW_L	LOENABLE: This field controls the steering of read and write cycles that address the BIOS area from 0D0h to 0D3FFFh. 00: DRAM Disabled. All reads are sent to DRAM. All writes are forwarded to DMI. 01: Read Only. All reads are sent to DRAM. All writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM. All reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.

3.25 Programmable Attribute Map 4 (PAM4)—Offset 84h

This register controls the read, write and shadowing attributes of the BIOS range from D8000h to DFFFFh. The Uncore allows programmable memory attributes on 13 legacy memory segments of various sizes in the 768KB to 1MB address range. Seven Programmable Attribute Map (PAM) registers are used to support these features. Cacheability of these areas is controlled via the MTRR register in the core.

Two bits are used to specify memory attributes for each memory segment. These bits apply to host accesses to the PAM areas. These attributes are:

RE - Read Enable. When RE=1, the host read accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when RE=0, the host read accesses are directed to DMI.

WE - Write Enable. When WE=1, the host write accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when WE=0, the host read accesses are directed to DMI.

The RE and WE attributes permit a memory segment to be Read Only, Write Only, Read/Write or Disabled. For example, if a memory segment has RE=1 and WE=0, the segment is Read Only.


Access Method



Type: CFG
(Size: 8 bits)

Offset: [B:0, D:0, F:0] + 84h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
7:6	0h RO	Reserved (RSVD): Reserved.
5:4	0h RW_L	HIENABLE: This field controls the steering of read and write cycles that address the BIOS area from 0DC000h to 0DFFFFh. 00: DRAM Disabled. All accesses are directed to DMI. 01: Read Only. All reads are sent to DRAM, all writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM, all reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.
3:2	0h RO	Reserved (RSVD): Reserved.
1:0	0h RW_L	LOENABLE: This field controls the steering of read and write cycles that address the BIOS area from 0D8000h to 0DBFFFh. 00: DRAM Disabled. All reads are sent to DRAM. All writes are forwarded to DMI. 01: Read Only. All reads are sent to DRAM. All writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM. All reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.

3.26 Programmable Attribute Map 5 (PAM5)—Offset 85h


This register controls the read, write and shadowing attributes of the BIOS range from E_0h to E_7FFFh. The Uncore allows programmable memory attributes on 13 legacy memory segments of various sizes in the 768KB to 1MB address range. Seven Programmable Attribute Map (PAM) registers are used to support these features. Cacheability of these areas is controlled via the MTRR register in the core.

Two bits are used to specify memory attributes for each memory segment. These bits apply to host accesses to the PAM areas. These attributes are:

RE - Read Enable. When RE=1, the host read accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when RE=0, the host read accesses are directed to DMI.

WE - Write Enable. When WE=1, the host write accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when WE=0, the host read accesses are directed to DMI.

The RE and WE attributes permit a memory segment to be Read Only, Write Only, Read/Write or Disabled. For example, if a memory segment has RE=1 and WE=0, the segment is Read Only.



Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:0, F:0] + 85h

Default: 0h

7	4	0
0 0 0 0	0 0	0 0
RSVD	HIENABLE	LOENABLE

Bit Range	Default and Access	Field Name (ID): Description
7:6	0h RO	Reserved (RSVD): Reserved.
5:4	0h RW_L	HIENABLE: This field controls the steering of read and write cycles that address the BIOS area from 0E4000h to 0E7FFFh. 00: DRAM Disabled. All accesses are directed to DMI. 01: Read Only. All reads are sent to DRAM, all writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM, all reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.
3:2	0h RO	Reserved (RSVD): Reserved.
1:0	0h RW_L	LOENABLE: This field controls the steering of read and write cycles that address the BIOS area from 0E0h to 0E3FFFh. 00: DRAM Disabled. All reads are sent to DRAM. All writes are forwarded to DMI. 01: Read Only. All reads are sent to DRAM. All writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM. All reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.

3.27 Programmable Attribute Map 6 (PAM6)—Offset 86h

This register controls the read, write and shadowing attributes of the BIOS range from E_8000h to E_FFFFh. The Uncore allows programmable memory attributes on 13 legacy memory segments of various sizes in the 768KB to 1MB address range. Seven Programmable Attribute Map (PAM) registers are used to support these features. Cacheability of these areas is controlled via the MTRR register in the core.

Two bits are used to specify memory attributes for each memory segment. These bits apply to host accesses to the PAM areas. These attributes are:

RE - Read Enable. When RE=1, the host read accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when RE=0, the host read accesses are directed to DMI.

WE - Write Enable. When WE=1, the host write accesses to the corresponding memory segment are claimed by the Uncore and directed to main memory. Conversely, when WE=0, the host read accesses are directed to DMI.




The RE and WE attributes permit a memory segment to be Read Only, Write Only, Read/Write or Disabled. For example, if a memory segment has RE=1 and WE=0, the segment is Read Only.

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:0, F:0] + 86h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
7:6	0h RO	Reserved (RSVD): Reserved.
5:4	0h RW_L	HIENABLE: This field controls the steering of read and write cycles that address the BIOS area from 0EC000h to 0EFFFFh. 00: DRAM Disabled. All accesses are directed to DMI. 01: Read Only. All reads are sent to DRAM, all writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM, all reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.
3:2	0h RO	Reserved (RSVD): Reserved.
1:0	0h RW_L	LOENABLE: This field controls the steering of read and write cycles that address the BIOS area from 0E8000h to 0EBFFFh. 00: DRAM Disabled. All reads are sent to DRAM. All writes are forwarded to DMI. 01: Read Only. All reads are sent to DRAM. All writes are forwarded to DMI. 10: Write Only. All writes are sent to DRAM. All reads are serviced by DMI. 11: Normal DRAM Operation. All reads and writes are serviced by DRAM.

3.28 Legacy Access Control (LAC)—Offset 87h

This 8-bit register controls steering of MDA cycles and a fixed DRAM hole from 15-16MB.

There can only be at most one MDA device in the system.

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:0, F:0] + 87h

Default: 0h




7	0	0	0	4	0	0	0	0
HEN			RSVD		MDAP60	MDAP12	MDAP11	MDAP10

Bit Range	Default and Access	Field Name (ID): Description															
7	0h RW	<p>HEN: This field enables a memory hole in DRAM space. The DRAM that lies "behind" this space is not remapped. 0: No memory hole. 1: Memory hole from 15MB to 16MB. This bit is Intel TXT lockable.</p>															
6:4	0h RO	<p>Reserved (RSVD): Reserved.</p>															
3	0h RW	<p>MDAP60: This bit works with the VGA Enable bits in the BCTRL register of Device 1 Function 2 to control the routing of Processor initiated transactions targeting MDA compatible I/O and memory address ranges. This bit should not be set if device 1 function 2 VGA Enable bit is not set. If device 1 function 2 VGA enable bit is not set, then accesses to IO address range x3BCh-x3BFh remain on the backbone. If the VGA enable bit is set and MDA is not present, then accesses to IO address range x3BCh-x3BFh are forwarded to PCI Express through device 1 function 2 if the address is within the corresponding IOBASE and IOLIMIT, otherwise they remain on the backbone. MDA resources are defined as the following: Memory: 0B0h - 0B7FFFh I/O: 3B4h, 3B5h, 3B8h, 3B9h, 3BAh, 3BFh, (including ISA address aliases, A[15:10] are not used in decode) Any I/O reference that includes the I/O locations listed above, or their aliases, will remain on the backbone even if the reference also includes I/O locations not listed above. The following table shows the behavior for all combinations of MDA and VGA:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">VGAEN</th> <th style="text-align: center;">MDAP</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">0</td> <td style="text-align: center;">0</td> <td>All References to MDA and VGA space are not claimed by Device 1 Function 2.</td> </tr> <tr> <td style="text-align: center;">0</td> <td style="text-align: center;">1</td> <td>Illegal combination</td> </tr> <tr> <td style="text-align: center;">1</td> <td style="text-align: center;">0</td> <td>All VGA and MDA references are routed to PCI Express Graphics Attach device 1 function 2.</td> </tr> <tr> <td style="text-align: center;">1</td> <td style="text-align: center;">1</td> <td>All VGA references are routed to PCI Express Graphics Attach device 1 function 2. MDA references are not claimed by device 1 function 2. VGA and MDA memory cycles can only be routed across PEG12 when MAE (PCICMD12[1]) is set. VGA and MDA I/O cycles can only be routed across PEG12 if IOAE (PCICMD12[0]) is set.</td> </tr> </tbody> </table>	VGAEN	MDAP	Description	0	0	All References to MDA and VGA space are not claimed by Device 1 Function 2.	0	1	Illegal combination	1	0	All VGA and MDA references are routed to PCI Express Graphics Attach device 1 function 2.	1	1	All VGA references are routed to PCI Express Graphics Attach device 1 function 2. MDA references are not claimed by device 1 function 2. VGA and MDA memory cycles can only be routed across PEG12 when MAE (PCICMD12[1]) is set. VGA and MDA I/O cycles can only be routed across PEG12 if IOAE (PCICMD12[0]) is set.
VGAEN	MDAP	Description															
0	0	All References to MDA and VGA space are not claimed by Device 1 Function 2.															
0	1	Illegal combination															
1	0	All VGA and MDA references are routed to PCI Express Graphics Attach device 1 function 2.															
1	1	All VGA references are routed to PCI Express Graphics Attach device 1 function 2. MDA references are not claimed by device 1 function 2. VGA and MDA memory cycles can only be routed across PEG12 when MAE (PCICMD12[1]) is set. VGA and MDA I/O cycles can only be routed across PEG12 if IOAE (PCICMD12[0]) is set.															



Bit Range	Default and Access	Field Name (ID): Description															
2	0h RW	<p>MDAP12: This bit works with the VGA Enable bits in the BCTRL register of Device 1 Function 2 to control the routing of Processor initiated transactions targeting MDA compatible I/O and memory address ranges. This bit should not be set if device 1 function 2 VGA Enable bit is not set.</p> <p>If device 1 function 2 VGA enable bit is not set, then accesses to IO 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 function 2 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: 0B0h - 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" data-bbox="623 768 1398 1045"> <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 Function 2.</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 PCI Express Graphics Attach device 1 function 2.</td> </tr> <tr> <td>1</td> <td>1</td> <td>All VGA references are routed to PCI Express Graphics Attach device 1 function 2. MDA references are not claimed by device 1 function 2. VGA and MDA memory cycles can only be routed across PEG12 when MAE (PCICMD12[1]) is set. VGA and MDA I/O cycles can only be routed across PEG12 if IOAE (PCICMD12[0]) is set.</td> </tr> </tbody> </table>	VGAEN	MDAP	Description	0	0	All References to MDA and VGA space are not claimed by Device 1 Function 2.	0	1	Illegal combination	1	0	All VGA and MDA references are routed to PCI Express Graphics Attach device 1 function 2.	1	1	All VGA references are routed to PCI Express Graphics Attach device 1 function 2. MDA references are not claimed by device 1 function 2. VGA and MDA memory cycles can only be routed across PEG12 when MAE (PCICMD12[1]) is set. VGA and MDA I/O cycles can only be routed across PEG12 if IOAE (PCICMD12[0]) is set.
VGAEN	MDAP	Description															
0	0	All References to MDA and VGA space are not claimed by Device 1 Function 2.															
0	1	Illegal combination															
1	0	All VGA and MDA references are routed to PCI Express Graphics Attach device 1 function 2.															
1	1	All VGA references are routed to PCI Express Graphics Attach device 1 function 2. MDA references are not claimed by device 1 function 2. VGA and MDA memory cycles can only be routed across PEG12 when MAE (PCICMD12[1]) is set. VGA and MDA I/O cycles can only be routed across PEG12 if IOAE (PCICMD12[0]) is set.															
1	0h RW	<p>MDAP11: This bit works with the VGA Enable bits in the BCTRL register of Device 1 Function 1 to control the routing of Processor initiated transactions targeting MDA compatible I/O and memory address ranges. This bit should not be set if device 1 function 1 VGA Enable bit is not set.</p> <p>If device 1 function 1 VGA enable bit is not set, then accesses to IO 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 function 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: 0B0h - 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" data-bbox="623 1472 1398 1749"> <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 Function 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 PCI Express Graphics Attach device 1 function 1.</td> </tr> <tr> <td>1</td> <td>1</td> <td>All VGA references are routed to PCI Express Graphics Attach device 1 function 1. MDA references are not claimed by device 1 function 1. VGA and MDA memory cycles can only be routed across PEG11 when MAE (PCICMD11[1]) is set. VGA and MDA I/O cycles can only be routed across PEG11 if IOAE (PCICMD11[0]) is set.</td> </tr> </tbody> </table>	VGAEN	MDAP	Description	0	0	All References to MDA and VGA space are not claimed by Device 1 Function 1.	0	1	Illegal combination	1	0	All VGA and MDA references are routed to PCI Express Graphics Attach device 1 function 1.	1	1	All VGA references are routed to PCI Express Graphics Attach device 1 function 1. MDA references are not claimed by device 1 function 1. VGA and MDA memory cycles can only be routed across PEG11 when MAE (PCICMD11[1]) is set. VGA and MDA I/O cycles can only be routed across PEG11 if IOAE (PCICMD11[0]) is set.
VGAEN	MDAP	Description															
0	0	All References to MDA and VGA space are not claimed by Device 1 Function 1.															
0	1	Illegal combination															
1	0	All VGA and MDA references are routed to PCI Express Graphics Attach device 1 function 1.															
1	1	All VGA references are routed to PCI Express Graphics Attach device 1 function 1. MDA references are not claimed by device 1 function 1. VGA and MDA memory cycles can only be routed across PEG11 when MAE (PCICMD11[1]) is set. VGA and MDA I/O cycles can only be routed across PEG11 if IOAE (PCICMD11[0]) is set.															



Bit Range	Default and Access	Field Name (ID): Description															
0	0h RW	<p>MDAP10: This bit works with the VGA Enable bits in the BCTRL register of Device 1 Function 0 to control the routing of Processor initiated transactions targeting MDA compatible I/O and memory address ranges. This bit should not be set if device 1 function 0 VGA Enable bit is not set.</p> <p>If device 1 function 0 VGA enable bit is not set, then accesses to IO 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 function 0 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: 0B0h - 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 Function 0.</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 PCI Express Graphics Attach device 1 function 0.</td> </tr> <tr> <td>1</td> <td>1</td> <td>All VGA references are routed to PCI Express Graphics Attach device 1 function 0. MDA references are not claimed by device 1 function 0. VGA and MDA memory cycles can only be routed across PEG10 when MAE (PCICMD10[1]) is set. VGA and MDA I/O cycles can only be routed across PEG10 if IOAE (PCICMD10[0]) is set.</td> </tr> </tbody> </table>	VGAEN	MDAP	Description	0	0	All References to MDA and VGA space are not claimed by Device 1 Function 0.	0	1	Illegal combination	1	0	All VGA and MDA references are routed to PCI Express Graphics Attach device 1 function 0.	1	1	All VGA references are routed to PCI Express Graphics Attach device 1 function 0. MDA references are not claimed by device 1 function 0. VGA and MDA memory cycles can only be routed across PEG10 when MAE (PCICMD10[1]) is set. VGA and MDA I/O cycles can only be routed across PEG10 if IOAE (PCICMD10[0]) is set.
VGAEN	MDAP	Description															
0	0	All References to MDA and VGA space are not claimed by Device 1 Function 0.															
0	1	Illegal combination															
1	0	All VGA and MDA references are routed to PCI Express Graphics Attach device 1 function 0.															
1	1	All VGA references are routed to PCI Express Graphics Attach device 1 function 0. MDA references are not claimed by device 1 function 0. VGA and MDA memory cycles can only be routed across PEG10 when MAE (PCICMD10[1]) is set. VGA and MDA I/O cycles can only be routed across PEG10 if IOAE (PCICMD10[0]) is set.															

3.29 System Management RAM Control (SMRAMC)—Offset 88h

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

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:0, F:0] + 88h

Default: 2h

7				4				0
0	0	0	0	0	0	0	1	0
RSVD	D_OPEN	D_CLS	D_LCK	G_SMRAME			C_BASE_SEG	



Bit Range	Default and Access	Field Name (ID): Description
7	0h RO	Reserved (RSVD): Reserved.
6	0h RW_LV	D_OPEN: When D_OPEN = 1 and D_LCK = 0, the SMM DRAM space is made visible even when SMM decode is not active. This is intended to help BIOS initialize SMM space. Software should ensure that D_OPEN = 1 and D_CLS = 1 are not set at the same time.
5	0h RW_L	D_CLS: When D_CLS = 1, SMM DRAM space is not accessible to data references, even if SMM decode is active. Code references may still access SMM DRAM space. This will allow SMM software to reference through SMM space to update the display even when SMM is mapped over the VGA range. Software should ensure that D_OPEN = 1 and D_CLS = 1 are not set at the same time.
4	0h RW_KL	D_LCK: When D_LCK=1, then D_OPEN is reset to 0 and all writeable fields in this register are locked (become RO). D_LCK can be set to 1 via a normal configuration space write but can only be cleared by a Full Reset. The combination of D_LCK and D_OPEN provide convenience with security. The BIOS can use the D_OPEN function to initialize SMM space and then use D_LCK to "lock down" SMM space in the future so that no application software (or even BIOS itself) can violate the integrity of SMM space, even if the program has knowledge of the D_OPEN function.
3	0h RW_L	G_SMRAME: If set to '1', then Compatible SMRAM functions are enabled, providing 128KB of DRAM accessible at the A_0h address while in SMM. Once D_LCK is set, this bit becomes RO.
2:0	2h RO	C_BASE_SEG: This field indicates the location of SMM space. SMM DRAM is not remapped. It is simply made visible if the conditions are right to access SMM space, otherwise the access is forwarded to DMI. Only SMM space between A_0h and B_FFFFh is supported, so this field is hardwired to 010b.


3.30 Remap Base Address Register (REMAPBASE)—Offset 90h

Access Method

Type: CFG
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 90h

Default: 7FFFF00h





Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:20	7FFFFh RW_L	REMAPBASE: The value in this register defines the lower boundary of the Remap window. The Remap window is inclusive of this address. In the decoder A[19:0] of the Remap Base Address are assumed to be 0's. Thus the bottom of the defined memory range will be aligned to a 1MB boundary. When the value in this register is greater than the value programmed into the Remap Limit register, the Remap window is disabled. These bits are Intel TXT lockable.
19:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW_KL	LOCK: This bit will lock all writeable settings in this register, including itself.


3.31 Remap Limit Address Register (REMAPLIMIT)—Offset 98h

Access Method

Type: CFG
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 98h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:20	0h RW_L	REMAPLMT: The value in this register defines the upper boundary of the Remap window. The Remap window is inclusive of this address. In the decoder A[19:0] of the remap limit address are assumed to be F's. Thus the top of the defined range will be one byte less than a 1MB boundary. When the value in this register is less than the value programmed into the Remap Base register, the Remap window is disabled. These Bits are Intel TXT lockable.
19:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW_KL	LOCK: This bit will lock all writeable settings in this register, including itself.



3.32 Top of Memory (TOM)—Offset A0h


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

Access Method

Type: CFG
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + A0h

Default: 7FFFF00h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:20	7FFFFh RW_L	TOM: This register reflects the total amount of populated physical memory. This is NOT necessarily the highest main memory address (holes may exist in main memory address map due to addresses allocated for memory mapped IO). These bits correspond to address bits 38:20 (1MB granularity). Bits 19:0 are assumed to be 0. All the bits in this register are locked in Intel TXT mode.
19:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW_KL	LOCK: This bit will lock all writeable settings in this register, including itself.

3.33 Top of Upper Usable DRAM (TOUUD)—Offset A8h



This 64 bit register defines the Top of Upper Usable DRAM. Configuration software should set this value to TOM minus all ME stolen memory if reclaim is disabled. If reclaim is enabled, this value should be set to reclaim limit + 1byte, 1MB aligned, since reclaim limit is 1MB aligned. Address bits 19:0 are assumed to be 000_0h for the purposes of address comparison. The Host interface positively decodes an address towards DRAM if the incoming address is less than the value programmed in this register and greater than or equal to 4GB. BIOS Restriction: Minimum value for TOUUD is 4GB. These bits are Intel TXT lockable.

Access Method

Type: CFG
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + A8h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:20	0h RW_L	TOUUD: This register contains bits 38 to 20 of an address one byte above the maximum DRAM memory above 4G that is usable by the operating system. Configuration software should set this value to TOM minus all ME stolen memory if reclaim is disabled. If reclaim is enabled, this value should be set to reclaim limit 1MB aligned since reclaim limit + 1byte is 1MB aligned. Address bits 19:0 are assumed to be 000_0h for the purposes of address comparison. The Host interface positively decodes an address towards DRAM if the incoming address is less than the value programmed in this register and greater than 4GB. All the bits in this register are locked in Intel TXT mode.
19:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW_KL	LOCK: This bit will lock all writeable settings in this register, including itself.

3.34 Base Data of Stolen Memory (BDSM)—Offset B0h


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

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + B0h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31:20	0h RW_L	BDSM: This register contains bits 31 to 20 of the base address of stolen DRAM memory. BIOS determines the base of graphics stolen memory by subtracting the graphics stolen memory size (PCI Device 0 offset 50 bits 15:8) from TOLUD (PCI Device 0 offset BC bits 31:20).
19:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW_KL	LOCK: This bit will lock all writeable settings in this register, including itself.

3.35 Base of GTT stolen Memory (BGSM)—Offset B4h


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

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + B4h

Default: 100h



Bit Range	Default and Access	Field Name (ID): Description
31:20	1h RW_L	BGSM: This register contains the base address of stolen DRAM memory for the GTT. BIOS determines the base of GTT stolen memory by subtracting the GTT graphics stolen memory size (PCI Device 0 offset 50 bits 7:6) from the Graphics Base of Data Stolen Memory (PCI Device 0 offset B0 bits 31:20).
19:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW_KL	LOCK: This bit will lock all writeable settings in this register, including itself.



3.36 TSEG Memory Base (TSEGMB)—Offset B8h


This register contains the base address of TSEG DRAM memory. BIOS determines the base of TSEG memory which should be at or below Graphics Base of GTT Stolen Memory (PCI Device 0 Offset B4 bits 31:20).

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + B8h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:20	0h RW_L	TSEGMB: This register contains the base address of TSEG DRAM memory. BIOS determines the base of TSEG memory which should be at or below Graphics Base of GTT Stolen Memory (PCI Device 0 Offset B4 bits 31:20). BIOS should program the value of TSEGMB to be the same as BGSM when TSEG is disabled.
19:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW_KL	LOCK: This bit will lock all writeable settings in this register, including itself.

3.37 Top of Low Usable DRAM (TOLUD)—Offset BCh

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

Programming Example:

- C1DRB3 is set to 4 GB
- TSEG is enabled and TSEG size is set to 1 MB
- Processor 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_0h 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: 4GB = 1_0000_0h

The system memory requirements are: 4GB (max addressable space) - 1GB (pci space) - 35MB (lost memory) = 3GB - 35MB (minimum granularity) = 0_ECB0_0h

Since 0_ECB0_0h (PCI and other system requirements) is less than 1_0000_0h, TOLUD should be programmed to ECBh.


These bits are Intel TXT lockable.

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + BCh

Default: 100h



Bit Range	Default and Access	Field Name (ID): Description
31:20	1h RW_L	<p>TOLUD: This register contains bits 31 to 20 of an address one byte above the maximum DRAM memory below 4G that is usable by the operating system. Address bits 31 down to 20 programmed to 01h implies a minimum memory size of 1 MB. Configuration software should set this value to the smaller of the following 2 choices: maximum amount memory in the system minus ME stolen memory plus one byte or the minimum address allocated for PCI memory. Address bits 19:0 are assumed to be 0_0h 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 TXT mode.</p> <p>This register should be 1 MB aligned when reclaim is enabled.</p>
19:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW_KL	LOCK: This bit will lock all writeable settings in this register, including itself.



3.38 Error Status (ERRSTS)—Offset C8h

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.

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + C8h

Default: 0h

15	12	8	4	0
0	0	0	0	0
RSVD				DMERR
				DSERR

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



3.39 Error Command (ERRCMD)—Offset CAh

This register controls the Host Bridge responses to various system errors. Since the Host Bridge 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 will be generated on DMI whenever the corresponding flag is set in the ERRSTS register. The actual generation of the SERR message is globally enabled for Device #0 via the PCI Command register.

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + CAh

Default: 0h

15	12	8	4	0
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0
RSVD				DMERR DSERR

Bit Range	Default and Access	Field Name (ID): Description
15:2	0h RO	Reserved (RSVD): Reserved.
1	0h RW	DMERR: 1: The Host Bridge generates an SERR message over DMI when it detects a multiple-bit error reported by the DRAM controller. 0: Reporting of this condition via SERR messaging is disabled. For systems not supporting ECC this bit should be disabled.
0	0h RW	DSERR: 1: The Host Bridge generates an SERR special cycle over DMI when the DRAM controller detects a single bit error. 0: Reporting of this condition via SERR messaging is disabled. For systems that do not support ECC, this bit should be disabled.

3.40 SMI Command (SMICMD)—Offset CCh


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.

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + CCh

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
15:2	0h RO	Reserved (RSVD): Reserved.
1	0h RW	DMESMI: 1: The Host generates an SMI DMI message when it detects a multiple-bit error reported by the DRAM controller. 0: Reporting of this condition via SMI messaging is disabled. For systems not supporting ECC this bit should be disabled.
0	0h RW	DSESMI: 1: The Host generates an SMI DMI special cycle when the DRAM controller detects a single bit error. 0: Reporting of this condition via SMI messaging is disabled. For systems that do not support ECC this bit should be disabled.

3.41 SCI Command (SCICMD)—Offset CEh


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.

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + CEh

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
15:2	0h RO	Reserved (RSVD): Reserved.
1	0h RW	DMESCI: 1: The Host generates an SCI DMI message when it detects a multiple-bit error reported by the DRAM controller. 0: Reporting of this condition via SCI messaging is disabled. For systems not supporting ECC this bit should be disabled.
0	0h RW	DSMESCI: 1: The Host generates an SCI DMI special cycle when the DRAM controller detects a single bit error. 0: Reporting of this condition via SCI messaging is disabled. For systems that do not support ECC this bit should be disabled.

3.42 Scratchpad Data (SKPD)—Offset DCh


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

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + DCh

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:0	0h RW	SKPD: 1 DWORD of data storage.



3.43 Capabilities A (CAPID0)—Offset E4h


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

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + E4h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:26	0h RO	Reserved (RSVD): Reserved.
25	0h RO	ECCDIS: 0: ECC capable 1: Not ECC capable
24	0h RO	Reserved (RSVD): Reserved.
23	0h RO_KFW	VTDD: 0: Enable VTd 1: Disable VTd
22:15	0h RO	Reserved (RSVD): Reserved.
14	0h RO	DDPCD: Allows Dual Channel operation but only supports 1 DIMM per channel. 0: 2 DIMMs per channel enabled 1: 2 DIMMs per channel disabled. This setting hardwires bits 2 and 3 of the rank population field for each channel to zero. (MCHBAR offset 260h, bits 22-23 for channel 0 and MCHBAR offset 660h, bits 22-23 for channel 1)
13	0h RO	X2APIC_EN: Extended Interrupt Mode. 0: Hardware does not support Extended APIC mode. 1: Hardware supports Extended APIC mode.
12	0h RO	PDCD: 0: Capable of Dual Channels 1: Not Capable of Dual Channel - only single channel capable.
11:0	0h RO	Reserved (RSVD): Reserved.



3.44 Capabilities B (CAPID0)—Offset E8h

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

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + E8h

Default: 0h

3	2	2	2	1	1	8	4	0										
1	8	4	0	6	2													
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0 0 0 0 0	0 0 0 0	0 0 0 0										
IMGU_DIS	RSVD	SMT	CACHESZ	RSVD	PLL_REF100_CFG	PEGG3_DIS	RSVD	ADDGFXEN	ADDGFXCAP	RSVD	DMIG3DIS	RSVD	GMM_DIS	RSVD	DMFC_DDR3	RSVD	LPDDR3_EN	RSVD

Bit Range	Default and Access	Field Name (ID): Description
31	0h RO_KFW	IMGU_DIS: 0: Device 5 associated memory spaces are accessible. 1: Device 5 associated memory and IO spaces are disabled by hardwiring the D1F2EN field, bit 1 of the Device Enable register, (DEVEN Dev 0 Offset 54h) to '0'.
30:29	0h RO	Reserved (RSVD): Reserved.
28	0h RO	SMT: This setting indicates whether or not the Processor is SMT capable.
27:25	0h RO	CACHESZ: This setting indicates the supporting cache sizes.
24:21	0h RO	Reserved (RSVD): Reserved.
20	0h RO	PEGG3_DIS: the processor: PCIe Gen 3 Disable. 0: Capable of running any of the Gen 3-compliant PEG controllers in Gen 3 mode (Devices 0/1/0, 0/1/1, 0/1/2) 1: Not capable of running any of the PEG controllers in Gen 3 mode
19	0h RO	Reserved (RSVD): Reserved.
18	0h RO	ADDGFXEN: 0: Additive Graphics Disabled 1: Additive Graphics Enabled
17	0h RO	ADDGFXCAP: 0: Capable of Additive Graphics 1: Not capable of Additive Graphics
16	0h RO	Reserved (RSVD): Reserved.
15	0h RO	DMIG3DIS: DMI Gen 3 Disable.



Bit Range	Default and Access	Field Name (ID): Description
14:9	0h RO	Reserved (RSVD): Reserved.
8	0h RO_KFW	GMM_DIS: 0: Device 8 associated memory spaces are accessible. 1: Device 8 associated memory and IO spaces are disabled by hardwiring the D8EN field, bit 1 of the Device Enable register, (DEVEN Dev 0 Offset 54h) to '0'.
7	0h RO	Reserved (RSVD): Reserved.
6:4	0h RO	DMFC_DDR3: This field controls which values may be written to the Memory Frequency Select field 6:4 of the Clocking Configuration registers (MCHBAR Offset C00h). Any attempt to write an unsupported value will be ignored. 000: MC capable of DDR3 2667 (2667 is the upper limit) 001: MC capable of up to DDR3 2667 010: MC capable of up to DDR3 2400 011: MC capable of up to DDR3 2133 100: MC capable of up to DDR3 1867 101: MC capable of up to DDR3 1600 110: MC capable of up to DDR3 1333 111: MC capable of up to DDR3 1067
3	0h RO	Reserved (RSVD): Reserved.
2	0h RO	LPDDR3_EN: Allow LPDDR3 operation
1:0	0h RO	Reserved (RSVD): Reserved.

3.45 Capabilities C (CAPID0)—Offset ECh

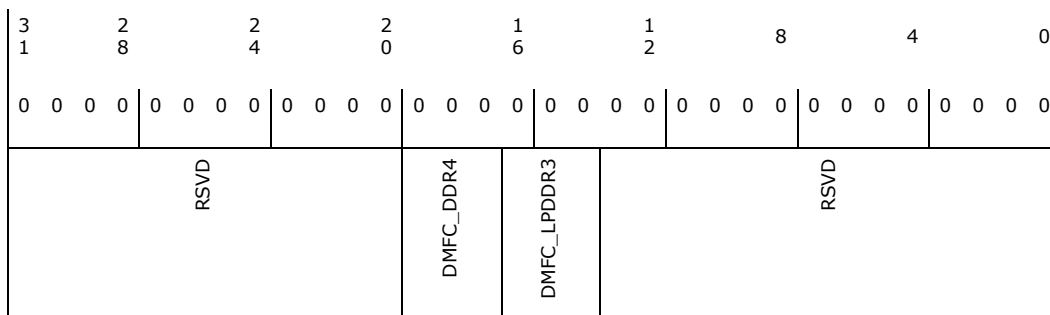
Control of bits in this register are only required for customer visible SKU differentiation.

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + ECh

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31:20	0h RO	Reserved (RSVD): Reserved.
19:17	0h RO	DMFC_DDR4: hardware will update this field with the value of maximum frequency for DDR4.
16:14	0h RO	DMFC_LPDDR3: hardware will update this field with the value of maximum frequency for LPDDR3.
13:0	0h RO	Reserved (RSVD): Reserved.

§ §



4 Processor Graphics Registers

Table 4-1. Summary of Bus: 0, Device: 2, Function: 0 (CFG)

Offset	Size (Bytes)	Register Name (Register Symbol)	Default Value
0-1h	2	Vendor Identification (VID2)—Offset 0h	8086h
2-3h	2	Device Identification (DID2)—Offset 2h	3EXXh
4-5h	2	PCI Command (PCICMD)—Offset 4h	0h
6-7h	2	PCI Status (PCISTS2)—Offset 6h	10h
8-8h	1	Revision Identification (RID2)—Offset 8h	0h
9-8h	3	Class Code (CC)—Offset 9h	30h
C-Ch	1	Cache Line Size (CLS)—Offset Ch	0h
D-Dh	1	Master Latency Timer (MLT2)—Offset Dh	0h
E-Eh	1	Header Type (HDR2)—Offset Eh	0h
10-17h	8	Graphics Translation Table, Memory Mapped Range Address (GTTMMADR)—Offset 10h	4h
18-1Fh	8	Graphics Memory Range Address (GMADR)—Offset 18h	Ch
20-23h	4	I/O Base Address (IOBAR)—Offset 20h	1h
2C-2Dh	2	Subsystem Vendor Identification (SVID2)—Offset 2Ch	0h
2E-2Fh	2	Subsystem Identification (SID2)—Offset 2Eh	0h
30-33h	4	Video BIOS ROM Base Address (ROMADR)—Offset 30h	0h
34-34h	1	Capabilities Pointer (CAPPOINT)—Offset 34h	40h
3C-3Ch	1	Interrupt Line (INTRLINE)—Offset 3Ch	0h
3D-3Dh	1	Interrupt Pin (INTRPIN)—Offset 3Dh	1h
3E-3Eh	1	Minimum Grant (MINGNT)—Offset 3Eh	0h
3F-3Fh	1	Maximum Latency (MAXLAT)—Offset 3Fh	0h
44-47h	4	Capabilities A (CAPID0)—Offset 44h	0h
48-4Bh	4	Capabilities B (CAPID0)—Offset 48h	0h
54-57h	4	Device Enable (DEVEN0)—Offset 54h	84BFh
5C-5Fh	4	Base Data of Stolen Memory (BDSM)—Offset 5Ch	0h
62-62h	1	Multi Size Aperture Control (MSAC)—Offset 62h	1h
70-71h	2	PCI Express Capability Header (PCIECAPHDR)—Offset 70h	AC10h
AC-ADh	2	Message Signaled Interrupts Capability ID (MSI)—Offset ACh	D005h
AE-AFh	2	Message Control (MC)—Offset AEh	0h
B0-B3h	4	Message Address (MA)—Offset B0h	0h
B4-B5h	2	Message Data (MD)—Offset B4h	0h
D0-D1h	2	Power Management Capabilities ID (PMCAPID)—Offset D0h	1h
D2-D3h	2	Power Management Capabilities (PMCAP)—Offset D2h	22h
D4-D5h	2	Power Management Control/Status (PMCS)—Offset D4h	0h



4.1 Vendor Identification (VID2)—Offset 0h


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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:2, F:0] + 0h

Default: 8086h



Bit Range	Default and Access	Field Name (ID): Description
15:0	8086h RO	VID: PCI standard identification for Intel.

4.2 Device Identification (DID2)—Offset 2h


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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:2, F:0] + 2h

Default: 3EXXh



Bit Range	Default and Access	Field Name (ID): Description
15:8	3Eh RO	DID_MSB: This is the upper part of a 16 bit value assigned to the Graphics device. Bits 5 and 4 are updated based on the GFX level by hardware.
7:0	XXh ROV	DID_SKU: These are lower bits of the 16 bit value assigned to the processor Graphics device.



4.3 PCI Command (PCICMD)—Offset 4h

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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:2, F:0] + 4h

Default: 0h

15	12	8	4	0										
0	0	0	0	0										
RSVD				INTDIS	FB2B	SEN	WCC	PER	VPS	MWIE	SCE	BME	MAE	IOAE

Bit Range	Default and Access	Field Name (ID): Description
15:11	0h RO	Reserved (RSVD): Reserved.
10	0h RW	INTDIS: This bit disables the device from asserting INTx#. 0: Enable the assertion of this device's INTx# signal. 1: Disable the assertion of this device's INTx# signal. DO_INTx messages will not be sent to DMI.
9	0h RO	FB2B: Not Implemented. Hardwired to 0.
8	0h RO	SEN: Not Implemented. Hardwired to 0.
7	0h RO	WCC: Not Implemented. Hardwired to 0.
6	0h RO	PER: Not Implemented. Hardwired to 0. Since the Processor Graphics belongs to the category of devices that does not corrupt programs or data in system memory or hard drives, the Processor Graphics ignores any parity error that it detects and continues with normal operation.
5	0h RO	VPS: This bit is hardwired to 0 to disable snooping.
4	0h RO	MWIE: Hardwired to 0. The Processor Graphics does not support memory write and invalidate commands.
3	0h RO	SCE: This bit is hardwired to 0. The Processor Graphics ignores Special cycles.



Bit Range	Default and Access	Field Name (ID): Description
2	0h RW	BME: 0: Disable Processor Graphics bus mastering. 1: Enable the Processor Graphics to function as a PCI compliant master.
1	0h RW	MAE: This bit controls the Processor Graphics's response to memory space accesses. 0: Disable. 1: Enable.
0	0h RW	IOAE: This bit controls the Processor Graphics's response to I/O space accesses. 0: Disable. 1: Enable.

4.4 PCI Status (PCISTS2)—Offset 6h

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 Processor Graphics.

Access Method


Type: CFG
(Size: 16 bits)

Offset: [B:0, D:2, F:0] + 6h

Default: 10h

15		12		8		4		0				
0	0	0	0	0	0	0	0	0				
DPE	SSE	RMAS	RTAS	STAS	DEVT	DPD	FB2B	UDF	C66	CLIST	INTSTS	RSVD

Bit Range	Default and Access	Field Name (ID): Description
15	0h RO	DPE: Since the Processor Graphics does not detect parity, this bit is always hardwired to 0.
14	0h RO	SSE: The Processor Graphics never asserts SERR#, therefore this bit is hardwired to 0.
13	0h RO	RMAS: The Processor Graphics never gets a Master Abort, therefore this bit is hardwired to 0.
12	0h RO	RTAS: The Processor Graphics never gets a Target Abort, therefore this bit is hardwired to 0.
11	0h RO	STAS: Hardwired to 0. The Processor Graphics does not use target abort semantics.
10:9	0h RO	DEVT: N/A. These bits are hardwired to "00".
8	0h RO	DPD: Since Parity Error Response is hardwired to disabled (and the Processor Graphics does not do any parity detection), this bit is hardwired to 0.



Bit Range	Default and Access	Field Name (ID): Description
7	0h RO	FB2B: Hardwired to 0.
6	0h RO	UDF: Hardwired to 0.
5	0h RO	C66: N/A - Hardwired to 0.
4	1h RO	CLIST: This bit is set to 1 to indicate that the register at 34h provides an offset into the function's PCI Configuration Space containing a pointer to the location of the first item in the list.
3	0h RO_V	INTSTS: This bit reflects the state of the interrupt in the device. Only when the Interrupt Disable bit in the command register is a 0 and this Interrupt Status bit is a 1, will the devices INTx# signal be asserted.
2:0	0h RO	Reserved (RSVD): Reserved.

4.5 Revision Identification (RID2)—Offset 8h

This register contains the revision number for Device #2 Functions 0.


These bits are read only and writes to this register have no effect.

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:2, F:0] + 8h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
7:4	0h RO	RID_MSB: Four MSB of RID
3:0	0h RO	RID: Four LSB of RID



4.6 Class Code (CC)—Offset 9h


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

Access Method

Type: CFG
(Size: 24 bits)

Offset: [B:0, D:2, F:0] + 9h

Default: 30h



Bit Range	Default and Access	Field Name (ID): Description
23:16	3h RO_V	BCC: This is an 8-bit value that indicates the base class code. When MGGC0[VAMEN] is 0 this code has the value 03h, indicating a Display Controller. When MGGC0[VAMEN] is 1 this code has the value 04h, indicating a Multimedia Device.
15:8	0h RO_V	SUBCC: When MGGC0[VAMEN] is 0 this value will be determined based on Device 0 GGC register, GMS and IVD fields. 00h: VGA compatible 80h: Non VGA (GMS = "00h" or IVD = "1b") When MGGC0[VAMEN] is 1, this value is 80h, indicating other multimedia device.
7:0	0h RO	PI: When MGGC0[VAMEN] is 0 this value is 00h, indicating a Display Controller. When MGGC0[VAMEN] is 1 this value is 00h, indicating a NOP.

4.7 Cache Line Size (CLS)—Offset Ch


This register is implemented by PCI Express devices as a read-write field for legacy compatibility purposes but has no effect on any PCI Express device behavior.

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:2, F:0] + Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
7:0	0h RW	CLS: This field is implemented by PCI Express devices as a read-write field for legacy compatibility purposes but has no effect on any PCI Express device behavior.

4.8 Master Latency Timer (MLT2)—Offset Dh


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

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:2, F:0] + Dh

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
7:0	0h RO	MLTCV: Hardwired to 0s.

4.9 Header Type (HDR2)—Offset Eh


This register contains the Header Type of the Processor Graphics.

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:2, F:0] + Eh

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
7	0h RO	MFUNC: Indicates if the device is a Multi-Function Device. The Value of this register is hardwired to 0, indicating the processor graphics is a single function.
6:0	0h RO	H: This is a 7-bit value that indicates the Header Code for the Processor Graphics. This code has the value 00h, indicating a type 0 configuration space format.

4.10 Graphics Translation Table, Memory Mapped Range Address (GTTMMADR)—Offset 10h

This register requests allocation for the combined Graphics Translation Table Modification Range and Memory Mapped Range. The range requires 16 MB combined for MMIO and Global GTT aperture, with 2MB of that used by MMIO and 8MB used by GTT. GTTADR will begin at (GTTMMADR + 8 MB) while the MMIO base address will be the same as GTTMMADR. The region between (GTTMMADR + 2MB) - (GTTMMADR + 8MB) is reserved.

For the Global GTT, this range is defined as a memory BAR in graphics device config space. It is an alias into which software is required to write Page Table Entry values (PTEs). Software may read PTE values from the global Graphics Translation Table (GTT). PTEs cannot be written directly into the global GTT memory area.

The device snoops writes to this region in order to invalidate any cached translations within the various TLBs implemented on-chip.


The allocation is for 16 MB and the base address is defined by bits [38:24].

Access Method

Type: CFG
(Size: 64 bits)

Offset: [B:0, D:2, F:0] + 10h

Default: 4h





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

4.11 Graphics Memory Range Address (GMADR)—Offset 18h

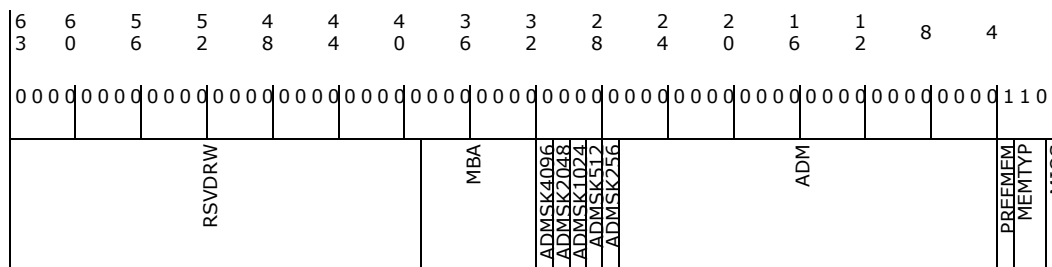
GMADR is the PCI aperture used by S/W to access tiled GFX surfaces in a linear fashion.

Access Method

Type: CFG
(Size: 64 bits)

Offset: [B:0, D:2, F:0] + 18h

Default: Ch





Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RW	RSVDRW: should be set to 0 since addressing above 512GB is not supported.
38:32	0h RW	MBA: Memory Base Address (MBA): Set by the OS, these bits correspond to address signals [38:32].
31	0h RW_L	ADMSK4096: 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[4:0]. Refer MSAC (Dev2, Func 0, offset 62h) for details.
30	0h RW_L	ADMSK2048: 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[4:0]. Refer MSAC (Dev2, Func 0, offset 62h) for details.
29	0h RW_L	ADMSK1024: 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[4:0]. Refer MSAC (Dev2, Func 0, offset 62h) for details.
28	0h RW_L	ADMSK512: 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[4:0]. Refer MSAC (Dev2, Func 0, offset 62h) for details.
27	0h RW_L	ADMSK256: 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[4:0]. Refer MSAC (Dev 2, Func 0, offset 62h) for details.
26:4	0h RO	ADM: Hardwired to 0s to indicate at least 128MB address range.
3	1h RO	PREFMEM: Hardwired to 1 to enable prefetching.
2:1	2h RO	MEMTYP: Memory Type (MEMTYP): 00: indicate 32-bit address. 10: Indicate 64-bit address
0	0h RO	MIOS: Hardwired to 0 to indicate memory space.

4.12 I/O Base Address (IOBAR)—Offset 20h

This register provides the Base offset of the I/O registers within Device #2. Bits 15:6 are programmable allowing the I/O Base to be located anywhere in 16bit I/O Address Space. Bits 2:1 are fixed and return zero; bit 0 is hardwired to a one indicating that 8 bytes of I/O space are decoded. Access to the 8Bs of IO space is allowed in PM state D0 when IO Enable (PCICMD bit 0) set. Access is disallowed in PM states D1-D3 or if IO Enable is clear or if Device #2 is turned off or if Processor Graphics is disabled.

Note that access to this IO BAR is independent of VGA functionality within Device #2.


If accesses to this IO bar is allowed then all 8, 16 or 32 bit IO cycles from IA cores that falls within the 8B are claimed. This IO BAR can be disabled and hidden from system software via DEV2CTL[0] IOBARDIS at offset 0

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:2, F:0] + 20h

Default: 1h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:6	0h RW	IOBASE: Set by the OS, these bits correspond to address signals [15:6]. Note: This field is RO 0s if DEV2CTL[0] IOBARDIS is 1b.
5:3	0h RO	Reserved (RSVD): Reserved.
2:1	0h RO	MEMTYPE: Hardwired to 0s to indicate 32-bit address.
0	1h RO	MIOS: Hardwired to "1" to indicate IO space. Note: This field is RO 0s if DEV2CTL[0] IOBARDIS is 1b.

4.13 Subsystem Vendor Identification (SVID2)—Offset 2Ch


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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:2, F:0] + 2Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
15:0	0h RW_O	SUBVID: This value is used to identify the vendor of the subsystem. This register should be programmed by BIOS during boot-up. Once written, this register becomes Read_Only. This register can only be cleared by a Reset.



4.14 Subsystem Identification (SID2)—Offset 2Eh


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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:2, F:0] + 2Eh

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
15:0	0h RW_O	SUBID: This value is used to identify a particular subsystem. This field should be programmed by BIOS during boot-up. Once written, this register becomes Read_Only. This register can only be cleared by a Reset.

4.15 Video BIOS ROM Base Address (ROMADR)—Offset 30h


The Processor Graphics does not use a separate BIOS ROM. Therefore, this register is hardwired to 0s.

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:2, F:0] + 30h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31:18	0h RO	RBA: Hardwired to 0's.
17:11	0h RO	ADMSK: Hardwired to 0s to indicate 256 KB address range.
10:1	0h RO	Reserved (RSVD): Reserved.
0	0h RO	RBE: 0: ROM not accessible.

4.16 Capabilities Pointer (CAPPOINT)—Offset 34h


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

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:2, F:0] + 34h

Default: 40h



Bit Range	Default and Access	Field Name (ID): Description
7:0	40h RO	CPV: This field contains an offset into the function's PCI Configuration Space for the first item in the New Capabilities Linked List, the CAPID0 register at offset 40h.

4.17 Interrupt Line (INTRLINE)—Offset 3Ch

This 8-bit register is used to communicate interrupt line routing information. It is read/write and should be implemented by the device. POST software will write the routing information into this register as it initializes and configures the system.

The value in this register tells which input of the system interrupt controller(s) the device's interrupt pin is connected to. The device itself does not use this value, rather it is used by device drivers and operating systems to determine priority and vector information.


Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:2, F:0] + 3Ch



Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
7:0	0h RW	INTCON: Used to communicate interrupt line routing information. POST software writes the routing information into this register as it initializes and configures the system. The value in this register indicates to which input of the system interrupt controller the device's interrupt pin is connected.

4.18 Interrupt Pin (INTRPIN)—Offset 3Dh


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

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:2, F:0] + 3Dh

Default: 1h



Bit Range	Default and Access	Field Name (ID): Description
7:0	1h RO	INTRPIN: As a single function device, the Processor Graphics specifies INTA# as its interrupt pin. 01h:INTA#.



4.19 Minimum Grant (MINGNT)—Offset 3Eh

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

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:2, F:0] + 3Eh

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
7:0	0h RO	MGV: The Processor Graphics does not burst as a PCI compliant master.

4.20 Maximum Latency (MAXLAT)—Offset 3Fh


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

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:2, F:0] + 3Fh

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
7:0	0h RO	MLV: The Processor Graphics has no specific requirements for how often it needs to access the PCI bus.



4.21 Capabilities A (CAPID0)—Offset 44h

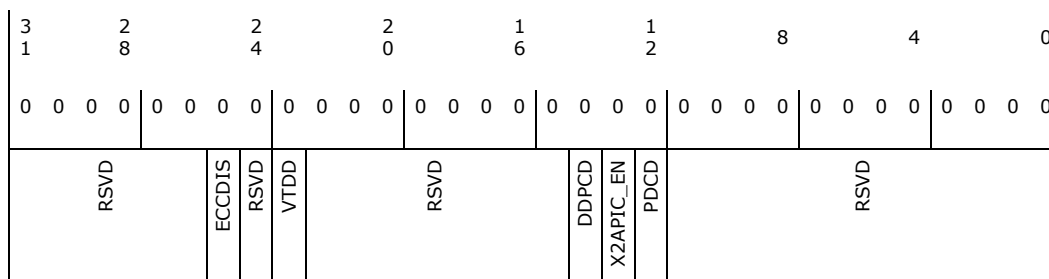
Control of bits in this register are only required for customer visible SKU differentiation.

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:2, F:0] + 44h

Default: 0h



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



4.22 Capabilities B (CAPID0)—Offset 48h

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

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:2, F:0] + 48h

Default: 0h

3	2	2	2	1	1	8	4	0										
1	8	4	0	6	2													
0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0 0 0 0 0	0 0 0 0	0 0 0 0										
IMGU_DIS	RSVD	SMT	CACHESZ	RSVD	PLL_REF100_CFG	PEGG3_DIS	RSVD	ADDGFEN	ADDGFXCAP	RSVD	DMIG3DIS	RSVD	GMM_DIS	RSVD	DMFC_DDR3	RSVD	LPDDR3_EN	RSVD

Bit Range	Default and Access	Field Name (ID): Description
31	0h RO_V	IMGU_DIS: 0: Device 5 associated memory spaces are accessible. 1: Device 5 associated memory and IO spaces are disabled by hardwiring the D1F2EN field, bit 1 of the Device Enable register, (DEVEN Dev 0 Offset 54h) to '0'.
30:29	0h RO	Reserved (RSVD): Reserved.
28	0h RO_V	SMT: This setting indicates whether or not the Processor is SMT capable.
27:25	0h RO_V	CACHESZ: This setting indicates the supporting cache sizes.
24:21	0h RO	Reserved (RSVD): Reserved.
20	0h RO_V	PEGG3_DIS: the processor: PCIe Gen 3 Disable. 0: Capable of running any of the Gen 3-compliant PEG controllers in Gen 3 mode (Devices 0/1/0, 0/1/1, 0/1/2) 1: Not capable of running any of the PEG controllers in Gen 3 mode
19	0h RO	Reserved (RSVD): Reserved.
18	0h RO_V	ADDGFEN: 0: Additive Graphics Disabled 1:-Additive Graphics Enabled
17	0h RO_V	ADDGFXCAP: 0: Capable of Additive Graphics 1: Not capable of Additive Graphics
16	0h RO	Reserved (RSVD): Reserved.



Bit Range	Default and Access	Field Name (ID): Description
15	0h RO_V	DMIG3DIS: DMI Gen 3 Disable.
14:9	0h RO	Reserved (RSVD): Reserved.
8	0h RO_V	GMM_DIS: 0: Device 8 associated memory spaces are accessible. 1: Device 8 associated memory and IO spaces are disabled by hardwiring the D8EN field, bit 1 of the Device Enable register, (DEVEN Dev 0 Offset 54h) to '0'.
7	0h RO	Reserved (RSVD): Reserved.
6:4	0h RO_V	DMFC_DDR3: This field controls which values may be written to the Memory Frequency Select field 6:4 of the Clocking Configuration registers (MCHBAR Offset C00h). Any attempt to write an unsupported value will be ignored. 000: MC capable of DDR3 2667 (2667 is the upper limit) 001: MC capable of up to DDR3 2667 010: MC capable of up to DDR3 2400 011: MC capable of up to DDR3 2133 100: MC capable of up to DDR3 1867 101: MC capable of up to DDR3 1600 110: MC capable of up to DDR3 1333 111: MC capable of up to DDR3 1067
3	0h RO	Reserved (RSVD): Reserved.
2	0h RO_V	LPDDR3_EN: Allow LPDDR3 operation
1:0	0h RO	Reserved (RSVD): Reserved.

4.23 Device Enable (DEVEN0)—Offset 54h

Allows for enabling/disabling of PCI devices and functions that are within the Processor package. The table below the bit definitions describes the behavior of all combinations of transactions to devices controlled by this register.

All the bits in this register are Intel TXT Lockable.

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:2, F:0] + 54h

Default: 84BFh

3	2	2	2	1	1	8	4	0									
1	8	4	0	6	2												
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	1 0 0 0	0 1 0 0	1 0 1 1	1 1 1 1										
		RSVD		D8EN	D7EN	D6EN	RSVD	D5EN	RSVD	D4EN	RSVD	D3EN	D2EN	D1F0EN	D1F1EN	D1F2EN	D0EN



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15	1h RO_V	D8EN: 0: Bus 0 Device 8 is disabled and not visible. 1: Bus 0 Device 8 is enabled and visible. This bit will be set to 0b and remain 0b if Device 8 capability is disabled.
14	0h RO_V	D7EN: 0: Bus 0 Device 7 is disabled and not visible. 1: Bus 0 Device 7 is enabled and visible. Non-production BIOS code should provide a setup option to enable Bus 0 Device 7. When enabled, Bus 0 Device 7 should be initialized in accordance to standard PCI device initialization procedures.
13	0h RO_V	D6EN: Reserved (RSVD):
12:11	0h RO	Reserved (RSVD): Reserved.
10	1h RO_V	D5EN: 0: Bus 0 Device 5 is disabled and not visible. 1: Bus 0 Device 5 is enabled and visible. This bit will be set to 0b and remain 0b if Device 5 capability is disabled.
9:8	0h RO	Reserved (RSVD): Reserved.
7	1h RO_V	D4EN: 0: Bus 0 Device 4 is disabled and not visible. 1: Bus 0 Device 4 is enabled and visible. This bit will be set to 0b and remain 0b if Device 4 capability is disabled.
6	0h RO	Reserved (RSVD): Reserved.
5	1h RO_V	D3EN: 0: Bus 0 Device 3 is disabled and hidden 1: Bus 0 Device 3 is enabled and visible This bit will be set to 0b and remain 0b if Device 3 capability is disabled.
4	1h RO_V	D2EN: 0: Bus 0 Device 2 is disabled and hidden 1: Bus 0 Device 2 is enabled and visible This bit will be set to 0b and remain 0b if Device 2 capability is disabled.
3	1h RO_V	D1F0EN: 0: Bus 0 Device 1 Function 0 is disabled and hidden. 1: Bus 0 Device 1 Function 0 is enabled and visible. This bit will be set to 0b and remain 0b if PEG10 capability is disabled.
2	1h RO_V	D1F1EN: 0: Bus 0 Device 1 Function 1 is disabled and hidden. 1: Bus 0 Device 1 Function 1 is enabled and visible.
1	1h RO_V	D1F2EN: 0: Bus 0 Device 1 Function 2 is disabled and hidden. 1: Bus 0 Device 1 Function 2 is enabled and visible.
0	1h RO	DOEN: Bus 0 Device 0 Function 0 may not be disabled and is therefore hardwired to 1.



4.24 Base Data of Stolen Memory (BDSM)—Offset 5Ch


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

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:2, F:0] + 5Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:20	0h RO_V	BDSM: This register contains bits 31 to 20 of the base address of stolen DRAM memory. BIOS determines the base of graphics stolen memory by subtracting the graphics stolen memory size (PCI Device 0 offset 50 bits 15:8) from TOLUD (PCI Device 0, offset BC, bits 31:20).
19:1	0h RO	Reserved (RSVD): Reserved.
0	0h RO_V	LOCK: This bit will lock all writeable settings in this register, including itself.

4.25 Multi Size Aperture Control (MSAC)—Offset 62h

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

This register is Intel TXT locked, becomes read-only when trusted environment is launched.

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:2, F:0] + 62h

Default: 1h



7	0	0	0	4	0	0	0	0	
								1	
			RSVDRW		APSZ4	APSZ3	APSZ2	APSZ1	APSZ0

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



Bit Range	Default and Access	Field Name (ID): Description
1	0h RW_KV	<p>APSZ1: This field is used in conjunction with other APSZ* fields to determine the size of Aperture (GMADR) and affects certain bits of GMADR register. The description below is for all APSZ* fields 4:0 -</p> <p>00000 = 128MB => GMADR.B[26:4] is hardwired to 0 00001 = 256MB => GMADR.B[27] = 0, RO 00010 = illegal (hardware will treat this as 00011) 00011 = 512MB => GMADR.B[28:27] = 0, RO 0100-00110 = illegal (hardware will treat this as 00111) 00111 = 1024MB => GMADR.B[29:27] = 0, RO 000-01110 = illegal (hardware will treat this as 01111) 01111 = 2048MB => GMADR.B[30:27] = 0, RO 10000-11110 = illegal (hardware will treat this as 11111) 11111 = 4096MB => GMADR.B[31:27] = 0, RO</p>
0	1h RW_KV	<p>APSZ0: This field is used in conjunction with other APSZ* fields to determine the size of Aperture (GMADR) and affects certain bits of GMADR register. The description below is for all APSZ* fields 4:0 -</p> <p>00000 = 128MB => GMADR.B[26:4] is hardwired to 0 00001 = 256MB => GMADR.B[27] = 0, RO 00010 = illegal (hardware will treat this as 00011) 00011 = 512MB => GMADR.B[28:27] = 0, RO 0100-00110 = illegal (hardware will treat this as 00111) 00111 = 1024MB => GMADR.B[29:27] = 0, RO 000-01110 = illegal (hardware will treat this as 01111) 01111 = 2048MB => GMADR.B[30:27] = 0, RO 10000-11110 = illegal (hardware will treat this as 11111) 11111 = 4096MB => GMADR.B[31:27] = 0, RO</p>

4.26 PCI Express Capability Header (PCIECAPHDR)—Offset 70h

This is the header register for the PCI Express Capability Structure, allowing the exposure of PCI Express Extended Capabilities which are required for SVM OS support.


Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:2, F:0] + 70h

Default: AC10h

15	12	8	4	0
1 0 1 0	1 1 0 0	0 0 0 1	0 0 0 0	0 0 0 0
NEXT_CAP				CAP_ID



Bit Range	Default and Access	Field Name (ID): Description
15:8	ACh RO	NEXT_CAP: This field contains the offset to the next PCI Capability structure, the MSI Capabilities at ACh
7:0	10h RO	CAP_ID: Indicates the PCI Express Capability structure. This field should return a Capability ID of 10h indicating that this is a PCI Express Capability structure

4.27 Message Signaled Interrupts Capability ID (MSI)—Offset ACh


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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:2, F:0] + ACh

Default: D005h



Bit Range	Default and Access	Field Name (ID): Description
15:8	D0h RO	POINTNEXT: This contains a pointer to the next item in the capabilities list which is the Power Management capability.
7:0	5h RO	CAPID: Value of 05h identifies this linked list item (capability structure) as being for MSI registers.

4.28 Message Control (MC)—Offset AEh

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


Access Method



Type: CFG
(Size: 16 bits)

Offset: [B:0, D:2, F:0] + AEh

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
15:8	0h RO	Reserved (RSVD): Reserved.
7	0h RO	CAP64B: Hardwired to 0 to indicate that the function does not implement the upper 32 bits of the Message address register and is incapable of generating a 64-bit memory address.
6:4	0h RW	MME: System software programs this field to indicate the actual number of messages allocated to this device. This number will be equal to or less than the number actually requested. The encoding is the same as for the MMC field below.
3:1	0h RO	MMC: System Software reads this field to determine the number of messages being requested by this device. 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	0h RW	MSIEN: Controls the ability of this device to generate MSIs.

4.29 Message Address (MA)—Offset B0h


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

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:2, F:0] + B0h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:2	0h RW	MESSADD: Used by system software to assign an MSI address to the device. The device handles an MSI by writing the padded contents of the MD register to this address.
1:0	0h RO	FDWORD: Hardwired to 0 so that addresses assigned by system software are always aligned on a DWORD address boundary.

4.30 Message Data (MD)—Offset B4h


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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:2, F:0] + B4h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
15:0	0h RW	MESSDATA: Base message data pattern assigned by system software and used to handle an MSI from the device. When the device should 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.

4.31 Power Management Capabilities ID (PMCAPID)—Offset D0h

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



Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:2, F:0] + D0h

Default: 1h

15	12	8	4	0
0	0	0	0	1
NEXT_PTR				CAP_ID

Bit Range	Default and Access	Field Name (ID): Description
15:8	0h RO	NEXT_PTR: This contains a pointer to the next item in the capabilities list. This is the final capability in the list and should be set to 00h.
7:0	1h RO	CAP_ID: SIG defines this ID is 01h for power management.

4.32 Power Management Capabilities (PMCAP)—Offset D2h

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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:2, F:0] + D2h

Default: 22h

15	12	8	4	0
0	0	0	0	0
PMES		D2	D1	RSVD
			DSI	RSVD
			PMECLK	VER



Bit Range	Default and Access	Field Name (ID): Description
15:11	0h RO	PMES: This field indicates the power states in which the Processor Graphics may assert PME#. Hardwired to 0 to indicate that the Processor Graphics does not assert the PME# signal.
10	0h RO	D2: The D2 power management state is not supported. This bit is hardwired to 0.
9	0h RO	D1: Hardwired to 0 to indicate that the D1 power management state is not supported.
8:6	0h RO	Reserved (RSVD): Reserved.
5	1h RO	DSI: Hardwired to 1 to indicate that special initialization of the Processor Graphics is required before generic class device driver is to use it.
4	0h RO	Reserved (RSVD): Reserved.
3	0h RO	PMECLK: Hardwired to 0 to indicate Processor Graphics does not support PME# generation.
2:0	2h RO	VER: Hardwired to 010b to indicate that there are 4 bytes of power management registers implemented and that this device complies with revision 1.1 of the PCI Power Management Interface Specification.


4.33 Power Management Control/Status (PMCS)—Offset D4h

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:2, F:0] + D4h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
15	0h RO	PMESTS: This bit is 0 to indicate that Processor Graphics does not support PME# generation from D3 (cold).
14:13	0h RO	DSCALE: The Processor Graphics does not support data register. This bit always returns 00 when read, write operations have no effect.
12:9	0h RO	DSEL: The Processor Graphics does not support data register. This bit always returns 0h when read, write operations have no effect.



Bit Range	Default and Access	Field Name (ID): Description										
8	0h RO	PMEEN: This bit is 0 to indicate that PME# assertion from D3 (cold) is disabled.										
7:2	0h RO	Reserved (RSVD): Reserved.										
1:0	0h RO_V	<p>PWRSTAT: This field indicates the current power state of the Processor Graphics and can be used to set the Processor Graphics into a new power state. If software attempts to write an unsupported state to this field, write operation should 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.</p> <table border="0"> <tr> <td>Bits[1:0]</td> <td>Power state</td> </tr> <tr> <td>00:</td> <td>D0 Default</td> </tr> <tr> <td>01:</td> <td>D1 Not Supported</td> </tr> <tr> <td>10:</td> <td>D2 Not Supported</td> </tr> <tr> <td>11:</td> <td>D3</td> </tr> </table>	Bits[1:0]	Power state	00:	D0 Default	01:	D1 Not Supported	10:	D2 Not Supported	11:	D3
Bits[1:0]	Power state											
00:	D0 Default											
01:	D1 Not Supported											
10:	D2 Not Supported											
11:	D3											

§ §



5 Dynamic Power Performance Management (DPPM) Registers

Table 5-1. Summary of Bus: 0, Device: 4, Function: 0 (CFG)

Offset	Size (Bytes)	Register Name (Register Symbol)	Default Value
54–57h	4	Device Enable (DEVEN)—Offset 54h	84BFh
E4–E7h	4	Capabilities A (CAPID0)—Offset E4h	0h
E8–EBh	4	Capabilities B (CAPID0)—Offset E8h	0h

5.1 Device Enable (DEVEN)—Offset 54h

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

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:4, F:0] + 54h

Default: 84BFh

3	2	2	2	1	1	8	4	0
1	8	4	0	6	2			
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
		RSVD			D8EN	D7EN	D6EN	RSVD
					D5EN	RSVD	D4EN	RSVD
							D3EN	D2EN
							D1F0EN	D1F1EN
							D1F2EN	D0EN



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15	1h RO_V	D8EN: 0: Bus 0 Device 8 is disabled and not visible. 1: Bus 0 Device 8 is enabled and visible. This bit will be set to 0b and remain 0b if Device 8 capability is disabled.
14	0h RO_V	D7EN: 0: Bus 0 Device 7 is disabled and not visible. 1: Bus 0 Device 7 is enabled and visible. Non-production BIOS code should provide a setup option to enable Bus 0 Device 7. When enabled, Bus 0 Device 7 should be initialized in accordance to standard PCI device initialization procedures.
13	0h RO_V	D6EN: Reserved (RSVD):
12:11	0h RO	Reserved (RSVD): Reserved.
10	1h RO_V	D5EN: 0: Bus 0 Device 5 is disabled and not visible. 1: Bus 0 Device 5 is enabled and visible. This bit will be set to 0b and remain 0b if Device 5 capability is disabled.
9:8	0h RO	Reserved (RSVD): Reserved.
7	1h RO_V	D4EN: 0: Bus 0 Device 4 is disabled and not visible. 1: Bus 0 Device 4 is enabled and visible. This bit will be set to 0b and remain 0b if Device 4 capability is disabled.
6	0h RO	Reserved (RSVD): Reserved.
5	1h RO_V	D3EN: 0: Bus 0 Device 3 is disabled and hidden 1: Bus 0 Device 3 is enabled and visible This bit will be set to 0b and remain 0b if Device 3 capability is disabled.
4	1h RO_V	D2EN: 0: Bus 0 Device 2 is disabled and hidden 1: Bus 0 Device 2 is enabled and visible This bit will be set to 0b and remain 0b if Device 2 capability is disabled.
3	1h RO_V	D1F0EN: 0: Bus 0 Device 1 Function 0 is disabled and hidden. 1: Bus 0 Device 1 Function 0 is enabled and visible. This bit will be set to 0b and remain 0b if PEG10 capability is disabled.
2	1h RO_V	D1F1EN: 0: Bus 0 Device 1 Function 1 is disabled and hidden. 1: Bus 0 Device 1 Function 1 is enabled and visible.
1	1h RO_V	D1F2EN: 0: Bus 0 Device 1 Function 2 is disabled and hidden. 1: Bus 0 Device 1 Function 2 is enabled and visible.
0	1h RO	DOEN: Bus 0 Device 0 Function 0 may not be disabled and is therefore hardwired to 1.



5.2 Capabilities A (CAPID0)—Offset E4h

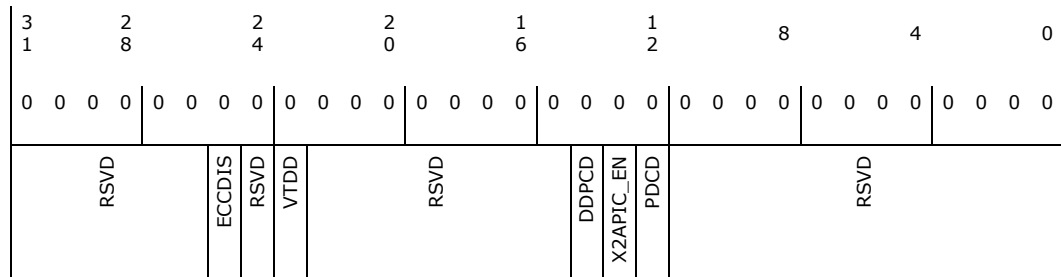
Control of bits in this register are only required for customer visible SKU differentiation.

Access Method


Type: CFG
(Size: 32 bits)

Offset: [B:0, D:4, F:0] + E4h

Default: 0h



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



5.3 Capabilities B (CAPID0)—Offset E8h

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

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:4, F:0] + E8h

Default: 0h

3	2	2	2	1	1	8	4	0
1	8	4	0	6	2			
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0 0 0 0 0	0 0 0 0	0 0 0 0
IMGU_DIS RSVD	SMT CACHESZ	RSVD PLL_REF100_CFG	PEGG3_DIS RSVD ADDGFEN ADDGFXCAP	RSVD DMIG3DIS	RSVD	GMM_DIS RSVD	DMFC_DDR3 RSVD LPDDR3_EN	RSVD

Bit Range	Default and Access	Field Name (ID): Description
31	0h RO_V	IMGU_DIS: 0: Device 5 associated memory spaces are accessible. 1: Device 5 associated memory and IO spaces are disabled by hardwiring the D1F2EN field, bit 1 of the Device Enable register, (DEVEN Dev 0 Offset 54h) to '0'.
30:29	0h RO	Reserved (RSVD): Reserved.
28	0h RO_V	SMT: This setting indicates whether or not the Processor is SMT capable.
27:25	0h RO_V	CACHESZ: This setting indicates the supporting cache sizes.
24:21	0h RO	Reserved (RSVD): Reserved.
20	0h RO_V	PEGG3_DIS: the processor: PCIe Gen 3 Disable. 0: Capable of running any of the Gen 3-compliant PEG controllers in Gen 3 mode (Devices 0/1/0, 0/1/1, 0/1/2) 1: Not capable of running any of the PEG controllers in Gen 3 mode
19	0h RO	Reserved (RSVD): Reserved.
18	0h RO_V	ADDGFEN: 0: Additive Graphics Disabled 1: Additive Graphics Enabled
17	0h RO_V	ADDGFXCAP: 0: Capable of Additive Graphics 1: Not capable of Additive Graphics
16	0h RO	Reserved (RSVD): Reserved.



Bit Range	Default and Access	Field Name (ID): Description
15	0h RO_V	DMIG3DIS: DMI Gen 3 Disable.
14:9	0h RO	Reserved (RSVD): Reserved.
8	0h RO_V	GMM_DIS: 0: Device 8 associated memory spaces are accessible. 1: Device 8 associated memory and IO spaces are disabled by hardwiring the D8EN field, bit 1 of the Device Enable register, (DEVEN Dev 0 Offset 54h) to '0'.
7	0h RO	Reserved (RSVD): Reserved.
6:4	0h RO_V	DMFC_DDR3: This field controls which values may be written to the Memory Frequency Select field 6:4 of the Clocking Configuration registers (MCHBAR Offset C00h). Any attempt to write an unsupported value will be ignored. 000: MC capable of DDR3 2667 (2667 is the upper limit) 001: MC capable of up to DDR3 2667 010: MC capable of up to DDR3 2400 011: MC capable of up to DDR3 2133 100: MC capable of up to DDR3 1867 101: MC capable of up to DDR3 1600 110: MC capable of up to DDR3 1333 111: MC capable of up to DDR3 1067
3	0h RO	Reserved (RSVD): Reserved.
2	0h RO_V	LPDDR3_EN: Allow LPDDR3 operation
1:0	0h RO	Reserved (RSVD): Reserved.

§ §



6 DMIBAR Registers

Table 6-1. Summary of Bus: 0, Device: 0, Function: 0 (MEM)

Offset	Size (Bytes)	Register Name (Register Symbol)	Default Value
0-3h	4	DMI Virtual Channel Enhanced Capability (DMIVCECH)—Offset 0h	4010002h
4-7h	4	DMI Port VC Capability Register 1 (DMIPVCCAP1)—Offset 4h	0h
8-8h	4	DMI Port VC Capability Register 2 (DMIPVCCAP2)—Offset 8h	0h
C-Dh	2	DMI Port VC Control (DMIPVCCTL)—Offset Ch	0h
10-13h	4	DMI VC0 Resource Capability (DMIVCORCAP)—Offset 10h	1h
14-17h	4	DMI VC0 Resource Control (DMIVCORCTL)—Offset 14h	8000017Fh
1A-1Bh	2	DMI VC0 Resource Status (DMIVCORSTS)—Offset 1Ah	2h
1C-1Fh	4	DMI VC1 Resource Capability (DMIVC1RCAP)—Offset 1Ch	8001h
20-23h	4	DMI VC1 Resource Control (DMIVC1RCTL)—Offset 20h	1000100h
26-27h	2	DMI VC1 Resource Status (DMIVC1RSTS)—Offset 26h	2h
34-37h	4	DMI VCm Resource Capability (DMIVCMRCAP)—Offset 34h	8000h
38-3Bh	4	DMI VCm Resource Control (DMIVCMRCTL)—Offset 38h	7000180h
3E-3Fh	2	DMI VCm Resource Status (DMIVCMRSTS)—Offset 3Eh	2h
40-43h	4	DMI Root Complex Link Declaration (DMIRCLDECH)—Offset 40h	8010005h
44-47h	4	DMI Element Self Description (DMIESD)—Offset 44h	1000202h
50-53h	4	DMI Link Entry 1 Description (DMILE1D)—Offset 50h	0h
58-5Bh	4	DMI Link Entry 1 Address (DMILE1A)—Offset 58h	0h
5C-5Fh	4	DMI Link Upper Entry 1 Address (DMILUE1A)—Offset 5Ch	0h
60-63h	4	DMI Link Entry 2 Description (DMILE2D)—Offset 60h	0h
68-6Bh	4	DMI Link Entry 2 Address (DMILE2A)—Offset 68h	0h
84-87h	4	Link Capabilities (LCAP)—Offset 84h	41AC43h
88-89h	2	Link Control (LCTL)—Offset 88h	0h
8A-8Bh	2	DMI Link Status (LSTS)—Offset 8Ah	1h
98-99h	2	Link Control 2 (LCTL2)—Offset 98h	1h
9A-9Bh	2	Link Status 2 (LSTS2)—Offset 9Ah	0h
1C4-1C7h	4	DMI Uncorrectable Error Status (DMIUESTS)—Offset 1C4h	0h
1C8-1CBh	4	DMI Uncorrectable Error Mask (DMIUEMSK)—Offset 1C8h	0h
1CC-1CFh	4	DMI Uncorrectable Error Severity (DMIUESEV)—Offset 1CCh	60010h
1D0-1D3h	4	DMI Correctable Error Status (DMICESTS)—Offset 1D0h	0h
1D4-1D7h	4	DMI Correctable Error Mask (DMICEMSK)—Offset 1D4h	2000h



6.1 DMI Virtual Channel Enhanced Capability (DMIVCECH)—Offset 0h


Indicates DMI Virtual Channel capabilities.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 0h

Default: 4010002h



Bit Range	Default and Access	Field Name (ID): Description
31:20	40h RO	PNC: Pointer to Next Capability: This field contains the offset to the next PCI Express capability structure in the linked list of capabilities (Link Declaration Capability).
19:16	1h RO	PCIEVCCV: PCI Express Virtual Channel Capability Version: Hardwired to 1 to indicate compliances with the 1.1 version of the PCI Express specification. Note: This version does not change for 2.0 compliance.
15:0	2h RO	ECID: Extended Capability ID: Value of 0002h identifies this linked list item (capability structure) as being for PCI Express Virtual Channel registers.

6.2 DMI Port VC Capability Register 1 (DMIPVCCAP1)—Offset 4h


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

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:7	0h RO	Reserved (RSVD): Reserved.
6:4	0h RO	LPEVCC: Low Priority Extended VC Count: Indicates the number of (extended) Virtual Channels in addition to the default VC belonging to the low-priority VC (LPVC) group that has the lowest priority with respect to other VC resources in a strict-priority VC Arbitration. The value of 0 in this field implies strict VC arbitration.
3	0h RO	Reserved (RSVD): Reserved.
2:0	0h RW_O	EVCC: Extended VC Count: Indicates the number of (extended) Virtual Channels in addition to the default VC supported by the device. The Private Virtual Channel, VC1 and the Manageability Virtual Channel are not included in this count.

6.3 DMI Port VC Capability Register 2 (DMIPVCCAP2)—Offset 8h


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

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 8h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31:24	0h RO	VCATO: Reserved for VC Arbitration Table Offset:
23:8	0h RO	Reserved (RSVD): Reserved.
7:0	0h RO	VCAC: Reserved for VC Arbitration Capability:


6.4 DMI Port VC Control (DMIPVCCTL)—Offset Ch

Access Method

Type: MEM
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
15:4	0h RO	Reserved (RSVD): Reserved.
3:1	0h RW	VCAS: VC Arbitration Select: This field will be programmed by software to the only possible value as indicated in the VC Arbitration Capability field. The value 000b when written to this field will indicate the VC arbitration scheme is hardware fixed (in the root complex). This field cannot be modified when more than one VC in the LPVC group is enabled. 000:Hardware fixed arbitration scheme. E.G. Round Robin Others:Reserved Refer the PCI express specification for more details.
0	0h RO	LVCAT: Reserved for Load VC Arbitration Table:



6.5 DMI VC0 Resource Capability (DMIVCORCAP)— Offset 10h

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 10h

Default: 1h

3 1	2 8	2 4	2 0	1 6	1 2	8	4	0
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 1
PATO		RSVD	MTS		REJSNPT	RSVD		PAC

Bit Range	Default and Access	Field Name (ID): Description
31:24	0h RO	PATO: Reserved for Port Arbitration Table Offset:
23	0h RO	Reserved (RSVD): Reserved.
22:16	0h RO	MTS: Reserved for Maximum Time Slots:
15	0h RO	REJSNPT: Reject Snoop Transactions: 0: Transactions with or without the No Snoop bit set within the TLP header are allowed on this VC. 1: Any transaction for which the No Snoop attribute is applicable but is not set within the TLP Header will be rejected as an Unsupported Request.
14:8	0h RO	Reserved (RSVD): Reserved.
7:0	1h RO	PAC: Port Arbitration Capability: Having only bit 0 set indicates that the only supported arbitration scheme for this VC is non-configurable hardware-fixed.



6.6 DMI VC0 Resource Control (DMIVC0RCTL)—Offset 14h

Controls the resources associated with PCI Express Virtual Channel 0.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 14h

Default: 8000017Fh

3 1	2 8	2 4	2 0	1 6	1 2	8	4	0	
1	0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0 1	0 1 1 1	1 1 1 1	
VC0E	RSVD	VC0ID	RSVD	PAS	RSVD	FC_FSM_STATE	TCMVC0M	TCVC0M	TCOV0M

Bit Range	Default and Access	Field Name (ID): Description
31	1h RO	VC0E: Virtual Channel 0 Enable: For VC0 this is hardwired to 1 and read only as VC0 can never be disabled.
30:27	0h RO	Reserved (RSVD): Reserved.
26:24	0h RO	VC0ID: Virtual Channel 0 ID: Assigns a VC ID to the VC resource. For VC0 this is hardwired to 0 and read only.
23:20	0h RO	Reserved (RSVD): Reserved.
19:17	0h RW	PAS: Port Arbitration Select: Configures the VC resource to provide a particular Port Arbitration service. Valid value for this field is a number corresponding to one of the asserted bits in the Port Arbitration Capability field of the VC resource. Because only bit 0 of that field is asserted. This field will always be programmed to '1'.
16:13	0h RO	Reserved (RSVD): Reserved.
12:8	1h ROV	FC_FSM_STATE: This register is for Save Restore to restore the FC fsm



Bit Range	Default and Access	Field Name (ID): Description
7	0h RO	TCMVCOM: Traffic Class m / Virtual Channel 0 Map:
6:1	3Fh RW	TCVCOM: Traffic Class / Virtual Channel 0 Map: Indicates the TCs (Traffic Classes) that are mapped to the VC resource. Bit locations within this field correspond to TC values. For example, when bit 7 is set in this field, TC7 is mapped to this VC resource. When more than one bit in this field is set, it indicates that multiple TCs are mapped to the VC resource. In order to remove one or more TCs from the TC/VC Map of an enabled VC, software should ensure that no new or outstanding transactions with the TC labels are targeted at the given Link.
0	1h RO	TCOVCOM: Traffic Class 0 / Virtual Channel 0 Map: Traffic Class 0 is always routed to VC0.

6.7 DMI VC0 Resource Status (DMIVC0RSTS)—Offset 1Ah

Reports the Virtual Channel specific status.

Access Method

Type: MEM
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + 1Ah

Default: 2h

15	12	8	4	0
0	0	0	0	0
RSVD				VC0NP
				RSVD

Bit Range	Default and Access	Field Name (ID): Description
15:2	0h RO	Reserved (RSVD): Reserved.
1	1h RO_V	VC0NP: Virtual Channel 0 Negotiation Pending: 0: The VC negotiation is complete. 1: The VC resource is still in the process of negotiation (initialization or disabling). This bit indicates the status of the process of Flow Control initialization. It is set by default on Reset, as well as whenever the corresponding Virtual Channel is Disabled or the Link is in the DL_Down state. It is cleared when the link successfully exits the FC_INIT2 state. BIOS Requirement: Before using a Virtual Channel, software should check whether the VC Negotiation Pending fields for that Virtual Channel are cleared in both Components on a Link.
0	0h RO	Reserved (RSVD): Reserved.




6.8 DMI VC1 Resource Capability (DMIVC1RCAP) – Offset 1Ch

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 1Ch

Default: 8001h



Bit Range	Default and Access	Field Name (ID): Description
31:24	0h RO	PATO: Reserved for Port Arbitration Table Offset:
23	0h RO	Reserved (RSVD): Reserved.
22:16	0h RO	MTS: Reserved for Maximum Time Slots:
15	1h RO	REJSNPT: Reject Snoop Transactions: 0: Transactions with or without the No Snoop bit set within the TLP header are allowed on this VC. 1: When Set, any transaction for which the No Snoop attribute is applicable but is not Set within the TLP Header will be rejected as an Unsupported Request.
14:8	0h RO	Reserved (RSVD): Reserved.
7:0	1h RO	PAC: Port Arbitration Capability: Having only bit 0 set indicates that the only supported arbitration scheme for this VC is non-configurable hardware-fixed.



6.9 DMI VC1 Resource Control (DMIVC1RCTL)—Offset 20h

Controls the resources associated with PCI Express Virtual Channel 1.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 20h

Default: 1000100h

3	2	2	2	1	1	8	4	0	
1	8	4	0	6	2				
0	0	0	0	0	0	0	0	0	
0	0	0	1	0	0	0	0	0	
0	0	0	0	0	0	0	0	0	
0	0	0	0	0	0	0	1	0	
0	0	0	0	0	0	0	0	0	
0	0	0	0	0	0	0	0	0	
VC1E	RSVD	VC1ID	RSVD	PAS	RSVD	FC_FSM_STATE	TCMVC1M	TCVC1M	TCOV1M

Bit Range	Default and Access	Field Name (ID): Description
31	0h RW	<p>VC1E: Virtual Channel 1 Enable: 0: Virtual Channel is disabled. 1: Virtual Channel is enabled. Refer exceptions below.</p> <p>Software should 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.</p> <p>BIOS Requirement:</p> <ol style="list-style-type: none"> To enable a Virtual Channel, the VC Enable bits for that Virtual Channel should be set in both Components on a Link. To disable a Virtual Channel, the VC Enable bits for that Virtual Channel should be cleared in both Components on a Link. Software should ensure that no traffic is using a Virtual Channel at the time it is disabled. Software should fully disable a Virtual Channel in both Components on a Link before re-enabling the Virtual Channel.
30:27	0h RO	Reserved (RSVD): Reserved.
26:24	1h RW	VC1ID: Virtual Channel 1 ID: Assigns a VC ID to the VC resource. Assigned value should be non-zero. This field can not be modified when the VC is already enabled.
23:20	0h RO	Reserved (RSVD): Reserved.
19:17	0h RW	PAS: Port Arbitration Select: Configures the VC resource to provide a particular Port Arbitration service. Valid value for this field is a number corresponding to one of the asserted bits in the Port Arbitration Capability field of the VC resource.
16:13	0h RO	Reserved (RSVD): Reserved.
12:8	1h ROV	FC_FSM_STATE: This register is for Save Restore to restore the FC fsm



Bit Range	Default and Access	Field Name (ID): Description
7	0h RO	TCMVC1M: Traffic Class m / Virtual Channel 1:
6:1	0h RW	<p>TCVC1M: Traffic Class / Virtual Channel 1 Map: Indicates the TCs (Traffic Classes) that are mapped to the VC resource. Bit locations within this field correspond to TC values. For example, when bit 6 is set in this field, TC6 is mapped to this VC resource. When more than one bit in this field is set, it indicates that multiple TCs are mapped to the VC resource. In order to remove one or more TCs from the TC/VC Map of an enabled VC, software should ensure that no new or outstanding transactions with the TC labels are targeted at the given Link.</p> <p>BIOS Requirement: Program this field with the value 010001b, which maps TC1 and TC5 to VC1.</p>
0	0h RO	TC0VC1M: Traffic Class 0 / Virtual Channel 1 Map: Traffic Class 0 is always routed to VC0.

6.10 DMI VC1 Resource Status (DMIVC1RSTS)—Offset 26h


Reports the Virtual Channel specific status.

Access Method

Type: MEM
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + 26h

Default: 2h



Bit Range	Default and Access	Field Name (ID): Description
15:2	0h RO	Reserved (RSVD): Reserved.
1	1h RO_V	<p>VC1NP: Virtual Channel 1 Negotiation Pending: 0: The VC negotiation is complete. 1: The VC resource is still in the process of negotiation (initialization or disabling). Software may use this bit when enabling or disabling the VC. This bit indicates the status of the process of Flow Control initialization. It is set by default on Reset, as well as whenever the corresponding Virtual Channel is Disabled or the Link is in the DL_Down state. It is cleared when the link successfully exits the FC_INIT2 state. Before using a Virtual Channel, software should check whether the VC Negotiation Pending fields for that Virtual Channel are cleared in both Components on a Link.</p>
0	0h RO	Reserved (RSVD): Reserved.




6.11 DMI VCm Resource Capability (DMIVCMRCAP)— Offset 34h

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 34h

Default: 8000h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15	1h RO	REJSNPT: Reject Snoop Transactions: 0: Transactions with or without the No Snoop bit set within the TLP header are allowed on the VC. 1: When Set, any transaction for which the No Snoop attribute is applicable but is not Set within the TLP Header will be rejected as an Unsupported Request
14:0	0h RO	Reserved (RSVD): Reserved.



6.12 DMI VCm Resource Control (DMIVCMRCTL)— Offset 38h

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 38h

Default: 7000180h

3 1	2 8	2 4	2 0	1 6	1 2	8	4	0
0 0 0 0	0 1 1 1	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 1	1 0 0 0	0 0 0 0
VCMEN	RSVD	VCID	RSVD		FC_FSM_STATE		TCVCM MAP	

Bit Range	Default and Access	Field Name (ID): Description
31	0h RW	<p>VCMEN: Virtual Channel enable: 0: Virtual Channel is disabled. 1: Virtual Channel is enabled. Refer exceptions below. Software should 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.</p> <p>BIOS Requirement:</p> <ol style="list-style-type: none"> To enable a Virtual Channel, the VC Enable bits for that Virtual Channel should be set in both Components on a Link. To disable a Virtual Channel, the VC Enable bits for that Virtual Channel should be cleared in both Components on a Link. Software should ensure that no traffic is using a Virtual Channel at the time it is disabled. Software should fully disable a Virtual Channel in both Components on a Link before re-enabling the Virtual Channel.
30:27	0h RO	Reserved (RSVD): Reserved.
26:24	7h RW	VCID: Virtual Channel ID: Assigns a VC ID to the VC resource. Assigned value should be non-zero. This field can not be modified when the VC is already enabled.
23:13	0h RO	Reserved (RSVD): Reserved.
12:8	1h ROV	FC_FSM_STATE: This register is for Save Restore to restore the FC fsm
7:0	80h RO	<p>TCVCM MAP: Traffic Class/Virtual Channel Map: Indicates the TCs (Traffic Classes) that are mapped to the VC resource. Bit locations within this field correspond to TC values.</p> <p>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 should ensure that no new or outstanding transactions with the TC labels are targeted at the given Link.</p>



6.13 DMI VCm Resource Status (DMIVCMRSTS)—Offset 3Eh

Access Method

Type: MEM
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + 3Eh

Default: 2h

15	12	8	4	0
0	0	0	0	0
0	0	0	0	0
RSVD				VCNEGPND
				RSVD

Bit Range	Default and Access	Field Name (ID): Description
15:2	0h RO	Reserved (RSVD): Reserved.
1	1h RO_V	VCNEGPND: Virtual Channel Negotiation Pending: 0: The VC negotiation is complete. 1: The VC resource is still in the process of negotiation (initialization or disabling). Software may use this bit when enabling or disabling the VC. This bit indicates the status of the process of Flow Control initialization. It is set by default on Reset, as well as whenever the corresponding Virtual Channel is Disabled or the Link is in the DL_Down state. It is cleared when the link successfully exits the FC_INIT2 state. Before using a Virtual Channel, software should check whether the VC Negotiation Pending fields for that Virtual Channel are cleared in both Components on a Link.
0	0h RO	Reserved (RSVD): Reserved.



6.14 DMI Root Complex Link Declaration (DMIRCLDECH)—Offset 40h


This capability declares links from the respective element to other elements of the root complex component to which it belongs and to an element in another root complex component. Refer PCI Express specification for link/topology declaration requirements.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 40h

Default: 8010005h



Bit Range	Default and Access	Field Name (ID): Description
31:20	80h RO	PNC: Pointer to Next Capability: This field contains the offset to the next PCI Express capability structure in the linked list of capabilities (Internal Link Control Capability).
19:16	1h RO	LDCV: Link Declaration Capability Version: Hardwired to 1 to indicate compliances with the 1.1 version of the PCI Express specification. Note: This version does not change for 2.0 compliance.
15:0	5h RO	ECID: Extended Capability ID: Value of 0005h identifies this linked list item (capability structure) as being for PCI Express Link Declaration Capability.



6.15 DMI Element Self Description (DMIESD)—Offset 44h


Provides information about the root complex element containing this Link Declaration Capability.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 44h

Default: 1000202h



Bit Range	Default and Access	Field Name (ID): Description
31:24	1h RO	PORTNUM: Port Number: Specifies the port number associated with this element with respect to the component that contains this element. This port number value is utilized by the egress port of the component to provide arbitration to this Root Complex Element.
23:16	0h RW_O	CID: Component ID: Identifies the physical component that contains this Root Complex Element. BIOS Requirement: should be initialized according to guidelines in the PCI Express* Isochronous/Virtual Channel Support Hardware Programming Specification (HPS).
15:8	2h RO	NLE: Number of Link Entries: Indicates the number of link entries following the Element Self Description. This field reports 2 (one for MCH egress port to main memory and one to egress port belonging to ICH on other side of internal link).
7:4	0h RO	Reserved (RSVD): Reserved.
3:0	2h RO	ETYP: Element Type: Indicates the type of the Root Complex Element. Value of 2h represents an Internal Root Complex Link (DMI).



6.16 DMI Link Entry 1 Description (DMILE1D)—Offset 50h


First part of a Link Entry which declares an internal link to another Root Complex Element.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 50h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:24	0h RW_O	TPN: Target Port Number: Specifies the port number associated with the element targeted by this link entry (egress port of PCH). The target port number is with respect to the component that contains this element as specified by the target component ID. This can be programmed by BIOS, but the default value will likely be correct because the DMI RCRB in the PCH will likely be associated with the default egress port for the PCH meaning it will be assigned port number 0.
23:16	0h RW_O	TCID: Target Component ID: Identifies the physical component that is targeted by this link entry. BIOS Requirement: should be initialized according to guidelines in the PCI Express* Isochronous/Virtual Channel Support Hardware Programming Specification (HPS).
15:2	0h RO	Reserved (RSVD): Reserved.
1	0h RO	LTP: Link Type: Indicates that the link points to memory-mapped space (for RCRB). The link address specifies the 64-bit base address of the target RCRB.
0	0h RW_O	LV: Link Valid: 0: Link Entry is not valid and will be ignored. 1: Link Entry specifies a valid link.



6.17 DMI Link Entry 1 Address (DMILE1A)—Offset 58h


Second part of a Link Entry which declares an internal link to another Root Complex Element.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 58h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:12	0h RW_O	LA: Link Address: Memory mapped base address of the RCRB that is the target element (egress port of PCH) for this link entry.
11:0	0h RO	Reserved (RSVD): Reserved.

6.18 DMI Link Upper Entry 1 Address (DMILUE1A)—Offset 5Ch


Second part of a Link Entry which declares an internal link to another Root Complex Element.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5Ch

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31:8	0h RO	Reserved (RSVD): Reserved.
7:0	0h RW_O	ULA: Upper Link Address: Memory mapped base address of the RCRB that is the target element (egress port of PCH) for this link entry.

6.19 DMI Link Entry 2 Description (DMILE2D)—Offset 60h


First part of a Link Entry which declares an internal link to another Root Complex Element.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 60h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:24	0h RO	TPN: Target Port Number: Specifies the port number associated with the element targeted by this link entry (Egress Port). The target port number is with respect to the component that contains this element as specified by the target component ID.
23:16	0h RW_O	TCID: Target Component ID: Identifies the physical or logical component that is targeted by this link entry. BIOS Requirement: should be initialized according to guidelines in the PCI Express* Isochronous/Virtual Channel Support Hardware Programming Specification (HPS).
15:2	0h RO	Reserved (RSVD): Reserved.
1	0h RO	LTYP: Link Type: Indicates that the link points to memory-mapped space (for RCRB). The link address specifies the 64-bit base address of the target RCRB.
0	0h RW_O	LV: Link Valid: 0: Link Entry is not valid and will be ignored. 1: Link Entry specifies a valid link.



6.20 DMI Link Entry 2 Address (DMILE2A)—Offset 68h


Second part of a Link Entry which declares an internal link to another Root Complex Element.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 68h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:12	0h RW_O	LA: Link Address: Memory mapped base address of the RCRB that is the target element (Egress Port) for this link entry.
11:0	0h RO	Reserved (RSVD): Reserved.

6.21 Link Capabilities (LCAP)—Offset 84h

Indicates DMI specific capabilities.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 84h

Default: 41AC43h



3 1	2 8	2 4	2 0	1 6	1 2	8	4	0
0 0 0 0	0 0 0 0	0 1 0 0	0 0 0 1	1 0 1 0	1 1 0 0	0 1 0 0	0 0 1 1	
RSVD		ASPM_OPT_COMPLIANCE	RSVD	L1SELAT	LOSELAT	ASLPMS	MLW	MLS

Bit Range	Default and Access	Field Name (ID): Description
31:23	0h RO	Reserved (RSVD): Reserved.
22	1h RO	ASPM_OPT_COMPLIANCE: ASPM Optionality Compliance. This bit should be set to 1b in all Functions. Components implemented against certain earlier versions of this specification will have this bit set to 0b. Software is permitted to use the value of this bit to help determine whether to enable ASPM or whether to run ASPM compliance tests.
21:18	0h RO	Reserved (RSVD): Reserved.
17:15	3h RW_O	L1SELAT: L1 Exit Latency: Indicates the length of time this Port requires to complete the transition from L1 to L0. The value 010b indicates the range of 2 us to less than 4 us. 000: Less than 1 us 001: 1 us to less than 2 us 010: 2 us to less than 4 us 011: 4 us to less than 8 us 100: 8 us to less than 16 us 101: 16 us to less than 32 us 110: 32 us-64 us 111: More than 64 us Both bytes of this register that contain a portion of this field should be written simultaneously in order to prevent an intermediate (and undesired) value from ever existing.
14:12	2h RW_O	LOSELAT: L0s Exit Latency: Indicates the length of time this Port requires to complete the transition from L0s to L0. 000: Less than 64 ns 001: 64 ns to less than 128 ns 010: 128 ns to less than 256 ns 011: 256 ns to less than 512 ns 100: 512 ns to less than 1 us 101: 1 us to less than 2 us 110: 2 us-4 us 111: More than 4 us



Bit Range	Default and Access	Field Name (ID): Description
11:10	3h RO	ASLPMS: Active State Link PM Support: L0s & L1 entry supported.
9:4	4h RO	MLW: Indicates the maximum number of lanes supported for this link.
3:0	3h RW_OV	MLS: This default value reflects gen1. Later the field may be changed by BIOS to allow gen2. Defined encodings are: 0001b: 2.5 GT/s Link speed supported 0010b: 5.0 GT/s and 2.5 GT/s Link speeds supported 0011b: 8.0 GT/s and 5.0 GT/s and 2.5 GT/s Link speeds supported

6.22 Link Control (LCTL)—Offset 88h

Allows control of PCI Express link.

Access Method

Type: MEM
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + 88h

Default: 0h

15	12	8	4	0						
0	0	0	0	0						
RSVD				HAWD	RSVD	ES	RSVD	RL	RSVD	ASPM

Bit Range	Default and Access	Field Name (ID): Description
15:10	0h RO	Reserved (RSVD): Reserved.
9	0h RO	HAWD: OPI - N/A Hardware Autonomous Width Disable: Hardware Autonomous Width Disable - When Set, this bit disables hardware from changing the Link width for reasons other than attempting to correct unreliable Link operation by reducing Link width. Devices that do not implement the ability autonomously to change Link width are permitted to hardwire this bit to 0b.
8	0h RO	Reserved (RSVD): Reserved.
7	0h RW	ES: OPI - N/A Extended Synch: Extended synch 0: Standard Fast Training Sequence (FTS). 1: Forces the transmission of additional ordered sets when exiting the L0s state and when in the Recovery state. This mode provides external devices (e.g., logic analyzers) monitoring the Link time to achieve bit and symbol lock before the link enters L0 and resumes communication. This is a test mode only and may cause other undesired side effects such as buffer overflows or underruns.
6	0h RO	Reserved (RSVD): Reserved.



Bit Range	Default and Access	Field Name (ID): Description
5	0h RW	RL: Retrain Link: 0: Normal operation. 1: Full Link retraining is initiated by directing the Physical Layer LTSSM from L0, L0s, or L1 states to the Recovery state. This bit always returns 0 when read. This bit is cleared automatically (no need to write a 0).
4:2	0h RO	Reserved (RSVD): Reserved.
1:0	0h RO	ASPM: Active State PM: Controls the level of active state power management supported on the given link. 00: Disabled 01: L0s Entry Supported 10: L1 Entry Supported 11: L0s and L1 Entry Supported

6.23 DMI Link Status (LSTS)—Offset 8Ah

Indicates DMI status.

Access Method

Type: MEM
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + 8Ah

Default: 1h

15	12	8	4	0
0	0	0	0	1
RSVD		LTRN	RSVD	NWID
				NSPD

Bit Range	Default and Access	Field Name (ID): Description
15:12	0h RO	Reserved (RSVD): Reserved.
11	0h ROV	LTRN: Link Training: Indicates that the Physical Layer LTSSM is in the Configuration or Recovery state, or that 1b was written to the Retrain Link bit but Link training has not yet begun. Hardware clears this bit when the LTSSM exits the Configuration/Recovery state once Link training is complete.



Bit Range	Default and Access	Field Name (ID): Description
10	0h RO	Reserved (RSVD): Reserved.
9:4	0h ROV	NWID: Negotiated Width: Indicates negotiated link width. This field is valid only when the link is in the L0, L0s, or L1 states (after link width negotiation is successfully completed). 00h: Reserved 01h: X1 02h: X2 04h: X4 All other encodings are reserved.
3:0	1h ROV	NSPD: Negotiated Speed: Indicates negotiated link speed. 1h: 2.5 Gb/s 2h: 5.0 Gb/s All other encodings are reserved. The value in this field is undefined when the Link is not up.

6.24 Link Control 2 (LCTL2)—Offset 98h

Access Method

Type: MEM
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + 98h

Default: 1h

15	12	8	4	0
0	0	0	0	1
ComplianceDeemphasis	compos	txmargin	selectabledeemphasis	TLS
	entermodcompliance		HASD	
			EC	



Bit Range	Default and Access	Field Name (ID): Description
15:12	0h RWS	<p>ComplianceDeemphasis: Compliance De-emphasis: For 8 GT/s Data Rate: This field sets the Transmitter Preset level in Polling.Compliance state if the entry occurred due to the Enter Compliance bit being 1b. This bit sets the de-emphasis level in Polling.Compliance state if the entry occurred due to the Enter Compliance bit being 1b. Defined encodings are: 0001b -3.5 dB 0000b -6 dB When the Link is operating at 2.5 GT/s, the setting of this bit has no effect. Components that support only 2.5 GT/s speed are permitted to hardwire this bit to 0b. For a Multi-Function device associated with an Upstream Port, the bit in Function 0 is of type RWS, and only Function 0 controls the component's Link behavior. In all other Functions of that device, this bit is of type RsvdP. The default value of this bit is 0000b. This bit is intended for debug, compliance testing purposes. System firmware and software is allowed to modify this bit only during debug or compliance testing.</p>
11	0h RWS	<p>compos: Compliance SOS: When set to 1b, the LTSSM is required to send SKP Ordered Sets periodically in between the (modified) compliance patterns. For a Multi-Function device associated with an Upstream Port, the bit in Function 0 is of type RWS, and only Function 0 controls the component's Link behavior. In all other Functions of that device, this bit is of type RsvdP. The default value of this bit is 0b. This bit is applicable when the Link is operating at 2.5 GT/s or 5 GT/s data rates only. Components that support only the 2.5 GT/s speed are permitted to hardwire this field to 0b.</p>
10	0h RWS	<p>entermodcompliance: Enter Modified Compliance: When this bit is set to 1b, the device transmits modified compliance pattern if the LTSSM enters Polling.Compliance state. Components that support only the 2.5GT/s speed are permitted to hardwire this bit to 0b. Default value of this field is 0b.</p>
9:7	0h RWS_V	<p>txmargin: Transmit Margin: This field controls the value of the non-deemphasized voltage level at the Transmitter pins. This field is reset to 000b on entry to the LTSSM Polling.Configuration substrate. 000: Normal operating range 001: 800-1200 mV for full swing and 400-700 mV for half-swing 010 - (n-1): Values should be monotonic with a non-zero slope. The value of n should be greater than 3 and less than 7. At least two of these should be below the normal operating range n: 200-400 mV for full-swing and 100-200 mV for half-swing n -111: reserved Default value is 000b. Components that support only the 2.5GT/s speed are permitted to hardwire this bit to 0b. When operating in 5GT/s mode with full swing, the deemphasis ratio should be maintained within +/- 1dB from the specification defined operational value (either -3.5 or -6 dB).</p>
6	0h RWS	<p>selectabledeemphasis: Selectable De-emphasis: When the Link is operating at 5GT/s speed, selects the level of de-emphasis. Encodings: 1b: -3.5 dB 0b: -6 dB Default value is implementation specific, unless a specific value is required for a selected form factor or platform. When the Link is operating at 2.5GT/s speed, the setting of this bit has no effect. Components that support only the 2.5GT/s speed are permitted to hardwire this bit to 0b.</p>



Bit Range	Default and Access	Field Name (ID): Description
5	0h RWS	HASD: Hardware Autonomous Speed Disable: When set to 1b this bit disables hardware from changing the link speed for reasons other than attempting to correct unreliable link operation by reducing link speed.
4	0h RWS	EC: Enter Compliance: Software is permitted to force a link to enter Compliance mode at the speed indicated in the Target Link Speed field by setting this bit to 1b in both components on a link and then initiating a hot reset on the link.
3:0	1h RWS	<p>TLS: Target Link Speed: For Downstream Ports, this field sets an upper limit on Link operational speed by restricting the values advertised by the Upstream component in its training sequences. The encoding is the binary value of the bit in the Supported Link Speeds Vector (in the Link Capabilities 2 register) that corresponds to the desired target Link speed. All other encodings are reserved. For example, 5.0 GT/s corresponds to bit 2 in the Supported Link Speeds Vector, so the encoding for a 5.0 GT/s target Link speed in this field is 0010b.</p> <p>If a value is written to this field that does not correspond to a supported speed (as indicated by the Max Link Speed Vector), the result is undefined. The default value of this field is the highest Link speed supported by the component (as reported in the Max Link Speed field of the Link Capabilities register) unless the corresponding platform/form factor requires a different default value.</p> <p>For both Upstream and Downstream Ports, this field is used to set the target compliance mode speed when software is using the Enter Compliance bit to force a Link into compliance mode. For a Multi-Function device associated with an Upstream Port, the field in Function 0 is of type RWS, and only Function 0 controls the components Link behavior. In all other Functions of that device, this field is of type RsvdP.</p>

6.25 Link Status 2 (LSTS2)—Offset 9Ah

Access Method

Type: MEM
(Size: 16 bits)

Offset: [B:0, D:0, F:0] + 9Ah

Default: 0h

15	12	8	4	0					
0	0	0	0	0					
RSVD				LNKEQREQ	EQPH3SUCC	EQPH2SUCC	EQPH1SUCC	EQCOMPLETE	CURDELVL

Bit Range	Default and Access	Field Name (ID): Description
15:6	0h RO	Reserved (RSVD): Reserved.
5	0h RW1C	LNKEQREQ: This bit is Set by hardware to request the Link equalization process to be performed on the Link.
4	0h ROV	EQPH3SUCC: Equalization Phase 3 Successful When set to 1b, this bit indicates that Phase 3 of the Transmitter Equalization procedure has successfully completed.
3	0h ROV	EQPH2SUCC: Equalization Phase 2 Successful When set to 1b, this bit indicates that Phase 2 of the Transmitter Equalization procedure has successfully completed.



Bit Range	Default and Access	Field Name (ID): Description
2	0h ROV	EQPH1SUCC: Equalization Phase 1 Successful When set to 1b, this bit indicates that Phase 1 of the Transmitter Equalization procedure has successfully completed.
1	0h ROV	EQCOMPLETE: Equalization Complete When set to 1b, this bit indicates that the Transmitter Equalization procedure has completed.
0	0h RO	CURDELVL: Current De-emphasis Level: Current De-emphasis Level - When the Link is operating at 5 GT/s speed, this reflects the level of de-emphasis. 1: -3.5 dB 0: -6 dB When the Link is operating at 2.5 GT/s speed, this bit is 0b.

6.26 DMI Uncorrectable Error Status (DMIUESTS)—Offset 1C4h


DMI Uncorrectable Error Status register. This register is for test and debug purposes only.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 1C4h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:21	0h RO	Reserved (RSVD): Reserved.
20	0h RW1CS	URES: Unsupported Request Error Status:
19	0h RO	Reserved (RSVD): Reserved.
18	0h RW1CS	MTLPS: Malformed TLP Status:
17	0h RW1CS	ROS: Receiver Overflow Status:
16	0h RW1CS	UCS: Unexpected Completion Status:
15	0h RO	Reserved (RSVD): Reserved.



Bit Range	Default and Access	Field Name (ID): Description
14	0h RW1CS	CTS: Completion Timeout Status:
13	0h RO	Reserved (RSVD): Reserved.
12	0h RW1CS	PTLPS: Poisoned TLP Status:
11:5	0h RO	Reserved (RSVD): Reserved.
4	0h RW1CS	DLPEs: Data Link Protocol Error Status:
3:0	0h RO	Reserved (RSVD): Reserved.

6.27 DMI Uncorrectable Error Mask (DMIUEMSK)—Offset 1C8h

DMI Uncorrectable Error Mask register. This register is for test and debug purposes only.

Access Method


Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 1C8h

Default: 0h

3	2	2	2	1	1	8	4	0							
1	8	4	0	6	2										
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0							
RSVD		ECCERRM	RSVD	UREM	RSVD	MTLPM	ROM	UCM	RSVD	CPLTM	RSVD	PTLPM	RSVD	DLPEM	RSVD

Bit Range	Default and Access	Field Name (ID): Description
31:23	0h RO	Reserved (RSVD): Reserved.
22	0h RWS	ECCERRM: 2 Bit Error Mask:
21	0h RO	Reserved (RSVD): Reserved.
20	0h RWS	UREM: Unsupported Request Error Mask:



Bit Range	Default and Access	Field Name (ID): Description
19	0h RO	Reserved (RSVD): Reserved.
18	0h RWS	MTLPM: Malformed TLP Mask:
17	0h RWS	ROM: Receiver Overflow Mask:
16	0h RWS	UCM: Unexpected Completion Mask:
15	0h RO	Reserved (RSVD): Reserved.
14	0h RWS	CPLTM: Completion Timeout Mask:
13	0h RO	Reserved (RSVD): Reserved.
12	0h RWS	PTLPM: Poisoned TLP Mask:
11:5	0h RO	Reserved (RSVD): Reserved.
4	0h RWS	DLPEM: Data Link Protocol Error Mask:
3:0	0h RO	Reserved (RSVD): Reserved.

6.28 DMI Uncorrectable Error Severity (DMIUESEV)—Offset 1CCh

DMI Uncorrectable Error Severity register. This register controls whether an individual error is reported as a non-fatal or fatal error. An error is reported as fatal when the corresponding error bit in the severity register is set. If the bit is cleared, the corresponding error is considered nonfatal. It is for test and debug purposes only.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 1CCh

Default: 60010h

3 1	2 8	2 4	2 0	1 6	1 2	8	4	0
0 0 0 0	0 0 0 0	0 0 0 0	0 1 1 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 1	0 0 0 0
	RSVD	ECCERS RSVD URES	ECRCES MTLPES ROEV UCES	CAES CTES FCPES PTLPES		RSVD	DLPEM	RSVD



Bit Range	Default and Access	Field Name (ID): Description
31:23	0h RO	Reserved (RSVD): Reserved.
22	0h RWS	ECCERRS: 2 Bit Error Mask:
21	0h RO	Reserved (RSVD): Reserved.
20	0h RWS	URES: Unsupported Request Error Severity:
19	0h RO	ECRCES: Reserved for ECRC Error Severity:
18	1h RWS	MTLPES: Malformed TLP Error Severity:
17	1h RWS	ROEV: Receiver Overflow Error Severity:
16	0h RWS	UCES: Unexpected Completion Error Severity:
15	0h RO	CAES: Reserved for Completer Abort Error Severity:
14	0h RWS	CTES: Completion Timeout Error Severity:
13	0h RO	FCPES: Reserved for Flow Control Protocol Error Severity:
12	0h RWS	PTLPES: Poisoned TLP Error Severity:
11:5	0h RO	Reserved (RSVD): Reserved.
4	1h RWS	DLPEs: Data Link Protocol Error Severity:
3:0	0h RO	Reserved (RSVD): Reserved.

6.29 DMI Correctable Error Status (DMICESTS)—Offset 1D0h


DMI Correctable Error Status Register. This register is for test and debug purposes only.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 1D0h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:14	0h RO	Reserved (RSVD): Reserved.
13	0h RW1CS	ANFES: Advisory Non-Fatal Error Status: When set, indicates that an Advisory Non-Fatal Error occurred.
12	0h RW1CS	RTTS: Replay Timer Timeout Status:
11:9	0h RO	Reserved (RSVD): Reserved.
8	0h RW1CS	RNRS: REPLAY_NUM Rollover Status:
7	0h RW1CS	BDLLPS: Bad DLLP Status:
6	0h RW1CS	BTLPS: Bad TLP Status:
5:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW1CS	RES: Receiver Error Status: Physical layer receiver Error occurred. These errors include: elastic Buffer Collision, 8b/10b error, De-skew Timeout Error.

6.30 DMI Correctable Error Mask (DMICEMSK)—Offset 1D4h


DMI Correctable Error Mask register. This register is for test and debug purposes only.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 1D4h

Default: 2000h





Bit Range	Default and Access	Field Name (ID): Description
31:14	0h RO	Reserved (RSVD): Reserved.
13	1h RWS	ANFEM: Advisory Non-Fatal Error Mask: When set, masks Advisory Non-Fatal errors from (a) signaling ERR_COR to the device control register, and (b) updating the Uncorrectable Error Status register. This register is set by default to enable compatibility with software that does not comprehend Role-Based Error Reporting.
12:0	0h RO	Reserved (RSVD): Reserved.

§ §



7 MCHBAR Registers

Table 7-1. Summary of Bus: 0, Device: 0, Function: 0 (MEM) (Sheet 1 of 3)

Offset	Size (Bytes)	Register Name (Register Symbol)	Default Value
4000h	4	MCHBAR_CH0_CR_TC_PRE_0_0_0_MCHBAR—Offset 4000h	0h
401Ch	4	MCHBAR_CH0_CR_SC_GS_CFG_0_0_0_MCHBAR—Offset 401Ch	0h
4048–404Bh	4	ECCERRLOG0—Offset 4048h	0h
404C–404Fh	4	ECC Error Log 1 (ECCERRLOG1)—Offset 404Ch	0h
4070h	4	MCHBAR_CH0_CR_TC_ODT_0_0_0_MCHBAR—Offset 4070h	0h
4238–423Bh	4	Refresh Parameters (TC)—Offset 4238h	4600980Fh
423C–423Fh	4	Refresh Timing Parameters (TC)—Offset 423Ch	B41004h
4260–4263h	4	Power Management DIMM Idle Energy (PM)—Offset 4260h	0h
4264–4267h	4	Power Management DIMM Power Down Energy (PM)—Offset 4264h	0h
4268–426Bh	4	Power Management DIMM Activate Energy (PM)—Offset 4268h	0h
426C–426Fh	4	Power Management DIMM RdCas Energy (PM)—Offset 426Ch	0h
4270–4273h	4	Power Management DIMM WrCas Energy (PM)—Offset 4270h	0h
4400h	4	MCHBAR_CH1_CR_TC_PRE_0_0_0_MCHBAR—Offset 4400h	0h
441Ch	4	MCHBAR_CH0_CR_SC_GS_CFG_0_0_0_MCHBAR—Offset 441Ch	0h
4448–444Bh	4	ECCERRLOG0—Offset 4448h	0h
444C–444Fh	4	ECC Error Log 1 (ECCERRLOG1)—Offset 444Ch	0h
4470h	4	MCHBAR_CH0_CR_TC_ODT_0_0_0_MCHBAR—Offset 4470h	0h
4638–463Bh	4	Refresh Parameters (TC)—Offset 4638h	4600980Fh
463C–463Fh	4	Refresh Timing Parameters (TC)—Offset 463Ch	B41004h
4660–4663h	4	Power Management DIMM Idle Energy (PM)—Offset 4660h	0h
4664–4667h	4	Power Management DIMM Power Down Energy (PM)—Offset 4664h	0h
4668–466Bh	4	Power Management DIMM Activate Energy (PM)—Offset 4668h	0h
466C–466Fh	4	Power Management DIMM RdCas Energy (PM)—Offset 466Ch	0h
4670–4673h	4	Power Management DIMM WrCas Energy (PM)—Offset 4670h	0h
4C1Ch	4	MCSCHEDS_CR_SC_GS_CFG_0_0_0_MCHBAR—Offset 4C1Ch	0h
4C40–4C43h	4	PM—Offset 4C40h	0h
4C48–4C4Bh	4	ECCERRLOG0—Offset 4C48h	0h
4C4C–4C4Fh	4	ECC Error Log 1 (ECCERRLOG1)—Offset 4C4Ch	0h
4C70h	4	MCSCHEDS_CR_TC_ODT_0_0_0_MCHBAR—Offset 4C70h	0h
4E38–4E3Bh	4	Refresh Parameters (TC)—Offset 4E38h	4600980Fh
4E3C–4E3Fh	4	Refresh Timing Parameters (TC)—Offset 4E3Ch	B41004h
4E60–4E63h	4	Power Management DIMM Idle Energy (PM)—Offset 4E60h	0h
4E64–4E67h	4	Power Management DIMM Power Down Energy (PM)—Offset 4E64h	0h
4E68–4E6Bh	4	Power Management DIMM Activate Energy (PM)—Offset 4E68h	0h
4E6C–4E6Fh	4	Power Management DIMM RdCas Energy (PM)—Offset 4E6Ch	0h



Table 7-1. Summary of Bus: 0, Device: 0, Function: 0 (MEM) (Continued) (Sheet 2 of 3)

Offset	Size (Bytes)	Register Name (Register Symbol)	Default Value
4E70–4E73h	4	Power Management DIMM WrCas Energy (PM)—Offset 4E70h	0h
5000–5003h	4	Address Decoder Inter Channel Configuration Register (MAD)—Offset 5000h	0h
5004–5007h	4	Address Decoder Intra Channel Configuration Register (MAD)—Offset 5004h	0h
5008–500Bh	4	Address Decoder Intra Channel Configuration Register (MAD)—Offset 5008h	0h
500C–500Fh	4	Address Decode DIMM Parameters. (MAD)—Offset 500Ch	0h
5010–5013h	4	Address Decode DIMM Parameters (MAD)—Offset 5010h	0h
5034h	4	MCDECS_CR_MRC_REVISION_0_0_0_MCHBAR_MCMAIN—Offset 5034h	0h
5040–5043h	4	Request Count from GT (DRAM)—Offset 5040h	0h
5044–5047h	4	Request Count from IA (DRAM)—Offset 5044h	0h
5048–504Bh	4	Request Count from IO (DRAM)—Offset 5048h	0h
5050–5053h	4	RD Data Count (DRAM)—Offset 5050h	0h
5054–5057h	4	WR Data Count (DRAM)—Offset 5054h	0h
5060–5063h	4	Self Refresh Configuration Register (PM)—Offset 5060h	10200h
5090–5093h	4	Address Compare for ECC Error Injection (ECC)—Offset 5090h	0h
5094–5097h	4	Address Mask for ECC Error Injection (ECC)—Offset 5094h	FFFFFFFFh
5400h	4	NCDECS_CR_GFXVTBAR_0_0_0_MCHBAR_NCU—Offset 5400h	0h
5410h	4	NCDECS_CR_VTDPVC0BAR_0_0_0_MCHBAR_NCU—Offset 5410h	0h
5820–5823h	4	PACKAGE—Offset 5820h	0h
5828–582Fh	8	PKG—Offset 5828h	0h
5830–5837h	8	PKG—Offset 5830h	0h
5838–583Fh	8	PKG—Offset 5838h	0h
5840–5847h	8	PKG—Offset 5840h	0h
5848–584Fh	8	PKG—Offset 5848h	0h
5858–585Fh	8	PKG—Offset 5858h	0h
5880–5883h	4	DDR—Offset 5880h	0h
5884–5887h	4	DRAM—Offset 5884h	3h
5888–588Bh	4	DRAM—Offset 5888h	0h
588C–588Fh	4	DDR—Offset 588Ch	0h
5890–5893h	4	DDR—Offset 5890h	FFFFh
5894–5897h	4	DDR—Offset 5894h	FFFFh
5898–589Bh	4	DDR—Offset 5898h	FFFFh
589C–589Fh	4	DDR—Offset 589Ch	FFFFh
58A0–58A3h	4	DDR—Offset 58A0h	0h
58A8–58ABh	4	PACKAGE—Offset 58A8h	7F00h
58B0–58B3h	4	DDR—Offset 58B0h	0h
58B4–58B7h	4	DDR—Offset 58B4h	0h
58C0–58C7h	8	DDR—Offset 58C0h	0h
58C8–58CFh	8	DDR—Offset 58C8h	0h


Table 7-1. Summary of Bus: 0, Device: 0, Function: 0 (MEM) (Continued) (Sheet 3 of 3)

Offset	Size (Bytes)	Register Name (Register Symbol)	Default Value
58D0–58D3h	4	DDR—Offset 58D0h	FFFFh
58D4–58D7h	4	DDR—Offset 58D4h	FFFFh
58D8–58DBh	4	DDR—Offset 58D8h	FFFFh
58DC–58DFh	4	DDR—Offset 58DCh	FFFFh
58F0–58F3h	4	PACKAGE—Offset 58F0h	0h
58FC–58FFh	4	IA—Offset 58FCh	0h
5900–5903h	4	GT—Offset 5900h	0h
5918–591Bh	4	SA—Offset 5918h	0h
5948–594Bh	4	GT—Offset 5948h	0h
594C–594Fh	4	EDRAM—Offset 594Ch	0h
5978–597Bh	4	Package—Offset 5978h	0h
597C–597Fh	4	PP0—Offset 597Ch	0h
5980–5983h	4	PP1—Offset 5980h	0h
5994–5997h	4	RP—Offset 5994h	FFh
5998–599Bh	4	RP—Offset 5998h	0h
5D10–5D17h	8	SSKPD—Offset 5D10h	0h
5DA8–5DABh	4	BIOS_RESET_CPL—Offset 5DA8h	0h
5E00h	4	PCU_CR_MC_BIOS_REQ_0_0_0_MCHBAR_PCU—Offset 5E00h	0h
5F3C–5F3Fh	4	CONFIG—Offset 5F3Ch	0h
5F40–5F47h	8	CONFIG—Offset 5F40h	0h
5F48–5F4Fh	8	CONFIG—Offset 5F48h	0h
5F50–5F53h	4	CONFIG—Offset 5F50h	0h
5F54–5F57h	4	TURBO—Offset 5F54h	0h
6200–6203h	4	Package Thermal DPPM Status (PKG)—Offset 6200h	8000h
6204–6207h	4	Memory Thermal DPPM Status (DDR)—Offset 6204h	0h



7.1 MCHBAR_CHO_CR_TC_PRE_0_0_0_MCHBAR— Offset 4000h


DDR timing

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4000h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
30:24	0h RO	tWRPRE: Holds DDR timing parameter tWRPRE. WR to PRE same bank minimum delay in DCLK cycles. Note: tWRRD_sg+tRDPRE should be greater than or equal to tWRPRE Supported range is 23-95.
19:16	0h RO	tRDPRE: Holds DDR timing parameter tRDPRE. RD to PRE same bank minimum delay in DCLK cycles. Supported range is 6-15.
14:8	0h RO	tRAS: Holds DDR timing parameter tRAS. ACT to PRE same bank minimum delay in DCLK cycles. Supported range is 28-64.
7:6	0h RO	tRPab_ext: Holds the value of tRPab-tRPpb for LPDDR3 in DCLK cycles LPDDR3 requires a longer time from PREAL to ACT vs. PRE to ACT, the offset between the two should be programmed to this field. When using DDR3/DDR4 this field should be programmed to 0. Supported range is 0-3.
5:0	0h RO	tRP: Holds DDR timing parameter tRP (and tRCD). PRE to ACT same bank minimum delay in DCLK cycles. ACT to CAS (RD or WR) same bank minimum delay in DCLK cycles. For LPDDR3 this field should hold tRPpb (and tRCD) values. Supported range is 8-63.



7.2 MCHBAR_CHO_CR_SC_GS_CFG_0_0_0_MCHBAR— Offset 401Ch

Scheduler configuration

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 401Ch

Default: 0h

3 1	2 8	2 4	2 0	1 6	1 2	8	4	0
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0 0 0 0 0	0 0 0 0	0 0 0 0
tCAL	ddr_probeless_low_frequency enable_odt_matrix	ck_to_cke	cmd_3st reset_delay	reset_on_command	LPDDR_2N_CS_MRW tCPDED	x8_device Address_mirror	RSVD N_to_1_ratio	CMD_stretch DRAM_technology

Bit Range	Default and Access	Field Name (ID): Description
31:29	0h RO	tCAL: For DDR4, holds tCAL value. Supported values: 0 (CAL mode disabled), 3-5 (CAL mode enabled, value is the delay in DCLK cycles from CSb to command). Updating this field is required only after sending MRS to MR4 enabling/disabling CAL mode before any other command is sent to DRAM. TC_MR4_shaddow_0_0_0_MCHBAR should be updated with the correct value of tCAL once its value changes.
28	0h RO	ddr_probeless_low_frequency: This bit controls whether the DDR probeless logic uses DDR_TX_DELAY_LOW or DDR_TX_DELAY_HIGH for the internal delay of the write data. If MRC supports two training frequencies, this bit should be set when training at the low frequency.
27	0h RO	enable_odt_matrix: When bit is set, the ranks that are used for terminating when giving read/write requests are selected according to SC_ODT_MATRIX control register and not according to the default behavior.
26:24	0h RO	ck_to_cke: When working with LPDDR when CKE is low we also turn off the CKE buffers. The LPDDR specification requires starting the CK toggling two DCLK cycles before re-asserting CKE. The field defines the number of DCLK cycles from CKoutputEnable assert on power down exit to CKE assert as the DDRIO can delay the CK pins differently than CKE so a different value is required to get two DCLK cycles of CK toggling before CKE rise. Typically this field should be programmed to 3 if (CLK_pi+CLK_logicdelay)-(CKE_pi+CKE_logicdelay) is less than 1 QCLK. Otherwise, it should be programmed to 4 supported range is 2-7.
23	0h RO	cmd_3st: Defines when command and address bus is driving. 0: Drive when channel is active. Tri-stated when all ranks are in CKE-off or when memory is in SR or deeper. 1: Command bus is always driving. When no new valid command is driven, previous command and address is driven



Bit Range	Default and Access	Field Name (ID): Description
22:20	0h RO	reset_delay: Inserts an N Dclk delay ranging from 0 to 7 after the N to 1 Reset on Cmd is triggered.
19:16	0h RO	reset_on_command: The N:1 logic can be triggered to insert a bubble and reset the N:1 logic after a programmable delay from a command after a PRE/ACT/RD/WR CMD. This allows one to synchronize the N:1 logic periodically to ensure the correct worst case pattern between victim and aggressor occurs when training the command bus. Reset N to 1 Logic on a WR (bit 16) Reset N to 1 Logic on a RD (bit 17) Reset N to 1 Logic on a ACT (bit 18) Reset N to 1 Logic on a PRE (bit 19)
15	0h RO	LPDDR_2N_CS_MRW: When sending an MRW command via the MRH for LPDDR, drive the CSb for two DCLK cycles
14:12	0h RO	tCPDED: Holds DDR timing parameter tCPDED. Power down to command bus tri-state delay in DCLK cycles. Supported range is 1-7 in 1N mode.
11:10	0h RO	x8_device: DIMM is made out of X8 devices LSB is for DIMM 0, MSB is for DIMM 1.
9:8	0h RO	Address_mirror: DIMM routing causes address mirroring LSB is for DIMM 0, MSB is for DIMM 1.
6:4	0h RO	N_to_1_ratio: When using N:1 command stretch mode, every how many B2B valid command cycles a bubble is required Supported range is 1 to 7
3:2	0h RO	CMD_stretch: Command stretch mode: 00: 1N 01: 2N 10: 3N 11: N:1
1:0	0h RO	DRAM_technology: DRAM technology: 00: DDR4 01: DDR3 10: LPDDR3 11: Illegal

7.3 ECCERRLOG0—Offset 4048h

This register logs ECC error information. Represents channel 0.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4048h

Default: 0h

3	2	2	2	1	1	8	4	0
1	8	4	0	6	2			
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
ERRBANK	ERRRANK	ERRCHUNK	ERRSYND			RSVD		MERRSTS CERRSTS



Bit Range	Default and Access	Field Name (ID): Description
31:29	0h ROSV	ERRBANK: This field holds the Bank Address of the read transaction that had the ECC error.
28:27	0h ROSV	ERRRANK: This field holds the Rank ID of the read transaction that had the ECC error.
26:24	0h ROSV	ERRCHUNK: Holds the chunk number of the error stored in the register.
23:16	0h ROSV	ERRSYND: This field contains the error syndrome. A value of 0xFF indicates that the error is due to poisoning.
15:2	0h RO	Reserved (RSVD): Reserved.
1	0h ROSV	MERRSTS: This bit is set when an uncorrectable multiple-bit error occurs on a memory read data transfer. When this bit is set, the address that caused the error and the error syndrome are also logged and they are locked until this bit is cleared. This bit is cleared when the corresponding bit in 0.0.0.PCI.ERRSTS is cleared.
0	0h ROSV	CERRSTS: This bit is set when a correctable single-bit error occurs on a memory read data transfer. When this bit is set, the address that caused the error and the error syndrome are also logged and they are locked to further single bit errors, until this bit is cleared. A multiple bit error that occurs after this bit is set will override the address/error syndrome information. This bit is cleared when the corresponding bit in 0.0.0.PCI.ERRSTS is cleared.

7.4 ECC Error Log 1 (ECCERRLOG1)—Offset 404Ch

This register logs ECC error information.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 404Ch

Default: 0h

3 1	2 8	2 4	2 0	1 6	1 2	8	4	0
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0
RSVD	ERRBANKGROUP	ERRCOL		ERROR				



Bit Range	Default and Access	Field Name (ID): Description
31:30	0h RO	Reserved (RSVD): Reserved.
29:28	0h ROSV	ERRBANKGROUP: This field holds the DRAM bank group (for DDR4) address of the read transaction that had the ECC error.
27:17	0h ROSV	ERRCOL: This field holds the DRAM column address of the read transaction that had the ECC error.
16:0	0h ROSV	ERRROW: This field holds the DRAM row (page) address of the read transaction that had the ECC error.

7.5 MCHBAR_CHO_CR_TC_ODT_0_0_0_MCHBAR— Offset 4070h

ODT timing related parameters

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4070h

Default: 0h

3 1	2 8	2 4	2 0	1 6	1 2	8	4	0			
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0 0 0	0 0 0 0	0 0 0 0			
ODT_Always_Rank0	tAONPD	tCWL	tCL	Write_Early_ODT	ODT_Write_Delay	RSVD	ODT_write_duration	RSVD	ODT_Read_Delay	RSVD	ODT_read_duration

Bit Range	Default and Access	Field Name (ID): Description
31	0h RO	ODT_Always_Rank0: Indicate that ODT should always be multiplexed out on ODT[0], to be used for LPDDR3 only.
30:26	0h RO	tAONPD: Holds DDR timing parameter tAONPD. Supported range is 4-31.
25:21	0h RO	tCWL: Holds DDR timing parameter tCWL (sometimes referred to as tWCL). Write command to data delay in DCLK cycles. Supported range is 4-20 (maximum is for 1N mode and tCAL=0) For LPDDR3 the minimum supported value is 4 if Dec_WRD=0 and 5 if Dec_WRD=1. For DDR3/4 the minimum supported value is 5 if Dec_WRD=0 and 6 if Dec_WRD=1.
20:16	0h RO	tCL: Holds DDR timing parameter tCL. Read command to data delay in DCLK cycles. Supported range is 5-31.



Bit Range	Default and Access	Field Name (ID): Description
15	0h RO	Write_Early_ODT: When this bit is set, the MC is will send one extra cycle of ODT prior to the write command. In this mode the ranks that will be terminated on this early cycle are selected according to the SC_ODT_MATRIX_0_0_0_MCHBAR control register.
14:12	0h RO	ODT_Write_Delay: Controls delay from WR-CAS to ODT assertion in DCLK cycles (Typical Programming = 0).
10:8	0h RO	ODT_write_duration: Controls the length of the ODT pulse for write commands. Default is 6 DCLK cycles (BL/2 + 2) 000: 6 DCLK cycles 001: 7 DCLK cycles 010: 8 DCLK cycles 011: 9 DCLK cycles 100: 10 DCLK cycles 101: 11 DCLK cycles 110: 12 DCLK cycles 111: 13 DCLK cycles
6:4	0h RO	ODT_Read_Delay: Controls delay from RD-CAS to ODT assertion in DCLK cycles (Typical Programming = tCL-tCWL). Note 1: All RD->RD and RD->WR restrictions should be greater than or equal to this field value. Note 2: odt_read_delay + odt_read_duration should not be programmed to less than tCL-
2:0	0h RO	ODT_read_duration: Controls the length of the ODT pulse for read commands. Default is 6 DCLK cycles (BL/2 + 2) 000: 6 DCLK cycles 001: 7 DCLK cycles 010: 8 DCLK cycles 011: 9 DCLK cycles 100: 10 DCLK cycles 101: 11 DCLK cycles 110: 12 DCLK cycles 111: 13 DCLK cycles

7.6 Refresh Parameters (TC)—Offset 4238h


Refresh parameters

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4238h

Default: 4600980Fh



Bit Range	Default and Access	Field Name (ID): Description
31:25	23h RW_L	tREFIX9: Maximum time allowed between refreshes to a rank (in intervals of 1024 DCLK cycles). Should be programmed to 8.9*tREFI/1024 (to allow for possible delays from ZQ or isoc).
24:16	0h RO	Reserved (RSVD): Reserved.
15:12	9h RW_L	Refresh_panic_wm: tREFI count level in which the refresh priority is panic (default is 9). The Maximum value for this field is 9.
11:8	8h RW_L	Refresh_HP_WM: tREFI count level that turns the refresh priority to high (default is 8)
7:0	Fh RW_L	OREF_RI: Rank idle period that defines an opportunity for refresh, in DCLK cycles

7.7 Refresh Timing Parameters (TC)—Offset 423Ch


Refresh timing parameters

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 423Ch

Default: B41004h





Bit Range	Default and Access	Field Name (ID): Description
31:26	0h RO	Reserved (RSVD): Reserved.
25:16	B4h RW_L	tRFC: Time of refresh - from beginning of refresh until next ACT or refresh is allowed (in DCLK cycles, default is 180).
15:0	1004h RW_L	tREFI: defines the average period between refreshes, and the rate that tREFI counter is incremented (in DCLK cycles, default is 4100).

7.8 Power Management DIMM Idle Energy (PM)—Offset 4260h


This register defines the energy of an idle DIMM with CKE on. Each 6-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are two 6-bit fields, one per DIMM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4260h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:14	0h RO	Reserved (RSVD): Reserved.
13:8	0h RW_L	DIMM1_IDLE_ENERGY: This register defines the energy consumed by DIMM1 for one clock cycle when the DIMM is idle with cke on.
7:6	0h RO	Reserved (RSVD): Reserved.
5:0	0h RW_L	DIMM0_IDLE_ENERGY: This register defines the energy consumed by DIMM0 for one clock cycle when the DIMM is idle with cke on.



7.9 Power Management DIMM Power Down Energy (PM)—Offset 4264h


This register defines the energy of an idle DIMM with CKE off. Each 6-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are two 6-bit fields, one per DIMM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4264h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:14	0h RO	Reserved (RSVD): Reserved.
13:8	0h RW_L	DIMM1_PD_ENERGY: This register defines the energy consumed by DIMM1 for one clock cycle when the DIMM is idle with cke off.
7:6	0h RO	Reserved (RSVD): Reserved.
5:0	0h RW_L	DIMM0_PD_ENERGY: This register defines the energy consumed by DIMM0 for one clock cycle when the DIMM is idle with cke off.



7.10 Power Management DIMM Activate Energy (PM)—Offset 4268h


This register defines the combined energy contribution of activate and precharge commands. Each 8-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 8-bit fields, one per DIMM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4268h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	0h RW_L	DIMM1_ACT_ENERGY: This register defines the combined energy contribution of activate and precharge commands.
7:0	0h RW_L	DIMM0_ACT_ENERGY: This register defines the combined energy contribution of activate and precharge commands.



7.11 Power Management DIMM RdCas Energy (PM)—Offset 426Ch


This register defines the energy contribution of a read CAS command. Each 8-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 8-bit fields, one per DIMM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 426Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	0h RW_L	DIMM1_RD_ENERGY: This register defines the energy contribution of a read CAS command.
7:0	0h RW_L	DIMM0_RD_ENERGY: This register defines the energy contribution of a read CAS command.



7.12 Power Management DIMM WrCas Energy (PM)—Offset 4270h


This register defines the energy contribution of a write CAS command. Each 8-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 8-bit fields, one per DIMM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4270h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	0h RW_L	DIMM1_WR_ENERGY: This register defines the energy contribution of a write CAS command.
7:0	0h RW_L	DIMM0_WR_ENERGY: This register defines the energy contribution of a write CAS command.



7.13 MCHBAR_CH1_CR_TC_PRE_0_0_0_MCHBAR— Offset 4400h

DDR timing constraints related to PRE commands

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4400h

Default: 0h

3 1	2 8	2 4	2 0	1 6	1 2	8	4	0
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0 0	0 0 0 0	0 0 0 0
RSVD	tWRPRE	RSVD	tRDPRE	RSVD	tRAS	tRPab	tRP	

Bit Range	Default and Access	Field Name (ID): Description
30:24	0h RO	tWRPRE: Holds DDR timing parameter tWRPRE. WR to PRE same bank minimum delay in DCLK cycles. Note: tWRRD_sg+tRDPRE should be greater than or equal to tWRPRE Supported range is 23-95.
19:16	0h RO	tRDPRE: Holds DDR timing parameter tRDPRE. RD to PRE same bank minimum delay in DCLK cycles. Supported range is 6-15.
14:8	0h RO	tRAS: Holds DDR timing parameter tRAS. ACT to PRE same bank minimum delay in DCLK cycles. Supported range is 28-64.
7:6	0h RO	tRPab: Holds the value of tRPab-tRPpb for LPDDR3 in DCLK cycles LPDDR3 requires a longer time from PREAL to ACT vs. PRE to ACT, the offset between the two should be programmed to this field. When using DDR3/DDR4 this field should be programmed to 0. Supported range is 0-3.
5:0	0h RO	tRP: Holds DDR timing parameter tRP (and tRCD). PRE to ACT same bank minimum delay in DCLK cycles. ACT to CAS (RD or WR) same bank minimum delay in DCLK cycles. For LPDDR3 this field should hold tRPpb (and tRCD) values. Supported range is 8-63.



7.14 MCHBAR_CHO_CR_SC_GS_CFG_0_0_0_MCHBAR— Offset 441Ch

Scheduler configuration

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 441Ch

Default: 0h

3 1	2 8	2 4	2 0	1 6	1 2	8	4	0
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0 0 0	0 0 0 0	0 0 0 0
tCAL	ddr_probeless_low_frequency enable_odt_matrix	ck_to_cke	cmd_3st reset_delay	reset_on_command	LPDDR_2N_CS_MRW tCPDED	x8_device Address_mirror	RSVD N_to_1_ratio	CMD_stretch DRAM_technology

Bit Range	Default and Access	Field Name (ID): Description
31:29	0h RO	tCAL: For DDR4, holds tCAL value. Supported values: 0 (CAL mode disabled), 3-5 (CAL mode enabled, value is the delay in DCLK cycles from CSb to command). Updating this field is required only after sending MRS to MR4 enabling/disabling CAL mode before any other command is sent to DRAM. TC_MR4_shaddow_0_0_0_MCHBAR should be updated with the correct value of tCAL once its value changes.
28	0h RO	ddr_probeless_low_frequency: This bit controls whether the DDR probeless logic uses DDR_TX_DELAY_LOW or DDR_TX_DELAY_HIGH for the internal delay of the write data. If MRC supports two training frequencies, this bit should be set when training at the low frequency.
27	0h RO	enable_odt_matrix: When bit is set, the ranks that are used for terminating when giving read/write requests are selected according to SC_ODT_MATRIX control register and not according to the default behavior.
26:24	0h RO	ck_to_cke: When working with LPDDR when CKE is low we also turn off the CKE buffers. The LPDDR specification requires starting the CK toggling two DCLK cycles before re-asserting CKE. The field defines the number of DCLK cycles from CKoutputEnable assert on power down exit to CKE assert as the DDRIO can delay the CK pins differently than CKE so a different value is required to get two DCLK cycles of CK toggling before CKE rise. Typically this field should be programmed to 3 if (CLK_pi+CLK_logicdelay)-(CKE_pi+CKE_logicdelay) is less than 1 QCLK. Otherwise, it should be programmed to 4 supported range is 2-7.
23	0h RO	cmd_3st: Defines when command and address bus is driving. 0: Drive when channel is active. Tri-stated when all ranks are in CKE-off or when memory is in SR or deeper. 1: Command bus is always driving. When no new valid command is driven, previous command and address is driven



Bit Range	Default and Access	Field Name (ID): Description
22:20	0h RO	reset_delay: Inserts an N Dclk delay ranging from 0 to 7 after the N to 1 Reset on Cmd is triggered.
19:16	0h RO	reset_on_command: The N:1 logic can be triggered to insert a bubble and reset the N:1 logic after a programmable delay from a command after a PRE/ACT/RD/WR CMD. This allows one to synchronize the N:1 logic periodically to ensure the correct worst case pattern between victim and aggressor occurs when training the command bus. Reset N to 1 Logic on a WR (bit 16) Reset N to 1 Logic on a RD (bit 17) Reset N to 1 Logic on a ACT (bit 18) Reset N to 1 Logic on a PRE (bit 19)
15	0h RO	LPDDR_2N_CS_MRW: When sending an MRW command via the MRH for LPDDR drive the CSb for two DCLK cycles
14:12	0h RO	tCPDED: Holds DDR timing parameter tCPDED. Power down to command bus tri-state delay in DCLK cycles. Supported range is 1-7 in 1N mode.
11:10	0h RO	x8_device: DIMM is made out of X8 devices LSB is for DIMM 0, MSB is for DIMM 1.
9:8	0h RO	Address_mirror: DIMM routing causes address mirroring LSB is for DIMM 0, MSB is for DIMM 1.
6:4	0h RO	N_to_1_ratio: When using N:1 command stretch mode, every how many B2B valid command cycles a bubble is required Supported range is 1 to 7
3:2	0h RO	CMD_stretch: Command stretch mode: 00: 1N 01: 2N 10: 3N 11: N:1
1:0	0h RO	DRAM_technology: DRAM technology: 00: DDR4 01: DDR3 10: LPDDR3 11: Illegal

7.15 ECCERRLOG0—Offset 4448h

This register logs ECC error information. Represents Channel 1.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4448h

Default: 0h

3	2	2	2	1	1	8	4	0
1	8	4	0	6	2			
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
ERRBANK	ERRRANK	ERRCHUNK	ERRSYND			RSVD		MERRSTS CERRSTS



Bit Range	Default and Access	Field Name (ID): Description
31:29	0h ROSV	ERRBANK: This field holds the Bank Address of the read transaction that had the ECC error.
28:27	0h ROSV	ERRRANK: This field holds the Rank ID of the read transaction that had the ECC error.
26:24	0h ROSV	ERRCHUNK: Holds the chunk number of the error stored in the register.
23:16	0h ROSV	ERRSYND: This field contains the error syndrome. A value of 0xFF indicates that the error is due to poisoning.
15:2	0h RO	Reserved (RSVD): Reserved.
1	0h ROSV	MERRSTS: This bit is set when an uncorrectable multiple-bit error occurs on a memory read data transfer. When this bit is set, the address that caused the error and the error syndrome are also logged and they are locked until this bit is cleared. This bit is cleared when the corresponding bit in 0.0.0.PCI.ERRSTS is cleared.
0	0h ROSV	CERRSTS: This bit is set when a correctable single-bit error occurs on a memory read data transfer. When this bit is set, the address that caused the error and the error syndrome are also logged and they are locked to further single bit errors, until this bit is cleared. A multiple bit error that occurs after this bit is set will override the address/error syndrome information. This bit is cleared when the corresponding bit in 0.0.0.PCI.ERRSTS is cleared.

7.16 ECC Error Log 1 (ECCERRLOG1)—Offset 444Ch


This register logs ECC error information.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 444Ch

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31:30	0h RO	Reserved (RSVD): Reserved.
29:28	0h ROSV	ERRBANKGROUP: This field holds the DRAM bank group (for DDR4) address of the read transaction that had the ECC error.
27:17	0h ROSV	ERRCOL: This field holds the DRAM column address of the read transaction that had the ECC error.
16:0	0h ROSV	ERRROW: This field holds the DRAM row (page) address of the read transaction that had the ECC error.

7.17 MCHBAR_CHO_CR_TC_ODT_0_0_0_MCHBAR— Offset 4470h

ODT timing related parameters

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4470h

Default: 0h

3	2	2	2	1	1	8	4	0			
1	8	4	0	6	2						
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0 0 0	0 0 0 0	0 0 0 0			
ODT_Always_Rank0	tAONPD	tCWL	tCL	Write_Early_ODT	ODT_Write_Delay	RSVD	ODT_write_duration	RSVD	ODT_Read_Delay	RSVD	ODT_read_duration

Bit Range	Default and Access	Field Name (ID): Description
31	0h RO	ODT_Always_Rank0: Indicate that ODT should always be multiplexed out on ODT[0], to be used for LPDDR3 only.
30:26	0h RO	tAONPD: Holds DDR timing parameter tAONPD. Supported range is 4-31.
25:21	0h RO	tCWL: Holds DDR timing parameter tCWL (sometimes referred to as tWCL). Write command to data delay in DCLK cycles. Supported range is 4-20 (maximum is for 1N mode and tCAL=0) For LPDDR3 the minimum supported value is 4 if Dec_WRD=0 and 5 if Dec_WRD=1. For DDR3/4, the minimum supported value is 5 if Dec_WRD=0 and 6 if Dec_WRD=1.
20:16	0h RO	tCL: Holds DDR timing parameter tCL. Read command to data delay in DCLK cycles. Supported range is 5-31.



Bit Range	Default and Access	Field Name (ID): Description
15	0h RO	Write_Early_ODT: When this bit is set, the MC is will send one extra cycle of ODT prior to the write command. In this mode the ranks that will be terminated on this early cycle are selected according to the SC_ODT_MATRIX_0_0_0_MCHBAR control register.
14:12	0h RO	ODT_Write_Delay: Controls delay from WR-CAS to ODT assertion in DCLK cycles (Typical Programming = 0).
10:8	0h RO	ODT_write_duration: Controls the length of the ODT pulse for write commands. Default is 6 DCLK cycles (BL/2 + 2) 000: 6 DCLK cycles 001: 7 DCLK cycles 010: 8 DCLK cycles 011: 9 DCLK cycles 100: 10 DCLK cycles 101: 11 DCLK cycles 110: 12 DCLK cycles 111: 13 DCLK cycles
6:4	0h RO	ODT_Read_Delay: Controls delay from RD-CAS to ODT assertion in DCLK cycles (Typical Programming = tCL-tCWL). Note 1: All RD->RD and RD->WR restrictions should be greater than or equal to this field value. Note 2: odt_read_delay + odt_read_duration should not be programmed to less than tCL-
2:0	0h RO	ODT_read_duration: Controls the length of the ODT pulse for read commands. Default is 6 DCLK cycles (BL/2 + 2) 000: 6 DCLK cycles 001: 7 DCLK cycles 010: 8 DCLK cycles 011: 9 DCLK cycles 100: 10 DCLK cycles 101: 11 DCLK cycles 110: 12 DCLK cycles 111: 13 DCLK cycles

7.18 Refresh Parameters (TC)—Offset 4638h

Refresh parameters

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4638h

Default: 4600980Fh



Bit Range	Default and Access	Field Name (ID): Description
31:26	0h RO	Reserved (RSVD): Reserved.
25:16	B4h RW_L	tRFC: Time of refresh - from beginning of refresh until next ACT or refresh is allowed (in DCLK cycles, default is 180)
15:0	1004h RW_L	tREFI: defines the average period between refreshes, and the rate that tREFI counter is incremented (in DCLK cycles, default is 4100)

7.20 Power Management DIMM Idle Energy (PM)—Offset 4660h

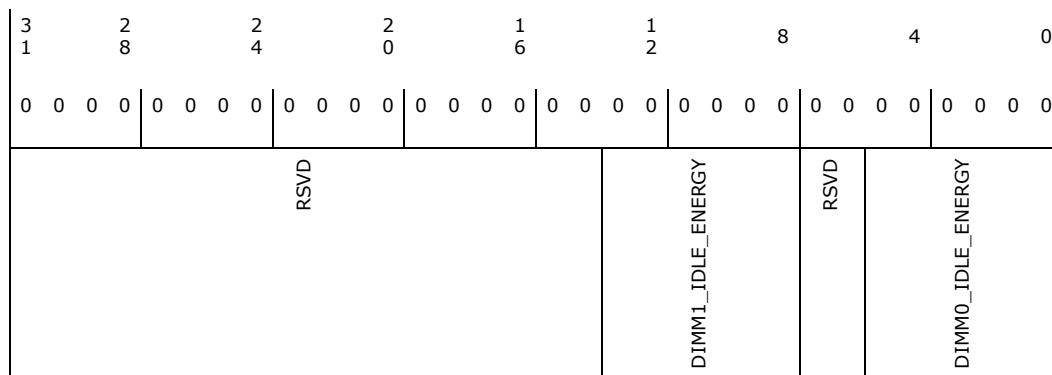
This register defines the energy of an idle DIMM with CKE on. Each 6-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 6-bit fields, one per DIMM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4660h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:14	0h RO	Reserved (RSVD): Reserved.
13:8	0h RW_L	DIMM1_IDLE_ENERGY: This register defines the energy consumed by DIMM1 for one clock cycle when the DIMM is idle with cke on.
7:6	0h RO	Reserved (RSVD): Reserved.
5:0	0h RW_L	DIMM0_IDLE_ENERGY: This register defines the energy consumed by DIMM0 for one clock cycle when the DIMM is idle with cke on.



7.21 Power Management DIMM Power Down Energy (PM)—Offset 4664h


This register defines the energy of an idle DIMM with CKE off. Each 6-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 6-bit fields, one per DIMM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4664h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:14	0h RO	Reserved (RSVD): Reserved.
13:8	0h RW_L	DIMM1_PD_ENERGY: This register defines the energy consumed by DIMM1 for one clock cycle when the DIMM is idle with cke off.
7:6	0h RO	Reserved (RSVD): Reserved.
5:0	0h RW_L	DIMM0_PD_ENERGY: This register defines the energy consumed by DIMM0 for one clock cycle when the DIMM is idle with cke off.



7.22 Power Management DIMM Activate Energy (PM)—Offset 4668h


This register defines the combined energy contribution of activate and precharge commands. Each 8-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 8-bit fields, one per DIMM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4668h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	0h RW_L	DIMM1_ACT_ENERGY: This register defines the combined energy contribution of activate and precharge commands.
7:0	0h RW_L	DIMM0_ACT_ENERGY: This register defines the combined energy contribution of activate and precharge commands.



7.23 Power Management DIMM RdCas Energy (PM)—Offset 466Ch


This register defines the energy contribution of a read CAS command. Each 8-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 8-bit fields, one per DIMM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 466Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	0h RW_L	DIMM1_RD_ENERGY: This register defines the energy contribution of a read CAS command.
7:0	0h RW_L	DIMM0_RD_ENERGY: This register defines the energy contribution of a read CAS command.



7.24 Power Management DIMM WrCas Energy (PM)—Offset 4670h


This register defines the energy contribution of a write CAS command. Each 8-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 8-bit fields, one per DIMM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4670h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	0h RW_L	DIMM1_WR_ENERGY: This register defines the energy contribution of a write CAS command.
7:0	0h RW_L	DIMM0_WR_ENERGY: This register defines the energy contribution of a write CAS command.



7.25 MCSCHEDS_CR_SC_GS_CFG_0_0_0_MCHBAR— Offset 4C1Ch


Scheduler configuration.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4C1Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:4	0h RO	Reserved (RSVD): Reserved.
3:2	0h RW_L	CMD_strech: Command stretch mode: 00: 1N 01: 2N 10: 3N 11: N:1
1:0	0h RO	Reserved (RSVD): Reserved.



7.26 PM—Offset 4C40h

This register defines the power-down (CKE-off) operation - power-down mode, and idle timers associated with power down entry

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4C40h

Default: 0h

3	2	2	2	1	1	8	4	0
1	8	4	0	6	2			
0	0	0	0	0	0	0	0	0
	RSVD	dis_cke_tt	TT_idle_counter	Global_PD	Slow_Exit	PPD	APD	PDWN_idle_counter

Bit Range	Default and Access	Field Name (ID): Description
31:25	0h RO	Reserved (RSVD): Reserved.
24	0h RW_L	dis_cke_tt: 1'b0: CKE TT is enabled. When throttling is asserted TT_idle_counter value is loaded into the CKE counter. The CKE FSM will be forced to the countdown state upon activation of throttling. When a rank becomes non-isoch-empty the CKE FSM will turn on if currently off. 1'b1: CKE TT is defeatured.
23:16	0h RW_L	TT_idle_counter: Amount of cycles to wait before going to PD when thermal throttling is enabled
15	0h RW_L	Global_PD: Power down entry and exit conditions are determined globally for the whole channel and not on a per rank basis
14	0h RW_L	Slow_Exit: Indicate if DDR (applicable only for DDR3/DDR4) that it is in slow exit mode so when exiting PPD the MC should wait tXPDLL before sending a CAS command and not tXP
13	0h RW_L	PPD: When rank is idle close all pages and go to PPD. If both APD and PPD are set and not all banks are closed when idle first go to APD then once all page idle timers expire go out of APD, issue a PREALL and then power down to PPD.
12	0h RW_L	APD: Put rank in APD when idle. This field is controller by hardware unless DDR_PTM_CTL_0_0_0_MCHBAR_PCU.PDWN_CONFIG_CTL is set
11:0	0h RW_L	PDWN_idle_counter: This defines the rank idle period in DCLK cycles that causes power-down entrance. This field is controller by hardware unless DDR_PTM_CTL_0_0_0_MCHBAR_PCU.PDWN_CONFIG_CTL is set



7.27 ECCERRLOG0—Offset 4C48h

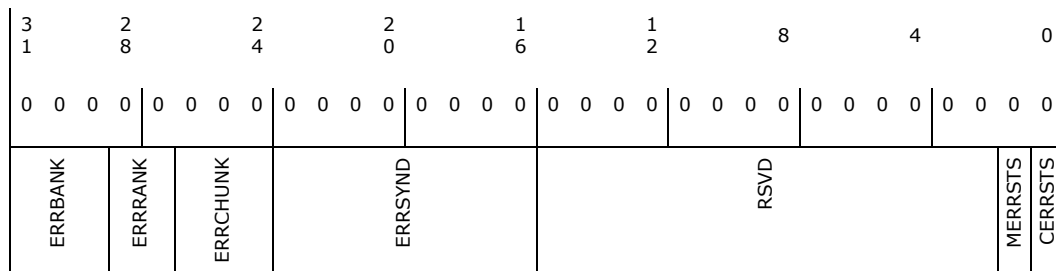
This register logs ECC error information. Represents both Channel 0 and Channel 1.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4C48h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:29	0h ROSV	ERRBANK: This field holds the Bank Address of the read transaction that had the ECC error.
28:27	0h ROSV	ERRRANK: This field holds the Rank ID of the read transaction that had the ECC error.
26:24	0h ROSV	ERRCHUNK: Holds the chunk number of the error stored in the register.
23:16	0h ROSV	ERRSYND: This field contains the error syndrome. A value of 0xFF indicates that the error is due to poisoning.
15:2	0h RO	Reserved (RSVD): Reserved.
1	0h ROSV	MERRSTS: This bit is set when an uncorrectable multiple-bit error occurs on a memory read data transfer. When this bit is set, the address that caused the error and the error syndrome are also logged and they are locked until this bit is cleared. This bit is cleared when the corresponding bit in 0.0.0.PCI.ERRSTS is cleared.
0	0h ROSV	CERRSTS: This bit is set when a correctable single-bit error occurs on a memory read data transfer. When this bit is set, the address that caused the error and the error syndrome are also logged and they are locked to further single bit errors, until this bit is cleared. A multiple bit error that occurs after this bit is set will override the address/error syndrome information. This bit is cleared when the corresponding bit in 0.0.0.PCI.ERRSTS is cleared.



7.28 ECC Error Log 1 (ECCERRLOG1)—Offset 4C4Ch


This register logs ECC error information.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4C4Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:30	0h RO	Reserved (RSVD): Reserved.
29:28	0h ROSV	ERRBANKGROUP: This field holds the DRAM bank group (for DDR4) address of the read transaction that had the ECC error.
27:17	0h ROSV	ERRCOL: This field holds the DRAM column address of the read transaction that had the ECC error.
16:0	0h ROSV	ERRORW: This field holds the DRAM row (page) address of the read transaction that had the ECC error.



7.29 MCSCHEDS_CR_TC_ODT_0_0_0_MCHBAR—Offset 4C70h


ODT timing related parameters.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4C70h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:26	0h RO	Reserved (RSVD): Reserved.
25:21	6h RW_L	tCWL: Holds DDR timing parameter tCWL (sometimes referred to as tWCL). Write command to data delay in DCLK cycles Supported range is 4-20 (maximum is for 1N mode and tCAL=0) For LPDDR3 the minimum supported value is 4 if Dec_WRD=0 5 and if Dec_WRD=1 For DDR3/4 the minimum supported value is 5 if Dec_WRD=0 6 and if Dec_WRD=1
20:16	5h RW_L	tCL: Holds DDR timing parameter tCL. Read command to data delay in DCLK cycles. Supported range is 5-31.
15:0	0h RO	Reserved (RSVD): Reserved.



7.30 Refresh Parameters (TC)—Offset 4E38h


Refresh parameters

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4E38h

Default: 4600980Fh



Bit Range	Default and Access	Field Name (ID): Description
31:25	23h RW_L	tREFIX9: Maximum time allowed between refreshes to a rank (in intervals of 1024 DCLK cycles). Should be programmed to 8.9*tREFI/1024 (to allow for possible delays from ZQ or isoc).
24:16	0h RO	Reserved (RSVD): Reserved.
15:12	9h RW_L	Refresh_panic_wm: tREFI count level in which the refresh priority is panic (default is 9). The Maximum value for this field is 9.
11:8	8h RW_L	Refresh_HP_WM: tREFI count level that turns the refresh priority to high (default is 8)
7:0	Fh RW_L	OREF_RI: Rank idle period that defines an opportunity for refresh, in DCLK cycles



7.31 Refresh Timing Parameters (TC)—Offset 4E3Ch


Refresh timing parameters

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4E3Ch

Default: B41004h



Bit Range	Default and Access	Field Name (ID): Description
31:26	0h RO	Reserved (RSVD): Reserved.
25:16	B4h RW_L	tRFC: Time of refresh - from beginning of refresh until next ACT or refresh is allowed (in DCLK cycles, default is 180)
15:0	1004h RW_L	tREFI: defines the average period between refreshes, and the rate that tREFI counter is incremented (in DCLK cycles, default is 4100)

7.32 Power Management DIMM Idle Energy (PM)—Offset 4E60h


This register defines the energy of an idle DIMM with CKE on. Each 6-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 6-bit fields, one per DIMM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4E60h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:14	0h RO	Reserved (RSVD): Reserved.
13:8	0h RW_L	DIMM1_IDLE_ENERGY: This register defines the energy consumed by DIMM1 for one clock cycle when the DIMM is idle with cke on
7:6	0h RO	Reserved (RSVD): Reserved.
5:0	0h RW_L	DIMM0_IDLE_ENERGY: This register defines the energy consumed by DIMM0 for one clock cycle when the DIMM is idle with cke on.

7.33 Power Management DIMM Power Down Energy (PM)—Offset 4E64h

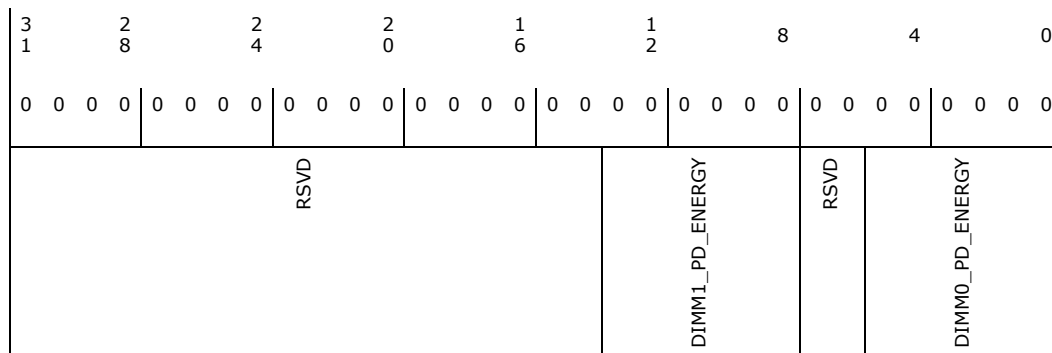
This register defines the energy of an idle DIMM with CKE off. Each 6-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 6-bit fields, one per DIMM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4E64h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31:14	0h RO	Reserved (RSVD): Reserved.
13:8	0h RW_L	DIMM1_PD_ENERGY: This register defines the energy consumed by DIMM1 for one clock cycle when the DIMM is idle with cke off
7:6	0h RO	Reserved (RSVD): Reserved.
5:0	0h RW_L	DIMM0_PD_ENERGY: This register defines the energy consumed by DIMM0 for one clock cycle when the DIMM is idle with cke off

7.34 Power Management DIMM Activate Energy (PM)—Offset 4E68h


This register defines the combined energy contribution of activate and precharge commands. Each 8-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 8-bit fields, one per DIMM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4E68h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	0h RW_L	DIMM1_ACT_ENERGY: This register defines the combined energy contribution of activate and precharge commands.
7:0	0h RW_L	DIMM0_ACT_ENERGY: This register defines the combined energy contribution of activate and precharge commands.



7.35 Power Management DIMM RdCas Energy (PM) – Offset 4E6Ch


This register defines the energy contribution of a read CAS command. Each 8-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 8-bit fields, one per DIMM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4E6Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	0h RW_L	DIMM1_RD_ENERGY: This register defines the energy contribution of a read CAS command.
7:0	0h RW_L	DIMM0_RD_ENERGY: This register defines the energy contribution of a read CAS command.



7.36 Power Management DIMM WrCas Energy (PM)— Offset 4E70h


This register defines the energy contribution of a write CAS command. Each 8-bit field corresponds to an integer multiple of the base DRAM command energy for that DIMM. There are 2 8-bit fields, one per DIMM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 4E70h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	0h RW_L	DIMM1_WR_ENERGY: This register defines the energy contribution of a write CAS command.
7:0	0h RW_L	DIMM0_WR_ENERGY: This register defines the energy contribution of a write CAS command.



7.37 Address Decoder Inter Channel Configuration Register (MAD)—Offset 5000h


This register holds parameters used by the channel decode stage. It defines virtual channel L mapping, as well as channel S size. Also defined is the DDR type installed in the system (DDR4 or DDR3).

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5000h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:19	0h RO	Reserved (RSVD): Reserved.
18:12	0h RW_L	CH_S_SIZE: Channel S size in multiplies of 1GB (min. rank size in the processor). Needed for channel decode stage. Supports range of 0GB - 64GB.
11:5	0h RO	Reserved (RSVD): Reserved.
4	0h RW_L	CH_L_MAP: Channel L mapping to physical channel. 0: Channel0 1: Channel1
3:2	0h RO	Reserved (RSVD): Reserved.
1:0	0h RW_L	DDR_TYPE: DDR_TYPE - defines the DDR type in system: 00: DDR4 01: DDR3 10: LPDDR3



7.38 Address Decoder Intra Channel Configuration Register (MAD)—Offset 5004h

This register holds parameters used by the DRAM decode stage.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5004h

Default: 0h

3	2	2	2	1	1	8	4	0			
1	8	4	0	6	2						
0	0	0	0	0	0	0	0	0			
0	0	0	0	0	0	0	0	0			
0	0	0	0	0	0	0	0	0			
0	0	0	0	0	0	0	0	0			
RSVD	HORI_ADDR	RSVD	HORI	RSVD	ECC	RSVD	EIM	RSVD	RI	RSVD	DIMM_L_MAP

Bit Range	Default and Access	Field Name (ID): Description
31	0h RO	Reserved (RSVD): Reserved.
30:28	0h RW_L	HORI_ADDR: High Order Rank Interleave Address. Specifies which address bit 20-27 to use as the rank interleave bit 000: bit 20 001: bit 21 ... 111: bit 27
27:25	0h RO	Reserved (RSVD): Reserved.
24	0h RW_L	HORI: High order rank interleaving enable bit 0: Disabled 1: Enabled High Order Rank Interleave (HORI) is mutually exclusive with Rank Interleave (RI)
23:14	0h RO	Reserved (RSVD): Reserved.
13:12	0h RW_L	Reserved (RSVD): Reserved.
13:12	0h RW_L	ECC: ECC configuration in the channel: 00: No ECC active in the channel. 11: ECC active in both IO and ECC logic. Note: This field should be programmed identically for all populated channels.
11:9	0h RO	Reserved (RSVD): Reserved.
8	0h RW_L	EIM: Enhanced mode enable bit 0: Disabled 1: Enabled



Bit Range	Default and Access	Field Name (ID): Description
7:5	0h RO	Reserved (RSVD): Reserved.
4	0h RW_L	RI: Rank interleaving enable bit 0: Disabled 1: Enabled
3:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW_L	DIMM_L_MAP: Virtual DIMM L mapping to physical DIMM 0: DIMM0 1: DIMM1

7.39 Address Decoder Intra Channel Configuration Register (MAD)—Offset 5008h

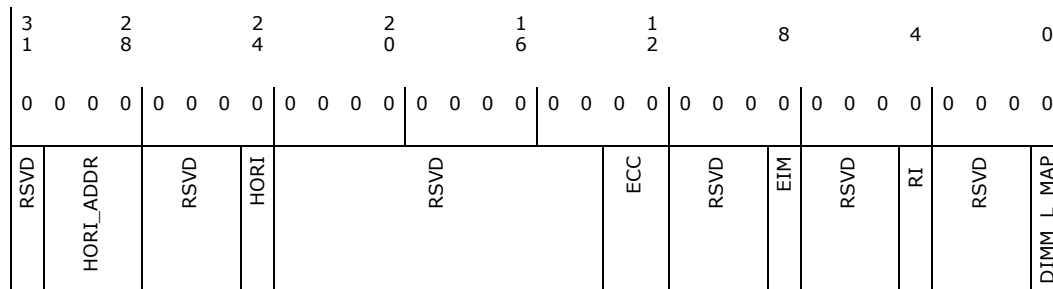
This register holds parameters used by the DRAM decode stage.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5008h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31	0h RO	Reserved (RSVD): Reserved.
30:28	0h RW_L	HORI_ADDR: High Order Rank Interleave Address. Specifies which address bit 20-27 to use as the rank interleave bit 000: bit 20 001: bit 21 ... 111: - bit 27
27:25	0h RO	Reserved (RSVD): Reserved.



Bit Range	Default and Access	Field Name (ID): Description
24	0h RW_L	HORI: High order rank interleaving enable bit 0: Disabled 1: Enabled High Order Rank Interleave (HORI) is mutually exclusive with Rank Interleave (RI)
23:14	0h RO	Reserved (RSVD): Reserved.
13:12	0h RW_L	Reserved (RSVD): Reserved.
13:12	0h RW_L	ECC: ECC configuration in the channel: 00: No ECC active in the channel. 11: ECC active in both IO and ECC logic. Note: This field should be programmed identically for all populated channels.
11:9	0h RO	Reserved (RSVD): Reserved.
8	0h RW_L	EIM: Enhanced mode enable bit 0: Disabled 1: Enabled
7:5	0h RO	Reserved (RSVD): Reserved.
4	0h RW_L	RI: Rank interleaving enable bit 0: Disabled 1: Enabled
3:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW_L	DIMM_L_MAP: Virtual DIMM L mapping to physical DIMM 0: DIMM0 1: DIMM1

7.40 Address Decode DIMM Parameters. (MAD)—Offset 500Ch


This register defines channel DIMM characteristics - number of DIMMs, number of ranks, size and type.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 500Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:28	0h RO	Reserved (RSVD): Reserved.
27	0h RW_L	DS8Gb: Defines whether DIMM S is built from 8Gb DRAM modules. 0: Not 8Gb 1: 8Gb
26	0h RW_L	DSNOR: DIMM S number of ranks 0: 1 Rank 1: 2 Ranks
25:24	0h RW_L	DSW: DSW: DIMM S width of DDR chips 00: X8 chips 01: X16 chips 10: X32 chips 11: Reserved
23:22	0h RO	Reserved (RSVD): Reserved.
21:16	0h RW_L	DIMM_S_SIZE: Size of DIMM S in 1GB multiples
15:12	0h RO	Reserved (RSVD): Reserved.
11	0h RW_L	DL8Gb: Defines for DDR3 whether DIMM L is built from 8Gb DRAM modules. 0: Not 8Gb 1: 8Gb For non DDR3, this field should be set to 0.
10	0h RW_L	DLNOR: DIMM L number of ranks 0: 1 Rank 1: 2 Ranks
9:8	0h RW_L	DLW: DLW: DIMM L width of DDR chips 00: X8 chips 01: X16 chips 10: X32 chips 11: Reserved
7:6	0h RO	Reserved (RSVD): Reserved.
5:0	0h RW_L	DIMM_L_SIZE: Size of DIMM L in 1GB multiples



7.41 Address Decode DIMM Parameters (MAD)—Offset 5010h

This register defines channel DIMM characteristics - number of DIMMs, number of ranks, size and type.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5010h

Default: 0h

3	2	2	2	1	1	8	4	0			
1	8	4	0	6	2						
0	0	0	0	0	0	0	0	0			
0	0	0	0	0	0	0	0	0			
RSVD	DS8Gb	DSNOR	DSW	RSVD	DIMM_S_SIZE	RSVD	DL8Gb	DLNOR	DLW	RSVD	DIMM_L_SIZE

Bit Range	Default and Access	Field Name (ID): Description
31:28	0h RO	Reserved (RSVD): Reserved.
27	0h RW_L	DS8Gb: Defines whether DIMM S is built from 8Gb DRAM modules. 0: Not 8Gb 1: 8Gb
26	0h RW_L	DSNOR: DIMM S number of ranks 0: 1 Rank 1: 2 Ranks
25:24	0h RW_L	DSW: DSW: DIMM S width of DDR chips 00: X8 chips 01: X16 chips 10: X32 chips 11: Reserved
23:22	0h RO	Reserved (RSVD): Reserved.
21:16	0h RW_L	DIMM_S_SIZE: Size of DIMM S in 1GB multiples
15:12	0h RO	Reserved (RSVD): Reserved.
11	0h RW_L	DL8Gb: Defines whether DIMM L is built from 8Gb DRAM modules. 0: Not 8Gb 1: 8Gb
10	0h RW_L	DLNOR: DIMM L number of ranks 0: 1 Rank 1: 2 Ranks



Bit Range	Default and Access	Field Name (ID): Description
9:8	0h RW_L	DLW: DLW: DIMM L width of DDR chips 00: X8 chips 01: X16 chips 10: X32 chips 11: Reserved
7:6	0h RO	Reserved (RSVD): Reserved.
5:0	0h RW_L	DIMM_L_SIZE: Size of DIMM L in 1GB multiples

7.42 MCDECS_CR_MRC_REVISION_0_0_0_MCHBAR_MCMA IN—Offset 5034h


Scheduler configuration.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5034h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:0	0h RW_L	REVISION: BIOS MRC Revision. 7:0 = Build # 15:8 = Revision 23:16 = Minor 31:24 = Major

7.43 Request Count from GT (DRAM)—Offset 5040h

Counts every read/write request entering the Memory Controller to DRAM (sum of all channels) from the GT engine. Each partial write request counts as a request incrementing this counter. However same-cache-line partial write requests are




combined to a single 64-byte data transfers from DRAM. Therefore multiplying the number of requests by 64-bytes will lead to inaccurate GT memory bandwidth. The inaccuracy is proportional to the number of same-cache-line partial writes combined.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5040h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:0	0h RW_LV	count: Number of accesses

7.44 Request Count from IA (DRAM)—Offset 5044h


Counts every read/write request (demand and HW prefetch) entering the Memory Controller to DRAM (sum of all channels) from IA. Each partial write request counts as a request incrementing this counter. However same-cache-line partial write requests are combined to a single 64-byte data transfers from DRAM. Therefore multiplying the number of requests by 64-bytes will lead to inaccurate IA memory bandwidth. The inaccuracy is proportional to the number of same-cache-line partial writes combined.


Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5044h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31:0	0h RW_LV	count: Number of accesses

7.45 Request Count from IO (DRAM)—Offset 5048h


Counts every read/write request entering the Memory Controller to DRAM (sum of all channels) from all IO sources (e.g. PCIe, Display Engine, USB audio, etc.). Each partial write request counts as a request incrementing this counter. However same-cache-line partial write requests are combined to a single 64-byte data transfers from DRAM. Therefore multiplying the number of requests by 64-bytes will lead to inaccurate IO memory bandwidth. The inaccuracy is proportional to the number of same-cache-line partial writes combined.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5048h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:0	0h RW_LV	count: Number of accesses

7.46 RD Data Count (DRAM)—Offset 5050h


Counts every read (RdCAS) issued by the Memory Controller to DRAM (sum of all channels). All requests result in 64-byte data transfers from DRAM. Use for accurate memory bandwidth calculations.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5050h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:0	0h RW_LV	count: Number of accesses

7.47 WR Data Count (DRAM)—Offset 5054h


Counts every write (WrCAS) issued by the Memory Controller to DRAM (sum of all channels). All requests result in 64-byte data transfers from DRAM. Use for accurate memory bandwidth calculations.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5054h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:0	0h RW_LV	count: Number of accesses



7.48 Self Refresh Configuration Register (PM)—Offset 5060h

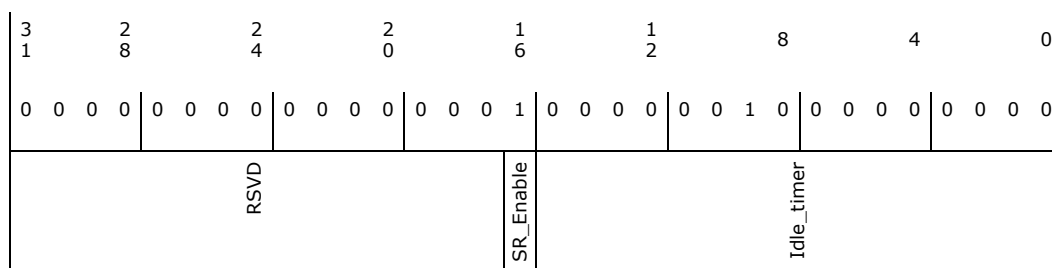
Self refresh mode control register - defines if and when DDR can go into SR

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5060h

Default: 10200h



Bit Range	Default and Access	Field Name (ID): Description
31:17	0h RO	Reserved (RSVD): Reserved.
16	1h RW_LV	SR_Enable: enables or disables self-refresh mechanism. In order to allow SR, both SREF_en bit should be set and SREF_exit signal should be cleared. PM_SREF_config may be updated in run-time
15:0	200h RW_LV	Idle_timer: This value is used when the SREF_enable field is set. It defines the number of cycles that there should not be any transaction in order to enter self-refresh. Supported range is 512 to 64K-1

7.49 Address Compare for ECC Error Injection (ECC)—Offset 5090h


Address compare for ECC error inject. Error injection is issued when
ECC_Inj_Addr_Compare[31:0] = ADDR[37:6] AND ECC_Inj_Addr_Mask[31:0]

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5090h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:0	0h RW_L	Address: Inject error when ECC_Inj_Addr_Compare[31:0] = ADDR[37:6]

7.50 Address Mask for ECC Error Injection (ECC)— Offset 5094h


Address compare for ECC error inject. Error injection is issued when
 $ECC_Inj_Addr_Compare[31:0] = ADDR[37:6] \text{ AND } ECC_Inj_Addr_Mask[31:0]$

Access Method

Type: MEM
 (Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5094h

Default: FFFFFFFFh



Bit Range	Default and Access	Field Name (ID): Description
31:0	FFFFFFFh RW_L	Mask: Inject error when $ECC_Inj_Addr_Compare[31:0] = ADDR[37:6] \text{ AND } ECC_Inj_Addr_Mask[31:0]$



7.51 NCDECS_CR_GFXVTBAR_0_0_0_MCHBAR_NCU—Offset 5400h

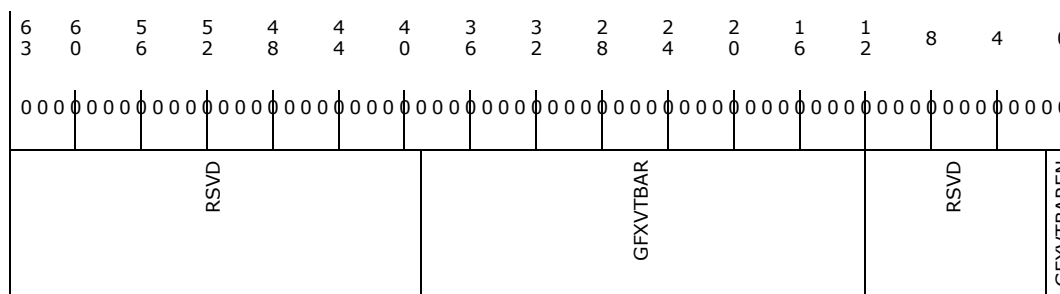
This is the base address for the Graphics VT configuration space. There is no physical memory within this 4KB window that can be addressed. The 4KB reserved by this register does not alias to any PCI 2.3 compliant memory mapped space. On reset, the GFX-VT configuration space is disabled and should be enabled by writing a 1 to GFX-VTBAREN.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 5400h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
38:12	0h RO	GFXVTBAR: This field corresponds to bits 38 to 12 of the base address GFX-VT configuration space. BIOS will program this register resulting in a base address for a 4KB block of contiguous memory address space. This register ensures that a naturally aligned 4KB space is allocated within the first 512GB of addressable memory space. System Software uses this base address to program the GFX-VT register set. All the Bits in this register are locked in LT mode.
0	0h RO	GFXVTBAREN: GFX-VTBAR is disabled and does not claim any memory 1: GFX-VTBAR memory mapped accesses are claimed and decoded appropriately This bit will remain 0 if VTd capability is disabled.

7.52 NCDECS_CR_VTDPVC0BAR_0_0_0_MCHBAR_NCU—Offset 5410h

This is the base address for the DMI/PEG VC0 configuration space. There is no physical memory within this 4KB window that can be addressed. The 4KB reserved by this register does not alias to any PCI 2.3 compliant memory mapped space. On reset, the DMI/PEG VC0 configuration space is disabled and should be enabled by writing a 1 to VC0BAREN.


Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 5410h



Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
38:12	0h RO	VTVCOBAR: This field corresponds to bits 38 to 12 of the base address DMI/PEG VC0 configuration space. BIOS will program this register resulting in a base address for a 4KB block of contiguous memory address space. This register ensures that a naturally aligned 4KB space is allocated within the first 512GB of addressable memory space. System Software uses this base address to program the DMI/PEG VC0 register set. All the Bits in this register are locked in LT mode.
0	0h RO	VTVCOBAREN: VC0BAR is disabled and does not claim any memory 1: VC0BAR memory mapped accesses are claimed and decoded appropriately. This bit will remain 0 if VTd capability is disabled.

7.53 PACKAGE—Offset 5820h

Thermal Limitation Interrupt Control. Hardware will read this information before generating a thermal interrupt.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5820h

Default: 0h



3	2	2	2	1	1	8	4	0
1	8	4	0	6	2			
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
TEMPERATURE_AVERAGING_TIME_WINDOW		POWER_INT_ENABLE	THRESHOLD_2_INT_ENABLE	THRESHOLD_2_REL_TEMP		THRESHOLD_1_INT_ENABLE	THRESHOLD_1_REL_TEMP	
							RSVD	
							OUT_OF_SPEC_INT_ENABLE	
							RSVD	
							PROCHOT_INT_ENABLE	
							LOW_TEMP_INT_ENABLE	
							HIGH_TEMP_INT_ENABLE	

Bit Range	Default and Access	Field Name (ID): Description
31:25	0h RW	TEMPERATURE_AVERAGING_TIME_WINDOW: averaging window for the running exponential average temperature. x = 2 msbs, that is [31:30] y = 5 lsbs, that is [29:25] The timing interval window is Floating Point number given by 1.x * power(2,y). The unit of measurement is defined in PACKAGE_POWER_SKU_UNIT_MSR[TIME_UNIT]. A value of zero means no averaging.
24	0h RW	POWER_INT_ENABLE: When this bit is set, a thermal interrupt will be sent upon throttling due to power limitations.
23	0h RW	THRESHOLD_2_INT_ENABLE: Controls the generation of a thermal interrupt whenever the Thermal Threshold 2 Temperature is crossed.
22:16	0h RW	THRESHOLD_2_REL_TEMP: This value indicates the offset in degrees below TJ Max Temperature that should trigger a Thermal Threshold 2 trip.
15	0h RW	THRESHOLD_1_INT_ENABLE: Controls the generation of a thermal interrupt whenever the Thermal Threshold 1 Temperature is crossed.
14:8	0h RW	THRESHOLD_1_REL_TEMP: This value indicates the offset in degrees below TJ Max Temperature that should trigger a Thermal Threshold 1 trip.
7:5	0h RO	Reserved (RSVD): Reserved.
4	0h RW	OUT_OF_SPEC_INT_ENABLE: Thermal interrupt enable for the Out Of Spec condition which is stored in the Out Of Spec status bit in PACKAGE_THERM_STATUS.
3	0h RO	Reserved (RSVD): Reserved.



Bit Range	Default and Access	Field Name (ID): Description
2	0h RW	PROCHOT_INT_ENABLE: Bidirectional PROCHOT# assertion interrupt enable. If set, a thermal interrupt is delivered on the rising edge of PROCHOT#.
1	0h RW	LOW_TEMP_INT_ENABLE: Enables a thermal interrupt to be generated on the transition from a high-temperature to a low-temperature when set, where 'high temperature' is dictated by the thermal monitor trip temperature minus offset as defined in IA32_TEMPERATURE_TARGET.
0	0h RW	HIGH_TEMP_INT_ENABLE: Enables a thermal interrupt to be generated on the transition from a low-temperature to a high-temperature when set, where 'high temperature' is dictated by the thermal monitor trip temperature minus offset as defined in IA32_TEMPERATURE_TARGET.

7.54 PKG—Offset 5828h


Sum the cycles per number of active cores

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 5828h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:0	0h ROV	DATA: RO: The counter value is incremented as a function of the number of cores that reside in C0 and active. If N cores are simultaneously in C0, then the number of "clock ticks" that are incremented is N. Counter rate is the Max Non-Turbo frequency (same as TSC)

7.55 PKG—Offset 5830h


C0.Any - Sum the cycles of any active cores.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 5830h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:0	0h ROV	DATA: RO, This counter increments whenever one or more IA cores are active and in C0 state. Counter rate is the Max Non-Turbo frequency (same as TSC)

7.56 PKG—Offset 5838h


Sum the cycles of active GT

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 5838h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:0	0h ROV	DATA: RO, This counter increments whenever GT slices or un slices are active and in C0 state. Counter rate is the Max Non-Turbo frequency (same as TSC)

7.57 PKG—Offset 5840h


Sum the cycles of overlap time between any IA cores and GT

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 5840h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:0	0h ROV	DATA: This counter increments whenever GT slices or un slices are active and in C0 state and in overlap with one of the IA cores that is active and in C0 state. Counter rate is the Max Non-Turbo frequency (same as TSC)

7.58 PKG—Offset 5848h


Sum the cycles of any active GT slice.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 5848h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:0	0h ROV	DATA: RO, This counter increments whenever any GT slice is active. Counter rate is in 24MHz

7.59 PKG—Offset 5858h



Sum the cycles of any media GT engine.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 5858h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:0	0h ROV	DATA: RO, This counter increments whenever any GT media engine is active. Counter rate is in 24 MHz

7.60 DDR—Offset 5880h


Mode control bits for DDR power and thermal management features.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5880h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31:9	0h RO	Reserved (RSVD): Reserved.
8	0h RW	DDR4_SKIP_REFRESH_EN: DDR4 DRAM supports temperature controlled refresh and self refresh. The temperature controlled refresh is essentially DRAM controls to skip some refresh issued by the host when temperature is low enough. When this bit is set and MAD_CHNL.DDR4=1, MC will enable DRAM's TC refresh mode aka skip refresh mode. hardware uses MAD_CHNL.DDR4 and PTM_CTL.DDR4_SKIP_REFRESH_EN to determine whether to support DDR thermal interrupt for refresh rate change. BIOS is responsible to configure this bit and is ZERO by default.
7	0h RW	DISABLE_DRAM_TS: When this bit is zero and MAD_CHNL.LPDDR=1, hardware will use DDR MR4 for DIMM thermal status purposes. Otherwise, hardware will ignore MR4 data and use the legacy CLTM/OLTM/EXTTS algorithms for computing DIMM thermal status.
6	0h RW	PDWN_CONFIG_CTL: This bit determined whether BIOS or hardware will control DDR powerdown modes and idle counter (via programming the PM_PDWN_config regs in iMC). When clear, hardware will manage the modes based on either core P-states or IA32_ENERGY_PERFORMANCE_BIAS MSR value (when enabled). When set, BIOS is in control of DDR CKE mode and idle timer value, and hardware algorithm does not run.
5	0h RW_KL	LOCK_PTM_REGS_PCU: When set, several PCU registers related to DDR power/thermal management all become unwritable (writes will be silently ignored). List of registered locked by this bit is: DDR_WARM_THRESHOLD_CH*, DDR_HOT_THRESHOLD_CH*, DDR_WARM_BUDGET_CH*, DDR_HOT_BUDGET_CH*, (note that RAPL regs, such as RAPL_LIMIT, are NOT included as those have separate lock bit). Note that BIOS should complete its writes to all of the locked registers prior to setting this bit, since it can only be reset via uncore reset.
4	0h RW	EXTTS_ENABLE: When clear (default), hardware ignores the EXTTS (external thermal status) indication which is obtained from the PCH (via PM_SYNC). When set, the value from EXTTS is used only when it is hotter than the thermal status reported by OLTM/CLTM algorithm (or used all of the time if neither of those modes is enabled).
3:2	0h RW	REFRESH_2X_MODE: These bits are read by reset hardware and later broadcast (together with the thermal status) into the iMC cregs that control 2x refresh modes. When DRAM is hot, it accumulates bits errors more quickly. The iMC refresh mechanism is how those errors get prevented and corrected (using ECC). Thus in order to maintain an acceptable overall error rate, the refresh rate needs to increase with temperature. This is a very coarse grain mechanism for accomplishing that. A value of 00 means the iMC 2x refresh is disabled. A value of 01 means that the iMC will enable 2x refresh whenever thermal status is WARM or HOT. A value of 10 means the iMC will enable 2x refresh only when HOT. The value 11 is illegal, and will trigger an assertion in the iMC (BIOS should not do this). This field is ignored for LPDDR when DISABLE_DRAM_TS is zero, in which case refresh rates in the MC are controlled by MR4 coming directly from DIMMs.
1	0h RW	CLTM_ENABLE: A value of 1 means CLTM (Closed Loop Thermal Management) hardware algorithm will be used to compute the memory thermal status (which will be written to the iMC). Note that OLTM and CLTM modes are mutex, so if both OLTM_ENABLE and CLTM_ENABLE are set, the OLTM_ENABLE will be ignored and CLTM mode will be active. BIOS should enable CLTM whenever DIMM thermal sensor data is available and memory thermal management is desired.
0	0h RW	OLTM_ENABLE: A value of 1 means OLTM (Open Loop Thermal Management) hardware algorithm will be used to compute the memory thermal status (which will be written to the iMC). Note that OLTM and CLTM modes are mutex, so if both OLTM_ENABLE and CLTM_ENABLE are set, the OLTM_ENABLE will be ignored and CLTM mode will be active. BIOS should enable OLTM in case of thermal sensor data absence, but memory thermal management is desired. Obviously lack of real temperature data means this mode will be somewhat conservative, and may result in the iMC throttling more often than necessary. Thus, for performance reasons, CLTM is preferred on systems with available DIMM thermal sensor data.



7.61 DRAM—Offset 5884h


Defines the base energy unit for DDR energy values in iMC command energy configuration regs, iMC rank energy counters (used for OLTM and Memory RAPL), OLTM thresholds, etc.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5884h

Default: 3h



Bit Range	Default and Access	Field Name (ID): Description
31:3	0h RO	Reserved (RSVD): Reserved.
2:0	3h RW	<p>SCALEFACTOR: Defines the base DDR energy unit of $2^{(-30-scalefactor)}$ Joules. The values are defined as follows:</p> <p>0d0 = 3'b000 = 931.3pJ, 0d1 = 3'b001 = 465.7pJ, 0d2 = 3'b010 = 232.8pJ, 0d3 = 3'b011 = 116.4pJ, 0d4 = 3'b100 = 58.2pJ, 0d5 = 3'b101 = 29.1pJ, 0d6 = 3'b110 = 14.6pJ, 0d7 = 3'b111 = 7.3pJ.</p> <p>The default reset value is 0d3 = 3'b011 = 116.4pJ.</p>

7.62 DRAM—Offset 5888h


Defines the minimum required power consumption of each DDR channel, in order to satisfy minimum memory bandwidth requirements for the platform. DDR RAPL should never throttle below the levels defined here. It is the responsibility of BIOS to comprehend the power consumption on each channel in order to write meaningful values into this register.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5888h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	0h RW	CH1: Minimum power level (in format of 5.3 W) used to clip DDR RAPL power budget for channel 1.
7:0	0h RW	CH0: Minimum power level (in format of 5.3 W) used to clip DDR RAPL power budget for channel 0.

7.63 DDR—Offset 588Ch


Per-DIMM thermal status values. The encoding of each DIMM thermal status is the same: 2'b00 = COLD, 2'b01 = WARM, 2'b11 = HOT, 2'b10 == Reserved.

Access Method


Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 588Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:12	0h RO	Reserved (RSVD): Reserved.
11:10	0h ROS	CH1_DIMM1: Thermal Status for Channel 1, DIMM1
9:8	0h ROS	CH1_DIMM0: Thermal Status for Channel 1, DIMM0



Bit Range	Default and Access	Field Name (ID): Description
7:4	0h RO	Reserved (RSVD): Reserved.
3:2	0h ROS	CHO_DIMM1: Thermal Status for Channel 0, DIMM1
1:0	0h ROS	CHO_DIMM0: Thermal Status for Channel 0, DIMM0

7.64 DDR—Offset 5890h


Per-DIMM temp/power thresholds used for CLTM/OLTM thermal status computation. These values can impact iMC throttling and memory thermal interrupts.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5890h

Default: FFFFh



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	FFh RWS_L	DIMM1: WARM_THRESHOLD for DIMM1 on this channel.
7:0	FFh RWS_L	DIMM0: WARM_THRESHOLD for DIMM0 on this channel.

7.65 DDR—Offset 5894h

Per-DIMM temp/power thresholds used for CLTM/OLTM thermal status computation. These values can impact iMC throttling and memory thermal interrupts.


Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5894h



Default: FFFFh



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	FFh RWS_L	DIMM1: WARM_THRESHOLD for DIMM1 on this channel.
7:0	FFh RWS_L	DIMM0: WARM_THRESHOLD for DIMM0 on this channel.

7.66 DDR—Offset 5898h


Per-DIMM temp/power thresholds used for CLTM/OLTM thermal status computation. These values can impact iMC throttling and memory thermal interrupts.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5898h

Default: FFFFh



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	FFh RWS_L	DIMM1: HOT_THRESHOLD for DIMM1 on this channel.
7:0	FFh RWS_L	DIMM0: HOT_THRESHOLD for DIMM0 on this channel.



7.67 DDR—Offset 589Ch


Per-DIMM temp/power thresholds used for CLTM/OLTM thermal status computation. These values can impact iMC throttling and memory thermal interrupts.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 589Ch

Default: FFFFh



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	FFh RWS_L	DIMM1: HOT_THRESHOLD for DIMM1 on this channel.
7:0	FFh RWS_L	DIMM0: HOT_THRESHOLD for DIMM0 on this channel.

7.68 DDR—Offset 58A0h

Enable bits and policy-free thresholds used for controlling memory thermal interrupt generation.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 58A0h

Default: 0h



3	2	2	2	1	1	8	4	0								
1	8	4	0	6	2											
0	0	0	0	0	0	0	0	0								
0	0	0	0	0	0	0	0	0								
POLICY_FREE_THRESHOLD2		POLICY_FREE_THRESHOLD1		RSVD		ENABLE_THRESHOLD2_INTERRUPT	RSVD	ENABLE_THRESHOLD1_INTERRUPT	RSVD	ENABLE_OOS_TEMP_INTERRUPT	RSVD	ENABLE_2X_REFRESH_INTERRUPT	RSVD	ENABLE_HOT_INTERRUPT	RSVD	ENABLE_WARM_INTERRUPT

Bit Range	Default and Access	Field Name (ID): Description
31:24	0h RW	POLICY_FREE_THRESHOLD2: A threshold temperature value used only for interrupt generation. No iMC throttling or other actions should be directly affected by this value. This only works when CLTM is enabled. This is an 8-bit unsigned value with variable units/format/resolution. THRESHOLD1 and THRESHOLD2 values and enables are fully independent from each other.
23:16	0h RW	POLICY_FREE_THRESHOLD1: A threshold temperature value used only for interrupt generation. No iMC throttling or other actions should be directly affected by this value. This only works when CLTM is enabled. This is an 8-bit unsigned value with variable units/format/resolution. THRESHOLD1 and THRESHOLD2 values and enables are fully independent from each other.
15:11	0h RO	Reserved (RSVD): Reserved.
10	0h RW	ENABLE_THRESHOLD2_INTERRUPT: When set, interrupts will be generated on both rising and falling transition of the hottest absolute DIMM temperature across the POLICY_FREE_THRESHOLD2 value. This interrupt will never get triggered by hardware in cases where CLTM is not enabled (i.e. does not work with OLTM). THRESHOLD1 and THRESHOLD2 values and enables are fully independent from each other.
9	0h RO	Reserved (RSVD): Reserved.
8	0h RW	ENABLE_THRESHOLD1_INTERRUPT: When set, interrupts will be generated on both rising and falling transition of the hottest absolute DIMM temperature across the POLICY_FREE_THRESHOLD1 value. This interrupt will never get triggered by hardware in cases where CLTM is not enabled (i.e. does not work with OLTM). THRESHOLD1 and THRESHOLD2 values and enables are fully independent from each other.
7	0h RO	Reserved (RSVD): Reserved.
6	0h RW	ENABLE_OOS_TEMP_INTERRUPT: When set, interrupts will be generated on a rising transition of hottest MR4 to 3'b111. This interrupt will never get triggered by hardware in cases where MAD_CHNL.LPDDR is zero or DISABLE_DRAM_TS is set.
5	0h RO	Reserved (RSVD): Reserved.



Bit Range	Default and Access	Field Name (ID): Description
4	0h RW	ENABLE_2X_REFRESH_INTERRUPT: When set, interrupts will be generated on a rising transition of the hottest DIMM thermal status across whichever threshold 2x refresh is configured for (WARM_THRESHOLD, HOT_THRESHOLD, or never, depending on DDR_PTM_CTL.REFRESH_2X_MODE). This interrupt will never be triggered by hardware in cases where 2X refresh is disabled OR when no thermal status updates are being performed because CLTM, OLTM, and EXTTS are all disabled.
3	0h RO	Reserved (RSVD): Reserved.
2	0h RW	ENABLE_HOT_INTERRUPT: When set, interrupts will be generated on a rising transition of the hottest DIMM thermal status from WARM to HOT (i.e. rise to or above HOT_THRESHOLD). This interrupt will never get triggered by hardware in cases where CLTM, OLTM, and EXTTS are all disabled.
1	0h RO	Reserved (RSVD): Reserved.
0	0h RW	ENABLE_WARM_INTERRUPT: When set, interrupts will be generated on a rising transition of the hottest DIMM thermal status from COLD to WARM (i.e. rise to or above WARM_THRESHOLD). This interrupt will never get triggered by hardware in cases where CLTM, OLTM, and EXTTS are all disabled.

7.69 PACKAGE—Offset 58A8h

Temperature margin in PECI temperature counts from the thermal profile specification. Platform fan control SW is expected to read therm_margin value to control fan or blower speed.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 58A8h

Default: 7F00h

3 1	2 8	2 4	2 0	1 6	1 2	8	4	0
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 1 1 1	1 1 1 1	0 0 0 0	0 0 0 0	0
RSVD				THERM_MARGIN				



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:0	7F00h RO_V	THERM_MARGIN: Temperature margin in PECI temperature counts from the thermal profile specification. THERM_MARGIN is in 2's complement format (8.8 format where MSB equals 1 Sign bit + 7 bits of integer temperature value and the LSB equals 8 precision bits of temperature value). A value of zero indicates the hottest Processor die temperature is on the thermal profile line. A negative value indicates gap to the thermal profile that platform SW should increase cooling capacity. A sustained negative value should be avoided as it may impact part reliability.

7.70 DDR—Offset 58B0h


Per-DIMM temperature values.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 58B0h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	0h ROS	DIMM1: Temperature of DIMM1 on this channel.
7:0	0h ROS	DIMM0: Temperature of DIMM0 on this channel.



7.71 DDR—Offset 58B4h


Per-DIMM temperature values.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 58B4h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	0h ROS	DIMM1: Temperature of DIMM1 on this channel.
7:0	0h ROS	DIMM0: Temperature of DIMM0 on this channel.

7.72 DDR—Offset 58C0h


Per-DIMM throttle duration counters. These accumulate the duration (in absolute wall clock time) that the iMC rank throttlers have been blocking memory traffic due to OLTM/CLTM/EXTTS thermal status. Note that RAPL throttling is done at the channel level, and thus is NOT included in these values.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 58C0h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
63:32	0h RO	Reserved (RSVD): Reserved.
31:16	0h ROS	DIMM1: Throttle duration of DIMM 1 on this channel, in units of 1/1024 seconds.
15:0	0h ROS	DIMM0: Throttle duration of DIMM 0 on this channel, in units of 1/1024 seconds.

7.73 DDR—Offset 58C8h


Per-DIMM throttle duration counters. These accumulate the duration (in absolute wall clock time) that the iMC rank throttlers have been blocking memory traffic due to OLTM/CLTM/EXTTS thermal status. Note that RAPL throttling is done at the channel level, and thus is NOT included in these values.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 58C8h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:32	0h RO	Reserved (RSVD): Reserved.
31:16	0h ROS	DIMM1: Throttle duration of DIMM 1 on this channel, in units of 1/1024 seconds.
15:0	0h ROS	DIMM0: Throttle duration of DIMM 0 on this channel, in units of 1/1024 seconds.



7.74 DDR—Offset 58D0h


Per-DIMM power budget for MC thermal throttling when thermal status is WARM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 58D0h

Default: FFFFh



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	FFh RWS_L	DIMM1: WARM_BUDGET for DIMM1 on this channel.
7:0	FFh RWS_L	DIMM0: WARM_BUDGET for DIMM0 on this channel.

7.75 DDR—Offset 58D4h


Per-DIMM power budget for MC thermal throttling when thermal status is WARM.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 58D4h

Default: FFFFh





Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	FFh RWS_L	DIMM1: WARM_BUDGET for DIMM1 on this channel.
7:0	FFh RWS_L	DIMM0: WARM_BUDGET for DIMM0 on this channel.

7.76 DDR—Offset 58D8h


Per-DIMM power budget for MC thermal throttling when thermal status is HOT.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 58D8h

Default: FFFFh



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	FFh RWS_L	DIMM1: HOT_BUDGET for DIMM1 on this channel.
7:0	FFh RWS_L	DIMM0: HOT_BUDGET for DIMM0 on this channel.



7.77 DDR—Offset 58DCh


Per-DIMM power budget for MC thermal throttling when thermal status is HOT.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 58DCh

Default: FFFFh



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	FFh RWS_L	DIMM1: HOT_BUDGET for DIMM1 on this channel.
7:0	FFh RWS_L	DIMM0: HOT_BUDGET for DIMM0 on this channel.

7.78 PACKAGE—Offset 58F0h


Package RAPL Performance Status Register. This register provides information on the performance impact of the RAPL power limit and indicates the duration for processor went below the requested P-state due to package power constraint.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 58F0h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31:0	0h ROS_V	COUNTS: Counter of the time units within which RAPL was limiting P-states. If limitation occurred anywhere within the time window of 1/1024 seconds, the count will be incremented (limitation on accuracy). This data can serve as a proxy for the potential performance impacts of RAPL on cores performance.

7.79 IA—Offset 58FCh

Interface to allow software to determine what is causing resolved frequency to be clamped below the requested frequency. Status bits are updated by hardware through the IO interface IO_IA_PERF_LIMIT, log bits are set by HW on a status bit edge detected and cleared by a SW write of '0'.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 58FCh

Default: 0h

3	2	2	2	1	1	8	4	0
1	8	4	0	6	2			
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
SPARE_IA_15_LOG	SPARE_IA_14_LOG	TURBO_ATTEN_LOG	MAX_TURBO_LIMIT_LOG	PBM_PL2_LOG	PBM_PL1_LOG	SPARE_IA_9_LOG	OTHER_LOG	VR_TDC_LOG
VR_THERMALERT_LOG	RATL_LOG	RSR_LIMIT_LOG	SPARE_IA_3_LOG	SPARE_IA_2_LOG	THERMAL_LOG	PROCHOT_LOG	SPARE_IA_15	SPARE_IA_14
TURBO_ATTEN	MAX_TURBO_LIMIT	PBM_PL2	PBM_PL1	SPARE_IA_9	OTHER	VR_TDC	VR_THERMALERT	RATL
RSR_LIMIT	SPARE_IA_3	SPARE_IA_2	THERMAL	PROCHOT				

Bit Range	Default and Access	Field Name (ID): Description
31	0h RW0C	SPARE_IA_15_LOG: Reserved
30	0h RW0C	SPARE_IA_14_LOG: Reserved
29	0h RW0C	TURBO_ATTEN_LOG: Turbo attenuation (multi core turbo) Log, RW, When set by hardware indicates that Turbo attenuation (multi core turbo) has cause IA frequency clipping. Software should write to this bit to clear the status in this bit
28	0h RW0C	MAX_TURBO_LIMIT_LOG: Max turbo limit Log, RW, When set by hardware indicates that Max turbo limit has cause IA frequency clipping. Software should write to this bit to clear the status in this bit
27	0h RW0C	PBM_PL2_LOG: PBM PL2, PL3 (pkg, platform) Log, RW, When set by hardware indicates that PBM PL2 or PL3(package or platform PL2 or PL3) has cause IA frequency clipping. Software should write to this bit to clear the status in this bit



Bit Range	Default and Access	Field Name (ID): Description
26	0h RW0C	PBM_PL1_LOG: PBM PL1 (pkg, platform) Log, RW, When set by hardware indicates that PBM PL1 (package or platform PL1) has cause IA frequency clipping. Software should write to this bit to clear the status in this bit
25	0h RW0C	SPARE_IA_9_LOG: Reserved
24	0h RW0C	OTHER_LOG: Other (IccMax, PL4, etc) Log, RW, When set by hardware indicates that other has cause reason IA frequency clipping. Software should write to this bit to clear the status in this bit
23	0h RW0C	VR_TDC_LOG: VR TDC (Thermal design current) Log, RW, When set by hardware indicates that VR TDC (Thermal design current) has cause IA frequency clipping. Software should write to this bit to clear the status in this bit
22	0h RW0C	VR_THERMALERT_LOG: Hot VR (any processor VR) Log, RW, When set by hardware indicates that Hot VR (any processor VR) has cause IA frequency clipping. Software should write to this bit to clear the status in this bit
21	0h RW0C	RATL_LOG: Running average thermal limit Log, RW, When set by hardware indicates that Running average thermal limit has cause IA frequency clipping. Software should write to this bit to clear the status in this bit
20	0h RW0C	RSR_LIMIT_LOG: Residency State Regulation Log, RW, When set by hardware indicates that Residency State Regulation has cause IA frequency clipping. Software should write to this bit to clear the status in this bit
19	0h RW0C	SPARE_IA_3_LOG: Reserved
18	0h RW0C	SPARE_IA_2_LOG: Reserved
17	0h RW0C	THERMAL_LOG: Thermal Log, RW, When set by hardware indicates that Thermal event has cause IA frequency clipping. Software should write to this bit to clear the status in this bit
16	0h RW0C	PROCHOT_LOG: PROCHOT# Log, RW, When set by hardware indicates that PROCHOT# has cause IA frequency clipping. Software should write to this bit to clear the status in this bit
15	0h ROV	SPARE_IA_15: Reserved
14	0h ROV	SPARE_IA_14: Reserved
13	0h ROV	TURBO_ATTEN: Turbo attenuation (multi core turbo) Status, RO, When set by hardware indicates that Turbo attenuation (multi core turbo) has cause IA frequency clipping
12	0h ROV	MAX_TURBO_LIMIT: Max turbo limit Status, RO, When set by hardware indicates that Max turbo limit has cause IA frequency clipping
11	0h ROV	PBM_PL2: PBM PL2, PL3 (pkg, platform) Status, RO, When set by hardware indicates that PBM PL2 or PL3(package or platform PL2 or PL3) has cause IA frequency clipping
10	0h ROV	PBM_PL1: PBM PL1 (pkg, platform), RO, When set by hardware indicates that PBM PL1 (package or platform PL1) has cause IA frequency clipping
9	0h ROV	SPARE_IA_9: Reserved
8	0h ROV	OTHER: Other (IccMax, PL4, etc) Status, RO, When set by hardware indicates that other has cause reason IA frequency clipping
7	0h ROV	VR_TDC: VR TDC (Thermal design current) Status, RO, When set by hardware indicates that VR TDC (Thermal design current) has cause IA frequency clipping
6	0h ROV	VR_THERMALERT: Hot VR (any processor VR) Status, RO, When set by hardware indicates that Hot VR (any processor VR) has cause IA frequency clipping



Bit Range	Default and Access	Field Name (ID): Description
5	0h ROV	RATL: Running average thermal limit Status, RO, When set by hardware indicates that Running average thermal limit has cause IA frequency clipping
4	0h ROV	RSR_LIMIT: Residency State Regulation Status, RO, When set by hardware indicates that Residency State Regulation has cause IA frequency clipping
3	0h ROV	SPARE_IA_3: Reserved
2	0h ROV	SPARE_IA_2: Reserved
1	0h ROV	THERMAL: Thermal Status, RO, When set by hardware indicates that Thermal event has cause IA frequency clipping
0	0h ROV	PROCHOT: PROCHOT# Status, RO, When set by hardware indicates that PROCHOT# has cause IA frequency clipping

7.80 GT—Offset 5900h

Interface to allow software to determine what is causing resolved frequency to be clamped below the requested frequency. Status bits are updated by hardware through the io interface IO_GT_PERF_LIMIT, log bits are set by HW on a status bit edge detected and cleared by a SW write of '0'.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5900h

Default: 0h

3	2	2	2	1	1	8	4	0
1	8	4	0	6	2			
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
SPARE_GT_15_LOG	SPARE_GT_14_LOG	SPARE_GT_13_LOG	INEFFICIENT_OPERATION_LOG	PBM_PL2_LOG	PBM_PL1_LOG	SPARE_GT_9_LOG	OTHER_LOG	VR_TDC_LOG
VR_THERMALERT_LOG	RATL_LOG	RSR_LIMIT_LOG	SPARE_GT_3_LOG	SPARE_GT_LOG_2	THERMAL_LOG	PROCHOT_LOG	SPARE_GT_15	SPARE_GT_14
SPARE_GT_13	INEFFICIENT_OPERATION	PBM_PL2	PBM_PL1	SPARE_GT_9	OTHER	VR_TDC	VR_THERMALERT	RATL
RSR_LIMIT	SPARE_GT_3	SPARE_GT_2	THERMAL	PROCHOT				



Bit Range	Default and Access	Field Name (ID): Description
31	0h RW0C	SPARE_GT_15_LOG: Reserved
30	0h RW0C	SPARE_GT_14_LOG: Reserved
29	0h RW0C	SPARE_GT_13_LOG: Reserved
28	0h RW0C	INEFFICIENT_OPERATION_LOG: Inefficient operation Log, RW, The current GT Frequency lower than the DCC target Frequency. Software should write to this bit to clear the status in this bit
27	0h RW0C	PBM_PL2_LOG: PBM PL2, PL3 (pkg, platform) Log, RW, When set by hardware indicates that PBM PL2 or PL3(package or platform PL2 or PL3) has cause GT frequency clipping. Software should write to this bit to clear the status in this bit
26	0h RW0C	PBM_PL1_LOG: PBM PL1 (pkg, platform) Log, RW, When set by hardware indicates that PBM PL1 (package or platform PL1) has cause GT frequency clipping. Software should write to this bit to clear the status in this bit
25	0h RW0C	SPARE_GT_9_LOG: Reserved
24	0h RW0C	OTHER_LOG: Other (IccMax, PL4, etc) Log, RW, When set by hardware indicates that other has cause reason GT frequency clipping. Software should write to this bit to clear the status in this bit
23	0h RW0C	VR_TDC_LOG: VR TDC (Thermal design current) Log, RW, When set by hardware indicates that VR TDC (Thermal design current) has cause GT frequency clipping. Software should write to this bit to clear the status in this bit
22	0h RW0C	VR_THERMALERT_LOG: Hot VR (any processor VR) Log, RW, When set by hardware indicates that Hot VR (any processor VR) has cause GT frequency clipping. Software should write to this bit to clear the status in this bit
21	0h RW0C	RATL_LOG: Running average thermal limit Log, RW, When set by hardware indicates that Running average thermal limit has cause GT frequency clipping. Software should write to this bit to clear the status in this bit
20	0h RW0C	RSR_LIMIT_LOG: Reserved
19	0h RW0C	SPARE_GT_3_LOG: Reserved
18	0h RW0C	SPARE_GT_LOG_2: Reserved
17	0h RW0C	THERMAL_LOG: Thermal Log, RW, When set by hardware indicates that Thermal event has cause GT frequency clipping. Software should write to this bit to clear the status in this bit
16	0h RW0C	PROCHOT_LOG: PROCHOT# Log, RW, When set by hardware indicates that PROCHOT# has cause GT frequency clipping. Software should write to this bit to clear the status in this bit
15	0h ROV	SPARE_GT_15: Reserved
14	0h ROV	SPARE_GT_14: Reserved
13	0h ROV	SPARE_GT_13: Reserved
12	0h ROV	INEFFICIENT_OPERATION: Inefficient operation Status, RO, The current GT Frequency lower than the DCC target Frequency
11	0h ROV	PBM_PL2: PBM PL2, PL3 (pkg, platform) Status, RO, When set by hardware indicates that PBM PL2 or PL3(package or platform PL2 or PL3) has cause GT frequency clipping



Bit Range	Default and Access	Field Name (ID): Description
10	0h ROV	PBM_PL1: PBM PL1 (pkg, platform), RO, When set by hardware indicates that PBM PL1 (package or platform PL1) has cause GT frequency clipping
9	0h ROV	SPARE_GT_9: Reserved
8	0h ROV	OTHER: Other (IccMax, PL4, etc) Status, RO, When set by hardware indicates that other has cause reason GT frequency clipping
7	0h ROV	VR_TDC: VR TDC (Thermal design current) Status, RO, When set by hardware indicates that VR TDC (Thermal design current has cause GT frequency clipping
6	0h ROV	VR_THERMALERT: Hot VR (any processor VR) Status, RO, When set by hardware indicates that Hot VR (any processor VR) has cause GT frequency clipping
5	0h ROV	RATL: Running average thermal limit Status, RO, When set by hardware indicates that Running average thermal limit has cause GT frequency clipping
4	0h ROV	RSR_LIMIT: Reserved
3	0h ROV	SPARE_GT_3: Reserved
2	0h ROV	SPARE_GT_2: Reserved
1	0h ROV	THERMAL: Thermal Status, RO, When set by hardware indicates that Thermal event has cause GT frequency clipping
0	0h ROV	PROCHOT: PROCHOT# Status, RO, When set by hardware indicates that PROCHOT# has cause GT frequency clipping

7.81 SA—Offset 5918h


System Agent Performance status. Indicates current SA PLLs ratios. Frequency to be calculated according to reference.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5918h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31:24	0h RO_V	UCLK_RATIO : RING UCLK RATIO. Reference=100Mhz
23:16	0h RO_V	ICLK_RATIO : IMGU ICLK RATIO. Reference=25Mhz
15:8	0h RO_V	FCLK_RATIO : SA FCLK RATIO. Reference=100Mhz
7	0h RO_V	QCLK_REFERENCE : DDR QCLK REFERENCE. 0=133Mhz, 1=100Mhz
6:0	0h RO_V	QCLK_RATIO : DDR QCLK RATIO. Reference determined by QCLK_REFERENCE

7.82 GT—Offset 5948h


P-state encoding for the Secondary Power Plane's current PLL frequency and the current VID.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5948h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:26	0h RO	Reserved (RSVD) : Reserved.



Bit Range	Default and Access	Field Name (ID): Description
25:17	0h RO_V	RP_STATE_RATIO_SLICE: GT Slices frequency, in granularity of 16.666Mhz. When GT is in RC6, or when all slices are disabled, this frequency is ZERO.
16:8	0h RO_V	RP_STATE_RATIO_UNSLICE: GT Unslice frequency, in granularity of 16.666 MHz. When GT is in RC6 this frequency is ZERO.
7:0	0h RO_V	RP_STATE_VOLTAGE: Voltage of the current RP-state.

7.83 EDRAM—Offset 594Ch


EDRAM die temperature in degrees (C). This field is updated by FW.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 594Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:8	0h RO	Reserved (RSVD): Reserved.
7:0	0h RO_V	DATA: Temperature in degrees (C).

7.84 Package—Offset 5978h


Package temperature in degrees (C). This field is updated by FW.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5978h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:8	0h RO	Reserved (RSVD): Reserved.
7:0	0h RO_V	DATA: Package temperature in degrees (C).

7.85 PP0—Offset 597Ch


PP0 (IA) temperature in degrees (C). This field is updated by FW.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 597Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:8	0h RO	Reserved (RSVD): Reserved.
7:0	0h RO_V	DATA: Temperature in degrees (C).

7.86 PP1—Offset 5980h

PP1 (GT) temperature in degrees (C). This field is updated by FW.


Access Method



Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5980h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:8	0h RO	Reserved (RSVD): Reserved.
7:0	0h RO_V	DATA: Temperature in degrees (C).

7.87 RP—Offset 5994h


This register allows SW to limit the maximum base frequency for the Integrated GFX Engine (GT) allowed during run-time.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5994h

Default: FFh



Bit Range	Default and Access	Field Name (ID): Description
31:8	0h RO	Reserved (RSVD): Reserved.
7:0	FFh RW	RPSTT_LIM: This field indicates the maximum base frequency limit for the Integrated GFX Engine (GT) allowed during run-time.



7.88 RP—Offset 5998h


This register contains the maximum base frequency capability for the Integrated GFX Engine (GT).

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5998h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:24	0h RO	Reserved (RSVD): Reserved.
23:16	0h ROS	RPN_CAP: This field indicates the maximum RPN base frequency capability for the Integrated GFX Engine (GT). Values are in units of 50 MHz.
15:8	0h ROS	RP1_CAP: This field indicates the maximum RP1 base frequency capability for the Integrated GFX Engine (GT). Values are in units of 50 MHz.
7:0	0h ROS	RPO_CAP: This field indicates the maximum RPO base frequency capability for the Integrated GFX Engine (GT). Values are in units of 50 MHz.



7.89 SSKPD—Offset 5D10h


This register holds 64 writable bits with no functionality behind them. It is for the convenience of BIOS and graphics drivers.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 5D10h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:0	0h RWS	SSKPD: 4 WORDs of data storage.

7.90 BIOS_RESET_CPL—Offset 5DA8h


This register is used as interface between BIOS and hardware. It is written by BIOS and read by hardware.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5DA8h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:4	0h RO	Reserved (RSVD): Reserved.
3	0h RW1S	ENABLE_PCIE_NDA_PG: This bit indicates if PCIe-NDA power-gating is enabled (disabled by default). Hardware looks at this bit after RST_CPL is set and decides whether or not to power-gate the PEG controllers and AFE. If it is asserted and all devices are disabled (post CPL), hardware will power-gate the devices. Note 1: This mode does not survive warm-reset, i.e. on a warm reset NDA mode is canceled and power to PEG controllers is resumed. Note 2: If checked only on CPL, no need to check also PCIE_ENUMERATION_DONE.
2	0h RO	Reserved (RSVD): Reserved.
1	0h RW	PCIE_ENUMERATION_DONE: This will be set after PCIe enumeration is done. This bit will be read by hardware. If it is set, hardware will look at the following register bits: DEVEN_0_0_0_PCI Bit: Bit Name 1: D1F2EN 2: D1F1EN 3: D1F0EN If all of these bits are set to a 0x0, this means that there is nothing connected to the PEG devices and the Gen3 PLL can be shut off. Note: implicit assumption - this bit is asserted prior to (or with) asserting RST_CPL.
0	0h RW	RST_CPL: This bit is set by BIOS to indicate to the Processor Power management function that it has completed to set up all PM relevant configuration and allow Processor Power management function to digest the configuration data and start active PM operation. It is expected that this bit will be set just before BIOS transfer of control to the OS. 0b: Not ready 1b: BIOS PM configuration complete

7.91 PCU_CR_MC_BIOS_REQ_0_0_0_MCHBAR_PCU—Offset 5E00h

This register allows BIOS to request Memory Controller clock frequency.


Access Method



Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5E00h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:4	0h RO	Reserved (RSVD): Reserved.
3:0	0h RW	<p>REQ_DATA: These 4 bits are the data for the request. The only possible request type is MC frequency request. The encoding of this field is the 133/266 MHz multiplier for DCLK/QCLK: Binary Dec DCLK Equation DCLK Freq QCLK Equation QCLK Freq</p> <p>000b 0d -----MC PLL – shutdown-----</p> <p>...</p> <p>0011b 3d 3*133.33 400.00 MHz 3*266.67 MHz 800.00 MHz</p> <p>0100b 4d 4*133.33 533.33 MHz 4*266.67 MHz 1066.67 MHz</p> <p>0101b 5d 5*133.33 666.67 MHz 5*266.67 MHz 1333.33 MHz</p> <p>0110b 6d 6*133.33 800.00 MHz 6*266.67 MHz 1600.00 MHz</p> <p>0111b 7d 7*133.33 933.33 MHz 7*266.67 MHz 1866.67 MHz</p> <p>1000b 8d 8*133.33 1066.67 MHz 8*266.67 MHz 2133.33 MHz</p> <p>...</p>

7.92 CONFIG—Offset 5F3Ch


This register is used to indicate the Nominal Configurable TDP ratio available for this specific SKU. System BIOS should use this value while building the _PSS table if the feature is enabled.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5F3Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:8	0h RO	Reserved (RSVD): Reserved.
7:0	0h RO_V	TDP_RATIO: Nominal TDP level ratio to be used for this specific processor (in units of 100 MHz). Note: A value of 0 in this field indicates invalid/undefined TDP point

7.93 CONFIG—Offset 5F40h

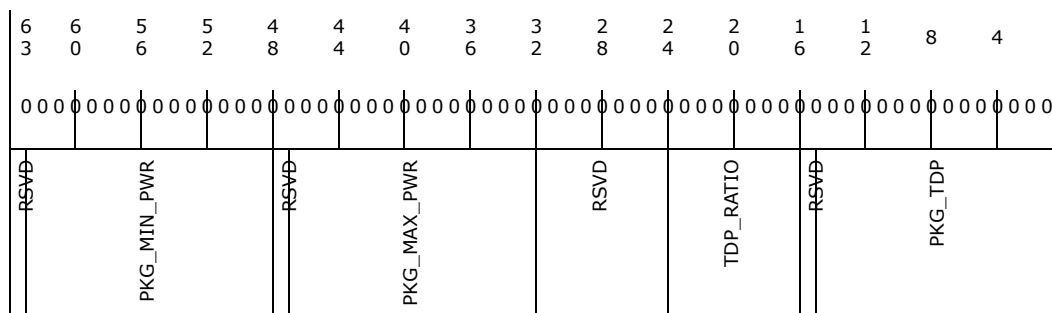
Level 1 configurable TDP settings

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 5F40h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
63	0h RO	Reserved (RSVD): Reserved.
62:48	0h RO_V	PKG_MIN_PWR: Min pkg power setting allowed for this config TDP level. Lower values will be clamped up to this value. Units defined in PACKAGE_POWER_SKU_MSR[PWR_UNIT]. Similar to PACKAGE_POWER_SKU[PKG_MIN_PWR].
47	0h RO	Reserved (RSVD): Reserved.
46:32	0h RO_V	PKG_MAX_PWR: Max pkg power setting allowed for this config TDP level1. Higher values will be clamped down to this value. Units defined in PACKAGE_POWER_SKU_MSR[PWR_UNIT]. Similar to PACKAGE_POWER_SKU[PKG_MAX_PWR].
31:24	0h RO	Reserved (RSVD): Reserved.
23:16	0h RO_V	TDP_RATIO: TDP ratio for config tdp level 1.
15	0h RO	Reserved (RSVD): Reserved.
14:0	0h RO_V	PKG_TDP: Power for this TDP level. Units defined in PACKAGE_POWER_SKU_MSR[PWR_UNIT] Similar to PACKAGE_POWER_SKU[PKG_TDP]

7.94 CONFIG—Offset 5F48h

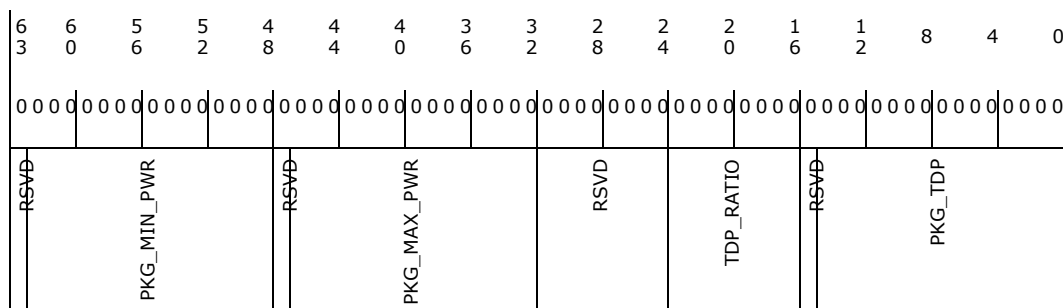
Level 2 configurable TDP settings

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 5F48h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
63	0h RO	Reserved (RSVD): Reserved.
62:48	0h RO_V	PKG_MIN_PWR: Min pkg power setting allowed for this config TDP level 2. Lower values will be clamped up to this value. Units defined in PACKAGE_POWER_SKU_MSR[PWR_UNIT]. Similar to PACKAGE_POWER_SKU[PKG_MIN_PWR].
47	0h RO	Reserved (RSVD): Reserved.
46:32	0h RO_V	PKG_MAX_PWR: Max pkg power setting allowed for config TDP level 2. Higher values will be clamped down to this value. Units defined in PACKAGE_POWER_SKU_MSR[PWR_UNIT]. Similar to PACKAGE_POWER_SKU[PKG_MAX_PWR].
31:24	0h RO	Reserved (RSVD): Reserved.
23:16	0h RO_V	TDP_RATIO: TDP ratio for level 2.
15	0h RO	Reserved (RSVD): Reserved.
14:0	0h RO_V	PKG_TDP: Power for this TDP level. Units defined in PACKAGE_POWER_SKU_MSR[PWR_UNIT] Similar to PACKAGE_POWER_SKU[PKG_TDP].

7.95 CONFIG—Offset 5F50h


Rd/Wr register to allow platform SW to select TDP point and set lock

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5F50h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31	0h RWS_KL	CONFIG_TDP_LOCK: Config TDP level select lock 0 - unlocked. 1 - locked till next reset.
30:2	0h RO	Reserved (RSVD): Reserved.
1:0	0h RWS_L	TDP_LEVEL: Config TDP level selected 0: nominal TDP level (default) 1: Level from CONFIG_TDP_LEVEL_1 2: Level from CONFIG_TDP_LEVEL_2 3: reserved

7.96 TURBO—Offset 5F54h

Read/write register to allow MSR/MMIO access to ACPI P-state notify (PCS 33).

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 5F54h

Default: 0h

3 1	2 8	2 4	2 0	1 6	1 2	8	4	0	
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0 0 0	0 0 0 0	0 0 0 0	
TURBO_ACTIVATION_RATIO_LOCK						RSVD			MAX_NON_TURBO_RATIO



Bit Range	Default and Access	Field Name (ID): Description
31	0h RWS_KL	TURBO_ACTIVATION_RATIO_LOCK: Lock this MSR until next reset 0: unlocked 1: locked
30:8	0h RO	Reserved (RSVD): Reserved.
7:0	0h RWS_L	MAX_NON_TURBO_RATIO: Processor will treat any P-state request above this ratio as a request for max turbo 0 is special encoding which disables the feature.

7.97 Package Thermal DPPM Status (PKG)—Offset 6200h

Thermal Status for Digital Thermometer

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 6200h

Default: 8000h

3	2	2	2	1	1	8	4	0
1	8	4	0	6	2			
0	0	0	0	0	1	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
Valid	Resolution	RSVD	Temperature	RSVD	POWER_LIMITATION_LOG	POWER_LIMITATION_STATUS	THRESHOLD2_LOG	THRESHOLD2_STATUS
					THRESHOLD1_LOG	THRESHOLD1_STATUS	OUT_OF_SPEC_LOG	OUT_OF_SPEC_STATUS
					PROCHOT_LOG	PROCHOT_STATUS	THERMAL_MONITOR_LOG	THERMAL_MONITOR_STATUS

Bit Range	Default and Access	Field Name (ID): Description
31	0h ROV	Valid: Set if temperature is within the valid thermal sensor range.
30:27	1h RO	Resolution: Supported resolution in degrees C. Hard-coded to '0001'. Currently reserved and not in use.
26:23	0h RO	Reserved (RSVD): Reserved.
22:16	0h ROV	Temperature: Temperature in degrees C, relative to the thermal monitor trip temperature.



Bit Range	Default and Access	Field Name (ID): Description
15:12	0h RO	Reserved (RSVD): Reserved.
11	0h RW0C	POWER_LIMITATION_LOG: Sticky log bit indicating that the power has transitioned out of its limitation status since the last time SW cleared this bit. Set by HW on a 0 to 1 transition of Power Limitation Status.
10	0h ROV	POWER_LIMITATION_STATUS: Indicates that either IA is running at P-state below the (max P-state - offset) or that GT is running at P-state below its P1 frequency.
9	0h RW0C	THRESHOLD2_LOG: Sticky log bit that asserts on a 0 to 1 or 1 to 0 transition of the Threshold2_Status bit. HW controls this transition.
8	0h ROV	THRESHOLD2_STATUS: Indicates that the current temperature (bits 23:16 of this register) is equal to or higher than the Threshold2 defined in the IA32_PACKAGE_THERM_INTERRUPT MSR. Note that because temperature and thresholds are defined as negative offsets, a higher number means a lower temperature.
7	0h RW0C	THRESHOLD1_LOG: Sticky log bit that asserts on a 0 to 1 or 1 to 0 transition of the Threshold1_Status bit. HW controls this transition.
6	0h ROV	THRESHOLD1_STATUS: Indicates that the current temperature (bits 23:16 in this register) is equal to or higher than the Threshold1 defined in the IA32_PACKAGE_THERM_INTERRUPT MSR. Note that because temperature and thresholds are defined as negative offsets, a higher number means a lower temperature.
5	0h RW0C	OUT_OF_SPEC_LOG: Sticky log bit indicating that the processor has operated out of its thermal specification since the last time SW cleared this bit. Set by HW on a 0 to 1 transition of Out_of_Spec_Status.
4	0h ROV	OUT_OF_SPEC_STATUS: Status bit indicating that the processor is operating out of its thermal specification.
3	0h RW0C	PROCHOT_LOG: Sticky log bit indicating that PROCHOT# has been asserted since the last time SW cleared this bit. Set by HW on a 0 to 1 transition of Prochot_Status.
2	0h ROV	PROCHOT_STATUS: Status bit indicating that PROCHOT# is currently being asserted.
1	0h RW0C	THERMAL_MONITOR_LOG: Sticky log bit indicating that the package has seen a thermal monitor event since the last time SW cleared this bit. Set by HW on a 0 to 1 transition of Thermal_Monitor_Status.
0	0h ROV	THERMAL_MONITOR_STATUS: Status bit indicating that any of the package thermal monitors have tripped and the package is currently thermally throttling.DPPM

7.98 Memory Thermal DPPM Status (DDR)—Offset 6204h

Status and log bits of memory thermal interrupt enabled through configuration of DDR_THERM_THRESHOLDS_CONFIG.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 6204h

Default: 0h



3	2	2	2	1	1	8	4	0
1	8	4	0	6	2			
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
RSVD						THRESHOLD2_LOG	THRESHOLD2_STATUS	THRESHOLD1_LOG
						THRESHOLD1_STATUS	OOS_TEMP_LOG	OOS_TEMP_STATUS
						REFRESH2X_LOG	REFRESH2X_STATUS	HOT_THRESHOLD_LOG
						HOT_THRESHOLD_STATUS	WARM_THRESHOLD_LOG	WARM_THRESHOLD_STATUS

Bit Range	Default and Access	Field Name (ID): Description
31:12	0h RO	Reserved (RSVD): Reserved.
11	0h RW0C	THRESHOLD2_LOG: Sticky log bit that asserts on a 0 to 1 transition of the THRESHOLD2_STATUS bit. HW controls this transition.
10	0h ROV	THRESHOLD2_STATUS: Status bit indicating that the hottest DIMM has crossed the THRESHOLD2 value programmed in bits 20:13 of DDR_THERM_DPPM_INTERRUPT.
9	0h RW0C	THRESHOLD1_LOG: Sticky log bit that asserts on a 0 to 1 transition of the THRESHOLD1_STATUS bit. HW controls this transition.
8	0h ROV	THRESHOLD1_STATUS: Status bit indicating that the hottest DIMM has crossed the THRESHOLD1 value programmed in bits 11:4 of DDR_THERM_DPPM_INTERRUPT.
7	0h RW0C	OOS_TEMP_LOG: Sticky log bit that asserts on a 0 to 1 transition of the OOS_TEMP_STATUS bit. HW controls this transition.
6	0h ROV	OOS_TEMP_STATUS: Status bit indicating that MR4 is currently indicating at least one DRAM with high temperature which is beyond the operating range. This can only occur currently when MAD_CHNL.LPDDR=1 and DDR_PTM_CTL.DISABLE_DRAM_TS=0.
5	0h RW0C	REFRESH2X_LOG: Sticky log bit that asserts on a 0 to 1 transition of the REFRESH2X_STATUS bit. HW controls this transition.
4	0h ROV	REFRESH2X_STATUS: Status bit indicating that the DIMM refresh rate has crossed the boundary (in either direction) between 1x or lower refresh rate, and higher than 1x refresh rate. The name is misleading for LPDDR where we may go above 2x refresh rate.
3	0h RW0C	HOT_THRESHOLD_LOG: Sticky log bit that asserts on a 0 to 1 transition of the HOT_THRESHOLD_STATUS bit. HW controls this transition.
2	0h ROV	HOT_THRESHOLD_STATUS: Status bit indicating that the DDR temperature is higher than or equal to the DDR Hot threshold defined in DDR_THERM_THRESHOLDS_CONFIG.
1	0h RW0C	WARM_THRESHOLD_LOG: Sticky log bit that asserts on a 0 to 1 transition of the WARM_THRESHOLD_STATUS bit. HW controls this transition.
0	0h ROV	WARM_THRESHOLD_STATUS: Status bit indicating that the DDR temperature is higher than or equal to the DDR Warm threshold defined in DDR_THERM_THRESHOLDS_CONFIG.





8 GFXVTBAR Registers

Table 8-1. Summary of Bus: 0, Device: 0, Function: 0 (MEM)

Offset	Size (Bytes)	Register Name (Register Symbol)	Default Value
0-3h	4	Version Register (VER)—Offset 0h	10h
8-Fh	8	Capability Register (CAP)—Offset 8h	1C0000C40660462h
10-17h	8	Extended Capability Register (ECAP)—Offset 10h	7E3FF0505Eh
18-1Bh	4	Global Command Register (GCMD)—Offset 18h	0h
1C-1Fh	4	Global Status Register (GSTS)—Offset 1Ch	0h
20-27h	8	Root-Entry Table Address Register (RTADDR)—Offset 20h	0h
28-2Fh	8	Context Command Register (CCMD)—Offset 28h	80h
34-37h	4	Fault Status Register (FSTS)—Offset 34h	0h
38-3Bh	4	Fault Event Control Register (FECTL)—Offset 38h	80000h
3C-3Fh	4	Fault Event Data Register (FEDATA)—Offset 3Ch	0h
40-43h	4	Fault Event Address Register (FEADDR)—Offset 40h	0h
44-47h	4	Fault Event Upper Address Register (FEUADDR)—Offset 44h	0h
58-5Fh	8	Advanced Fault Log Register (AFLOG)—Offset 58h	0h
64-67h	4	Protected Memory Enable Register (PMEN)—Offset 64h	0h
68-6Bh	4	Protected Low-Memory Base Register (PLMBASE)—Offset 68h	0h
6C-6Fh	4	Protected Low-Memory Limit Register (PLMLIMIT)—Offset 6Ch	0h
70-77h	8	Protected High-Memory Base Register (PHMBASE)—Offset 70h	0h
78-7Fh	8	Protected High-Memory Limit Register (PHMLIMIT)—Offset 78h	0h
80-87h	8	Invalidation Queue Head Register (IQH)—Offset 80h	0h
88-8Fh	8	Invalidation Queue Tail Register (IQT)—Offset 88h	0h
90-97h	8	Invalidation Queue Address Register (IQA)—Offset 90h	0h
9C-9Fh	4	Invalidation Completion Status Register (ICS)—Offset 9Ch	0h
A0-A3h	4	Invalidation Event Control Register (IECTL)—Offset A0h	80000h
A4-A7h	4	Invalidation Event Data Register (IEDATA)—Offset A4h	0h
A8-ABh	4	Invalidation Event Address Register (IEADDR)—Offset A8h	0h
AC-AFh	4	Invalidation Event Upper Address Register (IEUADDR)—Offset ACh	0h
B8-BFh	8	Interrupt Remapping Table Address Register (IRTA)—Offset B8h	0h
400-407h	8	Fault Recording Low Register (FRCDL)—Offset 400h	0h
408-40Fh	8	Fault Recording High Register (FRCDH)—Offset 408h	0h
500-507h	8	Invalidate Address Register (IVA)—Offset 500h	0h
508-50Fh	8	IOTLB Invalidate Register (IOTLB)—Offset 508h	20h
FF0-FF3h	4	DMA Remap Engine Policy Control (ARCHDIS)—Offset FF0h	1h
FF4-FF7h	4	DMA Remap Engine Policy Control (UARCHDIS)—Offset FF4h	100h




8.1 Version Register (VER)—Offset 0h

Register to report the architecture version supported. Backward compatibility for the architecture is maintained with new revision numbers, allowing software to load remapping hardware drivers written for prior architecture versions.

Access Method

Type: MEM
(Size: 32 bits) **Offset:** [B:0, D:0, F:0] + 0h

Default: 10h



Bit Range	Default and Access	Field Name (ID): Description
31:8	0h RO	Reserved (RSVD): Reserved.
7:4	1h RO	MAJOR: Indicates supported architecture version.
3:0	0h RO	MINOR: Indicates supported architecture minor version.


8.2 Capability Register (CAP)—Offset 8h

Register to report general remapping hardware capabilities

Access Method

Type: MEM
(Size: 64 bits) **Offset:** [B:0, D:0, F:0] + 8h

Default: 1C0000C40660462h





Bit Range	Default and Access	Field Name (ID): Description
63:59	0h RO	Reserved (RSVD): Reserved.
58	0h RO	SL64KP: A value of 1 in this field indicates 64-KByte page size is supported for second-level translation.
57	0h ROV	FL64KP: A value of 1 in this field indicates 64-KByte page size is supported for first-level translation.
56	1h ROV	FL1GP: A value of 1 in this field indicates 1-GByte page size is supported for first-level translation.
55	1h RO	DRD: 0: Hardware does not support draining of DMA read requests. 1: Hardware supports draining of DMA read requests.
54	1h RO	DWD: 0: Hardware does not support draining of DMA write requests. 1: Hardware supports draining of DMA write requests.
53:48	0h RO	MAMV: The value in this field indicates the maximum supported value for the Address Mask (AM) field in the Invalidation Address register (IVA_REG) and IOTLB Invalidation Descriptor (iotlb_inv_dsc). This field is valid only when the PSI field in Capability register is reported as Set.
47:40	0h RO	NFR: Number of fault recording registers is computed as N+1, where N is the value reported in this field. Implementations should support at least one fault recording register (NFR = 0) for each remapping hardware unit in the platform. The maximum number of fault recording registers per remapping hardware unit is 256.
39	0h RO	PSI: 0: Hardware supports only domain and global invalidates for IOTLB 1: Hardware supports page selective, domain and global invalidates for IOTLB. Hardware implementations reporting this field as set are recommended to support a Maximum Address Mask Value (MAMV) value of at least 9.
38	0h RO	Reserved (RSVD): Reserved.
37:34	3h ROV	SLLPS: This field indicates the super page sizes supported by hardware. A value of 1 in any of these bits indicates the corresponding super-page size is supported. The super-page sizes corresponding to various bit positions within this field are: 0: 21-bit offset to page frame (2MB) 1: 30-bit offset to page frame (1GB) 2: 39-bit offset to page frame (512GB) 3: 48-bit offset to page frame (1TB) Hardware implementations supporting a specific super-page size should support all smaller super-page sizes, i.e. only valid values for this field are 0001b, 0011b, 0111b, 1111b.
33:24	40h RO	FRO: This field specifies the location to the first fault recording register relative to the register base address of this remapping hardware unit. If the register base address is X, and the value reported in this field is Y, the address for the first fault recording register is calculated as X+(16*Y).
23	0h RO	Reserved (RSVD): Reserved.
22	1h RO	ZLR: 0: Indicates the remapping hardware unit blocks (and treats as fault) zero length DMA read requests to write-only pages. 1: Indicates the remapping hardware unit supports zero length DMA read requests to write-only pages. DMA remapping hardware implementations are recommended to report ZLR field as Set.



Bit Range	Default and Access	Field Name (ID): Description
21:16	26h RO	<p>MGAW: This field indicates the maximum DMA virtual addressability supported by remapping hardware. The Maximum Guest Address Width (MGAW) is computed as $(N+1)$, where N is the value reported in this field. For example, a hardware implementation supporting 48-bit MGAW reports a value of 47 (101111b) in this field. If the value in this field is X, untranslated and translated DMA requests to addresses above $2^{(X+1)}-1$ are always blocked by hardware. Translations requests to address above $2^{(X+1)}-1$ from allowed devices return a null Translation Completion Data Entry with R=W=0.</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	0h RO	<p>Reserved (RSVD): Reserved.</p>
12:8	4h RO	<p>SAGAW: This 5-bit field indicates the supported adjusted guest address widths (which in turn represents the levels of page-table walks for the 4KB base page size) supported by the hardware implementation.</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"> 0: 30-bit AGAW (2-level page table) 1: 39-bit AGAW (3-level page table) 2: 48-bit AGAW (4-level page table) 3: 57-bit AGAW (5-level page table) 4: 64-bit AGAW (6-level page table) <p>Software should 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	0h RO	<p>CM:</p> <ul style="list-style-type: none"> 0: Not-present and erroneous entries are not cached in any of the remapping caches. Invalidations are not required for modifications to individual not present or invalid entries. However, any modifications that result in decreasing the effective permissions or partial permission increases require invalidations for them to be effective. 1: Not-present and erroneous mappings may be cached in the remapping caches. Any software updates to the remapping structures (including updates to "not-present" or erroneous entries) require explicit invalidation. <p>Hardware implementations of this architecture should support a value of 0 in this field.</p>
6	1h RO	<p>PHMR:</p> <ul style="list-style-type: none"> 0: Indicates protected high-memory region is not supported. 1: Indicates protected high-memory region is supported.
5	1h RO	<p>PLMR:</p> <ul style="list-style-type: none"> 0: Indicates protected low-memory region is not supported. 1: Indicates protected low-memory region is supported.



Bit Range	Default and Access	Field Name (ID): Description
4	0h RO	RWBF: 0: Indicates no write-buffer flushing is needed to ensure changes to memory-resident structures are visible to hardware. 1: Indicates software should explicitly flush the write buffers to ensure updates made to memory-resident remapping structures are visible to hardware.
3	0h RO	AFL: 0: Indicates advanced fault logging is not supported. Only primary fault logging is supported. 1: Indicates advanced fault logging is supported.
2:0	2h RO	ND: 000b: Hardware supports 4-bit domain-ids with support for up to 16 domains. 001b: Hardware supports 6-bit domain-ids with support for up to 64 domains. 010b: Hardware supports 8-bit domain-ids with support for up to 256 domains. 011b: Hardware supports 10-bit domain-ids with support for up to 1024 domains. 100b: Hardware supports 12-bit domain-ids with support for up to 4K domains. 100b: Hardware supports 14-bit domain-ids with support for up to 16K domains. 110b: Hardware supports 16-bit domain-ids with support for up to 64K domains. 111b: Reserved.

8.3 Extended Capability Register (ECAP)—Offset 10h


Register to report remapping hardware extended capabilities

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 10h

Default: 7E3FF0505Eh



Bit Range	Default and Access	Field Name (ID): Description
63:41	0h RO	Reserved (RSVD): Reserved.
40	1h ROV	PASID: 0: Hardware does not support process address space IDs. 1: Hardware supports Process Address Space IDs.
39:35	Fh RO	PSS: This field reports the PASID size supported by the remapping hardware for requests with- PASID. A value of N in this field indicates hardware supports PASID field of N+1 bits (For example, value of 7 in this field, indicates 8-bit PASIDs are supported). Requests-with-PASID with PASID value beyond the limit specified by this field are treated as error by the remapping hardware. This field is valid only when PASID field is reported as Set.



Bit Range	Default and Access	Field Name (ID): Description
34	1h ROV	EAFS: 0: Hardware does not support the extended-accessed (EA) bit in first-level paging-structure entries. 1: Hardware supports the extendedaccessed (EA) bit in first-level paging-structure entries. This field is valid only when PASID field is reported as Set.
33	1h ROV	NWFS: 0: Hardware ignores the "No Write" (NW) flag in Device-TLB translation requests, and behaves as if NW is always 0. 1: Hardware supports the "No Write" (NW) flag in Device-TLB translation requests. This field is valid only when Device-TLB support (DT) field is reported as Set.
32	0h RO	POT: 0: Hardware does not support PASID-only Translation Type in extended-context-entries 1: Hardware supports PASID-only Translation Type in extended-context-entries
31	0h RO	SRS: 0: H/W does not support requests-with-PASID seeking supervisor privilege 1: H/W supports requests-with-PASID seeking supervisor privilege
30	0h RO	ERS: 0: H/W does not support requests seeking execute permission 1: H/W supports requests seeking execute permission
29	1h ROV	PRS: 0: Hardware does not support Page Requests 1: Hardware supports Page Requests
28	0h RO	Reserved (RSVD): Reserved.
27	1h ROV	DIS: 0: Hardware does not support deferred invalidations of IOTLB and Device-TLB. 1: Hardware supports deferred invalidations of IOTLB and Device-TLB.
26	1h ROV	NEST: 0: Hardware does not support nested translations. 1: Hardware supports nested translations.
25	1h ROV	MTS: 0: Hardware does not support Memory Type 1: Hardware supports Memory Type
24	1h ROV	ECS: 0: Hardware does not support extended-root-entries and Extended Context-Entries 1: Hardware supports extended-root-entries and Extended Context-Entries
23:20	Fh RO	MHMV: The value in this field indicates the maximum supported value for the Handle Mask (HM) field in the interrupt entry cache invalidation descriptor (iec_inv_dsc). This field is valid only when the IR field in Extended Capability register is reported as Set.
19:18	0h RO	Reserved (RSVD): Reserved.
17:8	50h RO	IRO: This field specifies the offset to the IOTLB registers relative to the register base address of this remapping hardware unit. If the register base address is X, and the value reported in this field is Y, the address for the first IOTLB invalidation register is calculated as $X + (16 * Y)$.
7	0h RO	SC: 0: Hardware does not support 1-setting of the SNP field in the page-table entries. 1: Hardware supports the 1-setting of the SNP field in the page-table entries.
6	1h ROV	PT: 0: Hardware does not support pass-through translation type in context entries. 1: Hardware supports pass-through translation type in context entries.



Bit Range	Default and Access	Field Name (ID): Description
5	0h RO	Reserved (RSVD): Reserved.
4	1h ROV	EIM: 0: On Intel®64 platforms, hardware supports only 8-bit APIC-IDs (xAPIC mode). 1: On Intel®64 platforms, hardware supports 32-bit APIC-IDs (x2APIC mode). This field is valid only on Intel®64 platforms reporting Interrupt Remapping support (IR field Set).
3	1h ROV	IR: 0: Hardware does not support interrupt remapping. 1: Hardware supports interrupt remapping. Implementations reporting this field as Set should also support Queued Invalidation (QI).
2	1h ROV	DT: 0: Hardware does not support device-IOTLBs. 1: Hardware supports Device-IOTLBs. Implementations reporting this field as Set should also support Queued Invalidation (QI).
1	1h ROV	QI: 0: Hardware does not support queued invalidations. 1: Hardware supports queued invalidations.
0	0h RO	C: This field indicates if hardware access to the root, context, page-table and interrupt-remap structures are coherent (snooped) or not. 0: Indicates hardware accesses to remapping structures are non-coherent. 1: Indicates hardware accesses to remapping structures are coherent. Hardware access to advanced fault log and invalidation queue are always coherent.

8.4 Global Command Register (GCMD)—Offset 18h

Register to control remapping hardware. If multiple control fields in this register need to be modified, software should serialize the modifications through multiple writes to this register.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 18h

Default: 0h

3	2	2	2	1	1	8	4	0
1	8	4	0	6	2			
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
TE	SRTIP	SFL	EAPL	WBF	QIE	IRE	SIRTP	CFI
							RSVD	



Bit Range	Default and Access	Field Name (ID): Description
31	0h RW_KV	<p>TE: Software writes to this field to request hardware to enable/disable DMA-remapping: 0: Disable DMA remapping 1: Enable DMA remapping Hardware reports the status of the translation enable operation through the TES field in the Global Status register. There may be active DMA requests in the platform when software updates this field. Hardware should 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. Hardware implementations supporting DMA draining should drain any in-flight DMA read/write requests queued within the Root-Complex before completing the translation enable command and reflecting the status of the command through the TES field in the Global Status register. The value returned on a read of this field is undefined.</p>
30	0h WO	<p>SRTP: Software sets this field to set/update the root-entry table pointer used by hardware. The root-entry table pointer is specified through the Root-entry Table Address (RTA_REG) register. Hardware reports the status of the "Set Root Table Pointer" operation through the RTPS field in the Global Status register. The "Set Root Table Pointer" operation should be performed before enabling or re-enabling (after disabling) DMA remapping through the TE field. After a "Set Root Table Pointer" operation, software should globally invalidate the context cache and then globally invalidate of IOTLB. This is required to ensure hardware uses only the remapping structures referenced by the new root table pointer, and not stale cached entries. While DMA remapping hardware is active, software may update the root table pointer through this field. However, to ensure valid in-flight DMA requests are deterministically remapped, software should ensure that the structures referenced by the new root table pointer are programmed to provide the same remapping results as the structures referenced by the previous root-table pointer. Clearing this bit has no effect. The value returned on read of this field is undefined.</p>
29	0h RO	<p>SFL: This field is valid only for implementations supporting advanced fault logging. Software sets this field to request hardware to set/update the fault-log pointer used by hardware. The fault-log pointer is specified through Advanced Fault Log register. Hardware reports the status of the 'Set Fault Log' operation through the FLS field in the Global Status register. The fault log pointer should be set before enabling advanced fault logging (through EAFL field). Once advanced fault logging is enabled, the fault log pointer may be updated through this field while DMA remapping is active. Clearing this bit has no effect. The value returned on read of this field is undefined.</p>
28	0h RO	<p>EAFL: This field is valid only for implementations supporting advanced fault logging. Software writes to this field to request hardware to enable or disable advanced fault logging: 0: Disable advanced fault logging. In this case, translation faults are reported through the Fault Recording registers. 1: Enable use of memory-resident fault log. When enabled, translation faults are recorded in the memory-resident log. The fault log pointer should be set in hardware (through the SFL field) before enabling advanced fault logging. Hardware reports the status of the advanced fault logging enable operation through the AFLS field in the Global Status register. The value returned on read of this field is undefined.</p>
27	0h RO	<p>WBF: This bit is valid only for implementations requiring write buffer flushing. Software sets this field to request that hardware flush the Root-Complex internal write buffers. This is done to ensure any updates to the memory-resident remapping structures are not held in any internal write posting buffers. Hardware reports the status of the write buffer flushing operation through the WBFS field in the Global Status register. Clearing this bit has no effect. The value returned on a read of this field is undefined.</p>



Bit Range	Default and Access	Field Name (ID): Description
26	0h WO	<p>QIE: This field is valid only for implementations supporting queued invalidations. Software writes to this field to enable or disable queued invalidations.</p> <p>0: Disable queued invalidations. 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>The value returned on a read of this field is undefined.</p>
25	0h WO	<p>IRE: 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 should 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 should 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	0h WO	<p>SIRTP: This field is valid only for implementations supporting interrupt-remapping. Software sets this field to set/update the interrupt remapping table pointer used by hardware. The interrupt remapping table pointer is specified through the Interrupt Remapping Table Address (IRTA_REG) register.</p> <p>Hardware reports the status of the 'Set Interrupt Remap Table Pointer' operation through the IRTPS field in the Global Status register.</p> <p>The 'Set Interrupt Remap Table Pointer' operation should be performed before enabling or re-enabling (after disabling) interrupt-remapping hardware through the IRE field.</p> <p>After a 'Set Interrupt Remap Table Pointer' operation, software should 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 should 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>
23	0h WO	<p>CFI: 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 Extended Interrupt Mode (x2APIC mode) is not enabled.</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>
22:0	0h RO	<p>Reserved (RSVD): Reserved.</p>



8.5 Global Status Register (GSTS)—Offset 1Ch

Register to report general remapping hardware status.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 1Ch

Default: 0h

3	2	2	2	1	1	8	4	0
1	8	4	0	6	2			
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
TES	RTPS	FLS	AFLS	WBFS	QIES	IRES	IRTPS	CFIS
							RSVD	

Bit Range	Default and Access	Field Name (ID): Description
31	0h ROV	TES: This field indicates the status of DMA-remapping hardware. 0: DMA-remapping hardware is not enabled 1: DMA-remapping hardware is enabled
30	0h ROV	RTPS: This field indicates the status of the root- table pointer in hardware. This field is cleared by hardware when software sets the SRTTP 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.
29	0h RO	FLS: This field: - Is cleared by hardware when software Sets the SFL field in the Global Command register. - Is Set by hardware when hardware completes the 'Set Fault Log Pointer' operation using the value provided in the Advanced Fault Log register.
28	0h RO	AFLS: This field is valid only for implementations supporting advanced fault logging. It indicates the advanced fault logging status: 0: Advanced Fault Logging is not enabled. 1: Advanced Fault Logging is enabled.
27	0h RO	WBFS: This field is valid only for implementations requiring write buffer flushing. This field indicates the status of the write buffer flush command. It is: - Set by hardware when software sets the WBF field in the Global Command register. - Cleared by hardware when hardware completes the write buffer flushing operation.
26	0h RO_V	QIES: This field indicates queued invalidation enable status. 0: queued invalidation is not enabled 1: queued invalidation is enabled
25	0h ROV	IRES: This field indicates the status of Interrupt-remapping hardware. 0: Interrupt-remapping hardware is not enabled 1: Interrupt-remapping hardware is enabled



Bit Range	Default and Access	Field Name (ID): Description
24	0h RO_V	IRTPS: This field indicates the status of the interrupt remapping table pointer in hardware. This field is cleared by hardware when software sets the SIRTP field in the Global Command register. This field is Set by hardware when hardware completes the set interrupt remap table pointer operation using the value provided in the Interrupt Remapping Table Address register.
23	0h RO_V	CFIS: This field indicates the status of Compatibility format interrupts on Intel®64 implementations supporting interrupt-remapping. The value reported in this field is applicable only when interrupt-remapping is enabled and Extended Interrupt Mode (x2APIC mode) is not enabled. 0: Compatibility format interrupts are blocked. 1: Compatibility format interrupts are processed as pass-through (bypassing interrupt remapping).
22:0	0h RO	Reserved (RSVD): Reserved.

8.6 Root-Entry Table Address Register (RTADDR)—Offset 20h


Register providing the base address of root-entry table.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 20h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.



Bit Range	Default and Access	Field Name (ID): Description
38:12	0h RW	RTA: This register points to base of page aligned, 4KB-sized root-entry table in system memory. Hardware ignores and not implements bits 63:HAW, where HAW is the host address width. Software specifies the base address of the root-entry table through this register, and programs it in hardware through the SRTP field in the Global Command register. Reads of this register returns value that was last programmed to it.
11	0h RW_V	RTT: This field specifies the type of root-table referenced by the Root Table Address (RTA) field; 0: Root Table 1: Extended Root Table
10:0	0h RO	Reserved (RSVD): Reserved.

8.7 Context Command Register (CCMD)—Offset 28h


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.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 28h

Default: 80h





Bit Range	Default and Access	Field Name (ID): Description
63	0h RW_V	<p>ICC: Software requests invalidation of context-cache by setting this field. Software should also set the requested invalidation granularity by programming the CIRG field. Software should read back and check the ICC field is Clear to confirm the invalidation is complete. Software should 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.</p> <p>Software should 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 should 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) should implicitly perform a write buffer flush before invalidating the context cache.</p>
62:61	0h RW	<p>CIRG: 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 should be specified in the DID field.</p> <p>11: Device-selective invalidation request. The target source-id(s) should be specified through the SID and FM fields, and the domain-id (that was programmed in the context-entry for these device(s)) should 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	1h ROV	<p>CAIG: Hardware reports the granularity at which an invalidation request was processed through the CAIG field at the time of reporting invalidation completion (by clearing the ICC field).</p> <p>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	0h RO	<p>Reserved (RSVD): Reserved.</p>
33:32	0h RW	<p>FM: Software may use the Function Mask to perform device-selective invalidations on behalf of devices supporting PCI Express Phantom Functions. This field specifies which bits of the function number portion (least significant three bits) of the SID field to mask when performing device-selective invalidations. The following encodings are defined for this field:</p> <p>00: No bits in the SID field masked.</p> <p>01: Mask most significant bit of function number in the SID field.</p> <p>10: Mask two most significant bit of function number in the SID field.</p> <p>11: Mask all three bits of function number in the SID field.</p> <p>The context-entries corresponding to all the source-ids specified through the FM and SID fields should have to the domain-id specified in the DID field.</p>
31:16	0h RW	<p>SID: Indicates the source-id of the device whose corresponding context-entry needs to be selectively invalidated. This field along with the FM field should be programmed by software for device-selective invalidation requests.</p>



Bit Range	Default and Access	Field Name (ID): Description
15:8	0h RO	Reserved (RSVD): Reserved.
7:0	0h RW	DID: Indicates the id of the domain whose context-entries need to be selectively invalidated. This field should 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 should ensure that the value written to this field is within this limit. Hardware may ignore and not implement bits15:N, where N is the supported domain-id width reported in the Capability register.

8.8 Fault Status Register (FSTS)—Offset 34h


Register indicating the various error status.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 34h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	0h RO	FRI: This field is valid only when the PPF field is Set. The FRI field indicates the index (from base) of the fault recording register to which the first pending fault was recorded when the PPF field was Set by hardware. The value read from this field is undefined when the PPF field is clear.
7	0h RW1CS	PRO: Hardware detected a Page Request Overflow error. Hardware implementations not supporting the Page Request Queue implement this bit as RsvdZ.
6	0h RO	ITE: Hardware detected a Device-IOTLB invalidation completion time-out. At this time, a fault event may be generated based on the programming of the Fault Event Control register. Hardware implementations not supporting device Device-IOTLBs implement this bit as RsvdZ.
5	0h RO	ICE: Hardware received an unexpected or invalid Device-IOTLB invalidation completion. This could be due to either an invalid ITag or invalid source-id in an invalidation completion response. At this time, a fault event may be generated based on the programming of the Fault Event Control register. Hardware implementations not supporting Device-IOTLBs implement this bit as RsvdZ.



Bit Range	Default and Access	Field Name (ID): Description
4	0h RW1CS	IQE: Hardware detected an error associated with the invalidation queue. This could be due to either a hardware error while fetching a descriptor from the invalidation queue, or hardware detecting an erroneous or invalid descriptor in the invalidation queue. At this time, a fault event may be generated based on the programming of the Fault Event Control register. Hardware implementations not supporting queued invalidations implement this bit as RsvdZ.
3	0h RO	APF: When this field is Clear, hardware sets this field when the first fault record (at index 0) is written to a fault log. At this time, a fault event is generated based on the programming of the Fault Event Control register. Software writing 1 to this field clears it. Hardware implementations not supporting advanced fault logging implement this bit as RsvdZ.
2	0h RO	AFO: Hardware sets this field to indicate advanced fault log overflow condition. At this time, a fault event is generated based on the programming of the Fault Event Control register. Software writing 1 to this field clears it. Hardware implementations not supporting advanced fault logging implement this bit as RsvdZ.
1	0h ROSV	PPF: This field indicates if there are one or more pending faults logged in the fault recording registers. Hardware computes this field as the logical OR of Fault (F) fields across all the fault recording registers of this remapping hardware unit. 0: No pending faults in any of the fault recording registers 1: One or more fault recording registers has pending faults. The FRI field is updated by hardware whenever the PPF field is set by hardware. Also, depending on the programming of Fault Event Control register, a fault event is generated when hardware sets this field.
0	0h RW1CS	PFO: Hardware sets this field to indicate overflow of fault recording registers. Software writing 1 clears this field. When this field is Set, hardware does not record any new faults until software clears this field.

8.9 Fault Event Control Register (FECTL)—Offset 38h

Register specifying the fault event interrupt message control bits.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 38h

Default: 80000h

3	2	2	2	1	1	8	4	0	
1	8	4	0	6	2				
1	0	0	0	0	0	0	0	0	
	0	0	0	0	0	0	0	0	
	0	0	0	0	0	0	0	0	
	0	0	0	0	0	0	0	0	
	0	0	0	0	0	0	0	0	
	0	0	0	0	0	0	0	0	
IM	IP	RSVD							



Bit Range	Default and Access	Field Name (ID): Description
31	1h RW	<p>IM: 0: No masking of interrupt. When an interrupt condition is detected, hardware issues an interrupt message (using the Fault Event Data and Fault Event Address register values). 1: This is the value on reset. Software may mask interrupt message generation by setting this field. Hardware is prohibited from sending the interrupt message when this field is set.</p>
30	0h ROV	<p>IP: Hardware sets the IP field whenever it detects an interrupt condition, which is defined as: When primary fault logging is active, an interrupt condition occurs when hardware records a fault through one of the Fault Recording registers and sets the PPF field in Fault Status register. When advanced fault logging is active, an interrupt condition occurs when hardware records a fault in the first fault record (at index 0) of the current fault log and sets the APF field in the Fault Status register. Hardware detected error associated with the Invalidation Queue, setting the IQE field in the Fault Status register. Hardware detected invalid Device-IOTLB invalidation completion, setting the ICE field in the Fault Status register. Hardware detected Device-IOTLB invalidation completion time-out, setting the ITE field in the Fault Status register. If any of the status fields in the Fault Status register was already Set at the time of setting any of these fields, it is not treated as a new interrupt condition. The IP field is kept set by hardware while the interrupt message is held pending. The interrupt message could be held pending due to interrupt mask (IM field) being Set or other transient hardware conditions. The IP field is cleared by hardware as soon as the interrupt message pending condition is serviced. This could be due to either: Hardware issuing the interrupt message due to either change in the transient hardware condition that caused interrupt message to be held pending, or due to software clearing the IM field. Software servicing all the pending interrupt status fields in the Fault Status register as follows:</p> <ul style="list-style-type: none"> • When primary fault logging is active, software clearing the Fault (F) field in all the Fault Recording registers with faults, causing the PPF field in Fault Status register to be evaluated as clear. • Software clearing other status fields in the Fault Status register by writing back the value read from the respective fields.
29:0	0h RO	<p>Reserved (RSVD): Reserved.</p>



8.10 Fault Event Data Register (FEDATA)—Offset 3Ch


Register specifying the interrupt message data

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 3Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RW	EIMD: This field is valid only for implementations supporting 32-bit interrupt data fields. Hardware implementations supporting only 16-bit interrupt data may treat this field as RsvdZ.
15:0	0h RW	IMD: Data value in the interrupt request.

8.11 Fault Event Address Register (FEADDR)—Offset 40h


Register specifying the interrupt message address.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 40h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31:2	0h RW	MA: When fault events are enabled, the contents of this register specify the DWORD-aligned address (bits 31:2) for the interrupt request.
1:0	0h RO	Reserved (RSVD): Reserved.

8.12 Fault Event Upper Address Register (FEUADDR)—Offset 44h


Register specifying the interrupt message upper address.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 44h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:0	0h RW	MUA: Hardware implementations supporting Extended Interrupt Mode are required to implement this register. Hardware implementations not supporting Extended Interrupt Mode may treat this field as RsvdZ.

8.13 Advanced Fault Log Register (AFLOG)—Offset 58h

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).

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 58h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31	0h RW	<p>EPM: This field controls DMA accesses to the protected low-memory and protected high-memory regions.</p> <p>0: Protected memory regions are disabled.</p> <p>1: Protected memory regions are enabled. DMA requests accessing protected memory regions are handled as follows:</p> <ul style="list-style-type: none"> When DMA remapping is not enabled, all DMA requests accessing protected memory regions are blocked. When DMA remapping is enabled: <ul style="list-style-type: none"> DMA requests processed as pass-through (Translation Type value of 10b in Context-Entry) and accessing the protected memory regions are blocked. DMA requests with translated address (AT=10b) and accessing the protected memory regions are blocked. DMA requests that are subject to address remapping, and accessing the protected memory regions may or may not be blocked by hardware. For such requests, software should 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. <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 should 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	0h RO	Reserved (RSVD): Reserved.
0	0h ROV	<p>PRS: This field indicates the status of protected memory region(s):</p> <p>0: Protected memory region(s) disabled.</p> <p>1: Protected memory region(s) enabled.</p>

8.15 Protected Low-Memory Base Register (PLMBASE)—Offset 68h

Register to set up the base address of DMA-protected low-memory region below 4GB. This register should be set up before enabling protected memory through PMEN_REG, and should not be updated when protected memory regions are enabled.

This register is always treated as RO for implementations not supporting protected low memory region (PLMR field reported as Clear in the Capability register).

The alignment of the protected low memory region base depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1s to this register, and finding the most significant zero bit position with 0 in the value read back from the register. Bits N:0 of this register is decoded by hardware as all 0s.

Software should setup the protected low memory region below 4GB.

Software should not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).


Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 68h



Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:20	0h RW	PLMB: This register specifies the base of protected low-memory region in system memory.
19:0	0h RO	Reserved (RSVD): Reserved.

8.16 Protected Low-Memory Limit Register (PLMLIMIT)—Offset 6Ch

Register to set up the limit address of DMA-protected low-memory region below 4GB. This register should be set up before enabling protected memory through PMEN_REG, and should not be updated when protected memory regions are enabled.

This register is always treated as RO for implementations not supporting protected low memory region (PLMR field reported as Clear in the Capability register).

The alignment of the protected low memory region limit depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1's to this register, and finding most significant zero bit position with 0 in the value read back from the register. Bits N:0 of the limit register is decoded by hardware as all 1s.

The Protected low-memory base and limit registers functions as follows:

- Programming the protected low-memory base and limit registers with the same value in bits 31:(N+1) specifies a protected low-memory region of size 2^(N+1) bytes.
- Programming the protected low-memory limit register with a value less than the protected low-memory base register disables the protected low-memory region.


Software should not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 6Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:20	0h RW	PLML: This register specifies the last host physical address of the DMA-protected low-memory region in system memory.
19:0	0h RO	Reserved (RSVD): Reserved.

8.17 Protected High-Memory Base Register (PHMBASE)—Offset 70h

Register to set up the base address of DMA-protected high-memory region. This register should be set up before enabling protected memory through PMEN_REG, and should not be updated when protected memory regions are enabled.

This register is always treated as RO for implementations not supporting protected high memory region (PHMR field reported as Clear in the Capability register).

The alignment of the protected high memory region base depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1's to this register, and finding most significant zero bit position below host address width (HAW) in the value read back from the register. Bits N:0 of this register are decoded by hardware as all 0s.

Software may setup the protected high memory region either above or below 4GB.


Software should not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 70h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:20	0h RW	PHMB: This register specifies the base of protected (high) memory region in system memory. Hardware ignores, and does not implement, bits 63:HAW, where HAW is the host address width.
19:0	0h RO	Reserved (RSVD): Reserved.

8.18 Protected High-Memory Limit Register (PHMLIMIT)—Offset 78h

Register to set up the limit address of DMA-protected high-memory region. This register should be set up before enabling protected memory through PMEN_REG, and should not be updated when protected memory regions are enabled.

This register is always treated as RO for implementations not supporting protected high memory region (PHMR field reported as Clear in the Capability register).

The alignment of the protected high memory region limit depends on the number of reserved bits (N:0) of this register. Software may determine the value of N by writing all 1s 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 should not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 78h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:20	0h RW	PHML: This register specifies the last host physical address of the DMA-protected high-memory region in system memory. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address width.
19:0	0h RO	Reserved (RSVD): Reserved.

8.19 Invalidation Queue Head Register (IQH)—Offset 80h


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.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 80h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:19	0h RO	Reserved (RSVD): Reserved.
18:4	0h ROV	QH: Specifies the offset (128-bit aligned) to the invalidation queue for the command that will be fetched next by hardware. Hardware resets this field to 0 whenever the queued invalidation is disabled (QIES field Clear in the Global Status register).
3:0	0h RO	Reserved (RSVD): Reserved.



8.20 Invalidation Queue Tail Register (IQT)—Offset 88h


Register indicating the invalidation tail head. This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 88h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:19	0h RO	Reserved (RSVD): Reserved.
18:4	0h RW_L	QT: Specifies the offset (128-bit aligned) to the invalidation queue for the command that will be written next by software.
3:0	0h RO	Reserved (RSVD): Reserved.



8.21 Invalidation Queue Address Register (IQA)— Offset 90h


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.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 90h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:12	0h RW_L	IQA: This field points to the base of 4KB aligned invalidation request queue. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address width. Reads of this field return the value that was last programmed to it.
11:3	0h RO	Reserved (RSVD): Reserved.
2:0	0h RW_L	QS: This field specifies the size of the invalidation request queue. A value of X in this field indicates an invalidation request queue of (2 ^X) 4KB pages. The number of entries in the invalidation queue is 2 ^(X + 8) .



8.22 Invalidation Completion Status Register (ICS)— Offset 9Ch

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.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 9Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW1CS	IWC: Indicates completion of Invalidation Wait Descriptor with Interrupt Flag (IF) field Set. Hardware implementations not supporting queued invalidations implement this field as RsvdZ.



8.23 Invalidation Event Control Register (IECTL)— Offset A0h

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.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + A0h

Default: 80000h



Bit Range	Default and Access	Field Name (ID): Description
31	1h RW_L	<p>IM: 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.</p>
30	0h ROV	<p>IP: Hardware sets the IP field whenever it detects an interrupt condition. Interrupt condition is defined as:</p> <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> — 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.
29:0	0h RO	Reserved (RSVD): Reserved.



8.24 Invalidation Event Data Register (IEDATA)— Offset A4h


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.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + A4h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RW_L	EIMD: This field is valid only for implementations supporting 32-bit interrupt data fields. Hardware implementations supporting only 16-bit interrupt data treat this field as Rsvd.
15:0	0h RW_L	IMD: Data value in the interrupt request.

8.25 Invalidation Event Address Register (IEADDR)— Offset A8h

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.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + A8h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:2	0h RW_L	MA: When fault events are enabled, the contents of this register specify the DWORD-aligned address (bits 31:2) for the interrupt request.
1:0	0h RO	Reserved (RSVD): Reserved.

8.26 Invalidation Event Upper Address Register (IEUADDR)—Offset ACh


Register specifying the Invalidation Event interrupt message upper address.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + ACh

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:0	0h RW_L	MUA: Hardware implementations supporting Queued Invalidation and Extended Interrupt Mode are required to implement this register. Hardware implementations not supporting Queued Invalidation or Extended Interrupt Mode may treat this field as RsvdZ.

8.27 Interrupt Remapping Table Address Register (IRTA)—Offset B8h


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.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + B8h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:12	0h RW_L	IRTA: This field points to the base of 4KB aligned interrupt remapping table. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address width. Reads of this field returns value that was last programmed to it.
11	0h ROV	EIME: This field is used by hardware on Intel®64 platforms as follows: 0: xAPIC mode is active. Hardware interprets only low 8-bits of Destination-ID field in the IRTEs. The high 24-bits of the Destination-ID field are treated as reserved. 1: x2APIC mode is active. Hardware interprets all 32-bits of Destination-ID field in the IRTEs. This field is implemented as RsvdZ on implementations reporting Extended Interrupt Mode (EIM) field as Clear in Extended Capability register.
10:4	0h RO	Reserved (RSVD): Reserved.
3:0	0h RW_L	S: This field specifies the size of the interrupt remapping table. The number of entries in the interrupt remapping table is $2^{(X+1)}$, where X is the value programmed in this field.



8.28 Fault Recording Low Register (FRCDL)—Offset 400h

Register to record fault information when primary fault logging is active. Hardware reports the number and location of fault recording registers through the Capability register. This register is relevant only for primary fault logging.


This register is sticky and can be cleared only through power good reset or by software clearing the RW1C fields by writing a 1.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 400h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:12	0h ROSV	FI: When the Fault Reason (FR) field indicates one of the DMA-remapping fault conditions, bits 63:12 of this field contain the page address in the faulted DMA request. Hardware treats bits 63:N as reserved (0), where N is the maximum guest address width (MGAW) supported. When the Fault Reason (FR) field indicates one of the interrupt-remapping fault conditions, bits 63:48 of this field indicate the interrupt_index computed for the faulted interrupt request, and bits 47:12 are cleared. This field is relevant only when the F field is Set.
11:0	0h RO	Reserved (RSVD): Reserved.

8.29 Fault Recording High Register (FRCDH)—Offset 408h

Register to record fault information when primary fault logging is active. Hardware reports the number and location of fault recording registers through the Capability register. This register is relevant only for primary fault logging.

This register is sticky and can be cleared only through power good reset or by software clearing the RW1C fields by writing a 1.


Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 408h



Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63	0h RW1CS	F: Hardware sets this field to indicate a fault is logged in this Fault Recording register. The F field is set by hardware after the details of the fault is recorded in other fields. When this field is Set, hardware may collapse additional faults from the same source-id (SID). Software writes the value read from this field to Clear it.
62	0h ROSV	T: Type of the faulted request: 0: Write request 1: Read request or AtomicOp request This field is relevant only when the F field is Set, and when the fault reason (FR) indicates one of the DMA-remapping fault conditions.
61:60	0h ROV	AT: This field captures the AT field from the faulted DMA request. Hardware implementations not supporting Device-IOTLBs (DI field Clear in Extended Capability register) treat this field as RsvdZ. When supported, this field is valid only when the F field is Set, and when the fault reason (FR) indicates one of the DMA-remapping fault conditions.
59:40	0h ROSV	PN: PASID value in the faulted request. This field is relevant only when the PP field is set. Hardware implementations not supporting PASID (PASID field Clear in Extended Capability register) implement this field as RsvdZ.
39:32	0h ROSV	FR: Reason for the fault. This field is relevant only when the F field is set.
31	0h ROSV	PP: When set, indicates the faulted request has a PASID tag. The value of the PASID field is reported in the PASID Value (PV) field. This field is relevant only when the F field is Set, and when the fault reason (FR) indicates one of the non-recoverable address translation fault conditions. Hardware implementations not supporting PASID (PASID field Clear in Extended Capability register) implement this field as RsvdZ.
30	0h ROSV	EXE: When set, indicates Execute permission was requested by the faulted read request. This field is relevant only when the PP field and T field are both Set. Hardware implementations not supporting PASID (PASID field Clear in Extended Capability register) implement this field as RsvdZ.
29	0h ROSV	PRIV: When set, indicates Supervisor privilege was requested by the faulted request. This field is relevant only when the PP field is Set. Hardware implementations not supporting PASID (PASID field Clear in Extended Capability register) implement this field as RsvdZ.
28:16	0h RO	Reserved (RSVD): Reserved.
15:0	0h ROSV	SID: Requester-id associated with the fault condition. This field is relevant only when the F field is set.



8.30 Invalidate Address Register (IVA)—Offset 500h


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.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 500h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:12	0h RW	ADDR: Software provides the DMA address that needs to be page-selectively invalidated. To make a page-selective invalidation request to hardware, software should first write the appropriate fields in this register, and then issue the appropriate page-selective invalidate command through the IOTLB_REG. Hardware ignores bits 63:N, where N is the maximum guest address width (MGAW) supported.
11:7	0h RO	Reserved (RSVD): Reserved.
6	0h RW	IH: The field provides hint to hardware about preserving or flushing the non-leaf (page-directory) entries that may be cached in hardware: 0: Software may have modified both leaf and non-leaf page-table entries corresponding to mappings specified in the ADDR and AM fields. On a page-selective invalidation request, hardware should 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 mappings specified by ADDR and AM fields.
5:0	0h RW	AM: The value in this field specifies the number of low order bits of the ADDR field that should be masked for the invalidation operation. This field enables software to request invalidation of contiguous mappings for size-aligned regions. For example: Mask ADDR bits Pages Value masked invalidated 0 None 1 1 12 2 2 13:12 4 3 14:12 8 4 15:12 16 When invalidating mappings for super-pages, software should specify the appropriate mask value. For example, when invalidating mapping for a 2MB page, software should specify an address mask value of at least 9. Hardware implementations report the maximum supported mask value through the Capability register.



8.31 IOTLB Invalidate Register (IOTLB)—Offset 508h


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.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 508h

Default: 20h



Bit Range	Default and Access	Field Name (ID): Description
63	0h RW_V	IVT: Software requests IOTLB invalidation by setting this field. Software should 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 should not submit another invalidation request through this register while the IVT field is Set, nor update the associated Invalidate Address register. Software should not submit IOTLB invalidation requests when there is a context-cache invalidation request pending at this remapping hardware unit. Hardware implementations reporting write-buffer flushing requirement (RWBF=1 in Capability register) should implicitly perform a write buffer flushing before invalidating the IOTLB.
62	0h RO	Reserved (RSVD): Reserved.
61:60	0h RW	IIRG: When requesting hardware to invalidate the IOTLB (by setting the IVT field), software writes the requested invalidation granularity through this field. The following are the encodings for the field. 00: Reserved. 01: Global invalidation request. 10: Domain-selective invalidation request. The target domain-id should be specified in the DID field. 11: Page-selective invalidation request. The target address, mask and invalidation hint should be specified in the Invalidate Address register, and the domain-id should be provided in the DID field. Hardware implementations may process an invalidation request by performing invalidation at a coarser granularity than requested. Hardware indicates completion of the invalidation request by clearing the IVT field. At this time, the granularity at which actual invalidation was performed is reported through the IAIG field
59	0h RO	Reserved (RSVD): Reserved.



Bit Range	Default and Access	Field Name (ID): Description
58:57	1h ROV	<p>IAIG: Hardware reports the granularity at which an invalidation request was processed through this field when reporting invalidation completion (by clearing the IVT field). The following are the encodings for this field.</p> <p>00: Reserved. This indicates hardware detected an incorrect invalidation request and ignored the request. Examples of incorrect invalidation requests include detecting an unsupported address mask value in Invalidate Address register for page-selective invalidation requests.</p> <p>01: Global Invalidation performed. This could be in response to a global, domain-selective, or page-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 a page-selective invalidation request.</p> <p>11: Domain-page-selective invalidation performed using the address, mask and hint specified by software in the Invalidate Address register and domain-id specified in DID field. This can be in response to a page-selective invalidation request.</p>
56:50	0h RO	Reserved (RSVD): Reserved.
49	0h RW	<p>DR: This field is ignored by hardware if the DRD field is reported as clear in the Capability register. When the DRD field is reported as Set in the Capability register, the following encodings are supported for this field:</p> <p>0: Hardware may complete the IOTLB invalidation without draining any translated DMA read requests.</p> <p>1: Hardware should drain DMA read requests.</p>
48	0h RW	<p>DW: This field is ignored by hardware if the DWD field is reported as Clear in the Capability register. When the DWD field is reported as Set in the Capability register, the following encodings are supported for this field:</p> <p>0: Hardware may complete the IOTLB invalidation without draining DMA write requests.</p> <p>1: Hardware should drain relevant translated DMA write requests.</p>
47:40	0h RO	Reserved (RSVD): Reserved.
39:32	0h RW	<p>DID: Indicates the ID of the domain whose IOTLB entries need to be selectively invalidated. This field should be programmed by software for domain-selective and page-selective invalidation requests.</p> <p>The Capability register reports the domain-id width supported by hardware. Software should ensure that the value written to this field is within this limit. Hardware ignores and not implements bits 47:(32+N), where N is the supported domain-id width reported in the Capability register.</p>
31:0	0h RO	Reserved (RSVD): Reserved.



8.32 DMA Remap Engine Policy Control (ARCHDIS)— Offset FF0h

This register contains all architectural disables and defeatures for the graphics DMA remap engine.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + FF0h

Default: 1h

3	2	2	2	1	1	8	4	0												
1	8	4	0	6	2															
0	0	0	0	0	0	0	0	1												
DMAR_LCKDN	DMA_RSRV_CTL	RSVD			NWFSCAPDIS	MTSCAPDIS	EAFSCAPDIS	FL64KPCAPCTRL	DTCAPDIS	PASIDCAPDIS	ECSCAPDIS	SCCAPDIS	PTCAPDIS	IRCAPDIS	QICAPDIS	NESTCAPDIS	DISCAPDIS	PRSCAPDIS	FL1GPCAPDIS	SLLPCAPCTRL

Bit Range	Default and Access	Field Name (ID): Description
31	0h RW_KL	DMAR_LCKDN: This register bit protects all the DMA remap engine specific policy configuration registers. Once this bit is set by software all the DMA remap engine registers within the range 0xF00 to 0xFFC will be read-only. This bit can only be clear through platform reset.
30	0h RW_L	DMA_RSRV_CTL: This bit indicates whether Reserved Bit checking is supported or not (i.e. support for Fault Reason 0xA, 0xB, or 0xC). 0: HW supports reserved field checking in root, context and page translation structures. 1: HW ignores reserved field checking in root, context, and page translation structures.
29:16	0h RO	Reserved (RSVD): Reserved.
15	0h RW_L	NWFSCAPDIS: This bit allows hiding the NWFS Capability. 0: ECAP_REG[NWFS] is determined by its own default value. 1: ECAP_REG[NWFS] is set to 0b.
14	0h RW_L	MTSCAPDIS: This bit allows hiding the MTS Capability. 0: ECAP_REG[MTS] is determined by its own default value. 1: ECAP_REG[MTS] is set to 0b.
13	0h RW_L	EAFSCAPDIS: This bit allows hiding the EAFS Capability. 0: ECAP_REG[EAFS] is determined by its own default value. 1: ECAP_REG[EAFS] is set to 0b.
12	0h RW_L	FL64KPCAPCTRL: This bit allows hiding the FL64KP Capability. 0: ECAP_REG[FL64KP] is determined by its own default value. 1: ECAP_REG[FL64KP] is set to 0b.



Bit Range	Default and Access	Field Name (ID): Description
11	0h RW_L	DTCAPDIS: This bit allows hiding the Device TLB Capability. 0: ECAP_REG[DT] is determined by its own default value. 1: ECAP_REG[DT] is set to 0b.
10	0h RW_L	PASIDCAPDIS: This bit allows hiding the PASID Capability. 0: ECAP_REG[PASID] is determined by its own default value. 1: ECAP_REG[PASID] is set to 0b.
9	0h RW_L	ECSCAPDIS: This bit allows hiding the Extended Context Capability. 0: ECAP_REG[ECS] is determined by its own default value. 1: ECAP_REG[ECS] is set to 0b. Additionally hardware will prevent writing of '1' to RTADDR_REG.b[11].
8	0h RO	SCCAPDIS: This bit allows hiding the Snoop Control Capability. 0: ECAP_REG[SC] is determined by its own default value. 1: ECAP_REG[SC] is set to 0b.
7	0h RW_L	PTCAPDIS: This bit allows hiding the Pass Through Capability. 0: ECAP_REG[PT] is determined by its own default value. 1: ECAP_REG[PT] is set to 0b.
6	0h RO_KFW	IRCAPDIS: This bit allows hiding the Interrupt Remapping Capability. 0: ECAP_REG[IR] is determined by its own default value. 1: ECAP_REG[IR] is set to 0b.
5	0h RO_KFW	QICAPDIS: This bit allows hiding the Queued Invalidation Capability. 0: ECAP_REG[QI] is determined by its own default value. 1: ECAP_REG[QI] is set to 0b.
4	0h RW_L	NESTCAPDIS: This bit allows hiding the Nested Translation Capability. 0: CAP_REG[NEST] is determined by its own default value. 1: CAP_REG[NEST] is set to 0b.
3	0h RW_L	DISCAPDIS: This bit allows hiding the Deferred Invalidation Support Capability. 0: CAP_REG[DIS] is determined by its own default value. 1: CAP_REG[DIS] is set to 0b.
2	0h RW_L	PRSCAPDIS: This bit allows hiding the Page Request Capability. 0: CAP_REG[PRS] is determined by its own default value. 1: CAP_REG[PRS] is set to 0b.
1	0h RW_L	FL1GPCAPDIS: This bit allows hiding the First Level 1G Page Capability. 0: CAP_REG[FL1GP] is determined by its own default value. 1: CAP_REG[FL1GP] is set to 0b.
0	1h RW_L	SLLPSCAPCTRL: This bit allows enabling/disabling the Super Page Capability. 0: CAP_REG[SLLPS] is set to 0x0 to disable superpages. 1: CAP_REG[SLLPS] is set to 0x3 to enable superpages. When SLLPSCAPCTRL is set to 0, CAP_REG[SLLPS]=0. If software ignores it and sets up Super Pages then IMPH will generate VT-d fault.



8.33 DMA Remap Engine Policy Control (UARCHDIS)—Offset FF4h

This register contains all micro-architectural disables and defeatures for the graphics DMA remap engine.

Access Method

Type: MEM
(Size: 32 bits) **Offset:** [B:0, D:0, F:0] + FF4h

Default: 100h

3 1	2 8	2 4	2 0	1 6	1 2	8	4	0																								
0	0	0	0	0	1	0	0	0																								
0	0	0	0	0	0	0	0	0																								
RSVD			NO_TLCLKUP_PEND		IQ_COH_DIS		L3_HIT2PEND_DIS		L2_HIT2PEND_DIS		L1_HIT2PEND_DIS		L0_HIT2PEND_DIS		CC_HIT2PEND_DIS		L3DIS		L2DIS		L1DIS		L0DIS		CCDIS		RSVD		GLBIOTLBINV		GLBCTXINV	

Bit Range	Default and Access	Field Name (ID): Description
31:23	0h RO	Reserved (RSVD): Reserved.
22	0h RW_L	NO_TLCLKUP_PEND: When this bit is set, all entries which hit to pending on another requests TLB allocation in the default engine are not allowed to look up peer aperture TLBs for a following graphics walk. They should do all page walks (including root and context) in the Processor Graphics engine.
21	0h RW_L	IQ_COH_DIS: When this bit is set to 1b, read requests from the Invalidation Queue are done in a non-coherent manner (no snoops are generated).
20	1h RW_L	L3_HIT2PEND_DIS: When set, this bit forces a lookup which matches an L3 TLB entry in PEND state to be treated as a miss without allocation.
19	0h RO	L2_HIT2PEND_DIS: When set, this bit forces a lookup which matches an L2 TLB entry in PEND state to be treated as a miss without allocation.
18	0h RW_L	L1_HIT2PEND_DIS: When set, this bit forces a lookup which matches an L1 TLB entry in PEND state to be treated as a miss without allocation.
17	0h RW_L	L0_HIT2PEND_DIS: When set, this bit forces a lookup which matches an L0 TLB entry in PEND state to be treated as a miss without allocation.
16	0h RW_L	CC_HIT2PEND_DIS: When set, this bit forces a lookup which matches a context cache entry in PEND state to be treated as a miss without allocation.
15	0h RW_L	L3DIS: 1: L3 TLB is disabled, and each GPA request that looks up the L3 will result in a miss. 0: Normal mode (default). L3 is enabled.



Bit Range	Default and Access	Field Name (ID): Description
14	0h RO	L2DIS: 1: L2 TLB is disabled, and each GPA request that looks up the L2 will result in a miss. 0: Normal mode (default). L2 is enabled.
13	0h RW_L	L1DIS: 1: L1 TLB is disabled, and each GPA request that looks up the L1 will result in a miss. 0: Normal mode (default). L1 is enabled.
12	0h RW_L	L0DIS: 1: L0 TLB is disabled, and each GPA request that looks up the L0 will result in a miss. 0: Normal mode (default). L0 is enabled.
11	0h RW_L	CCDIS: 1: Context Cache is disabled. Each GPA request results in a miss and will request a root walk. 0: Normal mode (default). Context Cache is enabled.
10:2	0h RO	Reserved (RSVD): Reserved.
1	0h RO	GLBIOTLBINV: This bit controls the IOTLB Invalidation behavior of the DMA remap engine. When this bit is set, any type of IOTLB Invalidation will be promoted to Global IOTLB Invalidation. This promotion applies to both register-based invalidation and queued invalidation.
0	0h RO	GLBCTXINV: This bit controls the Context Invalidation behavior of the DMA remap engine. When this bit is set, any type of Context Invalidation will be promoted to Global Context Invalidation. This promotion applies to both register-based invalidation and queued invalidation.

§ §

9 PXPEPBAR Registers

Table 9-1. Summary of Bus: 0, Device: 0, Function: 0 (MEM)

Offset	Size (Bytes)	Register Name (Register Symbol)	Default Value
14-17h	4	EP VC 0 Resource Control (EPVC0RCTL)—Offset 14h	80000FFh

9.1 EP VC 0 Resource Control (EPVC0RCTL)—Offset 14h

Controls the resources associated with Egress Port Virtual Channel 0.


Access Method

Type: MEM

Offset: [B:0, D:0, F:0] + 14h

(Size: 32 bits)

Default: 80000FFh



Bit Range	Default and Access	Field Name (ID): Description
31	1h RO	VC0E: VC0 Enable: For VC0 this is hardwired to 1 and read only as VC0 can never be disabled.
30:27	0h RO	Reserved (RSVD): Reserved.
26:24	0h RO	VC0ID: VC0 ID: Assigns a VC ID to the VC resource. For VC0 this is hardwired to 0 and read only.
23:20	0h RO	Reserved (RSVD): Reserved.
19:17	0h RW	PAS: Port Arbitration Select: This field configures the VC resource to provide a particular Port Arbitration service. The value of 0h corresponds to the bit position of the only asserted bit in the Port Arbitration Capability field.



Bit Range	Default and Access	Field Name (ID): Description
16:8	0h RO	Reserved (RSVD): Reserved.
7:1	7Fh RW	<p>TCVCOM: TC/VC0 Map: Indicates the TCs (Traffic Classes) that are mapped to the VC resource. Bit locations within this field correspond to TC values. For example, when bit 7 is set in this field, TC7 is mapped to this VC resource. When more than one bit in this field is set, it indicates that multiple TCs are mapped to the VC resource.</p> <p>In order to remove one or more TCs from the TC/VC Map of an enabled VC, software should ensure that no new or outstanding transactions with the TC labels are targeted at the given Link.</p>
0	1h RO	TCOVCOM: TC0/VC0 Map: Traffic Class 0 is always routed to VC0.

§ §



10 VCOPREMAP Registers

Table 10-1. Summary of Bus: 0, Device: 0, Function: 0 (MEM)

Offset	Size (Bytes)	Register Name (Register Symbol)	Default Value
0-3h	4	Version Register (VER)—Offset 0h	10h
8-Fh	8	Capability Register (CAP)—Offset 8h	D2008C40660462h
10-17h	8	Extended Capability Register (ECAP)—Offset 10h	F050DAh
18-1Bh	4	Global Command Register (GCMD)—Offset 18h	0h
1C-1Fh	4	Global Status Register (GSTS)—Offset 1Ch	0h
20-27h	8	Root-Entry Table Address Register (RTADDR)—Offset 20h	0h
28-2Fh	8	Context Command Register (CCMD)—Offset 28h	0h
34-37h	4	Fault Status Register (FSTS)—Offset 34h	0h
38-3Bh	4	Fault Event Control Register (FECTL)—Offset 38h	80000h
3C-3Fh	4	Fault Event Data Register (FEDATA)—Offset 3Ch	0h
40-43h	4	Fault Event Address Register (FEADDR)—Offset 40h	0h
44-47h	4	Fault Event Upper Address Register (FEUADDR)—Offset 44h	0h
58-5Fh	8	Advanced Fault Log Register (AFLOG)—Offset 58h	0h
64-67h	4	Protected Memory Enable Register (PMEN)—Offset 64h	0h
68-6Bh	4	Protected Low-Memory Base Register (PLMBASE)—Offset 68h	0h
6C-6Fh	4	Protected Low-Memory Limit Register (PLMLIMIT)—Offset 6Ch	0h
70-77h	8	Protected High-Memory Base Register (PHMBASE)—Offset 70h	0h
78-7Fh	8	Protected High-Memory Limit Register (PHMLIMIT)—Offset 78h	0h
80-87h	8	Invalidation Queue Head Register (IQH)—Offset 80h	0h
88-8Fh	8	Invalidation Queue Tail Register (IQT)—Offset 88h	0h
90-97h	8	Invalidation Queue Address Register (IQA)—Offset 90h	0h
9C-9Fh	4	Invalidation Completion Status Register (ICS)—Offset 9Ch	0h
A0-A3h	4	Invalidation Event Control Register (IECTL)—Offset A0h	80000h
A4-A7h	4	Invalidation Event Data Register (IEDATA)—Offset A4h	0h
A8-ABh	4	Invalidation Event Address Register (IEADDR)—Offset A8h	0h
AC-AFh	4	Invalidation Event Upper Address Register (IEUADDR)—Offset ACh	0h
B8-BFh	8	Interrupt Remapping Table Address Register (IRTA)—Offset B8h	0h
400-407h	8	Fault Recording Low Register (FRCDL)—Offset 400h	0h
408-40Fh	8	Fault Recording High Register (FRCDH)—Offset 408h	0h
500-507h	8	Invalidate Address Register (IVA)—Offset 500h	0h
508-50Fh	8	IOTLB Invalidate Register (IOTLB)—Offset 508h	0h



10.1 Version Register (VER)—Offset 0h


Register to report the architecture version supported. Backward compatibility for the architecture is maintained with new revision numbers, allowing software to load remapping hardware drivers written for prior architecture versions.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 0h

Default: 10h



Bit Range	Default and Access	Field Name (ID): Description
31:8	0h RO	Reserved (RSVD): Reserved.
7:4	1h RO	MAJOR: Indicates supported architecture version.
3:0	0h RO	MINOR: Indicates supported architecture minor version.

10.2 Capability Register (CAP)—Offset 8h


Register to report general remapping hardware capabilities


Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 8h

Default: D2008C40660462h





Bit Range	Default and Access	Field Name (ID): Description
63:59	0h RO	Reserved (RSVD): Reserved.
58	0h RO	SL64KP: A value of 1 in this field indicates 64-KByte page size is supported for second-level translation.
57	0h RO	FL64KP: A value of 1 in this field indicates 64-KByte page size is supported for first-level translation.
56	0h RO	FL1GP: A value of 1 in this field indicates 1-GByte page size is supported for first-level translation.
55	1h RO	DRD: 0: Hardware does not support draining of DMA read requests. 1: Hardware supports draining of DMA read requests.
54	1h RO	DWD: 0: Hardware does not support draining of DMA write requests. 1: Hardware supports draining of DMA write requests.
53:48	12h RO	MAMV: The value in this field indicates the maximum supported value for the Address Mask (AM) field in the Invalidation Address register (IVA_REG) and IOTLB Invalidation Descriptor (iotlb_inv_dsc). This field is valid only when the PSI field in Capability register is reported as Set.
47:40	0h RO	NFR: Number of fault recording registers is computed as N+1, where N is the value reported in this field. Implementations should support at least one fault recording register (NFR = 0) for each remapping hardware unit in the platform. The maximum number of fault recording registers per remapping hardware unit is 256.
39	1h ROV	PSI: 0: Hardware supports only domain and global invalidates for IOTLB 1: Hardware supports page selective, domain and global invalidates for IOTLB. Hardware implementations reporting this field as set are recommended to support a Maximum Address Mask Value (MAMV) value of at least 9.
38	0h RO	Reserved (RSVD): Reserved.
37:34	3h ROV	SLLPS: This field indicates the super page sizes supported by hardware. A value of 1 in any of these bits indicates the corresponding super-page size is supported. The super-page sizes corresponding to various bit positions within this field are: 0: 21-bit offset to page frame (2MB) 1: 30-bit offset to page frame (1GB) 2: 39-bit offset to page frame (512GB) 3: 48-bit offset to page frame (1TB) Hardware implementations supporting a specific super-page size should support all smaller super-page sizes, i.e. only valid values for this field are 0000b, 0001b, 0011b, 0111b, 1111b.
33:24	40h RO	FRO: This field specifies the location to the first fault recording register relative to the register base address of this remapping hardware unit. If the register base address is X, and the value reported in this field is Y, the address for the first fault recording register is calculated as X+(16*Y).
23	0h RO	Reserved (RSVD): Reserved.
22	1h RO	ZLR: 0: Indicates the remapping hardware unit blocks (and treats as fault) zero length DMA read requests to write-only pages. 1: Indicates the remapping hardware unit supports zero length DMA read requests to write-only pages. DMA remapping hardware implementations are recommended to report ZLR field as Set.



Bit Range	Default and Access	Field Name (ID): Description
21:16	26h RO	<p>MGAW: This field indicates the maximum DMA virtual addressability supported by remapping hardware. The Maximum Guest Address Width (MGAW) is computed as $(N+1)$, where N is the value reported in this field. For example, a hardware implementation supporting 48-bit MGAW reports a value of 47 (101111b) in this field. If the value in this field is X, untranslated and translated DMA requests to addresses above $2^{(X+1)}-1$ are always blocked by hardware. Translations requests to address above $2^{(X+1)}-1$ from allowed devices return a null Translation Completion Data Entry with R=W=0.</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	0h RO	<p>Reserved (RSVD): Reserved.</p>
12:8	4h RO	<p>SAGAW: This 5-bit field indicates the supported adjusted guest address widths (which in turn represents the levels of page-table walks for the 4KB base page size) supported by the hardware implementation.</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"> 0: 30-bit AGAW (2-level page table) 1: 39-bit AGAW (3-level page table) 2: 48-bit AGAW (4-level page table) 3: 57-bit AGAW (5-level page table) 4: 64-bit AGAW (6-level page table) <p>Software should 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	0h RO	<p>CM:</p> <ul style="list-style-type: none"> 0: Not-present and erroneous entries are not cached in any of the remapping caches. Invalidations are not required for modifications to individual not present or invalid entries. However, any modifications that result in decreasing the effective permissions or partial permission increases require invalidations for them to be effective. 1: Not-present and erroneous mappings may be cached in the remapping caches. Any software updates to the remapping structures (including updates to "not-present" or erroneous entries) require explicit invalidation. <p>Hardware implementations of this architecture should support a value of 0 in this field.</p>
6	1h RO	<p>PHMR:</p> <ul style="list-style-type: none"> 0: Indicates protected high-memory region is not supported. 1: Indicates protected high-memory region is supported.
5	1h RO	<p>PLMR:</p> <ul style="list-style-type: none"> 0: Indicates protected low-memory region is not supported. 1: Indicates protected low-memory region is supported.



Bit Range	Default and Access	Field Name (ID): Description
4	0h RO	RWBF: 0: Indicates no write-buffer flushing is needed to ensure changes to memory-resident structures are visible to hardware. 1: Indicates software should explicitly flush the write buffers to ensure updates made to memory-resident remapping structures are visible to hardware.
3	0h RO	AFL: 0: Indicates advanced fault logging is not supported. Only primary fault logging is supported. 1: Indicates advanced fault logging is supported.
2:0	2h RO	ND: 000b: Hardware supports 4-bit domain-ids with support for up to 16 domains. 001b: Hardware supports 6-bit domain-ids with support for up to 64 domains. 010b: Hardware supports 8-bit domain-ids with support for up to 256 domains. 011b: Hardware supports 10-bit domain-ids with support for up to 1024 domains. 100b: Hardware supports 12-bit domain-ids with support for up to 4K domains. 100b: Hardware supports 14-bit domain-ids with support for up to 16K domains. 110b: Hardware supports 16-bit domain-ids with support for up to 64K domains. 111b: Reserved.

10.3 Extended Capability Register (ECAP)—Offset 10h

Register to report remapping hardware extended capabilities

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 10h

Default: F050DAh



Bit Range	Default and Access	Field Name (ID): Description
63:40	0h RO	Reserved (RSVD): Reserved.
39:35	0h RO	PSS: This field reports the PASID size supported by the remapping hardware for requests with- PASID. A value of N in this field indicates hardware supports PASID field of N+1 bits (For example, value of 7 in this field, indicates 8-bit PASIDs are supported). Requests-with-PASID with PASID value beyond the limit specified by this field are treated as error by the remapping hardware. This field is valid only when PASID field is reported as Set.
34	0h RO	EAFS: 0: Hardware does not support the extended-accessed (EA) bit in first-level paging-structure entries. 1: Hardware supports the extendedaccessed (EA) bit in first-level paging-structure entries. This field is valid only when PASID field is reported as Set.



Bit Range	Default and Access	Field Name (ID): Description
33	0h RO	NWFS: 0: Hardware ignores the "No Write" (NW) flag in Device-TLB translation requests, and behaves as if NW is always 0. 1: Hardware supports the "No Write" (NW) flag in Device-TLB translation requests. This field is valid only when Device-TLB support (DT) field is reported as Set.
32	0h RO	POT: 0: Hardware does not support PASID-only Translation Type in extended-context-entries 1: Hardware supports PASID-only Translation Type in extended-context-entries
31	0h RO	SRS: 0: H/W does not support requests-with-PASID seeking supervisor privilege 1: H/W supports requests-with-PASID seeking supervisor privilege
30	0h RO	ERS: 0: H/W does not support requests seeking execute permission 1: H/W supports requests seeking execute permission
29	0h RO	PRS: 0: Hardware does not support Page Requests 1: Hardware supports Page Requests
28	0h RO	PASID: 0: Hardware does not support process address space IDs. 1: Hardware supports Process Address Space IDs.
27	0h RO	DIS: 0: Hardware does not support deferred invalidations of IOTLB and Device-TLB. 1: Hardware supports deferred invalidations of IOTLB and Device-TLB.
26	0h RO	NEST: 0: Hardware does not support nested translations. 1: Hardware supports nested translations.
25	0h RO	MTS: 0: Hardware does not support Memory Type 1: Hardware supports Memory Type
24	0h RO	ECS: 0: Hardware does not support extended-root-entries and Extended Context-Entries 1: Hardware supports extended-root-entries and Extended Context-Entries
23:20	Fh RO	MHMV: The value in this field indicates the maximum supported value for the Handle Mask (HM) field in the interrupt entry cache invalidation descriptor (ie_c_inv_dsc). This field is valid only when the IR field in Extended Capability register is reported as Set.
19:18	0h RO	Reserved (RSVD): Reserved.
17:8	50h RO	IRO: This field specifies the offset to the IOTLB registers relative to the register base address of this remapping hardware unit. If the register base address is X, and the value reported in this field is Y, the address for the first IOTLB invalidation register is calculated as X+(16*Y).
7	1h ROV	SC: 0: Hardware does not support 1-setting of the SNP field in the page-table entries. 1: Hardware supports the 1-setting of the SNP field in the page-table entries.
6	1h ROV	PT: 0: Hardware does not support pass-through translation type in context entries. 1: Hardware supports pass-through translation type in context entries.
5	0h RO	Reserved (RSVD): Reserved.



Bit Range	Default and Access	Field Name (ID): Description
4	1h ROV	EIM: 0: On Intel®64 platforms, hardware supports only 8-bit APIC-IDs (xAPIC mode). 1: On Intel®64 platforms, hardware supports 32-bit APIC-IDs (x2APIC mode). This field is valid only on Intel®64 platforms reporting Interrupt Remapping support (IR field Set).
3	1h ROV	IR: 0: Hardware does not support interrupt remapping. 1: Hardware supports interrupt remapping. Implementations reporting this field as Set should also support Queued Invalidation (QI).
2	0h RO	DT: 0: Hardware does not support device-IOTLBs. 1: Hardware supports Device-IOTLBs. Implementations reporting this field as Set should also support Queued Invalidation (QI).
1	1h ROV	QI: 0: Hardware does not support queued invalidations. 1: Hardware supports queued invalidations.
0	0h RO	C: This field indicates if hardware access to the root, context, page-table and interrupt-remap structures are coherent (snooped) or not. 0: Indicates hardware accesses to remapping structures are non-coherent. 1: Indicates hardware accesses to remapping structures are coherent. Hardware access to advanced fault log and invalidation queue are always coherent.

10.4 Global Command Register (GCMD)—Offset 18h

Register to control remapping hardware. If multiple control fields in this register need to be modified, software should serialize the modifications through multiple writes to this register.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 18h

Default: 0h

3	2	2	2	1	1	8	4	0
1	8	4	0	6	2			
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
TE	SRTP	SFL	EAPL	WBF	QIE	IRE	SIRTP	CFI
							RSVD	



Bit Range	Default and Access	Field Name (ID): Description
31	0h WO	<p>TE: Software writes to this field to request hardware to enable/disable DMA-remapping: 0: Disable DMA remapping 1: Enable DMA remapping Hardware reports the status of the translation enable operation through the TES field in the Global Status register. There may be active DMA requests in the platform when software updates this field. Hardware should 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. Hardware implementations supporting DMA draining should drain any in-flight DMA read/write requests queued within the Root-Complex before completing the translation enable command and reflecting the status of the command through the TES field in the Global Status register. The value returned on a read of this field is undefined.</p>
30	0h WO	<p>SRTP: Software sets this field to set/update the root-entry table pointer used by hardware. The root-entry table pointer is specified through the Root-entry Table Address (RTA_REG) register. Hardware reports the status of the "Set Root Table Pointer" operation through the RTPS field in the Global Status register. The "Set Root Table Pointer" operation should be performed before enabling or re-enabling (after disabling) DMA remapping through the TE field. After a "Set Root Table Pointer" operation, software should globally invalidate the context cache and then globally invalidate of IOTLB. This is required to ensure hardware uses only the remapping structures referenced by the new root table pointer, and not stale cached entries. While DMA remapping hardware is active, software may update the root table pointer through this field. However, to ensure valid in-flight DMA requests are deterministically remapped, software should ensure that the structures referenced by the new root table pointer are programmed to provide the same remapping results as the structures referenced by the previous root-table pointer. Clearing this bit has no effect. The value returned on read of this field is undefined.</p>
29	0h RO	<p>SFL: This field is valid only for implementations supporting advanced fault logging. Software sets this field to request hardware to set/update the fault-log pointer used by hardware. The fault-log pointer is specified through Advanced Fault Log register. Hardware reports the status of the 'Set Fault Log' operation through the FLS field in the Global Status register. The fault log pointer should be set before enabling advanced fault logging (through EAFL field). Once advanced fault logging is enabled, the fault log pointer may be updated through this field while DMA remapping is active. Clearing this bit has no effect. The value returned on read of this field is undefined.</p>
28	0h RO	<p>EAFL: This field is valid only for implementations supporting advanced fault logging. Software writes to this field to request hardware to enable or disable advanced fault logging: 0: Disable advanced fault logging. In this case, translation faults are reported through the Fault Recording registers. 1: Enable use of memory-resident fault log. When enabled, translation faults are recorded in the memory-resident log. The fault log pointer should be set in hardware (through the SFL field) before enabling advanced fault logging. Hardware reports the status of the advanced fault logging enable operation through the AFLS field in the Global Status register. The value returned on read of this field is undefined.</p>
27	0h RO	<p>WBF: This bit is valid only for implementations requiring write buffer flushing. Software sets this field to request that hardware flush the Root-Complex internal write buffers. This is done to ensure any updates to the memory-resident remapping structures are not held in any internal write posting buffers. Hardware reports the status of the write buffer flushing operation through the WBFS field in the Global Status register. Clearing this bit has no effect. The value returned on a read of this field is undefined.</p>



Bit Range	Default and Access	Field Name (ID): Description
26	0h WO	<p>QIE: This field is valid only for implementations supporting queued invalidations. Software writes to this field to enable or disable queued invalidations. 0: Disable queued invalidations. 1: Enable use of queued invalidations. Hardware reports the status of queued invalidation enable operation through QIES field in the Global Status register. The value returned on a read of this field is undefined.</p>
25	0h WO	<p>IRE: This field is valid only for implementations supporting interrupt remapping. 0: Disable interrupt-remapping hardware 1: Enable interrupt-remapping hardware Hardware reports the status of the interrupt remapping enable operation through the IRES field in the Global Status register. There may be active interrupt requests in the platform when software updates this field. Hardware should enable or disable interrupt-remapping logic only at deterministic transaction boundaries, so that any in-flight interrupts are either subject to remapping or not at all. Hardware implementations should drain any in-flight interrupts requests queued in the Root-Complex before completing the interrupt-remapping enable command and reflecting the status of the command through the IRES field in the Global Status register. The value returned on a read of this field is undefined.</p>
24	0h WO	<p>SIRTP: This field is valid only for implementations supporting interrupt-remapping. Software sets this field to set/update the interrupt remapping table pointer used by hardware. The interrupt remapping table pointer is specified through the Interrupt Remapping Table Address (IRTA_REG) register. Hardware reports the status of the 'Set Interrupt Remap Table Pointer' operation through the IRTPS field in the Global Status register. The 'Set Interrupt Remap Table Pointer' operation should be performed before enabling or re-enabling (after disabling) interrupt-remapping hardware through the IRE field. After a 'Set Interrupt Remap Table Pointer' operation, software should globally invalidate the interrupt entry cache. This is required to ensure hardware uses only the interrupt-remapping entries referenced by the new interrupt remap table pointer, and not any stale cached entries. While interrupt remapping is active, software may update the interrupt remapping table pointer through this field. However, to ensure valid in-flight interrupt requests are deterministically remapped, software should ensure that the structures referenced by the new interrupt remap table pointer are programmed to provide the same remapping results as the structures referenced by the previous interrupt remap table pointer. Clearing this bit has no effect. The value returned on a read of this field is undefined.</p>
23	0h WO	<p>CFI: This field is valid only for Intel®64 implementations supporting interrupt-remapping. Software writes to this field to enable or disable Compatibility Format interrupts on Intel®64 platforms. The value in this field is effective only when interrupt-remapping is enabled and Extended Interrupt Mode (x2APIC mode) is not enabled. 0: Block Compatibility format interrupts. 1: Process Compatibility format interrupts as pass-through (bypass interrupt remapping). Hardware reports the status of updating this field through the CFIS field in the Global Status register. The value returned on a read of this field is undefined.</p>
22:0	0h RO	<p>Reserved (RSVD): Reserved.</p>



10.5 Global Status Register (GSTS)—Offset 1Ch

Register to report general remapping hardware status.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 1Ch

Default: 0h

3	2	2	2	1	1	8	4	0
1	8	4	0	6	2			
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0
TES	RTPS	FLS	AFLS	WBFS	QIES	IRES	IRTPS	CFIS
							RSVD	

Bit Range	Default and Access	Field Name (ID): Description
31	0h ROV	TES: This field indicates the status of DMA-remapping hardware. 0: DMA-remapping hardware is not enabled 1: DMA-remapping hardware is enabled
30	0h RO_V	RTPS: This field indicates the status of the root- table pointer in hardware. This field is cleared by hardware when software sets the SRTP field in the Global Command register. This field is set by hardware when hardware completes the 'Set Root Table Pointer' operation using the value provided in the Root-Entry Table Address register.
29	0h RO	FLS: This field: - Is cleared by hardware when software Sets the SFL field in the Global Command register. - Is Set by hardware when hardware completes the 'Set Fault Log Pointer' operation using the value provided in the Advanced Fault Log register.
28	0h RO	AFLS: This field is valid only for implementations supporting advanced fault logging. It indicates the advanced fault logging status: 0: Advanced Fault Logging is not enabled. 1: Advanced Fault Logging is enabled.
27	0h RO	WBFS: This field is valid only for implementations requiring write buffer flushing. This field indicates the status of the write buffer flush command. It is: - Set by hardware when software sets the WBF field in the Global Command register. - Cleared by hardware when hardware completes the write buffer flushing operation.
26	0h RO_V	QIES: This field indicates queued invalidation enable status. 0: queued invalidation is not enabled 1: queued invalidation is enabled
25	0h ROV	IRES: This field indicates the status of Interrupt-remapping hardware. 0: Interrupt-remapping hardware is not enabled 1: Interrupt-remapping hardware is enabled



Bit Range	Default and Access	Field Name (ID): Description
24	0h RO_V	IRTPS: This field indicates the status of the interrupt remapping table pointer in hardware. This field is cleared by hardware when software sets the SIRTP field in the Global Command register. This field is Set by hardware when hardware completes the set interrupt remap table pointer operation using the value provided in the Interrupt Remapping Table Address register.
23	0h RO_V	CFIS: This field indicates the status of Compatibility format interrupts on Intel®64 implementations supporting interrupt-remapping. The value reported in this field is applicable only when interrupt-remapping is enabled and Extended Interrupt Mode (x2APIC mode) is not enabled. 0: Compatibility format interrupts are blocked. 1: Compatibility format interrupts are processed as pass-through (bypassing interrupt remapping).
22:0	0h RO	Reserved (RSVD): Reserved.

10.6 Root-Entry Table Address Register (RTADDR)—Offset 20h


Register providing the base address of root-entry table.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 20h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.



Bit Range	Default and Access	Field Name (ID): Description
38:12	0h RW	RTA: This register points to base of page aligned, 4KB-sized root-entry table in system memory. Hardware ignores and not implements bits 63:HAW, where HAW is the host address width. Software specifies the base address of the root-entry table through this register, and programs it in hardware through the SRTP field in the Global Command register. Reads of this register returns value that was last programmed to it.
11	0h RO	RTT: PLACEHOLDER: This field specifies the type of root-table referenced by the Root Table Address (RTA) field; 0: Root Table 1: Extended Root Table
10:0	0h RO	Reserved (RSVD): Reserved.

10.7 Context Command Register (CCMD)—Offset 28h


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.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 28h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
63	0h RW_V	<p>ICC: Software requests invalidation of context-cache by setting this field. Software should also set the requested invalidation granularity by programming the CIRG field. Software should read back and check the ICC field is Clear to confirm the invalidation is complete. Software should 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.</p> <p>Software should submit a context-cache invalidation request through this field only when there are no invalidation requests pending at this remapping hardware unit.</p> <p>Since information from the context-cache may be used by hardware to tag IOTLB entries, software should 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) should implicitly perform a write buffer flush before invalidating the context cache.</p>
62:61	0h RW	<p>CIRG: 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 should be specified in the DID field.</p> <p>11: Device-selective invalidation request. The target source-id(s) should be specified through the SID and FM fields, and the domain-id (that was programmed in the context-entry for these device(s)) should 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	0h ROV	<p>CAIG: Hardware reports the granularity at which an invalidation request was processed through the CAIG field at the time of reporting invalidation completion (by clearing the ICC field).</p> <p>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	0h RO	<p>Reserved (RSVD): Reserved.</p>
33:32	0h RW	<p>FM: 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. The following encodings are defined for this field:</p> <p>00: No bits in the SID field masked.</p> <p>01: Mask most significant bit of function number in the SID field.</p> <p>10: Mask two most significant bit of function number in the SID field.</p> <p>11: Mask all three bits of function number in the SID field.</p> <p>The context-entries corresponding to all the source-ids specified through the FM and SID fields should have to the domain-id specified in the DID field.</p>
31:16	0h RW	<p>SID: Indicates the source-id of the device whose corresponding context-entry needs to be selectively invalidated. This field along with the FM field should be programmed by software for device-selective invalidation requests.</p>



Bit Range	Default and Access	Field Name (ID): Description
15:8	0h RO	Reserved (RSVD): Reserved.
7:0	0h RW	DID: Indicates the id of the domain whose context-entries need to be selectively invalidated. This field should 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 should ensure that the value written to this field is within this limit. Hardware may ignore and not implement bits15:N, where N is the supported domain-id width reported in the Capability register.

10.8 Fault Status Register (FSTS)—Offset 34h


Register indicating the various error status.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 34h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15:8	0h RO	FRI: This field is valid only when the PPF field is Set. The FRI field indicates the index (from base) of the fault recording register to which the first pending fault was recorded when the PPF field was Set by hardware. The value read from this field is undefined when the PPF field is clear.
7	0h RO	PRO: Hardware detected a Page Request Overflow error. Hardware implementations not supporting the Page Request Queue implement this bit as RsvdZ.
6	0h RO	ITE: Hardware detected a Device-IOTLB invalidation completion time-out. At this time, a fault event may be generated based on the programming of the Fault Event Control register. Hardware implementations not supporting device Device-IOTLBs implement this bit as RsvdZ.
5	0h RO	ICE: Hardware received an unexpected or invalid Device-IOTLB invalidation completion. This could be due to either an invalid ITag or invalid source-id in an invalidation completion response. At this time, a fault event may be generated based on the programming of the Fault Event Control register. Hardware implementations not supporting Device-IOTLBs implement this bit as RsvdZ.



Bit Range	Default and Access	Field Name (ID): Description
4	0h RW1CS	IQE: Hardware detected an error associated with the invalidation queue. This could be due to either a hardware error while fetching a descriptor from the invalidation queue, or hardware detecting an erroneous or invalid descriptor in the invalidation queue. At this time, a fault event may be generated based on the programming of the Fault Event Control register. Hardware implementations not supporting queued invalidations implement this bit as RsvdZ.
3	0h RO	APF: When this field is Clear, hardware sets this field when the first fault record (at index 0) is written to a fault log. At this time, a fault event is generated based on the programming of the Fault Event Control register. Software writing 1 to this field clears it. Hardware implementations not supporting advanced fault logging implement this bit as RsvdZ.
2	0h RO	AFO: Hardware sets this field to indicate advanced fault log overflow condition. At this time, a fault event is generated based on the programming of the Fault Event Control register. Software writing 1 to this field clears it. Hardware implementations not supporting advanced fault logging implement this bit as RsvdZ.
1	0h ROSV	PPF: This field indicates if there are one or more pending faults logged in the fault recording registers. Hardware computes this field as the logical OR of Fault (F) fields across all the fault recording registers of this remapping hardware unit. 0: No pending faults in any of the fault recording registers 1: One or more fault recording registers has pending faults. The FRI field is updated by hardware whenever the PPF field is set by hardware. Also, depending on the programming of Fault Event Control register, a fault event is generated when hardware sets this field.
0	0h RW1CS	PFO: Hardware sets this field to indicate overflow of fault recording registers. Software writing 1 clears this field. When this field is Set, hardware does not record any new faults until software clears this field.

10.9 Fault Event Control Register (FECTL)—Offset 38h

Register specifying the fault event interrupt message control bits.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 38h

Default: 80000h

3	2	2	2	1	1	8	4	0	
1	8	4	0	6	2				
1	0	0	0	0	0	0	0	0	
	0	0	0	0	0	0	0	0	
	0	0	0	0	0	0	0	0	
	0	0	0	0	0	0	0	0	
	0	0	0	0	0	0	0	0	
	0	0	0	0	0	0	0	0	
IM	IP	RSVD							



Bit Range	Default and Access	Field Name (ID): Description
31	1h RW	<p>IM:</p> <p>0: No masking of interrupt. When an interrupt condition is detected, hardware issues an interrupt message (using the Fault Event Data and Fault Event Address register values).</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	0h ROV	<p>IP: Hardware sets the IP field whenever it detects an interrupt condition, which is defined as:</p> <p>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.</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 Fault Status register.</p> <p>Hardware detected error associated with the Invalidation Queue, setting the IQE field in the Fault Status register.</p> <p>Hardware detected invalid Device-IOTLB invalidation completion, setting the ICE field in the Fault Status register.</p> <p>Hardware detected Device-IOTLB invalidation completion time-out, setting the ITE field in the Fault Status register.</p> <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 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. This could be due to either:</p> <p>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>Software servicing all the pending interrupt status fields in the Fault Status register as follows:</p> <ul style="list-style-type: none"> - When primary fault logging is active, software clearing the Fault (F) field in all the Fault Recording registers with faults, causing the PPF field in Fault Status register to be evaluated as clear. - Software clearing other status fields in the Fault Status register by writing back the value read from the respective fields.
29:0	0h RO	Reserved (RSVD): Reserved.



10.10 Fault Event Data Register (FEDATA)—Offset 3Ch


Register specifying the interrupt message data

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 3Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RW	EIMD: This field is valid only for implementations supporting 32-bit interrupt data fields. Hardware implementations supporting only 16-bit interrupt data may treat this field as RsvdZ.
15:0	0h RW	IMD: Data value in the interrupt request.

10.11 Fault Event Address Register (FEADDR)—Offset 40h


Register specifying the interrupt message address.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 40h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31:2	0h RW	MA: When fault events are enabled, the contents of this register specify the DWORD-aligned address (bits 31:2) for the interrupt request.
1:0	0h RO	Reserved (RSVD): Reserved.

10.12 Fault Event Upper Address Register (FEUADDR)—Offset 44h


Register specifying the interrupt message upper address.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 44h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:0	0h RW	MUA: Hardware implementations supporting Extended Interrupt Mode are required to implement this register. Hardware implementations not supporting Extended Interrupt Mode may treat this field as RsvdZ.

10.13 Advanced Fault Log Register (AFLOG)—Offset 58h


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).

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 58h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:12	0h RO	FLA: This field specifies the base of 4KB aligned fault-log region in system memory. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address width. Software specifies the base address and size of the fault log region through this register, and programs it in hardware through the SFL field in the Global Command register. When implemented, reads of this field return the value that was last programmed to it.
11:9	0h RO	FLS: This field specifies the size of the fault log region pointed by the FLA field. The size of the fault log region is $2^X * 4KB$, where X is the value programmed in this register. When implemented, reads of this field return the value that was last programmed to it.
8:0	0h RO	Reserved (RSVD): Reserved.

10.14 Protected Memory Enable Register (PMEN)—Offset 64h

Register to enable the DMA-protected memory regions setup through the PLMBASE, PLMLIMIT, PHMBASE, PHMLIMIT registers. This register is always treated as RO for implementations not supporting protected memory regions (PLMR and PHMR fields reported as Clear in the Capability register).


Protected memory regions may be used by software to securely initialize remapping structures in memory. To avoid impact to legacy BIOS usage of memory, software is recommended to not overlap protected memory regions with any reserved memory regions of the platform reported through the Reserved Memory Region Reporting (RMRR) structures.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 64h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31	0h RW	<p>EPM: 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. DMA requests accessing protected memory regions are handled as follows:</p> <ul style="list-style-type: none"> • When DMA remapping is not enabled, all DMA requests accessing protected memory regions are blocked. • When DMA remapping is enabled: <ul style="list-style-type: none"> – DMA requests processed as pass-through (Translation Type value of 10b in Context-Entry) and accessing the protected memory regions are blocked. – DMA requests with translated address (AT=10b) and accessing the protected memory regions are blocked. – DMA requests that are subject to address remapping, and accessing the protected memory regions may or may not be blocked by hardware. For such requests, software should 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. <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 should 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	0h RO	Reserved (RSVD): Reserved.
0	0h ROV	<p>PRS: This field indicates the status of protected memory region(s):</p> <p>0: Protected memory region(s) disabled. 1: Protected memory region(s) enabled.</p>

10.15 Protected Low-Memory Base Register (PLMBASE)—Offset 68h

Register to set up the base address of DMA-protected low-memory region below 4GB. This register should be set up before enabling protected memory through PMEN_REG, and should not be updated when protected memory regions are enabled.

This register is always treated as RO for implementations not supporting protected low memory region (PLMR field reported as Clear in the Capability register).

The alignment of the protected low memory region base depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1s to this register, and finding the most significant zero bit position with 0 in the value read back from the register. Bits N:0 of this register is decoded by hardware as all 0s.

Software should setup the protected low memory region below 4GB.




Software should not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 68h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:20	0h RW	PLMB: This register specifies the base of protected low-memory region in system memory.
19:0	0h RO	Reserved (RSVD): Reserved.

10.16 Protected Low-Memory Limit Register (PLMLIMIT)—Offset 6Ch

Register to set up the limit address of DMA-protected low-memory region below 4GB. This register should be set up before enabling protected memory through PMEN_REG, and should not be updated when protected memory regions are enabled.

This register is always treated as RO for implementations not supporting protected low memory region (PLMR field reported as Clear in the Capability register).

The alignment of the protected low memory region limit depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1's to this register, and finding most significant zero bit position with 0 in the value read back from the register. Bits N:0 of the limit register is decoded by hardware as all 1s.

The Protected low-memory base and limit registers functions as follows:

- Programming the protected low-memory base and limit registers with the same value in bits 31:(N+1) specifies a protected low-memory region of size 2^(N+1) bytes.
- Programming the protected low-memory limit register with a value less than the protected low-memory base register disables the protected low-memory region.

Software should not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).




Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 6Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:20	0h RW	PLML: This register specifies the last host physical address of the DMA-protected low-memory region in system memory.
19:0	0h RO	Reserved (RSVD): Reserved.

10.17 Protected High-Memory Base Register (PHMBASE)—Offset 70h

Register to set up the base address of DMA-protected high-memory region. This register should be set up before enabling protected memory through PMEN_REG, and should not be updated when protected memory regions are enabled.

This register is always treated as RO for implementations not supporting protected high memory region (PHMR field reported as Clear in the Capability register).

The alignment of the protected high memory region base depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1's to this register, and finding most significant zero bit position below host address width (HAW) in the value read back from the register. Bits N:0 of this register are decoded by hardware as all 0s.

Software may setup the protected high memory region either above or below 4GB.


Software should not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 70h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:20	0h RW	PHMB: This register specifies the base of protected (high) memory region in system memory. Hardware ignores, and does not implement, bits 63:HAW, where HAW is the host address width.
19:0	0h RO	Reserved (RSVD): Reserved.

10.18 Protected High-Memory Limit Register (PHMLIMIT)—Offset 78h

Register to set up the limit address of DMA-protected high-memory region. This register should be set up before enabling protected memory through PMEN_REG, and should not be updated when protected memory regions are enabled.

This register is always treated as RO for implementations not supporting protected high memory region (PHMR field reported as Clear in the Capability register).

The alignment of the protected high memory region limit depends on the number of reserved bits (N:0) of this register. Software may determine the value of N by writing all 1's to this register, and finding most significant zero bit position below host address width (HAW) in the value read back from the register. Bits N:0 of the limit register is decoded by hardware as all 1s.

The protected high-memory base & limit registers functions as follows.

- Programming the protected low-memory base and limit registers with the same value in bits HAW:(N+1) specifies a protected low-memory region of size $2^{(N+1)}$ bytes.
- Programming the protected high-memory limit register with a value less than the protected high-memory base register disables the protected high-memory region.


Software should not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 78h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:20	0h RW	PHML: This register specifies the last host physical address of the DMA-protected high-memory region in system memory. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address width.
19:0	0h RO	Reserved (RSVD): Reserved.

10.19 Invalidation Queue Head Register (IQH)—Offset 80h


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.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 80h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
63:19	0h RO	Reserved (RSVD): Reserved.
18:4	0h ROV	QH: Specifies the offset (128-bit aligned) to the invalidation queue for the command that will be fetched next by hardware. Hardware resets this field to 0 whenever the queued invalidation is disabled (QIES field Clear in the Global Status register).
3:0	0h RO	Reserved (RSVD): Reserved.

10.20 Invalidation Queue Tail Register (IQT)—Offset 88h


Register indicating the invalidation tail head. This register is treated as RsvdZ by implementations reporting Queued Invalidation (QI) as not supported in the Extended Capability register.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 88h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:19	0h RO	Reserved (RSVD): Reserved.
18:4	0h RW_L	QT: Specifies the offset (128-bit aligned) to the invalidation queue for the command that will be written next by software.
3:0	0h RO	Reserved (RSVD): Reserved.



10.21 Invalidation Queue Address Register (IQA)— Offset 90h


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.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 90h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:12	0h RW_L	IQA: This field points to the base of 4KB aligned invalidation request queue. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address width. Reads of this field return the value that was last programmed to it.
11:3	0h RO	Reserved (RSVD): Reserved.
2:0	0h RW_L	QS: This field specifies the size of the invalidation request queue. A value of X in this field indicates an invalidation request queue of (2^X) 4KB pages. The number of entries in the invalidation queue is 2^(X + 8).

10.22 Invalidation Completion Status Register (ICS)— Offset 9Ch

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.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + 9Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:1	0h RO	Reserved (RSVD): Reserved.
0	0h RW1CS	IWC: Indicates completion of Invalidation Wait Descriptor with Interrupt Flag (IF) field Set. Hardware implementations not supporting queued invalidations implement this field as RsvdZ.

10.23 Invalidation Event Control Register (IECTL)—Offset A0h

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.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + A0h

Default: 80000h





Bit Range	Default and Access	Field Name (ID): Description
31	1h RW_L	IM: 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.
30	0h ROV	IP: Hardware sets the IP field whenever it detects an interrupt condition. Interrupt condition is defined as: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 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.
29:0	0h RO	Reserved (RSVD): Reserved.

10.24 Invalidation Event Data Register (IEDATA)—Offset A4h


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.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + A4h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RW_L	EIMD: This field is valid only for implementations supporting 32-bit interrupt data fields. Hardware implementations supporting only 16-bit interrupt data treat this field as Rsvd.
15:0	0h RW_L	IMD: Data value in the interrupt request.

10.25 Invalidation Event Address Register (IEADDR)—Offset A8h


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.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + A8h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:2	0h RW_L	MA: When fault events are enabled, the contents of this register specify the DWORD-aligned address (bits 31:2) for the interrupt request.
1:0	0h RO	Reserved (RSVD): Reserved.



10.26 Invalidation Event Upper Address Register (IEUADDR)—Offset ACh


Register specifying the Invalidation Event interrupt message upper address.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:0, F:0] + ACh

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:0	0h RW_L	MUA: Hardware implementations supporting Queued Invalidations and Extended Interrupt Mode are required to implement this register. Hardware implementations not supporting Queued Invalidations or Extended Interrupt Mode may treat this field as RsvdZ.

10.27 Interrupt Remapping Table Address Register (IRTA)—Offset B8h


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.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + B8h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:12	0h RW_L	IRTA: This field points to the base of 4KB aligned interrupt remapping table. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address width. Reads of this field returns value that was last programmed to it.
11	0h ROV	EIME: This field is used by hardware on Intel®64 platforms as follows: 0: xAPIC mode is active. Hardware interprets only low 8-bits of Destination-ID field in the IRTEs. The high 24-bits of the Destination-ID field are treated as reserved. 1: x2APIC mode is active. Hardware interprets all 32-bits of Destination-ID field in the IRTEs. This field is implemented as RsvdZ on implementations reporting Extended Interrupt Mode (EIM) field as Clear in Extended Capability register.
10:4	0h RO	Reserved (RSVD): Reserved.
3:0	0h RW_L	S: This field specifies the size of the interrupt remapping table. The number of entries in the interrupt remapping table is $2^{(X+1)}$, where X is the value programmed in this field.

10.28 Fault Recording Low Register (FRCDL)—Offset 400h

Register to record fault information when primary fault logging is active. Hardware reports the number and location of fault recording registers through the Capability register. This register is relevant only for primary fault logging.


This register is sticky and can be cleared only through power good reset or by software clearing the RW1C fields by writing a 1.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 400h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
63:12	0h ROSV	FI: When the Fault Reason (FR) field indicates one of the DMA-remapping fault conditions, bits 63:12 of this field contain the page address in the faulted DMA request. Hardware treats bits 63:N as reserved (0), where N is the maximum guest address width (MGAW) supported. When the Fault Reason (FR) field indicates one of the interrupt-remapping fault conditions, bits 63:48 of this field indicate the interrupt_index computed for the faulted interrupt request, and bits 47:12 are cleared. This field is relevant only when the F field is Set.
11:0	0h RO	Reserved (RSVD): Reserved.

10.29 Fault Recording High Register (FRCDH)—Offset 408h


Register to record fault information when primary fault logging is active. Hardware reports the number and location of fault recording registers through the Capability register. This register is relevant only for primary fault logging. This register is sticky and can be cleared only through power good reset or by software clearing the RW1C fields by writing a 1.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 408h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
63	0h RW1CS	F: Hardware sets this field to indicate a fault is logged in this Fault Recording register. The F field is set by hardware after the details of the fault is recorded in other fields. When this field is Set, hardware may collapse additional faults from the same source-id (SID). Software writes the value read from this field to Clear it.
62	0h ROSV	T: Type of the faulted request: 0: Write request 1: Read request or AtomicOp request This field is relevant only when the F field is Set, and when the fault reason (FR) indicates one of the DMA-remapping fault conditions.
61:60	0h RO	AT: This field captures the AT field from the faulted DMA request. Hardware implementations not supporting Device-IOTLBs (DI field Clear in Extended Capability register) treat this field as RsvdZ. When supported, this field is valid only when the F field is Set, and when the fault reason (FR) indicates one of the DMA-remapping fault conditions.
59:40	0h RO	PN: PASID value in the faulted request. This field is relevant only when the PP field is set. Hardware implementations not supporting PASID (PASID field Clear in Extended Capability register) implement this field as RsvdZ.
39:32	0h ROSV	FR: Reason for the fault. This field is relevant only when the F field is set.
31	0h RO	PP: When set, indicates the faulted request has a PASID tag. The value of the PASID field is reported in the PASID Value (PV) field. This field is relevant only when the F field is Set, and when the fault reason (FR) indicates one of the non-recoverable address translation fault conditions. Hardware implementations not supporting PASID (PASID field Clear in Extended Capability register) implement this field as RsvdZ.
30	0h RO	EXE: When set, indicates Execute permission was requested by the faulted read request. This field is relevant only when the PP field and T field are both Set. Hardware implementations not supporting PASID (PASID field Clear in Extended Capability register) implement this field as RsvdZ.
29	0h RO	PRIV: When set, indicates Supervisor privilege was requested by the faulted request. This field is relevant only when the PP field is Set. Hardware implementations not supporting PASID (PASID field Clear in Extended Capability register) implement this field as RsvdZ.
28:16	0h RO	Reserved (RSVD): Reserved.
15:0	0h ROSV	SID: Requester-id associated with the fault condition. This field is relevant only when the F field is set.

10.30 Invalidate Address Register (IVA)—Offset 500h


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.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 500h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:12	0h RW	ADDR: Software provides the DMA address that needs to be page-selectively invalidated. To make a page-selective invalidation request to hardware, software should first write the appropriate fields in this register, and then issue the appropriate page-selective invalidate command through the IOTLB_REG. Hardware ignores bits 63:N, where N is the maximum guest address width (MGAW) supported.
11:7	0h RO	Reserved (RSVD): Reserved.
6	0h RW	IH: The field provides hint to hardware about preserving or flushing the non-leaf (page-directory) entries that may be cached in hardware: 0: Software may have modified both leaf and non-leaf page-table entries corresponding to mappings specified in the ADDR and AM fields. On a page-selective invalidation request, hardware should 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 mappings specified by ADDR and AM fields.
5:0	0h RW	AM: The value in this field specifies the number of low order bits of the ADDR field that should be masked for the invalidation operation. This field enables software to request invalidation of contiguous mappings for size-aligned regions. For example: Mask ADDR bits Pages Value masked invalidated 0 None 1 1 12 2 2 13:12 4 3 14:12 8 4 15:12 16 When invalidating mappings for super-pages, software should specify the appropriate mask value. For example, when invalidating mapping for a 2MB page, software should specify an address mask value of at least 9. Hardware implementations report the maximum supported mask value through the Capability register.

10.31 IOTLB Invalidate Register (IOTLB)—Offset 508h


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.

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:0, F:0] + 508h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63	0h RW_V	IVT: Software requests IOTLB invalidation by setting this field. Software should 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 should not submit another invalidation request through this register while the IVT field is Set, nor update the associated Invalidate Address register. Software should not submit IOTLB invalidation requests when there is a context-cache invalidation request pending at this remapping hardware unit. Hardware implementations reporting write-buffer flushing requirement (RWBF=1 in Capability register) should implicitly perform a write buffer flushing before invalidating the IOTLB.
62	0h RO	Reserved (RSVD): Reserved.
61:60	0h RW	IIRG: When requesting hardware to invalidate the IOTLB (by setting the IVT field), software writes the requested invalidation granularity through this field. The following are the encodings for the field. 00: Reserved. 01: Global invalidation request. 10: Domain-selective invalidation request. The target domain-id should be specified in the DID field. 11: Page-selective invalidation request. The target address, mask and invalidation hint should be specified in the Invalidate Address register, and the domain-id should be provided in the DID field. Hardware implementations may process an invalidation request by performing invalidation at a coarser granularity than requested. Hardware indicates completion of the invalidation request by clearing the IVT field. At this time, the granularity at which actual invalidation was performed is reported through the IAIG field
59	0h RO	Reserved (RSVD): Reserved.
58:57	0h ROV	IAIG: Hardware reports the granularity at which an invalidation request was processed through this field when reporting invalidation completion (by clearing the IVT field). The following are the encodings for this field. 00: Reserved. This indicates hardware detected an incorrect invalidation request and ignored the request. Examples of incorrect invalidation requests include detecting an unsupported address mask value in Invalidate Address register for page-selective invalidation requests. 01: Global Invalidation performed. This could be in response to a global, domain-selective, or page-selective invalidation request. 10: Domain-selective invalidation performed using the domain-id specified by software in the DID field. This could be in response to a domain-selective or a page-selective invalidation request. 11: Domain-page-selective invalidation performed using the address, mask and hint specified by software in the Invalidate Address register and domain-id specified in DID field. This can be in response to a page-selective invalidation request.
56:50	0h RO	Reserved (RSVD): Reserved.



Bit Range	Default and Access	Field Name (ID): Description
49	0h RW	<p>DR: This field is ignored by hardware if the DRD field is reported as clear in the Capability register. When the DRD field is reported as Set in the Capability register, the following encodings are supported for this field:</p> <p>0: Hardware may complete the IOTLB invalidation without draining any translated DMA read requests.</p> <p>1: Hardware should drain DMA read requests.</p>
48	0h RW	<p>DW: This field is ignored by hardware if the DWD field is reported as Clear in the Capability register. When the DWD field is reported as Set in the Capability register, the following encodings are supported for this field:</p> <p>0: Hardware may complete the IOTLB invalidation without draining DMA write requests.</p> <p>1: Hardware should drain relevant translated DMA write requests.</p>
47:40	0h RO	<p>Reserved (RSVD): Reserved.</p>
39:32	0h RW	<p>DID: Indicates the ID of the domain whose IOTLB entries need to be selectively invalidated. This field should be programmed by software for domain-selective and page-selective invalidation requests.</p> <p>The Capability register reports the domain-id width supported by hardware. Software should ensure that the value written to this field is within this limit. Hardware ignores and not implements bits 47:(32+N), where N is the supported domain-id width reported in the Capability register.</p>
31:0	0h RO	<p>Reserved (RSVD): Reserved.</p>

§ §



11 PCI Express* Controllers Registers (D1:F0-2)

This chapter documents the registers of the processor PCIe Gen3 Controller devices. The processor contains multiple PCIe controller devices:

- Bus: 0, Device: 1, Function: 0 (PCI0 x16)
- Bus: 0, Device: 1, Function: 1 (PCI1 x8)
- Bus: 0, Device: 1, Function: 2 (PCI2 x4)

Note: Register default values are taken from device PCI0 only. Consult Volume 1 of this document for Device IDs

Table 11-1. Summary of Bus: 0, Device: 1, Function: 0 (CFG) (Sheet 1 of 3)

Offset	Size (Bytes)	Register Name (Register Symbol)	Default Value
0-1h	2	Vendor Identification (VID)—Offset 0h	8086h
2-3h	2	Device Identification (DID)—Offset 2h	3E81h
4-5h	2	PCI Command (PCICMD)—Offset 4h	0h
6-7h	2	PCI Status (PCISTS)—Offset 6h	10h
8-8h	1	Revision Identification (RID)—Offset 8h	0h
9-Bh	3	Class Code (CC)—Offset 9h	60400h
C-Ch	1	Cache Line Size (CL)—Offset Ch	0h
E-Eh	1	Header Type (HDR)—Offset Eh	81h
18-18h	1	Primary Bus Number (PBUSN)—Offset 18h	0h
19-19h	1	Secondary Bus Number (SBUSN)—Offset 19h	0h
1A-1Ah	1	Subordinate Bus Number (SUBUSN)—Offset 1Ah	0h
1C-1Ch	1	I/O Base Address (IOBASE)—Offset 1Ch	F0h
1D-1Dh	1	I/O Limit Address (IOLIMIT)—Offset 1Dh	0h
1E-1Fh	2	Secondary Status (SSTS)—Offset 1Eh	0h
20-21h	2	Memory Base Address (MBASE)—Offset 20h	FFF0h
22-23h	2	Memory Limit Address (MLIMIT)—Offset 22h	0h
24-25h	2	Prefetchable Memory Base Address (PMBASE)—Offset 24h	FFF1h
26-27h	2	Prefetchable Memory Limit Address (PMLIMIT)—Offset 26h	1h
28-2Bh	4	Prefetchable Memory Base Address Upper (PMBASEU)—Offset 28h	0h
2C-2Fh	4	Prefetchable Memory Limit Address Upper (PMLIMITU)—Offset 2Ch	0h
34-34h	1	Capabilities Pointer (CAPPTR)—Offset 34h	88h
3C-3Ch	1	Interrupt Line (INTRLINE)—Offset 3Ch	0h
3D-3Dh	1	Interrupt Pin (INTRPIN)—Offset 3Dh	1h
3E-3Fh	2	Bridge Control (BCTRL)—Offset 3Eh	0h
80-83h	4	Power Management Capabilities (PM)—Offset 80h	C8039001h
84-87h	4	Power Management Control/Status (PM)—Offset 84h	8h
88-8Bh	4	Subsystem ID and Vendor ID Capabilities (SS)—Offset 88h	800Dh



Table 11-1. Summary of Bus: 0, Device: 1, Function: 0 (CFG) (Continued) (Sheet 2 of 3)

Offset	Size (Bytes)	Register Name (Register Symbol)	Default Value
8C–8Fh	4	Subsystem ID and Subsystem Vendor ID (SS)—Offset 8Ch	8086h
90–91h	2	Message Signaled Interrupts Capability ID (MSI)—Offset 90h	A005h
92–93h	2	Message Control (MC)—Offset 92h	0h
94–97h	4	Message Address (MA)—Offset 94h	0h
98–99h	2	Message Data (MD)—Offset 98h	0h
A0–A1h	2	PCI Express-G Capability List (PEG)—Offset A0h	10h
A2–A3h	2	PCI Express-G Capabilities (PEG)—Offset A2h	142h
A4–A7h	4	Device Capabilities (DCAP)—Offset A4h	8001h
A8–A9h	2	Device Control (DCTL)—Offset A8h	0h
AA–ABh	2	Device Status (DSTS)—Offset AAh	0h
ACh	2	Link Capability (LCAP)—Offset ACh	33486h
B0–B1h	2	Link Control (LCTL)—Offset B0h	0h
B2–B3h	2	Link Status (LSTS)—Offset B2h	1000h
B4–B7h	4	Slot Capabilities (SLOTCAP)—Offset B4h	40h
B8–B9h	2	Slot Control (SLOTCTL)—Offset B8h	0h
BA–BBh	2	Slot Status (SLOTSTS)—Offset BAh	0h
BC–BFh	4	Root Control (RCTL)—Offset BCh	0h
C0–C3h	4	Root Status (RSTS)—Offset C0h	0h
C4–C7h	4	Device Capabilities 2 (DCAP2)—Offset C4h	B80h
C8–C9h	2	Device Control 2 (DCTL2)—Offset C8h	0h
D0–D1h	2	Link Control 2 (LCTL2)—Offset D0h	3h
D2–D3h	2	Link Status 2 (LSTS2)—Offset D2h	0h
104–107h	4	Port VC Capability Register 1 (PVCCAP1)—Offset 104h	0h
108–10Bh	4	Port VC Capability Register 2 (PVCCAP2)—Offset 108h	0h
10C–10Dh	2	Port VC Control (PVCCTL)—Offset 10Ch	0h
110–113h	4	VC0 Resource Capability (VC0RCAP)—Offset 110h	1h
114–117h	4	VC0 Resource Control (VC0RCTL)—Offset 114h	800000FFh
11A–11Bh	2	VC0 Resource Status (VC0RSTS)—Offset 11Ah	2h
1C0h	4	Advanced Error Extended Reporting Capability Header (AECH)—Offset 1C0h	0h
1C4–1C7h	4	PEG Uncorrectable Error Status—Offset 1C4h	0h
1C8–1CBh	4	PEG Uncorrectable Error Mask—Offset 1C8h	0h
1CC–1CFh	4	PEG Uncorrectable Error Severity—Offset 1CCh	0h
1D0–1D3h	4	PEG Correctable Error Status—Offset 1D0h	0h
1D4–1D7h	4	PEG Correctable Error Mask—Offset 1D4h	0h
1D8–1DBh	4	PEG Advanced Error Capabilities and Control—Offset 1D8h	0h
1DC–1DFh	4	PEG Header Log—Offset 1DCh, 1E0h, 1E4h, 1E8h	0h
1E0–1E3h	4		0h
1E4–1E7h	4		0h
1E8–1EBh	4		0h




Table 11-1. Summary of Bus: 0, Device: 1, Function: 0 (CFG) (Continued) (Sheet 3 of 3)

Offset	Size (Bytes)	Register Name (Register Symbol)	Default Value
1EC-1EFh	4	PEG Root Error Command—Offset 1ECh	0h
1F0-1F3h	4	PEG Root Error Status—Offset 1F0h	0h
1F4-1F7h	4	PEG Error Source Identification—Offset 1F4h	0h

11.1 Vendor Identification (VID)—Offset 0h


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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + 0h

Default: 8086h



Bit Range	Default and Access	Field Name (ID): Description
15:0	8086h RO	VID: Vendor Identification: PCI standard identification for Intel.

11.2 Device Identification (DID)—Offset 2h


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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + 2h

Default: 3E81h





Bit Range	Default and Access	Field Name (ID): Description
15:0	3E81h RO	DID_MSB: Device Identification Number MSB: Identifier assigned to the processor root port (virtual PCI-to-PCI bridge, PCI Express Graphics port).

11.3 PCI Command (PCICMD)—Offset 4h

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + 4h

Default: 0h

15	12	8	4	0
0	0	0	0	0
	RSVD	INTAAD	FB2B	SERRE
			RSVD	PERRE
			VGAPS	MWIE
			SCE	BME
			MAE	IOAE

Bit Range	Default and Access	Field Name (ID): Description
15:11	0h RO	Reserved (RSVD): Reserved.
10	0h RW	INTAAD: INTA Assertion Disable: 0: This device is permitted to generate INTA interrupt messages. 1: This device is prevented from generating interrupt messages. Any INTA emulation interrupts already asserted should be de-asserted when this bit is set. Only affects interrupts generated by the device (PCI INTA from a PME or Hot Plug event) controlled by this command register. It does not affect upstream MSIs, upstream PCI INTA-INTD assert and de-assert messages.
9	0h RO	FB2B: Fast Back-to-Back Enable: Not Applicable or Implemented. Hardwired to 0.
8	0h RW	SERRE: SERR# Message Enable: Controls the root port's SERR# messaging. The Processor communicates the SERR# condition by sending an SERR message to the PCH. This bit, when set, enables reporting of non-fatal and fatal errors detected by the device to the Root Complex. Note that errors are reported if enabled either through this bit or through the PCI-Express specific bits in the Device Control Register. In 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. 0: The SERR message is generated by the root port only under conditions enabled individually through the Device Control Register. 1: The root port is enabled to generate SERR messages which will be sent to the PCH for specific root port error conditions generated/detected or received on the secondary side of the virtual PCI to PCI bridge. The status of SERRs generated is reported in the PCISTS register.
7	0h RO	Reserved (RSVD): Reserved.



Bit Range	Default and Access	Field Name (ID): Description
6	0h RW	PERRE: Parity Error Response Enable: Controls whether or not the Master Data Parity Error bit in the PCI Status register can bet set. 0: Master Data Parity Error bit in PCI Status register can NOT be set. 1: Master Data Parity Error bit in PCI Status register CAN be set.
5	0h RO	VGAPS: VGA Palette Snoop: Not Applicable or Implemented. Hardwired to 0.
4	0h RO	MWIE: Memory Write and Invalidate Enable: Not Applicable or Implemented. Hardwired to 0.
3	0h RO	SCE: Special Cycle Enable: Not Applicable or Implemented. Hardwired to 0.
2	0h RW	BME: Bus Master Enable: Bus Master Enable (BME): Controls the ability of the PEG port to forward Memory Read/Write Requests in the upstream direction. 0: This device is prevented from making memory requests to its primary bus. Note that according to PCI Specification, as MSI interrupt messages are in-band memory writes, disabling the bus master enable bit prevents this device from generating MSI interrupt messages or passing them from its secondary bus to its primary bus. Upstream memory writes/reads, peer writes/reads, and MSIs will all be treated as illegal cycles. Writes are aborted. Reads are aborted and will return Unsupported Request status (or Master abort) in its completion packet. 1: This device is allowed to issue requests to its primary bus. Completions for previously issued memory read requests on the primary bus will be issued when the data is available. This bit does not affect forwarding of Completions from the primary interface to the secondary interface.
1	0h RW	MAE: Memory Access Enable: 0: All of device's memory space is disabled. 1: Enable the Memory and Pre-fetchable memory address ranges defined in the MBASE, MLIMIT, PMBASE, and PMLIMIT registers.
0	0h RW	IOAE: IO Access Enable: 0: All of device's I/O space is disabled. 1: Enable the I/O address range defined in the IOBASE, and IOLIMIT registers.

11.4 PCI Status (PCISTS)—Offset 6h

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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + 6h

Default: 10h

15		12		8		4		0
0	0	0	0	0	0	0	0	0
DPE	SSE	RMAS	RTAS	STAS	DEVT	PMDPE	FB2B	RSVD
							CAP66	CAPL
								INTAS
								RSVD



Bit Range	Default and Access	Field Name (ID): Description
15	0h RW1C	DPE: Detected Parity Error: This bit is Set by a Function whenever it receives a Poisoned TLP, regardless of the state the Parity Error Response bit in the Command register. On a Function with a Type 1 Configuration header, the bit is Set when the Poisoned TLP is received by its Primary Side. Default value of this bit is 0b. This bit will be set only for completions of requests encountering ECC error in DRAM. Poisoned Peer 2 peer posted forwarded will not set this bit. They are reported at the receiving port.
14	0h RW1C	SSE: Signaled System Error: This bit is set when this Device sends an SERR due to detecting an ERR_FATAL or ERR_NONFATAL condition and the SERR Enable bit in the Command register is '1'. Both received (if enabled by BCTRL1[1]) and internally detected error messages do not affect this field.
13	0h RO	RMAS: Received Master Abort Status: This bit is Set when a Requester receives a Completion with Unsupported Request Completion Status. On a Function with a Type 1 Configuration header, the bit is Set when the Unsupported Request is received by its Primary Side. Not applicable. We do not have UR on primary interface
12	0h RO	RTAS: Received Target Abort Status: This bit is Set when a Requester receives a Completion with Completer Abort Completion Status. On a Function with a Type 1 Configuration header, the bit is Set when the Completer Abort is received by its Primary Side. Default value of this bit is 0b. Not Applicable or Implemented. Hardwired to 0. The concept of a Completer abort does not exist on primary side of this device.
11	0h RO	STAS: Signaled Target Abort Status: This bit is Set when a Function completes a Posted or Non- Posted Request as a Completer Abort error. This applies to a Function with a Type 1 Configuration header when the Completer Abort was generated by its Primary Side. Default value of this bit is 0b. Not Applicable or Implemented. Hardwired to 0. The concept of a target abort does not exist on primary side of this device.
10:9	0h RO	DEVT: DEVSELB Timing: This device is not the subtractively decoded device on bus 0. This bit field is therefore hardwired to 00 to indicate that the device uses the fastest possible decode. Does not apply to PCI Express and should be hardwired to 00b.
8	0h RW1C	PMDPE: Master Data Parity Error: This bit is Set by a Requester (Primary Side for Type 1 Configuration Space header Function) if the Parity Error Response bit in the Command register is 1b and either of the following two conditions occur <ul style="list-style-type: none"> Requester receives a Completion marked poisoned Requester poisons a write Request If the Parity Error Response bit is 0b, this bit is never Set. Default value of this bit is 0b. This bit will be set only for completions of requests encountering ECC error in DRAM. Poisoned Peer 2 peer posted forwarded will not set this bit. They are reported at the receiving port.
7	0h RO	FB2B: Fast Back-to-Back: Not Applicable or Implemented. Hardwired to 0.
6	0h RO	Reserved (RSVD): Reserved.
5	0h RO	CAP66: 66/60MHz capability: Not Applicable or Implemented. Hardwired to 0.



Bit Range	Default and Access	Field Name (ID): Description
4	1h RO	CAPL: Capabilities List: Indicates that a capabilities list is present. Hardwired to 1.
3	0h ROV	INTAS: INTx Status: Indicates that an interrupt message is pending internally to the device. Only PME and Hot Plug sources feed into this status bit (not PCI INTA-INTD assert and de-assert 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	0h RO	Reserved (RSVD): Reserved.

11.5 Revision Identification (RID)—Offset 8h

This register contains the revision number of Device #1.


These bits are read only and writes to this register have no effect.

Access Method


Type: CFG
(Size: 8 bits)

Offset: [B:0, D:1, F:0] + 8h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
7:4	0h RO	RID_MSB: Revision Identification Number MSB: This is an 8-bit value that indicates the revision identification number for the root port.
3:0	0h RO	RID: Revision Identification Number: This is an 8-bit value that indicates the revision identification number for the root port.



11.6 Class Code (CC)—Offset 9h


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

Access Method

Type: CFG
(Size: 24 bits)

Offset: [B:0, D:1, F:0] + 9h

Default: 60400h



Bit Range	Default and Access	Field Name (ID): Description
23:16	6h RO	BCC: Base Class Code: Indicates the base class code for this device. This code has the value 06h, indicating a Bridge device.
15:8	4h RO	SUBCC: Sub-Class Code: Indicates the sub-class code for this device. The code is 04h indicating a PCI to PCI Bridge.
7:0	0h RO	PI: Programming Interface: Indicates the programming interface of this device. This value does not specify a particular register set layout and provides no practical use for this device.


11.7 Cache Line Size (CL)—Offset Ch

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:1, F:0] + Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
7:0	0h RW	CLS: Cache Line Size: Implemented by PCI Express devices as a read-write field for legacy compatibility purposes but has no impact on any PCI Express device functionality.



11.8 Header Type (HDR)—Offset Eh


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

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:1, F:0] + Eh

Default: 81h



Bit Range	Default and Access	Field Name (ID): Description
7:0	81h RO	HDR: Header Type Register: Device #1 returns 81 to indicate that this is a multi function device with bridge header layout. Device #6 returns 01 to indicate that this is a single function device with bridge header layout.

11.9 Primary Bus Number (PBUSN)—Offset 18h


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

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:1, F:0] + 18h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
7:0	0h RO	BUSN: Primary Bus Number: Configuration software typically programs this field with the number of the bus on the primary side of the bridge. Since the Processor root port is an internal device and its primary bus is always 0, these bits are read only and are hardwired to 0.



11.10 Secondary Bus Number (SBUSN)—Offset 19h


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.

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:1, F:0] + 19h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
7:0	0h RW	BUSN: Secondary Bus Number: This field is programmed by configuration software with the bus number assigned to PCI Express-G.

11.11 Subordinate Bus Number (SUBUSN)—Offset 1Ah


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.

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:1, F:0] + 1Ah

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
7:0	0h RW	BUSN: Subordinate Bus Number: This register is programmed by configuration software with the number of the highest subordinate bus that lies behind the processor root port bridge. When only a single PCI device resides on the PCI Express-G segment, this register will contain the same value as the SBUSN1 register.



11.12 I/O Base Address (IOBASE)—Offset 1Ch

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

$$IO_BASE = \< address = \< IO_LIMIT$$

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

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:1, F:0] + 1Ch

Default: F0h

7	4	0
1	1	0
1	1	0
1	1	0
IOBASE		RSVD

Bit Range	Default and Access	Field Name (ID): Description
7:4	Fh RW	IOBASE: I/O Address Base: Corresponds to A[15:12] of the I/O addresses passed by the root port to PCI Express-G.
3:0	0h RO	Reserved (RSVD): Reserved.

11.13 I/O Limit Address (IOLIMIT)—Offset 1Dh

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

$$IO_BASE = \< address = \< IO_LIMIT$$

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

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:1, F:0] + 1Dh

Default: 0h



7	4	0
0 0 0 0	0 0 0 0	0 0
IOLIMIT		RSVD

Bit Range	Default and Access	Field Name (ID): Description
7:4	0h RW	IOLIMIT: I/O Address Limit: Corresponds to A[15:12] of the I/O address limit of the root port. Devices between this upper limit and IOBASE1 will be passed to the PCI Express hierarchy associated with this device.
3:0	0h RO	Reserved (RSVD): Reserved.

11.14 Secondary Status (SSTS)—Offset 1Eh

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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + 1Eh

Default: 0h

15	12	8	4	0
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0
DPE	RSE RMA RTA	STA DEVT	SMDPE FB2B RSVD CAP66	RSVD

Bit Range	Default and Access	Field Name (ID): Description
15	0h RW1C	DPE: Detected Parity Error: This bit is set by the Secondary Side for a Type 1 Configuration Space header device whenever it receives a Poisoned TLP, regardless of the state of the Parity Error Response Enable bit in the Bridge Control Register.
14	0h RW1C	RSE: Received System Error: This bit is set when the Secondary Side for a Type 1 configuration space header device receives an ERR_FATAL or ERR_NONFATAL.
13	0h RW1C	RMA: Received Master Abort: This bit is set when the Secondary Side for Type 1 Configuration Space Header Device (for requests initiated by the Type 1 Header Device itself) receives a Completion with Unsupported Request Completion Status.
12	0h RW1C	RTA: Received Target Abort: This bit is set when the Secondary Side for Type 1 Configuration Space Header Device (for requests initiated by the Type 1 Header Device itself) receives a Completion with Completer Abort Completion Status.



Bit Range	Default and Access	Field Name (ID): Description
11	0h RO	STA: Signaled Target Abort: Not Applicable or Implemented. Hardwired to 0. The Processor does not generate Target Aborts (The root port will never complete a request using the Completer Abort Completion status). UR detected inside the Processor (such as in IMPH/MC will be reported in primary side status)
10:9	0h RO	DEVT: DEVSELB Timing: Not Applicable or Implemented. Hardwired to 0.
8	0h RW1C	SMDPE: Master Data Parity Error: When set indicates that the Processor received across the link (upstream) a Read Data Completion Poisoned TLP (EP=1). This bit can only be set when the Parity Error Enable bit in the Bridge Control register is set.
7	0h RO	FB2B: Fast Back-to-Back: Not Applicable or Implemented. Hardwired to 0.
6	0h RO	Reserved (RSVD): Reserved.
5	0h RO	CAP66: 66/60 MHz capability: Not Applicable or Implemented. Hardwired to 0.
4:0	0h RO	Reserved (RSVD): Reserved.

11.15 Memory Base Address (MBASE)—Offset 20h

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

$$\text{MEMORY_BASE} = \text{address} \ll \text{MEMORY_LIMIT}$$

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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + 20h

Default: FFF0h

15	12	8	4	0
1	1	1	1	0
1	1	1	1	0
1	1	1	1	0
MBASE				RSVD



Bit Range	Default and Access	Field Name (ID): Description
15:4	FFFh RW	MBASE: Memory Address Base: Corresponds to A[31:20] of the lower limit of the memory range that will be passed to PCI Express-G.
3:0	0h RO	Reserved (RSVD): Reserved.

11.16 Memory Limit Address (MLIMIT)—Offset 22h

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

$$\text{MEMORY_BASE} = \text{address} \&\text{lt;}; \text{MEMORY_LIMIT}$$

The upper 12 bits of the register are read/write and correspond to the upper 12 address bits A[31:20] of the 32 bit address. The bottom 4 bits of this register are read-only and return zeros when read. This register should 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 will be at the top of a 1MB aligned memory block. NOTE: Memory range covered by MBASE and MLIMIT registers are used to map non-prefetchable PCI Express-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.

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + 22h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
15:4	0h RW	MLIMIT: Memory Address Limit: Corresponds to A[31:20] of the upper limit of the address range passed to PCI Express-G.
3:0	0h RO	Reserved (RSVD): Reserved.

11.17 Prefetchable Memory Base Address (PMBASE)—Offset 24h

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

$$\text{PREFETCHABLE_MEMORY_BASE} = \&\text{lt; address} = \&\text{lt; PREFETCHABLE_MEMORY_LIMIT}$$

The upper 12 bits of this register are read/write and correspond to address bits A[31:20] of the 40-bit address. The lower 8 bits of the Upper Base Address register are read/write and correspond to address bits A[39:32] of the 40-bit address. This register should be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be 0. Thus, the bottom of the defined memory address range will be aligned to a 1MB boundary.

Access Method


Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + 24h

Default: FFF1h

15	12	8	4	0
1 1 1 1	1 1 1 1	1 1 1 1	1 1 1 1	0 0 0 1
PMBASE				AS64

Bit Range	Default and Access	Field Name (ID): Description
15:4	FFFh RW	PMBASE: Prefetchable Memory Base Address: Corresponds to A[31:20] of the lower limit of the memory range that will be passed to PCI Express-G.
3:0	1h RO	AS64: 64-bit Address Support: Indicates that the upper 32 bits of the prefetchable memory region base address are contained in the Prefetchable Memory base Upper Address register at 28h.



11.18 Prefetchable Memory Limit Address (PMLIMIT)—Offset 26h

This register in conjunction with the corresponding Upper Limit Address register controls the Processor 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 should 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 will be at the top of a 1MB aligned memory block. Note that prefetchable memory range is supported to allow segregation by the configuration software between the memory ranges that should be defined as UC and the ones that can be designated as a USWC (i.e. prefetchable) from the Processor perspective.

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + 26h

Default: 1h

15	12	8	4	0
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 1
PMLIMIT				AS64B

Bit Range	Default and Access	Field Name (ID): Description
15:4	0h RW	PMLIMIT: Prefetchable Memory Address Limit: Corresponds to A[31:20] of the upper limit of the address range passed to PCI Express-G.
3:0	1h RO	AS64B: 64-bit Address Support: Indicates that the upper 32 bits of the prefetchable memory region limit address are contained in the Prefetchable Memory Base Limit Address register at 2Ch

11.19 Prefetchable Memory Base Address Upper (PMBASEU)—Offset 28h

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

This register in conjunction with the corresponding Upper Base Address register controls the Processor to PCI Express-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 39-bit address. The lower 7 bits of the Upper Base Address register are read/write and correspond to address bits A[38:32] of the 39-bit address. This register should be initialized by the configuration software. For the purpose of address decode address bits A[19:0] are assumed to be 0. Thus, the bottom of the defined memory address range will be aligned to a 1MB boundary.

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:1, F:0] + 28h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:0	0h RW	PMBASEU: Prefetchable Memory Base Address: Corresponds to A[63:32] of the lower limit of the prefetchable memory range that will be passed to PCI Express-G.

11.20 Prefetchable Memory Limit Address Upper (PMLIMITU)—Offset 2Ch

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

This register in conjunction with the corresponding Upper Limit Address register controls the Processor to PCI Express-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 39-bit address. The lower 7 bits of the Upper Limit Address register are read/write and correspond to address bits A[39:32] of the 39-bit address. This register should 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 will be at the top of a 1MB aligned memory block.

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




Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:1, F:0] + 2Ch

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:0	0h RW	PMLIMITU: Prefetchable Memory Address Limit: Corresponds to A[63:32] of the upper limit of the prefetchable Memory range that will be passed to PCI Express-G.

11.21 Capabilities Pointer (CAPPTR)—Offset 34h


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

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:1, F:0] + 34h

Default: 88h



Bit Range	Default and Access	Field Name (ID): Description
7:0	88h RO	CAPPTR1: First Capability: The first capability in the list is the Subsystem ID and Subsystem Vendor ID Capability.



11.22 Interrupt Line (INTRLINE)—Offset 3Ch

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.

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:1, F:0] + 3Ch

Default: 0h

7	4	0
0	0	0
INTCON		

Bit Range	Default and Access	Field Name (ID): Description
7:0	0h RW	INTCON: Interrupt Connection: Used to communicate interrupt line routing information. BIOS Requirement: POST software writes the routing information into this register as it initializes and configures the system. The value indicates to which input of the system interrupt controller this device's interrupt pin is connected.

11.23 Interrupt Pin (INTRPIN)—Offset 3Dh

This register specifies which interrupt pin this device uses.

Access Method

Type: CFG
(Size: 8 bits)

Offset: [B:0, D:1, F:0] + 3Dh

Default: 1h

7	4	0
0	0	1
INTPINH		INTPINL



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

11.24 Bridge Control (BCTRL)—Offset 3Eh

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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + 3Eh

Default: 0h

15	12	8	4	0									
0	0	0	0	0									
RSVD		DTSERRE	DTSTS	SDT	PDT	FB2BEN	SRESET	MAMODE	VGA16D	VGAEN	ISAEN	SERREN	PEREN

Bit Range	Default and Access	Field Name (ID): Description
15:12	0h RO	Reserved (RSVD): Reserved.
11	0h RO	DTSERRE: Discard Timer SERR# Enable: Not Applicable or Implemented. Hardwired to 0.
10	0h RO	DTSTS: Discard Timer Status: Not Applicable or Implemented. Hardwired to 0.
9	0h RO	SDT: Secondary Discard Timer: Not Applicable or Implemented. Hardwired to 0.
8	0h RO	PDT: Primary Discard Timer: Not Applicable or Implemented. Hardwired to 0.



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15	0h RO	PMESTS: PME Status: Indicates that this device does not support PMEB generation from D3cold.
14:13	0h RO	DSCALE: Data Scale: Indicates that this device does not support the power management data register.
12:9	0h RO	DSEL: Data Select: Indicates that this device does not support the power management data register.
8	0h RW	PMEE: PME Enable: Indicates that this device does not generate PMEB assertion from any D-state. 0: PMEB generation not possible from any D State 1: PMEB generation enabled from any D State The setting of this bit has no effect on hardware. Refer PM_CAP[15:11]
7:4	0h RO	Reserved (RSVD): Reserved.
3	1h RO	NSR: No Soft Reset: No Soft Reset. When set to 1 this bit indicates that the device is transitioning from D3hot to D0 because the power state commands do not perform 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 uninitialized with only PME context preserved if PME is supported and enabled.
2	0h RO	Reserved (RSVD): Reserved.
1:0	0h RO_V	PS: Power State: Indicates the current power state of this device and can be used to set the device into a new power state. If software attempts to write an unsupported state to this field, write operation should complete normally on the bus, but the data is discarded and no state change occurs. 00: D0 01: D1 (Not supported in this device.) 10: D2 (Not supported in this device.) 11: D3 Support of D3cold does not require any special action. While in the D3hot state, this device can only act as the target of PCI configuration transactions (for power management control). This device also cannot generate interrupts or respond to MMR cycles in the D3 state. The device should return to the D0 state in order to be fully-functional. When the Power State is other than D0, the bridge will Master Abort (i.e. not claim) any downstream cycles (with exception of type 0 config cycles). Consequently, these unclaimed cycles will go down DMI and come back up as Unsupported Requests, which the Processor logs as Master Aborts in Device 0 PCISTS[13] There is no additional hardware functionality required to support these Power States.

11.27 Subsystem ID and Vendor ID Capabilities (SS)—Offset 88h

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 the operating system will test for its presence.


Access Method



Type: CFG
(Size: 32 bits)

Offset: [B:0, D:1, F:0] + 88h

Default: 800Dh



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

11.28 Subsystem ID and Subsystem Vendor ID (SS)– Offset 8Ch


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

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:1, F:0] + 8Ch

Default: 8086h





Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RW_O	SSID: Subsystem ID: Identifies the particular subsystem and is assigned by the vendor.
15:0	8086h RW_O	SSVID: Subsystem Vendor ID: Identifies the manufacturer of the subsystem and is the same as the vendor ID which is assigned by the PCI Special Interest Group.

11.29 Message Signaled Interrupts Capability ID (MSI)—Offset 90h

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.

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + 90h

Default: A005h

15	12	8	4	0
1 0 1 0	0 0 0 0	0 0 0 0	0 1 0 1	
PNC				CID

Bit Range	Default and Access	Field Name (ID): Description
15:8	A0h RO	PNC: Pointer to Next Capability: This contains a pointer to the next item in the capabilities list which is the PCI Express capability.
7:0	5h RO	CID: Capability ID: Value of 05h identifies this linked list item (capability structure) as being for MSI registers.

11.30 Message Control (MC)—Offset 92h

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 should be serviced, the device should not generate the same message again until the driver services the earlier one.


Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + 92h



Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
15:8	0h RO	Reserved (RSVD): Reserved.
7	0h RO	B64AC: 64-bit Address Capable: Hardwired to 0 to indicate that the function does not implement the upper 32 bits of the Message Address register and is incapable of generating a 64-bit memory address. This may need to change in future implementations when addressable system memory exceeds the 32b/4GB limit.
6:4	0h RW	MME: Multiple Message Enable: System software programs this field to indicate the actual number of messages allocated to this device. This number will be equal to or less than the number actually requested. The encoding is the same as for the MMC field below.
3:1	0h RO	MMC: Multiple Message Capable: System software reads this field to determine the number of messages being requested by this device. Value: Number of Messages Requested 000: 1 All of the following are reserved in this implementation: 001: 2 010: 4 011: 8 100: 16 101: 32 110: Reserved 111: Reserved
0	0h RW	MSIEN: MSI Enable: Controls the ability of this device to generate MSIs. 0: MSI will not be generated. 1: MSI will be generated when we receive PME messages. INTA will not be generated and INTA Status (PCISTS1[3]) will not be set.


11.31 Message Address (MA)—Offset 94h

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:1, F:0] + 94h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31:2	0h RW	MA: Message Address: Used by system software to assign an MSI address to the device. The device handles an MSI by writing the padded contents of the MD register to this address.
1:0	0h RO	FDWA: Force DWord Align: Hardwired to 0 so that addresses assigned by system software are always aligned on a DWord address boundary.

11.32 Message Data (MD)—Offset 98h

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + 98h

Default: 0h

15	12	8	4	0
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0
MD				

Bit Range	Default and Access	Field Name (ID): Description
15:0	0h RW	MD: Message Data: Base message data pattern assigned by system software and used to handle an MSI from the device. When the device should 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.

11.33 PCI Express-G Capability List (PEG)—Offset A0h

Enumerates the PCI Express capability structure.

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + A0h

Default: 10h

15	12	8	4	0
0 0 0 0	0 0 0 0	0 0 0 0	0 1 0 0	0 0 0 0
PNC			CID	



Bit Range	Default and Access	Field Name (ID): Description
15:8	0h RO	PNC: Pointer to Next Capability: This value terminates the capabilities list. The Virtual Channel capability and any other PCI Express specific capabilities that are reported via this mechanism are in a separate capabilities list located entirely within PCI Express Extended Configuration Space.
7:0	10h RO	CID: Capability ID: Identifies this linked list item (capability structure) as being for PCI Express registers.

11.34 PCI Express-G Capabilities (PEG)—Offset A2h

Indicates PCI Express device capabilities.

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + A2h

Default: 142h

15	12	8	4	0
0 0 0 0	0 0 0	1	0 1 0 0	0 0 1 0
RSVD	IMN	SI	DPT	PCIECV

Bit Range	Default and Access	Field Name (ID): Description
15:14	0h RO	Reserved (RSVD): Reserved.
13:9	0h RO	IMN: Interrupt Message Number: Not Applicable or Implemented. Hardwired to 0.
8	1h RW_O	SI: Slot Implemented: 0: The PCI Express Link associated with this port is connected to an integrated component or is disabled. 1: The PCI Express Link associated with this port is connected to a slot. BIOS Requirement: This field should be initialized appropriately if a slot connection is not implemented.
7:4	4h RO	DPT: Device/Port Type: Hardwired to 4h to indicate root port of PCI Express Root Complex.
3:0	2h RO	PCIECV: PCI Express Capability Version: PCI Express Capability Version (PCIECV): Hardwired to 2h to indicate compliance to the PCI Express Capabilities Register Expansion ECN.



11.35 Device Capabilities (DCAP)—Offset A4h


Indicates PCI Express device capabilities.

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:1, F:0] + A4h

Default: 8001h



Bit Range	Default and Access	Field Name (ID): Description
31:16	0h RO	Reserved (RSVD): Reserved.
15	1h RO	RBER: Role Based Error Reporting: Role Based Error Reporting (RBER): Indicates that this device implements the functionality defined in the Error Reporting ECN as required by the PCI Express 1.1 spec.
14:6	0h RO	Reserved (RSVD): Reserved.
5	0h RO	ETFS: Extended Tag Field Supported: Hardwired to indicate support for 5-bit Tags as a Requestor.
4:3	0h RO	PFS: Phantom Functions Supported: Not Applicable or Implemented. Hardwired to 0.
2:0	1h RW_O	MPS: Max Payload Size: Default indicates 256B max supported payload for Transaction Layer Packets (TLP).

11.36 Device Control (DCTL)—Offset A8h

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.

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + A8h

Default: 0h



15		12		8		4		0			
0	0	0	0	0	0	0	0	0			
RSVD		MRRS	NSE	RSVD		MPS	ROE	URRE	FERE	NERE	CERE

Bit Range	Default and Access	Field Name (ID): Description
15	0h RO	Reserved (RSVD): Reserved.
14:12	0h RO	MRRS: Reserved for Max Read Request Size:
11	0h RO	NSE: Reserved for Enable No Snoop:
10:8	0h RO	Reserved (RSVD): Reserved.
7:5	0h RW	MPS: Max Payload Size: 001:256B max supported payload for Transaction Layer Packets (TLP). As a receiver, the Device should handle TLPs as large as the set value; as transmitter, the Device should not generate TLPs exceeding the set value. BIOS should not set this field larger than the DCAP.MPS of the DSD.
4	0h RO	ROE: Reserved for Enable Relaxed Ordering:
3	0h RW	URRE: Unsupported Request Reporting Enable: Unsupported Request Reporting Enable (URRE): When set, allows signaling ERR_NONFATAL, ERR_FATAL, or ERR_CORR to the Root Control register when detecting an unmasked Unsupported Request (UR). An ERR_CORR is signaled when an unmasked Advisory Non-Fatal UR is received. An ERR_FATAL or ERR_NONFATAL is sent to the Root Control register when an uncorrectable non-Advisory UR is received with the severity bit set in the Uncorrectable Error Severity register.
2	0h RW	FERE: Fatal Error Reporting Enable: Fatal Error Reporting Enable (FERE): When set, enables signaling of ERR_FATAL to the Root Control register due to internally detected errors or error messages received across the link. Other bits also control the full scope of related error reporting.
1	0h RW	NERE: Non-Fatal Error Reporting Enable: Non-Fatal Error Reporting Enable (NERE): 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	0h RW	CERE: Correctable Error Reporting Enable: Correctable Error Reporting Enable (CERE): When set, enables signaling of ERR_CORR to the Root Control register due to internally detected errors or error messages received across the link. Other bits also control the full scope of related error reporting.



11.37 Device Status (DSTS)—Offset AAh

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.

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + AAh

Default: 0h

15	12	8	4	0					
0	0	0	0	0					
RSVD				TP	RSVD	URD	FED	NFED	CED

Bit Range	Default and Access	Field Name (ID): Description
15:6	0h RO	Reserved (RSVD): Reserved.
5	0h RO	TP: Transactions Pending: 0: All pending transactions (including completions for any outstanding non-posted requests on any used virtual channel) have been completed. 1: Indicates that the device has transaction(s) pending (including completions for any outstanding non-posted requests for all used Traffic Classes). Not Applicable or Implemented. Hardwired to 0.
4	0h RO	Reserved (RSVD): Reserved.
3	0h RW1C	URD: Unsupported Request Detected: When set this bit indicates that the Device received an Unsupported Request. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control Register. Additionally, the Non-Fatal Error Detected bit or the Fatal Error Detected bit is set according to the setting of the Unsupported Request Error Severity bit. In production systems setting the Fatal Error Detected bit is not an option as support for AER will not be reported.
2	0h RW1C	FED: Fatal Error Detected: When set this bit indicates that fatal error(s) were detected. Errors are logged in this register regardless of whether error reporting is enabled or not in 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	0h RW1C	NFED: Non-Fatal Error Detected: When set this bit indicates that non-fatal error(s) were detected. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. When Advanced Error Handling is enabled, errors are logged in this register regardless of the settings of the uncorrectable error mask register.
0	0h RW1C	CED: Correctable Error Detected: When set this bit indicates that correctable error(s) were detected. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. When Advanced Error Handling is enabled, errors are logged in this register regardless of the settings of the correctable error mask register.



11.38 Link Capability (LCAP)—Offset ACh

Indicates PCI link capabilities.

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:1, F:0] + ACh

Default: 33486h

3 1	2 8	2 4	2 0	1 6	1 2	8	4	0
0 0 0 0	0 0 0 0	0 0 0 0	0 0 1 1	0 0 1 1	0 1 0 0	1 0 0 0	0 1 1 0	0
RSVD		ASPM Optionality Compliance	RSVD	L1 Exit Latency	L0s Exit Latency	Active State Link PM Support	MLW	MLS

Bit Range	Default and Access	Field Name (ID): Description
31:23	0h RO	Reserved (RSVD): Reserved.
22	0h RO	ASPM Optionality Compliance: This bit should be set to 1b in all Functions. Components implemented against certain earlier versions of this specification will have this bit set to 0b. Software is permitted to use the value of this bit to help determine whether to enable ASPM or whether to run ASPM compliance tests.
21:18	0h RO	Reserved (RSVD): Reserved.
17:15	3h RW_O	L1 Exit Latency: Indicates the length of time this Port requires to complete the transition from L1 to L0. The value 010 b indicates the range of 2 us to less than 4 us. Both bytes of this register that contain a portion of this field should be written simultaneously in order to prevent an intermediate (and undesired) value from ever existing.
14:12	4h RO	L0s Exit Latency: Indicates the length of time this Port requires to complete the transition from L0s to L0. 000: Less than 64 ns 001: 64ns to less than 128ns 010: 128ns to less than 256 ns 011: 256ns to less than 512ns 100: 512ns to less than 1us 101: 1 us to less than 2 us 110: 2 us - 4 us 111: More than 4 us



Bit Range	Default and Access	Field Name (ID): Description
11:10	3h RW_O	Active State Link PM Support: Root port supports ASPM L0s and L1.
9:4	10h RW_OV	Max Link Width (MLW): Indicates the maximum number of lanes supported for this link.
3:0	3h RW_OV	Max Link Speed (MLS): The encoding is the binary value of the bit location in the Supported Link Speeds Vector (in the Link Capabilities 2 register) that corresponds to the maximum Link speed. For example, a value of 0010b in this field indicates that the maximum Link speed is that corresponding to bit 2 in the Supported Link Speeds Vector, which is 5.0 GT/s.

11.39 Link Control (LCTL)—Offset B0h

Allows control of PCI Express link.

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + B0h

Default: 0h

15		12		8		4		0					
0	0	0	0	0	0	0	0	0					
	RSVD		LABIE	LBMIE	HAWD	ECPM	ES	CCC	RL	LD	RCB	RSVD	ASPM

Bit Range	Default and Access	Field Name (ID): Description
15:12	0h RO	Reserved (RSVD): Reserved.
11	0h RW	LABIE: Link Autonomous Bandwidth Interrupt Enable: Link Autonomous Bandwidth Interrupt Enable - When Set, this bit enables the generation of an interrupt to indicate that the Link Autonomous Bandwidth Status bit has been Set. This bit is not applicable and is reserved for Endpoint devices, PCI Express to PCI/PCI-X bridges, and Upstream Ports of Switches. Devices that do not implement the Link Bandwidth Notification capability should hardwire this bit to 0b.
10	0h RW	LBMIE: Link Bandwidth Management Interrupt Enable: Link Bandwidth Management Interrupt Enable - When Set, this bit enables the generation of an interrupt to indicate that the Link Bandwidth Management Status bit has been Set. This bit is not applicable and is reserved for Endpoint devices, PCI Express to PCI/PCI-X bridges, and Upstream Ports of Switches.
9	0h RO	HAWD: Hardware Autonomous Width Disable: Hardware Autonomous Width Disable - When Set, this bit disables hardware from changing the Link width for reasons other than attempting to correct unreliable Link operation by reducing Link width. Devices that do not implement the ability autonomously to change Link width are permitted to hardwire this bit to 0b.



Bit Range	Default and Access	Field Name (ID): Description
8	0h RO	<p>ECPM: Enable Clock Power Management: Applicable only for form factors that support a "Clock Request" (CLKREQ#) mechanism, this enable functions as follows</p> <p>0: Clock power management is disabled and device should hold CLKREQ# signal low</p> <p>1: 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.</p> <p>Default value of this field is 0b.</p> <p>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) should hardwire this bit to 0b.</p>
7	0h RW	<p>ES: Extended Synch: Extended synch</p> <p>0: Standard Fast Training Sequence (FTS).</p> <p>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.</p> <p>This is a test mode only and may cause other undesired side effects such as buffer overflows or underruns.</p>
6	0h RW	<p>CCC: Common Clock Configuration:</p> <p>0: Indicates that this component and the component at the opposite end of this Link are operating with asynchronous reference clock.</p> <p>1: Indicates that this component and the component at the opposite end of this Link are operating with a distributed common reference clock.</p> <p>The state of this bit affects the L0s Exit Latency reported in LCAP[14:12] and the N_FTS value advertised during link training. Refer PEGLOSLAT at offset 22Ch.</p>
5	0h RW	<p>RL: Retrain Link:</p> <p>0: Normal operation.</p> <p>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	0h RW	<p>LD: Link Disable:</p> <p>0: Normal operation</p> <p>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	0h RO	<p>RCB: Read Completion Boundary: Hardwired to 0 to indicate 64 byte.</p>
2	0h RO	<p>Reserved (RSVD): Reserved.</p>
1:0	0h RO	<p>ASPM: Active State PM: Controls the level of active state power management supported on the given link.</p> <p>00: Disabled</p> <p>01: L0s Entry Supported</p> <p>10: L1 Entry Supported</p> <p>11: L0s and L1 Entry Supported</p>



11.40 Link Status (LSTS)—Offset B2h

Indicates PCI Express link status.

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + B2h

Default: 1000h

15		12		8		4		0
0	0	0	1	0	0	0	0	0
LABWS	LBWMS	DLLLA	SCC	LTRN	RSVD	NLW	CLS	

Bit Range	Default and Access	Field Name (ID): Description
15	0h RW1C	LABWS: Link Autonomous Bandwidth Status: This bit is set to 1b by hardware to indicate that hardware has autonomously changed link speed or width, without the port transitioning through DL_Down status, for reasons other than to attempt to correct unreliable link operation. This bit should 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.
14	0h RW1C	LBWMS: Link Bandwidth Management Status: This bit is set to 1b by hardware to indicate that either of the following has occurred without the port transitioning through DL_Down status: A link retraining initiated by a write of 1b to the Retrain Link bit has completed. Note: This bit is Set following any write of 1b to the Retrain Link bit, including when the Link is in the process of retraining for some other reason. Hardware has autonomously changed link speed or width to attempt to correct unreliable link operation, either through an LTSSM timeout or a higher level process This bit should 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.
13	0h ROV	DLLLA: Data Link Layer Link Active (Optional): This bit indicates the status of the Data Link Control and Management State Machine. It returns a 1b to indicate the DL_Active state, 0b otherwise. This bit should be implemented if the corresponding Data Link Layer Active Capability bit is implemented. Otherwise, this bit should be hardwired to 0b.
12	1h RO	SCC: Slot Clock Configuration: 0: The device uses an independent clock irrespective of the presence of a reference on the connector. 1: The device uses the same physical reference clock that the platform provides on the connector.
11	0h RO	LTRN: Link Training: Indicates that the Physical Layer LTSSM is in the Configuration or Recovery state, or that 1b was written to the Retrain Link bit but Link training has not yet begun. Hardware clears this bit when the LTSSM exits the Configuration/Recovery state once Link training is complete.



Bit Range	Default and Access	Field Name (ID): Description
10	0h RO	Reserved (RSVD): Reserved.
9:4	0h RO	NLW: Negotiated Link Width: Indicates negotiated link width. This field is valid only when the link is in the L0, L0s, or L1 states (after link width negotiation is successfully completed). 00h: Reserved 01h: X1 02h: X2 04h: X4 08h: X8 10h: X16 All other encodings are reserved.
3:0	0h RO	CLS: Current Link Speed: This field indicates the negotiated Link speed of the given PCI Express Link. The encoding is the binary value of the bit location in the Supported Link Speeds Vector (in the Link Capabilities 2 register) that corresponds to the current Link speed. For example, a value of 0010b in this field indicates that the current Link speed is that corresponding to bit 2 in the Supported Link Speeds Vector, which is 5.0 GT/s. The value in this field is undefined when the Link is not up.

11.41 Slot Capabilities (SLOTCAP)—Offset B4h


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

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:1, F:0] + B4h

Default: 40h



Bit Range	Default and Access	Field Name (ID): Description
31:19	0h RW_O	PSN: Physical Slot Number: Indicates the physical slot number attached to this Port. BIOS Requirement: This field should be initialized by BIOS to a value that assigns a slot number that is globally unique within the chassis.
18	1h RO	NCCS: No Command Completed Support: When set to 1b, this bit indicates that this slot does not generate software notification when an issued command is completed by the Hot-Plug Controller. This bit is only permitted to be set to 1b if the hotplug capable port is able to accept writes to all fields of the Slot Control register without delay between successive writes.
17	0h RO	EIP: Reserved for Electromechanical Interlock Present: When set to 1b, this bit indicates that an Electromechanical Interlock is implemented on the chassis for this slot.



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

11.42 Slot Control (SLOTCTL)—Offset B8h

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

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + B8h

Default: 0h

15	12	8	4	0
0 0 0	0 0 0	0 0 0	0 0 0	0 0 0
RSVD	DLLSCE	EIC PCC PIC	AIC HIPIE CCI	PDCE MSCE PFDE ABPE



Bit Range	Default and Access	Field Name (ID): Description
15:13	0h RO	Reserved (RSVD): Reserved.
12	0h RO	DLLSCE: Reserved for Data Link Layer State Changed Enable: Reserved for Data Link Layer State Changed Enable (DLLSCE): If the Data Link Layer Link Active capability is implemented, when set to 1b, this field enables software notification when Data Link Layer Link Active field is changed. If the Data Link Layer Link Active capability is not implemented, this bit is permitted to be read-only with a value of 0b.
11	0h RO	EIC: Reserved for Electromechanical Interlock Control: If an Electromechanical Interlock is implemented, a write of 1b to this field causes the state of the interlock to toggle. A write of 0b to this field has no effect. A read to this register always returns a 0.
10	0h RO	PCC: Reserved for Power Controller Control: If a Power Controller is implemented, this field when written sets the power state of the slot per the defined encodings. Reads of this field should reflect the value from the latest write, even if the corresponding hotplug command is not complete, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined. Depending on the form factor, the power is turned on/off either to the slot or within the adapter. Note that in some cases the power controller may autonomously remove slot power or not respond to a power-up request based on a detected fault condition, independent of the Power Controller Control setting. The defined encodings are: 0: Power On 1: Power Off If the Power Controller Implemented field in the Slot Capabilities register is set to 0b, then writes to this field have no effect and the read value of this field is undefined.
9:8	0h RO	PIC: Reserved Power Indicator Control: Reserved Power Indicator Control (PIC): If a Power Indicator is implemented, writes to this field set the Power Indicator to the written state. Reads of this field should reflect the value from the latest write, even if the corresponding hot-plug command is not complete, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined. 00:Reserved 01:On 10:Blink 11:Off If the Power Indicator Present bit in the Slot Capabilities register is 0b, this field is permitted to be read-only with a value of 00b.
7:6	0h RO	AIC: Reserved for Attention Indicator Control: Reserved for Attention Indicator Control (AIC): If an Attention Indicator is implemented, writes to this field set the Attention Indicator to the written state. Reads of this field should reflect the value from the latest write, even if the corresponding hot-plug command is not complete, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined. If the indicator is electrically controlled by chassis, the indicator is controlled directly by the downstream port through implementation specific mechanisms. 00:Reserved 01:On 10:Blink 11:Off If the Attention Indicator Present bit in the Slot Capabilities register is 0b, this field is permitted to be read only with a value of 00b.
5	0h RO	HPIE: Reserved for Hot-plug Interrupt Enable: When set to 1b, this bit enables generation of an interrupt on enabled hot-plug events Default value of this field is 0b. If the Hot Plug Capable field in the Slot Capabilities register is set to 0b, this bit is permitted to be read-only with a value of 0b.



Bit Range	Default and Access	Field Name (ID): Description
4	0h RO	CCI: Reserved for Command Completed Interrupt Enable: If Command Completed notification is supported (as indicated by No Command Completed Support field of Slot Capabilities Register), when set to 1b, this bit enables software notification when a hot-plug command is completed by the Hot-Plug Controller. Default value of this field is 0b. If Command Completed notification is not supported, this bit should be hardwired to 0b.
3	0h RO	PDCE: Presence Detect Changed Enable: When set to 1b, this bit enables software notification on a presence detect changed event.
2	0h RO	MSCE: Reserved for MRL Sensor Changed Enable: When set to 1b, this bit enables software notification on a MRL sensor changed event. Default value of this field is 0b. If the MRL Sensor Present field in the Slot Capabilities register is set to 0b, this bit is permitted to be read-only with a value of 0b.
1	0h RO	PFDE: Reserved for Power Fault Detected Enable: When set to 1b, this bit enables software notification on a power fault event. Default value of this field is 0b. If Power Fault detection is not supported, this bit is permitted to be read-only with a value of 0b
0	0h RO	ABPE: Reserved for Attention Button Pressed Enable: When set to 1b, this bit enables software notification on an attention button pressed event.

11.43 Slot Status (SLOTSTS)—Offset BAh

PCI Express Slot related registers.


Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + BAh

Default: 0h

15	12	8	4	0
0	0	0	0	0
	RSVD	DLLSC	EIS	PDS
			MSS	CC
			PDC	MSC
			PFD	ABP



Bit Range	Default and Access	Field Name (ID): Description
15:9	0h RO	Reserved (RSVD): Reserved.
8	0h RO	DLSC: Reserved for Data Link Layer State Changed: This bit is set when the value reported in the Data Link Layer Link Active field of the Link Status register is changed. In response to a Data Link Layer State Changed event, software should 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.
7	0h RO	EIS: Reserved for Electromechanical Interlock Status: If an Electromechanical Interlock is implemented, this bit indicates the current status of the Electromechanical Interlock. Defined encodings are: 0: Electromechanical Interlock Disengaged 1: Electromechanical Interlock Engaged
6	0h ROV	PDS: Presence Detect State: --In band presence detect state: 0: Slot Empty 1: Card present in slot This bit indicates the presence of an adapter in the slot, reflected by the logical "OR" of the Physical Layer in-band presence detect mechanism and, if present, any out-of-band presence detect mechanism defined for the slot's corresponding form factor. Note that the in-band presence detect mechanism requires that power be applied to an adapter for its presence to be detected. Consequently, form factors that require a power controller for hot-plug should implement a physical pin presence detect mechanism. Defined encodings are: 0: Slot Empty 1: Card Present in slot This register should 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 should return 1b.
5	0h RO	MSS: Reserved for MRL Sensor State: This register reports the status of the MRL sensor if it is implemented. Defined encodings are: 0: MRL Closed 1: MRL Open
4	0h RO	CC: Reserved for Command Completed: If Command Completed notification is supported (as indicated by No Command Completed Support field of Slot Capabilities Register), this bit is set when a hot-plug command has completed and the Hot-Plug Controller is ready to accept a subsequent command. The Command Completed status bit is set as an indication to host software that the Hot-Plug Controller has processed the previous command and is ready to receive the next command; it provides no guarantee that the action corresponding to the command is complete. If Command Completed notification is not supported, this bit should be hardwired to 0b.
3	0h RW1C	PDC: Presence Detect Changed: --A pulse indication that the inband presence detect state has changed This bit is set when the value reported in Presence Detect State is changed.
2	0h RO	MSC: Reserved for MRL Sensor Changed: If an MRL sensor is implemented, this bit is set when a MRL Sensor state change is detected. If an MRL sensor is not implemented, this bit should not be set.
1	0h RO	PFD: Reserved for Power Fault Detected: If a Power Controller that supports power fault detection is implemented, this bit is set when the Power Controller detects a power fault at this slot. Note that, depending on hardware capability, it is possible that a power fault can be detected at any time, independent of the Power Controller Control setting or the occupancy of the slot. If power fault detection is not supported, this bit should not be set.
0	0h RO	ABP: Reserved for Attention Button Pressed: If an Attention Button is implemented, this bit is set when the attention button is pressed. If an Attention Button is not supported, this bit should not be set.



11.44 Root Control (RCTL)—Offset BCh

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.

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:1, F:0] + BCh

Default: 0h

3	2	2	2	1	1	8	4	0			
1	8	4	0	6	2						
0	0	0	0	0	0	0	0	0	0	0	0
RSVD							CSVE	PMEIE	SEFEE	SENFUEE	SECEE

Bit Range	Default and Access	Field Name (ID): Description
31:5	0h RO	Reserved (RSVD): Reserved.
4	0h RO	CSVE: Reserved for CRS Software Visibility Enable: This bit, when set, enables the Root Port to return Configuration Request Retry Status (CRS) Completion Status to software. Root Ports that do not implement this capability should hardwire this bit to 0b.
3	0h RW	PMEIE: PME Interrupt Enable: 0: No interrupts are generated as a result of receiving PME messages. 1: Enables interrupt generation upon receipt of a PME message as reflected in the PME Status bit of the Root Status Register. A PME interrupt is also generated if the PME Status bit of the Root Status Register is set when this bit is set from a cleared state. If the bit change from 1 to 0 and interrupt is pending than interrupt is de-asserted
2	0h RW	SEFEE: System Error on Fatal Error Enable: Controls the Root Complex's response to fatal errors. 0: No SERR generated on receipt of fatal error. 1: Indicates that an SERR should be generated if a fatal error is reported by any of the devices in the hierarchy associated with this Root Port, or by the Root Port itself.
1	0h RW	SENFUEE: System Error on Non-Fatal Uncorrectable Error Enable: Controls the Root Complex's response to non-fatal errors. 0: No SERR generated on receipt of non-fatal error. 1: Indicates that an SERR should be generated if a non-fatal error is reported by any of the devices in the hierarchy associated with this Root Port, or by the Root Port itself.
0	0h RW	SECEE: System Error on Correctable Error Enable: Controls the Root Complex's response to correctable errors. 0: No SERR generated on receipt of correctable error. 1: Indicates that an SERR should be generated if a correctable error is reported by any of the devices in the hierarchy associated with this Root Port, or by the Root Port itself.



11.45 Root Status (RSTS)—Offset C0h


Provides information about PCI Express Root Complex specific parameters.

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:1, F:0] + C0h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:18	0h RO	Reserved (RSVD): Reserved.
17	0h RO	PMEP: PME Pending: Indicates that another PME is pending when the PME Status bit is set. When the PME Status bit is cleared by software; the PME is delivered by hardware by setting the PME Status bit again and updating the Requestor ID appropriately. The PME pending bit is cleared by hardware if no more PMEs are pending.
16	0h RW1C	PMES: PME Status: Indicates that PME was asserted by the requestor ID indicated in the PME Requestor ID field. Subsequent PMEs are kept pending until the status register is cleared by writing a 1 to this field. An interrupt is asserted If PMEIE is asserted and PMES is changing from 0 to 1 An interrupt is de-asserted If PMEIE is asserted and PMES is changing from 1 to 0 An Assert_PMEGPE is sent upstream If PMEGPEE in PEG Legacy cControl register (PEGLC) is asserted and PMES is changing from 0 to 1 An Deassert_PMEGPE is sent upstream If PMEGPEE in PEG Legacy cControl register (PEGLC) is asserted and PMES is changing from 1 to 0 An interrupt is de-asserted If PMEIE is asserted and PMES is changing from 1 to 0
15:0	0h ROV	PMERID: PME Requestor ID: Indicates the PCI requestor ID of the last PME requestor.



11.46 Device Capabilities 2 (DCAP2)—Offset C4h

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:1, F:0] + C4h

Default: B80h

3 1	2 8	2 4	2 0	1 6	1 2	8	4	0						
0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	1 0 1 1	1 0 0 0	0 0 0 0	0 0 0 0						
RSVD			OBFF_SUPPORTED	RSVD		LTRS	RSVD	ATOMIC128SUP	ATOMIC64SUP	ATOMIC32SUP	ATOMIC_OP_ROUTING_SUPPORT	ARIFS	CTODS	CTOR

Bit Range	Default and Access	Field Name (ID): Description
31:20	0h RO	Reserved (RSVD): Reserved.
19:18	0h RW_O	<p>OBFF_SUPPORTED: OBFF Supported</p> <p>00: OBFF Not Supported</p> <p>01: OBFF supported using Message signaling only</p> <p>10: OBFF supported using WAKE# signaling only</p> <p>11: OBFF supported using WAKE# and Message signaling</p> <p>The value reported in this field should indicate support for WAKE# signaling only if:</p> <ul style="list-style-type: none"> for a Downstream Port, driving the WAKE# signal for OBFF is supported and the connector or component connected Downstream is known to receive that same WAKE# signal for an Upstream Port, receiving the WAKE# signal for OBFF is supported and, if the component is on an add-in-card, that the component is connected to the WAKE# signal on the connector. <p>Root Ports, Switch Ports, and Endpoints are permitted to implement this capability. For a multi-Function device associated with an Upstream Port, each Function should report the same value for this field.</p> <p>For Bridges and Ports that do not implement this capability, this field should be hardwired to 00b.</p>
17:12	0h RO	Reserved (RSVD): Reserved.



Bit Range	Default and Access	Field Name (ID): Description
11	1h RO	LTRS: Latency Tolerance and BW reporting Mechanism Supported: A value of 1b indicates support for the optional Latency Tolerance & Bandwidth Requirement Reporting (LTBWR) mechanism capability. Root Ports, Switches and Endpoints are permitted to implement this capability. For Switches that implement LTBWR, this bit should be set only at the upstream port. For a multi-Function device, each Function should report the same value for this bit. For Bridges, Downstream Ports, and components that do not implement this capability, this bit should be hardwired to 0b.
10	0h RO	Reserved (RSVD): Reserved.
9	1h RO	ATOMIC128SUP: 128-bit CAS atomic operation completion support. This bit should be set to 1b if the Function supports this optional capability.
8	1h RO	ATOMIC64SUP: 64-bit atomic operation completion support. Includes FetchAdd, Swap, and CAS AtomicOps. This bit should be set to 1b if the Function supports this optional capability.
7	1h RO	ATOMIC32SUP: 32-bit atomic operation completion support. Includes FetchAdd, Swap, and CAS AtomicOps. This bit should be set to 1b if the Function supports this optional capability.
6	0h RO	ATOMIC_OP_ROUTING_SUPPORT: Atomic Operation Routing Supported. If set then atomic operations are supported.
5	0h RO	ARIFS: ARI Forwarding Supported: Applicable only to Switch Downstream Ports and Root Ports; should be 0b for other Function types. This bit should be set to 1b if a Switch Downstream Port or Root Port supports this optional capability.
4	0h RO	CTODS: Completion Timeout Disabled Supported: A value of 1b indicates support for the Completion Timeout Disable mechanism. The Completion Timeout Disable mechanism is required for Endpoints that issue Requests on their own behalf and PCI Express to PCI/PCI-X Bridges that take ownership of Requests issued on PCI Express. his mechanism is optional for Root Ports. The Root port does nopt support completion timeout disable
3:0	0h RO	CTOR: Completion Timer Ranges Supported: device Function support for the optional Completion Timeout programmability mechanism. This mechanism allows system software to modify the Completion Timeout value. This field is applicable only to Root Ports, Endpoints that issue Requests on their own behalf, and PCI Express to PCI/PCI-X Bridges that take ownership of Requests issued on PCI Express. For all other Functions this field is reserved and should be hardwired to 0000b. 0000b Completion Timeout programming not supported - the Function should implement a timeout value in the range 50 us to 50 ms.



11.47 Device Control 2 (DCTL2)—Offset C8h

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + C8h

Default: 0h

15	12	8	4	0
0	0	0	0	0
RSVD	OBFFEN	RSVD	LTREN	RSVD
			ATOMIC_OP_REQUESTER_EN	ARIFEN
				RSVD

Bit Range	Default and Access	Field Name (ID): Description
15	0h RO	Reserved (RSVD): Reserved.
14:13	0h RW	OBFFEN: Reserved.
12:11	0h RO	Reserved (RSVD): Reserved.
10	0h RW_V	LTREN: Latency Tolerance Reporting Mechanism Enable: When Set to 1b, this bit enables the Latency Tolerance & Reporting (LTR) mechanism. This bit is required for all Functions that support the LTR Capability. For a Multi-Function device associated with an upstream port of a device that implements LTBWR, the bit in Function 0 is of type RW, and only Function 0 controls the components Link behavior. In all other Functions of that device, this bit is of type RsvdP. Components that do not implement LTR are permitted to hardwire this bit to 0b. Default value of this bit is 0b. This bit is cleared when the port goes to DL_down state. HW ignores the value of this bit.
9:7	0h RO	Reserved (RSVD): Reserved.



Bit Range	Default and Access	Field Name (ID): Description
15:12	0h RWS	<p>ComplianceDeemphasis: Compliance De-emphasis: For 8 GT/s Data Rate: This field sets the Transmitter Preset level in Polling.Compliance state if the entry occurred due to the Enter Compliance bit being 1b.</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>Defined encodings are: 0001b -3.5 dB 0000b -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.</p> <p>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.</p> <p>The default value of this bit is 0000b.</p> <p>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	0h RWS	<p>compos: Compliance SOS: When set to 1b, the LTSSM is required to send SKP Ordered Sets periodically in between the (modified) compliance patterns.</p> <p>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.</p> <p>The default value of this bit is 0b.</p> <p>This bit is applicable when the Link is operating at 2.5 GT/s or 5 GT/s data rates only. Components that support only the 2.5 GT/s speed are permitted to hardwire this field to 0b.</p>
10	0h RWS	<p>entermocompliance: Enter Modified Compliance: When this bit is set to 1b, the device transmits modified compliance pattern if the LTSSM enters Polling.Compliance state.</p> <p>Components that support only the 2.5GT/s speed are permitted to hardwire this bit to 0b.</p> <p>Default value of this field is 0b.</p>
9:7	0h RWS_V	<p>txmargin: Transmit Margin: This field controls the value of the non-deemphasized voltage level at the Transmitter pins. This field is reset to 000b on entry to the LTSSM Polling.Configuration substrate (Refer 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 should be monotonic with a non-zero slope. The value of n should be greater than 3 and less than 7. At least two of these should be below the normal operating range n: 200-400 mV for full-swing and 100-200 mV for half-swing n -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> <p>When operating in 5GT/s mode with full swing, the deemphasis ratio should be maintained within +/- 1dB from the specification defined operational value (either -3.5 or -6 dB).</p>
6	0h RWS	<p>selectabledeemphasis: Selectable De-emphasis: When the Link is operating at 5GT/s speed, selects the level of de-emphasis. Encodings: 1b -3.5 dB 0b -6 dB</p> <p>Default value is implementation specific, unless a specific value is required for a selected form factor or platform.</p> <p>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>



Bit Range	Default and Access	Field Name (ID): Description
5	0h RWS	HASD: Hardware Autonomous Speed Disable: When set to 1b this bit disables hardware from changing the link speed for reasons other than attempting to correct unreliable link operation by reducing link speed.
4	0h RWS	EC: Enter Compliance: Software is permitted to force a link to enter Compliance mode at the speed indicated in the Target Link Speed field by setting this bit to 1b in both components on a link and then initiating a hot reset on the link.
3:0	3h RWS	TLS: Target Link Speed: For Downstream Ports, this field sets an upper limit on Link operational speed by restricting the values advertised by the Upstream component in its training sequences. The encoding is the binary value of the bit in the Supported Link Speeds Vector (in the Link Capabilities 2 register) that corresponds to the desired target Link speed. All other encodings are reserved. For example, 5.0 GT/s corresponds to bit 2 in the Supported Link Speeds Vector, so the encoding for a 5.0 GT/s target Link speed in this field is 0010b. If a value is written to this field that does not correspond to a supported speed (as indicated by the Max Link Speed Vector), the result is undefined. The default value of this field is the highest Link speed supported by the component (as reported in the Max Link Speed field of the Link Capabilities register) unless the corresponding platform/form factor requires a different default value. For both Upstream and Downstream Ports, this field is used to set the target compliance mode speed when software is using the Enter Compliance bit to force a Link into compliance mode. For a Multi-Function device associated with an Upstream Port, the field in Function 0 is of type RWS, and only Function 0 controls the components Link behavior. In all other Functions of that device, this field is of type RsvdP.

11.49 Link Status 2 (LSTS2)—Offset D2h

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + D2h

Default: 0h

15	12	8	4	0
0	0	0	0	0
RSVD			LNKEQREQ	EQPH3SUCC
			EQPH2SUCC	EQPH1SUCC
			EQCOMPLETE	CURDELVL

Bit Range	Default and Access	Field Name (ID): Description
15:6	0h RO	Reserved (RSVD): Reserved.
5	0h RW1C	LNKEQREQ: This bit is Set by hardware to request the Link equalization process to be performed on the Link.
4	0h ROV	EQPH3SUCC: Equalization Phase 3 Successful When set to 1b, this bit indicates that Phase 3 of the Transmitter Equalization procedure has successfully completed.



11.51 Port VC Capability Register 2 (PVCCAP2)—Offset 108h


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

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:1, F:0] + 108h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:24	0h RO	VCATO: VC Arbitration Table Offset: Indicates the location of the VC Arbitration Table. This field contains the zero-based offset of the table in DQWORDS (16 bytes) from the base address of the Virtual Channel Capability Structure. A value of 0 indicates that the table is not present (due to fixed VC priority).
23:8	0h RO	Reserved (RSVD): Reserved.
7:0	0h RO	VCAC: Reserved for VC Arbitration Capability:


11.52 Port VC Control (PVCCTL)—Offset 10Ch

Access Method

Type: CFG
(Size: 16 bits)

Offset: [B:0, D:1, F:0] + 10Ch

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
15:4	0h RO	Reserved (RSVD): Reserved.
3:1	0h RW	VCAS: VC Arbitration Select: This field will be programmed by software to the only possible value as indicated in the VC Arbitration Capability field. Since there is no other VC supported than the default, this field is reserved.
0	0h RO	VCARB: Reserved for Load VC Arbitration Table: Used for software to update the VC Arbitration Table when VC arbitration uses the VC Arbitration Table. As a VC Arbitration Table is never used by this component this field will never be used.

11.53 VC0 Resource Capability (VC0RCAP)—Offset 110h

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:1, F:0] + 110h

Default: 1h

3	2	2	2	1	1	8	4	0
1	8	4	0	6	2			
0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	1
PATO		RSVD	MTS		RSNPT	RSVD		PAC

Bit Range	Default and Access	Field Name (ID): Description
31:24	0h RO	PATO: Reserved for Port Arbitration Table Offset:
23	0h RO	Reserved (RSVD): Reserved.
22:16	0h RO	MTS: Reserved for Maximum Time Slots:



Bit Range	Default and Access	Field Name (ID): Description
15	0h RO	RSNPT: Reject Snoop Transactions: Reject Snoop Transactions (RSNPT): 0: Transactions with or without the No Snoop bit set within the TLP header are allowed on this VC. 1: When Set, any transaction for which the No Snoop attribute is applicable but is not Set within the TLP Header will be rejected as an Unsupported Request
14:8	0h RO	Reserved (RSVD): Reserved.
7:0	1h RO	PAC: Port Arbitration Capability: Port Arbitration Capability - Indicates types of Port Arbitration supported by the VC resource. This field is valid for all Switch Ports, Root Ports that support peer-to-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 (Refer 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 6-7Reserved Processor only supported arbitration indicates "Non-configurable hardware-fixed arbitration scheme".

11.54 VC0 Resource Control (VC0RCTL)—Offset 114h

Controls the resources associated with PCI Express Virtual Channel 0.

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:1, F:0] + 114h

Default: 800000FFh

3	2	2	2	1	1	8	4	0
1	8	4	0	6	2			
1	0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0 0	1 1 1 1	1 1 1 1
VC0E	RSVD	VC0ID	RSVD	PAS	RSVD	TCHVCOM	TCVCOM	TCOVCOM



Bit Range	Default and Access	Field Name (ID): Description
15:2	0h RO	Reserved (RSVD): Reserved.
1	1h RO_V	VCONP: VC0 Negotiation Pending: 0: The VC negotiation is complete. 1: The VC resource is still in the process of negotiation (initialization or disabling). This bit indicates the status of the process of Flow Control initialization. It is set by default on Reset, as well as whenever the corresponding Virtual Channel is Disabled or the Link is in the DL_Down state. It is cleared when the link successfully exits the FC_INIT2 state. Before using a Virtual Channel, software should check whether the VC Negotiation Pending fields for that Virtual Channel are cleared in both Components on a Link.
0	0h RO	Reserved (RSVD): Reserved.

11.56 Advanced Error Extended Reporting Capability Header (AECH)—Offset 1C0h

Size: 32 bits

The AER capability can optionally be included or excluded from the capabilities list. The full AER is supported.

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:1, F:0] + 1C0h

Default: 0h

Bit Range	Default & Access	Field Name (ID): Description
31:20	0h RW/O	Next Capability Offset (NCO): Set to 000h as this is the last capability in the list.
19:16	0h RW/O	Capability Version (CV): For systems that support AER, BIOS should write a 1h to this register else it should write 0
15:0	0h RW/O	Capability ID (CID): For systems that support AER, BIOS should write a 0001h to this register else it should write 0

11.57 PEG Uncorrectable Error Status—Offset 1C4h


Reports error status of individual error sources on a PCI Express device. An individual error status bit that is set indicates that a particular error occurred. Error status is cleared by writing a 1 to the respective bit. This register is for test and debug purposes only.

Access Method

Type: CFG

Offset: [B:0, D:1, F:0] + 1C4h (Size: 32 bits)

Default: 0h



Bit Range	Default & Access	Field Name (ID): Description
31:21	0h RW1CS	Reserved (RSVD): Reserved.
20	0h RW1CS	URES: Unsupported Request Error Status
19	0h RO	ECRCES: Reserved for ECRC Error Status
18	0h RW1CS	MTLPS: Malformed TLP Status
17	0h RW1CS	ROS: Receiver Overflow Status
16	0h RW1CS	UCS: Unexpected Completion Status
15	0h RO	CAS: Reserved for Completer Abort Status
14	0h RW1CS	CTS: Completion Timeout Status
13	0h RO	FCPES: Reserved for Flow Control Protocol Error Status
12	0h RW1CS	PTLPS: Poisoned TLP Status
11:5	0h RW1CS	Reserved (RSVD): Reserved.
4	0h RW1CS	DLPES: Data Link Protocol Error Status: The Data Link Layer Protocol Error that causes this bit to be set will also cause the Fatal Error Detected bit in Device Status[2] to be set if not already set.
3:0	0h RW1CS	Reserved (RSVD): Reserved.

11.58 PEG Uncorrectable Error Mask—Offset 1C8h

Controls reporting of individual errors by the device (or logic associated with this port) to the PCI Express Root Complex. As these errors are not originating on the other side of a PCI Express link, no PCI Express error message is sent, but the unmasked error is reported directly to the root control logic. A masked error (respective bit set to 1 in the mask register) has no action taken. There is a mask bit per error bit of the Uncorrectable Error Status register. This register is for test and debug purposes only.

Access Method

Type: CFG **Offset:** [B:0, D:1, F:0] + 1C8h
(Size: 32 bits)

Default: 0h

Bit Range	Default & Access	Field Name (ID): Description
31:21	0h RWS	Reserved (RSVD): Reserved.
20	0h RWS	UREM: Unsupported Request Error Mask
19	0h RO	ECRCM: Reserved for ECRC Error Mask
18	0h RWS	MTLPM: Malformed TLP Mask
17	0h RWS	ROM: Receiver Overflow Mask
16	0h RWS	UCM: Unexpected Completion Mask
15	0h RO	CAM: Reserved for Completer Abort Mask
14	0h RWS	CTM: Completion Timeout Mask
13	0h RO	FCPEM: Reserved for Flow Control Protocol Error Mask
12	0h RWS	PTLPM: Poisoned TLP Mask



Type: CFG **Offset:** [B:0, D:1, F:0] + 1D0h
 (Size: 32 bits)

Default: 0h

Bit Range	Default & Access	Field Name (ID): Description
31:14	0h RW1CS	Reserved (RSVD): Reserved.
13	0h RW1CS	ANFES: Advisory Non-Fatal Error Status: Advisory Non-Fatal Error Status When set, indicates that an Advisory Non-Fatal Error occurred
12	0h RW1CS	RTTS: Replay Timer Timeout Status
11:9	0h RW1CS	Reserved (RSVD): Reserved.
8	0h RW1CS	RNRS: Replay Number Rollover Status
7	0h RW1CS	BDLLPS: Bad DLLP Status
6	0h RW1CS	BTLPS: Bad TLP Status
5:1	0h RW1CS	Reserved (RSVD): Reserved.
0	0h RW1CS	RES: Receiver Error Status: Physical layer receiver error occurred.

11.61 PEG Correctable Error Mask—Offset 1D4h

Controls reporting of individual correctable errors by the device (or logic associated with this port) to the PCI Express Root Complex. As these errors are not originating on the other side of a PCI Express link, no PCI Express error message is sent, but the unmasked error is reported directly to the root control logic. A masked error (respective bit set to 1 in the mask register) has no action taken. There is a mask bit per error bit of the Correctable Error Status register. This register is for test and debug purposes only.

Access Method

Type: CFG **Offset:** [B:0, D:1, F:0] + 1D4h
 (Size: 32 bits)

Default: 0h

Bit Range	Default & Access	Field Name (ID): Description
31:14	0h RWS	Reserved (RSVD): Reserved.
13	0h RWS	ANFEM: Advisory Non-Fatal Error Mask: Advisory Non-Fatal Error Mask (ANFEM): When set, masks Advisory Non-Fatal errors from: a. signaling ERR_COR to the device control register, and a. updating the Uncorrectable Error Status register. This register is set by default to enable compatibility with software that does not comprehend Role-Based Error Reporting.
12	0h RWS	RTTM: Replay Timer Timeout Mask
11:9	0h RWS	Reserved (RSVD): Reserved.
8	0h RWS	RNRM: Replay Number Rollover Mask
7	0h RWS	BDLLPM: Bad DLLP Mask



Bit Range	Default & Access	Field Name (ID): Description
6	0h RWS	BTLPM: Bad TLP Mask
5:1	0h RWS	Reserved (RSVD): Reserved.
0	0h RWS	REM: Receiver Error Mask

11.62 PEG Advanced Error Capabilities and Control—Offset 1D8h

The PCI Express Advanced Error Reporting (AER) capabilities defined through this register. First Error is logged here.

Access Method

Type: CFG **Offset:** [B:0, D:1, F:0] + 1D8h
(Size: 32 bits)

Default: 0h

Bit Range	Default & Access	Field Name (ID): Description
31:11	0h RO	Reserved (RSVD): Reserved.
10	0h RO	MHRE: Multiple Header Recording Enable
9	0h RO	MHRC: Multiple Header Recording Capable
8	0h RO	ECRCCE: ECRC Check Enable
7	0h RO	ECRCCC: ECRC Check Capable
6	0h RO	ECRCGE: ECRC Generation Enable
5	0h RO	ECRCGC: ECRC Generation Capable
4:0	0h ROS	FEP: First Error Pointer

11.63 PEG Header Log—Offset 1DCh, 1E0h, 1E4h, 1E8h

The PCI Express Advanced Error Reporting (AER) capability for header logging will use this register for a header log.

Note: This description refers to all Header Log registers (HL0-HL3).

Access Method

Type: CFG **Offset:** [B:0, D:1, F:0] + 1DCh/1E0h/1E4h/1E8h
(Size: 32 bits)

Default: 0h

Bit Range	Default & Access	Field Name (ID): Description
31:0	0h ROS	HTAE: Header of TLP Associated with Error



11.64 PEG Root Error Command—Offset 1ECh

The Root Error Command register allows further control of Root Complex response to Correctable, Non-Fatal, and Fatal error Messages. Bit fields enable or disable generation of interrupts in addition to Do_SERR VDM sent to PCH.

Access Method

Type: CFG **Offset:** [B:0, D:1, F:0] + 1ECh
(Size: 32 bits)

Default: 0h

Bit Range	Default & Access	Field Name (ID): Description
31:3	0h RW	Reserved (RSVD): Reserved.
2	0h RW	FERE: Fatal Error Reporting Enable
1	0h RW	NFERE: Non-Fatal Error Reporting Enable
0	0h RW	CERE: Correctable Error Reporting Enable

11.65 PEG Root Error Status—Offset 1F0h

The Root Error Status register reports status of error messages received by the root complex, and of errors detected by the Root Port itself (which are treated conceptually as if the Root Port had sent an error message to itself). This register is updated regardless of the settings of the Root Control register and the Root Error Command register (which is not even implemented).

When an error is received by the Root Complex, the respective error received bit is set which indicates that a particular error category occurred. Software may clear an error status by writing a 1 to the respective bit.

Access Method

Type: CFG **Offset:** [B:0, D:1, F:0] + 1F0h
(Size: 32 bits)

Default: 0h

Bit Range	Default & Access	Field Name (ID): Description
31:27	0h RO	AEIMN: Reserved for Advanced Error Interrupt Message Number:
26:7	0h RW1CS	Reserved (RSVD): Reserved.
6	0h RW1CS	FEMR: Fatal Error Messages Received:
5	0h RW1CS	NFEMR: Non-Fatal Error Messages Received:
4	0h RW1CS	FUF: First Uncorrectable Fatal:
3	0h RW1CS	MEFNR: Multiple ERR_FATAL/NONFATAL Received:



Bit Range	Default & Access	Field Name (ID): Description
2	0h RW1CS	EFNFR: ERR_FATAL/NONFATAL Received: Set when either a fatal or a non-fatal error message is received and this bit is not already set.
1	0h RW1CS	MECR: Multiple ERR_COR Received:
0	0h RW1CS	ECR: ERR_COR Received: Set when a correctable error message is received and this bit is not already set.

11.66 PEG Error Source Identification—Offset 1F4h

The PCI Express Advanced Error Reporting (AER) register the requester ID of the first PCIE Correctable and Uncorrectable (Fatal/Non-Fatal) errors.

Access Method

Type: CFG
(Size: 32 bits)

Offset: [B:0, D:1, F:0] + 1F4h

Default: 0h

Bit Range	Default & Access	Field Name (ID): Description
31:16	0h ROS	EFNSI: Error Fatal/Non-Fatal Source Identification
15:0	0h ROS	ECSI: Error Correctable Source Identification

§ §

12 GTTMMADR Registers

Table 12-1. Summary of Bus: 0, Device: 2, Function: 0 (MEM)

Offset	Size (Bytes)	Register Name (Register Symbol)	Default Value
108000-108003h	4	Top of Low Usable DRAM (MTOLUD)—Offset 108000h	100h
108080-108087h	8	Top of Upper Usable DRAM (MTOUUD)—Offset 108080h	0h
1080C0-1080C3h	4	Base Data of Stolen Memory (MBDSM)—Offset 1080C0h	0h
108100-108103h	4	Base of GTT stolen Memory (MBGSM)—Offset 108100h	100h
108180-108183h	4	Protected Memory Enable Register (MPMEN)—Offset 108180h	0h
1081C0-1081C3h	4	Protected Low-Memory Base Register (MPLMBASE)—Offset 1081C0h	0h
108200-108203h	4	Protected Low-Memory Limit Register (MPLMLIMIT)—Offset 108200h	0h
108240-108247h	8	Protected High-Memory Base Register (MPHMBASE)—Offset 108240h	0h
108280-108287h	8	Protected High-Memory Limit Register (MPHMLIMIT)—Offset 108280h	0h
1082C0-1082C3h	4	Protected Audio Video Path Control (MPAVPC)—Offset 1082C0h	0h
108300-108303h	4	Global Command Register (MGCMD)—Offset 108300h	0h

12.1 Top of Low Usable DRAM (MTOLUD)—Offset 108000h

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

Programming Example:

C1DRB3 is set to 4GB

TSEG is enabled and TSEG size is set to 1MB

Processor Graphics is enabled, and Graphics Mode Select is set to 32MB

GTT Graphics Stolen Memory Size set to 2MB

BIOS knows the OS requires 1G of PCI space.

BIOS also knows the range from 0_FEC0_0h to 0_FFFF_FFFFh is not usable by the system. This 20MB range at the very top of addressable memory space is lost to APIC and Intel TXT.

According to the above equation, TOLUD is originally calculated to: 4GB = 1_0000_0h

The system memory requirements are: 4GB (max addressable space) - 1GB (pci space) - 35MB (lost memory) = 3GB - 35MB (minimum granularity) = 0_ECB0_0h

Since 0_ECB0_0h (PCI and other system requirements) is less than 1_0000_0h, TOLUD should be programmed to ECBh.

These bits are Intel TXT lockable.




Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:2, F:0] + 108000h

Default: 100h



Bit Range	Default and Access	Field Name (ID): Description
31:20	1h RO_V	<p>TOLUD: This register contains bits 31 to 20 of an address one byte above the maximum DRAM memory below 4G that is usable by the operating system. Address bits 31 down to 20 programmed to 01h implies a minimum memory size of 1MB. Configuration software should set this value to the smaller of the following 2 choices: maximum amount memory in the system minus ME stolen memory plus one byte or the minimum address allocated for PCI memory. Address bits 19:0 are assumed to be 0_0h 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 TXT mode.</p> <p>This register should be 1MB aligned when reclaim is enabled.</p>
19:1	0h RO	Reserved (RSVD): Reserved.
0	0h RO_V	LOCK: This bit will lock all writeable settings in this register, including itself.

12.2 Top of Upper Usable DRAM (MTOUUD)—Offset 108080h

This 64 bit register defines the Top of Upper Usable DRAM.

Configuration software should set this value to TOM minus all ME stolen memory if reclaim is disabled. If reclaim is enabled, this value should be set to reclaim limit + 1byte, 1MB aligned, since reclaim limit is 1MB aligned. Address bits 19:0 are assumed to be 000_0h for the purposes of address comparison. The Host interface positively decodes an address towards DRAM if the incoming address is less than the value programmed in this register and greater than or equal to 4GB.

BIOS Restriction: Minimum value for TOUUD is 4GB.

These bits are Intel TXT lockable.




Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:2, F:0] + 108080h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:20	0h RO_V	TOUUD: This register contains bits 38 to 20 of an address one byte above the maximum DRAM memory above 4G that is usable by the operating system. Configuration software should set this value to TOM minus all ME stolen memory if reclaim is disabled. If reclaim is enabled, this value should be set to reclaim limit 1MB aligned since reclaim limit + 1byte is 1MB aligned. Address bits 19:0 are assumed to be 000_0h for the purposes of address comparison. The Host interface positively decodes an address towards DRAM if the incoming address is less than the value programmed in this register and greater than 4GB. All the bits in this register are locked in Intel TXT mode.
19:1	0h RO	Reserved (RSVD): Reserved.
0	0h RO_V	LOCK: This bit will lock all writeable settings in this register, including itself.

12.3 Base Data of Stolen Memory (MBDSM)—Offset 1080C0h


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

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:2, F:0] + 1080C0h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:20	0h RO_V	BDSM: This register contains bits 31 to 20 of the base address of stolen DRAM memory. BIOS determines the base of graphics stolen memory by subtracting the graphics stolen memory size (PCI Device 0 offset 50 bits 15:8) from TOLUD (PCI Device 0 offset BC bits 31:20).
19:1	0h RO	Reserved (RSVD): Reserved.
0	0h RO_V	LOCK: This bit will lock all writeable settings in this register, including itself.

12.4 Base of GTT stolen Memory (MBGSM)—Offset 108100h


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

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:2, F:0] + 108100h

Default: 100h





Bit Range	Default and Access	Field Name (ID): Description
31:20	1h RO_V	BGSM: This register contains the base address of stolen DRAM memory for the GTT. BIOS determines the base of GTT stolen memory by subtracting the GTT graphics stolen memory size (PCI Device 0 offset 50 bits 7:6) from the Graphics Base of Data Stolen Memory (PCI Device 0 offset B0 bits 31:20).
19:1	0h RO	Reserved (RSVD): Reserved.
0	0h RO_V	LOCK: This bit will lock all writeable settings in this register, including itself.

12.5 Protected Memory Enable Register (MPMEN)— Offset 108180h

Register to enable the DMA-protected memory regions setup through the PLMBASE, PLMLIMIT, PHMBASE, PHMLIMIT registers. This register is always treated as RO for implementations not supporting protected memory regions (PLMR and PHMR fields reported as Clear in the Capability register).


Protected memory regions may be used by software to securely initialize remapping structures in memory. To avoid impact to legacy BIOS usage of memory, software is recommended to not overlap protected memory regions with any reserved memory regions of the platform reported through the Reserved Memory Region Reporting (RMRR) structures.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:2, F:0] + 108180h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31	0h RO_V	<p>EPM: This field controls DMA accesses to the protected low-memory and protected high-memory regions.</p> <p>0: Protected memory regions are disabled.</p> <p>1: Protected memory regions are enabled. DMA requests accessing protected memory regions are handled as follows:</p> <ul style="list-style-type: none"> When DMA remapping is not enabled, all DMA requests accessing protected memory regions are blocked. When DMA remapping is enabled: <ul style="list-style-type: none"> DMA requests processed as pass-through (Translation Type value of 10b in Context-Entry) and accessing the protected memory regions are blocked. DMA requests with translated address (AT=10b) and accessing the protected memory regions are blocked. DMA requests that are subject to address remapping, and accessing the protected memory regions may or may not be blocked by hardware. For such requests, software should 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. <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 should 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	0h RO	Reserved (RSVD): Reserved.
0	0h RO_V	<p>PRS: This field indicates the status of protected memory region(s):</p> <p>0: Protected memory region(s) disabled.</p> <p>1: Protected memory region(s) enabled.</p>

12.6 Protected Low-Memory Base Register (MPLMBASE)—Offset 1081C0h

Register to set up the base address of DMA-protected low-memory region below 4GB. This register should be set up before enabling protected memory through PMEN_REG, and should not be updated when protected memory regions are enabled.

This register is always treated as RO for implementations not supporting protected low memory region (PLMR field reported as Clear in the Capability register).

The alignment of the protected low memory region base depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1s to this register, and finding the most significant zero bit position with 0 in the value read back from the register. Bits N:0 of this register is decoded by hardware as all 0s.

Software should setup the protected low memory region below 4GB.

Software should not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).


Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:2, F:0] + 1081C0h



Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:20	0h RO_V	PLMB: This register specifies the base of protected low-memory region in system memory.
19:0	0h RO	Reserved (RSVD): Reserved.

12.7 Protected Low-Memory Limit Register (MPLMLIMIT)—Offset 108200h

Register to set up the limit address of DMA-protected low-memory region below 4GB. This register should be set up before enabling protected memory through PMEN_REG, and should not be updated when protected memory regions are enabled.

This register is always treated as RO for implementations not supporting protected low memory region (PLMR field reported as Clear in the Capability register).

The alignment of the protected low memory region limit depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1's to this register, and finding most significant zero bit position with 0 in the value read back from the register. Bits N:0 of the limit register is decoded by hardware as all 1s.

The Protected low-memory base and limit registers functions as follows:

- Programming the protected low-memory base and limit registers with the same value in bits 31:(N+1) specifies a protected low-memory region of size 2^(N+1) bytes.
- Programming the protected low-memory limit register with a value less than the protected low-memory base register disables the protected low-memory region.


Software should not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:2, F:0] + 108200h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
31:20	0h RO_V	PLML: This register specifies the last host physical address of the DMA-protected low-memory region in system memory.
19:0	0h RO	Reserved (RSVD): Reserved.

12.8 Protected High-Memory Base Register (MPHMBASE)—Offset 108240h

Register to set up the base address of DMA-protected high-memory region. This register should be set up before enabling protected memory through PMEN_REG, and should not be updated when protected memory regions are enabled.

This register is always treated as RO for implementations not supporting protected high memory region (PHMR field reported as Clear in the Capability register).

The alignment of the protected high memory region base depends on the number of reserved bits (N:0) of this register. Software may determine N by writing all 1's to this register, and finding most significant zero bit position below host address width (HAW) in the value read back from the register. Bits N:0 of this register are decoded by hardware as all 0s.

Software may setup the protected high memory region either above or below 4GB.


Software should not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:2, F:0] + 108240h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:20	0h RO_V	PHMB: This register specifies the base of protected (high) memory region in system memory. Hardware ignores, and does not implement, bits 63:HAW, where HAW is the host address width.
19:0	0h RO	Reserved (RSVD): Reserved.

12.9 Protected High-Memory Limit Register (MPHMLIMIT)—Offset 108280h

Register to set up the limit address of DMA-protected high-memory region. This register should be set up before enabling protected memory through PMEN_REG, and should not be updated when protected memory regions are enabled.

This register is always treated as RO for implementations not supporting protected high memory region (PHMR field reported as Clear in the Capability register).

The alignment of the protected high memory region limit depends on the number of reserved bits (N:0) of this register. Software may determine the value of N by writing all 1s 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 should not modify this register when protected memory regions are enabled (PRS field Set in PMEN_REG).

Access Method

Type: MEM
(Size: 64 bits)

Offset: [B:0, D:2, F:0] + 108280h

Default: 0h



Bit Range	Default and Access	Field Name (ID): Description
63:39	0h RO	Reserved (RSVD): Reserved.
38:20	0h RO_V	PHML: This register specifies the last host physical address of the DMA-protected high-memory region in system memory. Hardware ignores and does not implement bits 63:HAW, where HAW is the host address width.
19:0	0h RO	Reserved (RSVD): Reserved.

12.10 Protected Audio Video Path Control (MPAVPC)—Offset 1082C0h


All the bits in this register are locked by Intel TXT. When locked the R/W bits are RO.

Access Method

Type: MEM
(Size: 32 bits)

Offset: [B:0, D:2, F:0] + 1082C0h

Default: 0h





Bit Range	Default and Access	Field Name (ID): Description
31:20	0h RO_V	PCMBASE: Sizes supported in the processor: 1M, 2M, 4M and 8M. Base value programmed (from Top of Stolen Memory) itself defines the size of the WOPCM. Separate WOPCM size programming is redundant information and not required. Default 1M size programming. 4M recommended for the processor. This register is locked (becomes read-only) when PAVPE = 1b.
19:7	0h RO_V	RSVD2: These bits are reserved for future use.
6	0h RO_V	ASMFEN: ASMF method enabled 0b Disabled (default). 1b Enabled. This register is locked when PAVPLCK is set.
5	0h RO_V	RSVD1: These bits are reserved for future use.
4	0h RO_V	OVTATTACK: Override of Unsolicited Connection State Attack and Terminate. 0: Disable Override. Attack Terminate allowed. 1: Enable Override. Attack Terminate disallowed. This register bit is locked when PAVPE is set.
3	0h RO_V	HVYMODESEL: This bit is applicable only for PAVP2 operation mode. This bit is also applicable for PAVP3 mode only if the per-App memory config is disabled due to the clearing of bit 9 in the Crypto Function Control_1 register (address 0x320F0). 0: Lite Mode (Non-Serpent mode) 1: Serpent Mode For enabled PAVP3 mode, this one type boot time programming has been replaced by per-App programming (through the Media Crypto Copy command). Note that PAVP2 or PAVP3 mode selection is done by programming bit 8 of the MFX_MODE - Video Mode register.
2	0h RO_V	PAVPLCK: This bit locks all writeable contents in this register when set (including itself). Only a hardware reset can unlock the register again. This lock bit needs to be set only if PAVP is enabled (bit 1 of this register is asserted).
1	0h RO_V	PAVPE: 0: PAVP functionality is disabled. 1: PAVP functionality is enabled. This register is locked when PAVPLCK is set.
0	0h RO_V	PCME: This field enables Protected Content Memory within Graphics Stolen Memory. This memory is the same as the WOPCM area, whose size is defined by bit 5 of this register. This register is locked when PAVPLOCK is set. A value of 0 in this field indicates that Protected Content Memory is disabled, and cannot be programmed in this manner when PAVP is enabled. A value of 1 in this field indicates that Protected Content Memory is enabled, and is the only programming option available when PAVP is enabled. (Note that the processor legacy Lite mode programming of PCME bit = 0 is not supported. For non-PAVP3 Mode, even for Lite mode configuration, this bit should be programmed to 1 and HVYMODESEL = 0). This bit should always be programmed to 1 if bits 1 and 2 (PAVPE and PAVP lock bits) are both set. With per-App Memory configuration support, the range check for the WOPCM memory area should always happen when this bit is set, regardless of Lite or Serpent mode, or PAVP2 or PAVP3 mode programming.



Bit Range	Default and Access	Field Name (ID): Description
29	0h RO	<p>SFL: This field is valid only for implementations supporting advanced fault logging. Software sets this field to request hardware to set/update the fault-log pointer used by hardware. The fault-log pointer is specified through Advanced Fault Log register. Hardware reports the status of the 'Set Fault Log' operation through the FLS field in the Global Status register.</p> <p>The fault log pointer should be set before enabling advanced fault logging (through EAFI field). Once advanced fault logging is enabled, the fault log pointer may be updated through this field while DMA remapping is active. Clearing this bit has no effect. The value returned on read of this field is undefined.</p>
28	0h RO	<p>EAFI: This field is valid only for implementations supporting advanced fault logging. Software writes to this field to request hardware to enable or disable advanced fault logging:</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 should be set in hardware (through the SFL field) before enabling advanced fault logging. Hardware reports the status of the advanced fault logging enable operation through the AFLS field in the Global Status register.</p> <p>The value returned on read of this field is undefined.</p>
27	0h RO	<p>WBF: This bit is valid only for implementations requiring write buffer flushing. Software sets this field to request that hardware flush the Root-Complex internal write buffers. This is done to ensure any updates to the memory-resident remapping structures are not held in any internal write posting buffers.</p> <p>Hardware reports the status of the write buffer flushing operation through the WBF field in the Global Status register.</p> <p>Clearing this bit has no effect. The value returned on a read of this field is undefined.</p>
26	0h RO_V	<p>QIE: This field is valid only for implementations supporting queued invalidations. 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>The value returned on a read of this field is undefined.</p>
25	0h RO_V	<p>IRE: This field is valid only for implementations supporting interrupt remapping.</p> <p>0: Disable interrupt-remapping hardware</p> <p>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 should 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 should 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>



Bit Range	Default and Access	Field Name (ID): Description
24	0h WO	<p>SIRTP: This field is valid only for implementations supporting interrupt-remapping. Software sets this field to set/update the interrupt remapping table pointer used by hardware. The interrupt remapping table pointer is specified through the Interrupt Remapping Table Address (IRTA_REG) register.</p> <p>Hardware reports the status of the 'Set Interrupt Remap Table Pointer' operation through the IRTPS field in the Global Status register.</p> <p>The 'Set Interrupt Remap Table Pointer' operation should be performed before enabling or re-enabling (after disabling) interrupt-remapping hardware through the IRE field.</p> <p>After a 'Set Interrupt Remap Table Pointer' operation, software should 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 should 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>
23	0h RO_V	<p>CFI: 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 Extended Interrupt Mode (x2APIC mode) is not enabled.</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>
22:0	0h RO	<p>Reserved (RSVD): Reserved.</p>

§ §

13 EDRAMBAR Registers

Table 13-1. Summary of EDRAMBAR Registers

Offset	Size (Bytes)	Register Name (Register Symbol)
3250h	4	(3250h) GT_PMON_RD_EDRAM_HIT
3254h	4	(3254h) GT_PMON_RD_EDRAM_MISS
3258h	4	(3258h) GT_PMON_PTLWR_EDRAM_HIT
325Ch	4	(325Ch) GT_PMON_PTLWR_EDRAM_MISS
3260h	4	(3260h) GT_PMON_WR_EDRAM_HIT
3264h	4	(3264h) GT_PMON_WR_EDRAM_MISS
3268h	4	(3268h) GT_PMON_CNT_MISC_CTRL

13.1 (3250h) GT_PMON_RD_EDRAM_HIT

Counts read Requests that hit in EDRAM.

Bit Range	Default & Access	Field Name (ID): Description
31:0	0h RO_VFW	COUNTER: 32 bit counter.

13.2 (3254h) GT_PMON_RD_EDRAM_MISS

Counts read Requests that miss in EDRAM.

Bit Range	Default & Access	Field Name (ID): Description
31:0	0h RO_VFW	COUNTER: 32 bit counter.

13.3 (3258h) GT_PMON_PTLWR_EDRAM_HIT

Counts partial write requests that hit in EDRAM

Bit Range	Default & Access	Field Name (ID): Description
31:0	0h RO_VFW	COUNTER: 32 bit counter.



13.4 (325Ch) GT_PMON_PTLWR_EDRAM_MISS

Counts partial write Requests that miss in EDRAM

Bit Range	Default & Access	Field Name (ID): Description
31:0	0h RO_VFW	COUNTER: 32 bit counter.

13.5 (3260h) GT_PMON_WR_EDRAM_HIT

Counts full write requests that hit in EDRAM

Bit Range	Default & Access	Field Name (ID): Description
31:0	0h RO_VFW	COUNTER: 32 bit counter.

13.6 (3264h) GT_PMON_WR_EDRAM_MISS

Counts full write requests that miss in EDRAM

Bit Range	Default & Access	Field Name (ID): Description
31:0	0h RO_VFW	COUNTER: 32 bit counter.

13.7 (3268h) GT_PMON_CNT_MISC_CTRL

Holds GT_Pmon configuration and overflow bits.

Bit Range	Default & Access	Field Name (ID): Description
31:15	-	RSVD: Reserved
14:14	0h RO_VFW	WR_MISS_OVFLW: Full Write Miss counter was overflowed.
13:13	0h RO_VFW	WR_HIT_OVFLW: Full Write Hit counter was overflowed.
12:12	0h RO_VFW	PTLWR_MISS_OVFLW: Partial Write Miss counter was overflowed.
11:11	0h RO_VFW	PTLWR_HIT_OVFLW: Partial Write Hit counter was overflowed.
10:10	0h RO_VFW	RD_MISS_OVFLW: Read Miss counter was overflowed.



Bit Range	Default & Access	Field Name (ID): Description
9:9	0h RO_VFW	RD_HIT_OVFLW: Read Hit counter was overflowed.
8:1	2h RW	COS_ENABLE: Enable counting by 'Class Of Service' indications. CoS indicates the source of the transaction. There are 8 available sources, Put 1 on a specific bit location to enable counting requests from the specific source (bitwise OR). The sources are (according to CoS definition): bit 0 - NTR IA Read or uncacheable write bit 1 - NTR GT Read or uncacheable write bit 2 - NTR Coherent eviction (explicit + implicit) bit 3 - RSVD bit 4 - Default XTM bit 5 - XTM IMGU bit 6 - XTM Intel HD Audio bit 7 - XTM DE Default value: Count GT requests only (bit 1=1).
0:0	0h RW	ENABLE: GT_Pmon counters global enable.

§ §